Monthly Threat Intelligence Rollup

DEEP seas

12/01/23-12/31/23

Notable Cyberattacks

Summary of noteworthy cyberattacks in the last thirty days.

Incident	Activity Summary
AeroBlade Attackers Targeting U.S. Aviation Company	Details of a suspected Chinese nation-state actor targeting a U.S. aviation company have been released in a new paper by Blackberry Research & Intelligence Team. The campaign by the actor, dubbed AeroBlade by Blackberry, began with operational development in September 2022, culminating in a full-blown operation by the summer of 2023. The attackers attempted to phish the employees of an unnamed U.S. aviation supplier with a previously unseen malware, which remained in development for approximately a year. The malware is no simple tool, but rather a fully featured backdoor malware capable of enumerating and exfiltrating directory structures, establishing a reverse shell connection, and many other features. Also notable is the evasiveness of the malware, which can detect if it is running in a virtualized environment, as well as evading current detection logic. While Blackberry assesses that the malware is the result of a commercial cyberespionage campaign with an intent on exfiltrating and ransoming data, the operational tempo and rapid feature development of the malware, emphasis on evasion, and the choice of target is highly suggestive of Chinese nation-state involvement. Economic espionage against the aviation sector has been a feature of Chinese nation-state operations for nearly a decade now and taken as a whole this activity is likely being conducted at the behest of Beijing rather than a commercial cyberespionage group.[i]
JetBrains' TeamCity Servers Under Active Exploitation	DeepSeas is aware of reports that Russian nation-state actors, specifically APT29, have been exploiting an authentication bypass vulnerability in JetBrains' TeamCity servers. First identified and reported to JetBrains in September 2023, initial exploitation of CVE-2023-42793 was observed in October 2023 against a U.S.-based biomedical manufacturer. According to CISA the victims include, "[an] energy trade association; companies that provide software for billing, medical devices, customer care, employee monitoring, financial management, marketing, sales, […] video games, tools manufacturers, and small and large IT companies." DeepSeas has also observed chatter in cybercriminal and dark web forums regarding CVE-2023-42793, along with claims that working exploits are available for rent and/or purchase. Other attempts to exploit CVE-2023-42793, presumably by cybercriminal and financially motivated actors, have been observed in the wild with varying levels of sophistication and success. Based on the age of the exploit, and the fact that a patch has been available since September 2023, the number of vulnerable TeamCity servers is likely much lower than the number of vulnerable TeamCity servers in August and September 2023, and thus the severity is somewhat lessened. However, given that APT29 has been observed exploiting this vulnerability, patching and enforcement of two-factor authentication (2FA) is highly recommended, as are follow-up investigations to verify that no second-stage implants were delivered that may bypass a patched TeamCity server.[ii]

Emergent Malware and TTPs

Newly emerging malware and tactics, techniques, and procedures identified in the last thirty days.

Malware/Campaign	Activity Summary	TTP Analysis
U.S. Government Targeted Via Adobe ColdFusion Vulnerability CVE-2023-26360	A new report by CISA outlines recent observations of attempts by unidentified threat actors to compromise Adobe ColdFusion servers belonging to unspecified U.S. government agencies.	The vulnerability exploited is CVE-2023-26360, an improper access control issue impacting version 2018 Update 15 and earlier, as well as 2021 Update 5 and earlier. Other end of life (EoL) versions are likewise vulnerable, including ColdFusion 2016 and ColdFusion 11. Exploitation of ColdFusion is by no means a new or novel tactic. ColdFusion was heavily exploited beginning in 2014 and 2015, and exploitation has continued through to present day, primarily due to ColdFusion being used to develop web applications in an enterprise environment. After exploiting CVE-2023-26360, the attackers conducted reconnaissance, cleaned up artifacts of their compromise, and installed a web shell that appears to be a copy of a Chinese-language web shell dubbed ByPassGodzilla. Analysis of this web shell is ongoing and attribution to Chinese actors should be considered provisional at best.[iii]
Lazarus Group Utilizing Log4Shell in Operations Against U.S. Government Targets	In a recent find, Cisco Talos has uncovered a new malware campaign conducted by APT38, also known as the Lazarus threat group, that they have dubbed "Operation Blacksmith."	This campaign utilizes the well-known vulnerability CVE-2021-44228, commonly known as Log4Shell, which is a remote code execution vulnerability in the Apache Foundation's Log4j library. After exploiting this dangerous weakness in Log4j, Lazarus' campaign deployed three new malware families - NineRAT, DLRAT, and BottomLoader - each uniquely written in DLang and with unique purposes. NineRAT, after being dropped on a victim's device, will use Telegram for its C2 capabilities and establish persistence, which allows for interaction with the affected device and data exfiltration. DLRAT is like NineRAT, except for having added reconnaissance capabilities to also gather system information. BottomLoader is the simplest of the three, being a downloader used to infect a victim with additional malware, such as HazyLoad. BottomLoader, however, is still an integral factor in this case, as Cisco Talos was able to determine the threat group Andariel (Onyx Sleet, SILENT CHOLLIMA), which is known to be under Lazarus' umbrella, was involved in the campaign due to the use of HazyLoad, known to be used by Andariel and other North Korean threat actors.[iv]

Threat Actor Campaigns

New activity related to threat actor campaigns in the last thirty days.

Threat Actors	Activity Summary
Russian State-Aligned Groups Launch Wave of Phishing Campaigns	In recent weeks, a number of reports regarding Russian nation-state activity have been published by various security vendors. While it is common for vendors to be observing and analyzing the same activity, each report details different targets, tactics, and nations, presumably with the intent of espionage rather than disruption. APT28 has been observed using phishing lures regarding the Israeli-Hamas conflict to deliver new malware; while another report covers targeting of NATO members, as well as MENA nations; and finally, Microsoft issued updated guidance to a March 2023 report regarding APT28 exploitation of Microsoft Exchange, also for espionage. At a high level, this is indicative that APT28 has not suffered operationally following a protracted war in Ukraine. The development of a new malware, the so-called HeadLace malware identified by IBM X-Force, bears further investigation to determine whether the malware is a rewrite of a previous APT28 malware or something new entirely. Fortunately, APT28 has constrained itself to espionage rather than carrying out disruptive or destructive attacks. Though it cannot be ruled out that accesses may be shared with other Russian state-aligned groups like Turla, Sandworm, Gamaredon, and others.[v]
MUSTANG PANDA Targeting Taiwanese Government Entities	A recent report by Lab52 details a campaign by the Chinese nation-state group MUSTANG PANDA utilizing lure documents detailing recent developments in the Taiwanese government to deliver a new variant of the venerable PlugX malware, very similar to the SmugX malware previously identified being utilized by MUSTANG PANDA. The malware itself is delivered via a Microsoft Installer format file (*.MSI), which drops the lure document - a harmless PDF version of a document - as well as the PlugX variant. The lure theme hints at the targets, which are consistent with recent MUSTANG PANDA operations, suggesting that this is not necessarily a new campaign on the part of Beijing, but rather the next phase or a continuation of an existing campaign previously identified. DeepSeas is currently analyzing samples of the malware to determine the efficacy of existing detection logic against this new variant of PlugX. Chinese nation-state groups often share tools and malware, making it likely that other Chinese nation-state groups may be utilizing it.[vi]

Dark Web Markets

High-profile ransomware data dumps and dark web access sales identified in last thirty days.

Activity	Note
Tool Sale	An actor was observed selling a remote code execution exploit in Cisco IOS XE routers he claimed is not related to CVE-2023-20198 for USD 15,000. Later he claimed that there are “great accesses available with this exploit," naming G2 Satellite Solutions, TDS Telecom, Comcast, and, just in time for tax season, H&R Block as potentially vulnerable.
Access Sale	An actor was observed selling RDP local admin access to a U.S.-based healthcare software company with USD 2 billion in revenue and 6,000 employees for USD 6,000.
Access Sale	An actor was observed selling access to a Virginia-based government software contractor, REI Systems, an unnamed European stock exchange, an unnamed vendor that works with Colonial Pipeline and 20 more oil and gas pipeline companies, and the U.S.-based cloud software company CloudBees.
Access Sale	An actor was observed selling access to an unnamed Canadian energy company with USD 6.5 billion in revenue for USD 6,000, an unnamed American accounting company with USD 700 million in revenue for USD 3,500, and RDP local admin access to an unnamed Australian company with around USD 3 billion in revenue for USD 4,000.
Access Sale	An actor was observed selling RDP domain admin access to a New York City-based water transport company for USD 350. The New York City Department of Transportation lists at least eight commercial ferry and sightseeing services in the city.
Access Sale	A highly credible actor was observed selling access to U.S. Health and Human Services servers. This actor was responsible for several high-profile attacks in the past, including the breach of health care information belonging to members of Congress and their staffs.
Actor Developments	A longtime criminal actor from Venezuela, using the screen name Kelvin Security, was reportedly arrested in Spain in connection with the breach of multiple Spanish municipalities. Kelvin Security has been a presence in cyber criminal venues since at least 2013.
Access Sale	An access seller was observed selling access to more than two dozen victims, including what was likely a Spanish banking giant and a large Belgian logistics company.
Access Sale	An actor was selling RDP and web shell access to a French subdomain belonging to Bank of America. He did not specify a price.
Access Sale	An access broker sold domain admin RDP access to an eCommerce company in the fashion and clothing retail vertical with more than USD 1 billion in revenue for a buy now price of USD 7,331. In another post they said that there were 170 hosts on the network, suggesting that this is a pure eCommerce operation and not the web part of a bricks and mortar retailer.
Access Sale	An access broker was selling RDP local admin access to a U.S.-based software developer with USD 321.5 million in revenue for USD 11,000.
Access Sale	A user on an English language crime forum claimed to have "live access paths" into CVS Health. He claimed to have access to source code, developer environments, patient info, etc. He did not provide evidence of access but said he would provide proof of access to trusted forum

	members.
Access Sale	An access broker was observed selling access to a Canadian electronics retailer with USD 5 billion in revenue, as well as an unnamed U.S.-based energy company with USD 800 million in revenue.
Access Sale	An access broker was observed selling domain user Citrix access to an unnamed U.S. based technology company with USD 7 billion in revenue for a buy now price of USD 8,000.

By The Numbers

Summarizing incidents in graphical format

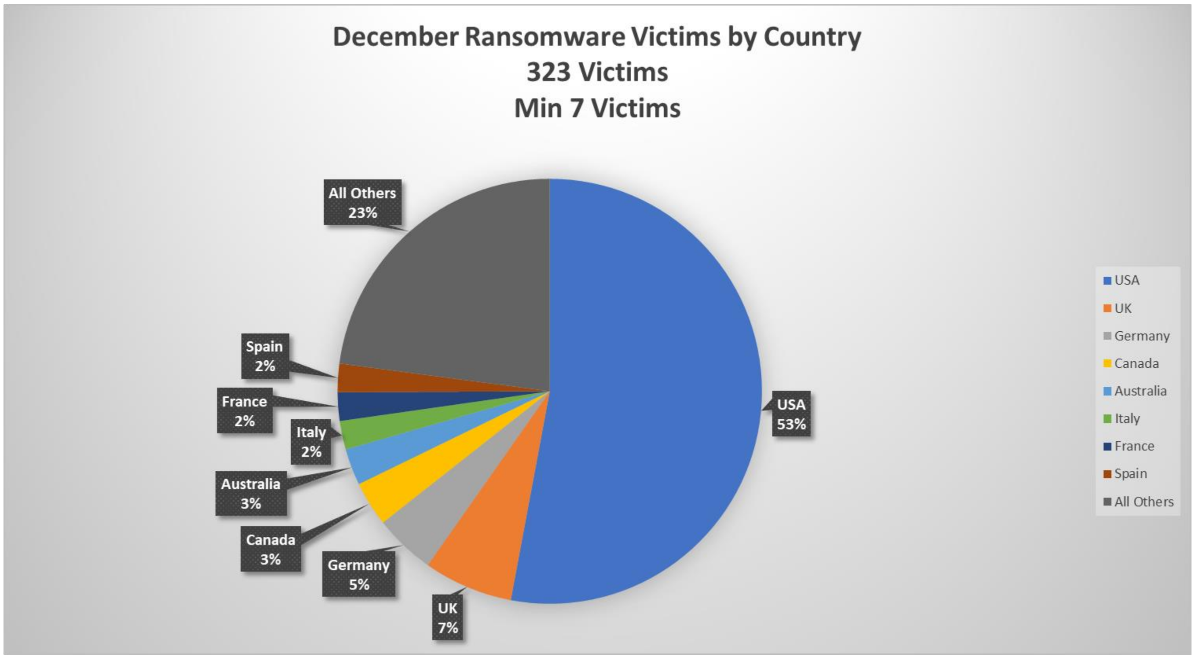

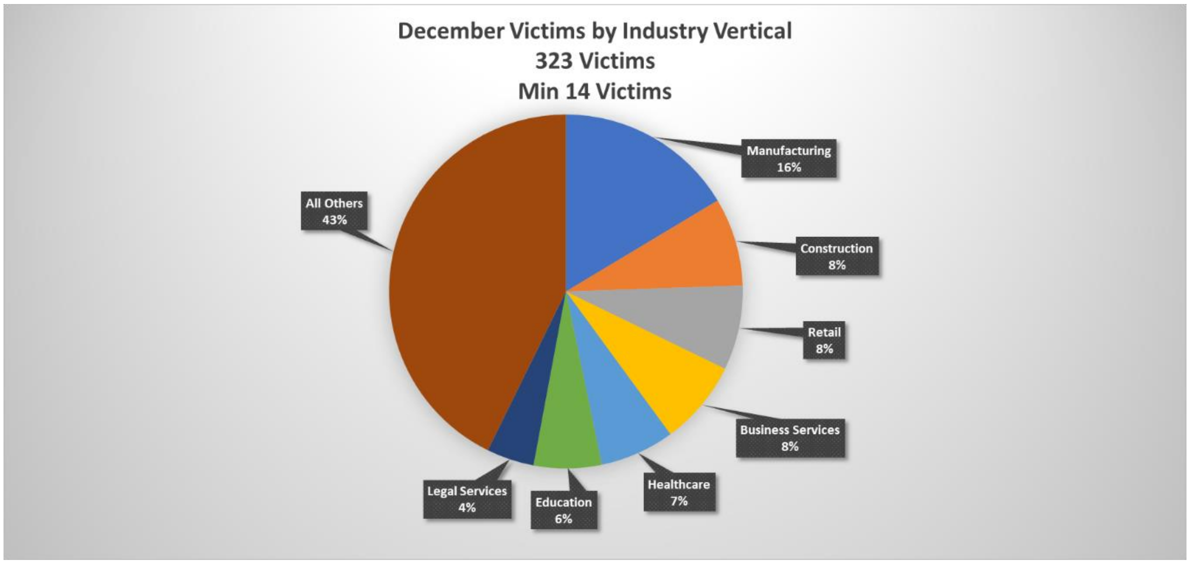

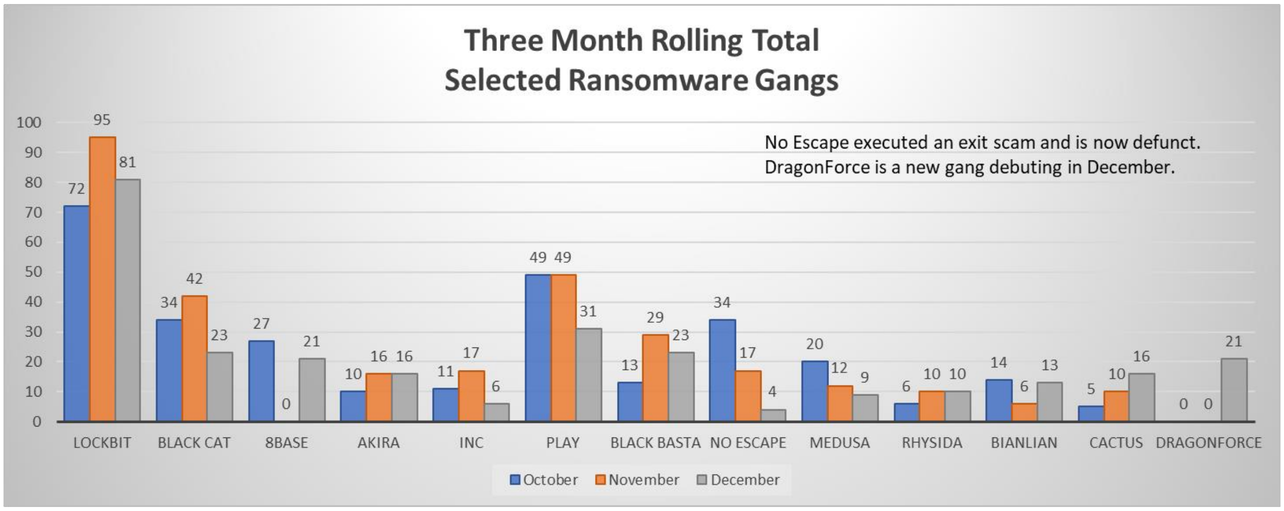

Three Month Rolling Total
Selected Ransomware Gangs

No Escape executed an exit scam and is now defunct.
DragonForce is a new gang debuting in December.

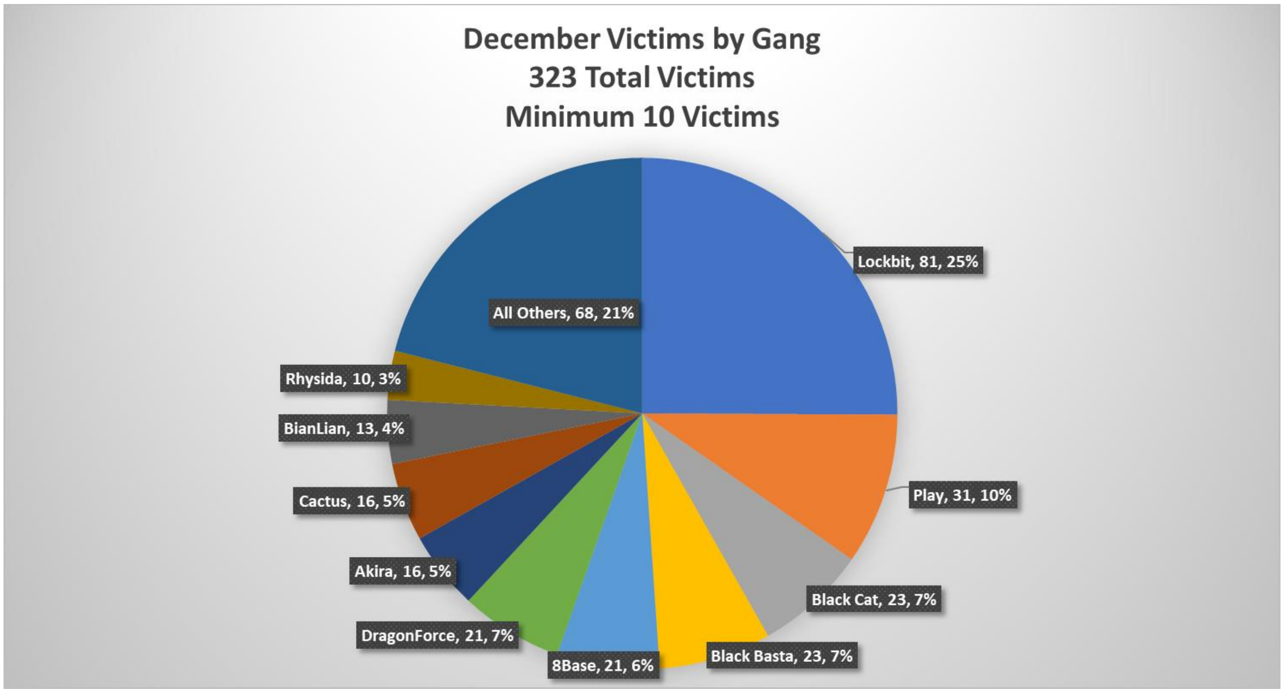

December Victims by Gang
323 Total Victims
Minimum 10 Victims

New Detection Content

Noteworthy new detection logic added in the last 30 days, excluding rule tuning.

- **DS_Known Cryptominers**
  - This rule identifies known cryptocurrency miners.

- **DS_BHUNT Malware**
  - "A novel modular crypto wallet stealing malware dubbed BHUNT; has been spotted targeting cryptocurrency wallet contents, passwords, and security phrases. To evade detection and triggering security warnings, BHUNT is packed and heavily encrypted using Themida and VMProtect, two virtual machine packers that hinder reverse-engineering and analysis by researcher."

- **DS_BlueKeep Crypto Miner**
  - "This query was originally published in the threat analytics report, Exploitation of CVE-2019-0708 (BlueKeep). CVE-2019-0708, also known as BlueKeep, is a critical remote code execution vulnerability involving RDP. Soon after its disclosure, the NSA issued a rare advisory about this vulnerability out of concern that it could be used to quickly spread malware. Attackers have since used this vulnerability to install cryptocurrency miners on targets. Microsoft has issued updates for this vulnerability, as well as guidance for protecting operating systems that we no longer support. Microsoft Defender ATP also contains behavioral detections for defending against this threat."

- **PingID Multiple Failed MFA Requests for User**
  - The following analytic identifies multiple failed multi-factor authentication (MFA) requests for a single user within a PingID (PingOne) environment. Specifically, the analytic triggers when 10 or more MFA user prompts fail within 10 minutes. PingID environments can be very different depending on the organization. Security teams should test this detection and customize these arbitrary thresholds. The detected behavior may represent an adversary who has obtained legitimate credentials for a user and continuously repeats login attempts to bombard users with MFA push notifications, SMS messages, and phone calls - potentially resulting in the user finally accepting the authentication request. Threat actors like the Lapsus team and APT29 have leveraged this technique to bypass multi-factor authentication controls, as reported by Mandiant and others.

- **PingID New MFA Method After Credential Reset**
  - A common social engineering technique used by threat actors is the impersonation of a valid user to organizational support staff for a password reset. During the same support call, or quickly afterwards, the threat actor will request provisioning of a new MFA device. This does not require malware or phishing infrastructure and has proven to be successful in numerous historical attacks. This detection looks for the pattern of password reset, followed by MFA device provisioning.

- **Kubernetes Unauthorized Access**
  - The following analytic detects unauthorized access to Kubernetes by monitoring Kubernetes audit logs. It identifies anomalies in access patterns by segmenting and analyzing the source of requests. Unauthorized access is worth identifying for a SOC, as it could indicate an attacker attempting to infiltrate the system. The impact of such an attack could be severe, potentially leading to unauthorized access to sensitive systems or data.

- **Kubernetes Abuse of Secret by Unusual User Group**
  - The following analytic detects unauthorized access or misuse of Kubernetes Secrets by unusual user groups. It identifies anomalies in access patterns by segmenting and analyzing the source of requests by user group. Kubernetes Secrets, which store sensitive information like passwords, OAuth tokens, and SSH keys, are critical assets, and their misuse can lead to significant security breaches. This behavior is worth identifying for a SOC as it could indicate an attacker attempting to exfiltrate or misuse these secrets. The impact of such an attack could be severe, potentially leading to unauthorized access to

sensitive systems or data.

- **PingID New MFA Method Registered for User**
  - The following analytic identifies the registration of a new multi-factor authentication method for a PingID (PingOne) account. Adversaries who have obtained unauthorized access to a user account may register a new MFA method to maintain persistence.

- **Kubernetes Abuse of Secret by Unusual Location**
  - The following analytic detects unauthorized access or misuse of Kubernetes Secrets from unusual locations. It identifies anomalies in access patterns by segmenting and analyzing the source of requests by country. Kubernetes Secrets, which store sensitive information like passwords, OAuth tokens, and SSH keys, are critical assets, and their misuse can lead to significant security breaches. This behavior is worth identifying for a SOC, as it could indicate an attacker attempting to exfiltrate or misuse these secrets. The impact of such an attack could be severe, potentially leading to unauthorized access to sensitive systems or data.

- **PingID Mismatch Auth Source and Verification Response**
  - The following analytic identifies variations in the authentication event IP address versus the verification response event IP address to identify suspicious sign-in behavior. Currently, this detection is configured to identify when the originating country of an authentication request is different than the verification country.

- **Kubernetes Abuse of Secret by Unusual Username**
  - The following analytic detects unauthorized access or misuse of Kubernetes Secrets by unusual usernames. It identifies anomalies in access patterns by segmenting and analyzing the source of requests by username. Kubernetes Secrets, which store sensitive information like passwords, OAuth tokens, and SSH keys, are critical assets, and their misuse can lead to significant security breaches. This behavior is worth identifying for a SOC, as it could indicate an attacker attempting to exfiltrate or misuse these secrets. The impact of such an attack could be severe, potentially leading to unauthorized access to sensitive systems or data.

- **Kubernetes Access Scanning**
  - The following analytic detects potential scanning activities within a Kubernetes environment. It identifies unauthorized access attempts, probing of public APIs, or attempts to exploit known vulnerabilities. The analytic detects this behavior by monitoring Kubernetes audit logs for patterns indicative of scanning, such as repeated failed access attempts or unusual API requests. This behavior is worth identifying for a SOC, as it could indicate an attacker’s preliminary step in an attack, aiming to gather information about the system to find potential vulnerabilities. The impact of such an attack could be severe, potentially leading to unauthorized access to sensitive systems or data.

- **Windows Modify System Firewall with Notable Process Path**
  - The following analytic detects a potential suspicious modification of firewall rule, allowing execution of specific applications in public, as well as suspicious, Windows process file paths. This technique was identified when an adversary and red teams bypassed firewall file execution restriction in a targeted host. Take note that this event or command can be run by an administrator during testing, allowing for legitimate tools or applications.

- **Kubernetes Abuse of Secret by Unusual User Agent**
  - The following analytic detects unauthorized access or misuse of Kubernetes Secrets by unusual User-Agents. It identifies anomalies in access patterns by segmenting and analyzing the source of requests by user agent. Kubernetes Secrets, which store sensitive information like passwords, OAuth tokens, and SSH keys, are critical assets, and their misuse can lead to significant security breaches. This behavior is worth identifying for a SOC, as it could indicate an attacker attempting to exfiltrate or misuse these secrets. The impact of such an attack could be severe, potentially leading to unauthorized access to sensitive systems or data.

- **Kubernetes Suspicious Image Pulling**
  - The following analytic detects instances of suspicious image pulling in Kubernetes. It identifies this behavior by monitoring Kubernetes audit logs for image pull requests that do not match a predefined list of allowed images. This behavior is worth identifying for a SOC, as it could indicate an attacker attempting to deploy malicious software or infiltrate the system. The impact of such an attack could be severe, potentially leading to unauthorized access to sensitive systems or data.

- **Stealing Credentials from Web Browsers**
  - "Command seen used to steal credentials from a user’s browser. Malicious command may look similar to:
    - cmd.exe /Q /c copy ""C:\Users\username\AppData\Local\Microsoft\Edge\User Data\Default\Login Data"" ""C:\Windows\Temp\"""

- **TeamCity Java Scheduling Task**
  - This query looks for an event where TeamCity process (java.exe) creates a process of Windows task management utility (schtasks.exe). This query may have false positives where there is an official need for Java applications to launch the schtasks.exe

- **Suspicious Child Processes of TeamCity Java**
  - Looking for suspicious child processes being invoked by Java.exe in the TeamCity install path. This may indicate an exploit of the TeamCity server seen in the references below.

- **EDRSandblast Indicators**
  - Looking for the EDRSandblast tool and vulnerable drivers used to bypass EDRs. Query is on the tool name, or file names and hashes of the vulnerable drivers. This signature may not indicate malicious behavior has been observed, but the vulnerable driver should be remediated.

- **O365 New MFA Method Registered**
  - This analytic detects the registration of a new MFA method associated with a user account within Office 365 by monitoring O365 audit logs and configurations. While adding a new MFA method can be a routine and legitimate action, it can also be indicative of an attacker's attempt to maintain persistence on a compromised account. By registering a new MFA method, attackers can potentially bypass existing security measures, allowing them to authenticate using stolen credentials without raising alarms. Monitoring for such changes is crucial, especially if the addition is not preceded by a user request or if it deviates from typical user behavior. If an attacker successfully registers a new MFA method on a compromised account, they can solidify their access, making it harder for legitimate users to regain control. The attacker can then operate with the privileges of the compromised account, potentially accessing sensitive data, making unauthorized changes, or even escalating their privileges further. Immediate action would be required to verify the legitimacy of the MFA change and, if malicious, to remediate and secure the affected account.

- **PowerShell Remote Services Add TrustedHost**
  - The following analytic identifies a suspicious PowerShell script execution via EventCode 4104 that contains command to add or modify the ‘Trustedhost’ configuration in Windows OS. This behavior raises concerns due to the nature of modifications made to the 'TrustedHost' configuration, which typically involve adjusting settings crucial for remote connections and security protocols. Alterations in this area could potentially indicate attempts to manipulate trusted hosts or systems for unauthorized remote access, a tactic commonly observed in various unauthorized access or compromise attempts.

- **O365 Multiple AppIDs and UserAgents Authentication Spike**
  - This analytic is crafted to identify unusual and potentially malicious authentication activity within an O365 environment. It triggers when a single user account is involved in more than eight (8) authentication attempts, using three (3) or more unique application IDs, and more

than five (5) unique user agents, within a short timeframe. This pattern is atypical for regular user behavior and may indicate an adversary's attempt to probe the environment, testing for multi-factor authentication requirements across different applications and platforms. The detection is based on analysis of O365 audit logs, specifically focusing on authentication events. It employs statistical thresholds to highlight instances where the volume of authentication attempts and the diversity of application IDs and user agents associated with a single user account exceed normal parameters. Identifying this behavior is crucial as it provides an early indication of potential account compromise. Adversaries, once in possession of user credentials, often conduct reconnaissance to understand the security controls in place, including multi-factor authentication configurations. Tools like Invoke-MFASweep are commonly used for this purpose, automating the process of testing different user agents and application IDs to bypass MFA. By detecting these initial probing attempts, security teams can swiftly respond, potentially stopping an attack in its early stages and preventing further unauthorized access. This proactive stance is vital for maintaining the integrity of the organization's security posture. If validated as a true positive, this detection points to a compromised account, signaling that an attacker is actively attempting to navigate security controls to maintain access and potentially escalate privileges. This could lead to further exploitation, lateral movement within the network, and eventual data exfiltration. Recognizing and responding to this early stage of an attack is vital for preventing substantial harm and safeguarding sensitive organizational data and systems.

- **Windows Modify Registry DisableRemoteDesktopAntiAlias**
  - The following analytic identifies a modification in the Windows registry to DisableRemoteDesktopAntiAlias. This registry setting might be intended to manage or control anti-aliasing behavior (smoothing of edges and fonts) within remote desktop sessions. DarkGate malware modifies this registry as part of its malicious installation in a targeted host for its remote desktop capabilities.

- **O365 Multi-Source Failed Authentications Spike**
  - This analytic detects potential distributed password spraying attacks within an Office 365 environment. It identifies a significant increase in failed authentication attempts characterized by diverse user-and-IP address combinations, originating from multiple source IP addresses and utilizing various user agents. These patterns may indicate an adversary's attempt to circumvent security controls by employing a spectrum of IP addresses to test commonly used passwords against a wide range of user accounts. The detection examines UserLoginFailed events from O365 Management Activity logs, with a particular focus on events with ErrorNumber 50126, which indicates a failed authentication due to incorrect credentials. By aggregating data over a five-minute interval, the analytic calculates the distinct counts of user-and-IP combinations and unique users and source IPs. It then applies a set of thresholds to these metrics to identify abnormal activities that could suggest a coordinated attack. The predefined thresholds within the analytic (such as unique IPs, unique users, etc.) serve as initial benchmarks and should be tailored to align with the organization's typical user behavior and risk tolerance. Early detection of such distributed activities is crucial for security operations centers (SOCs) to intercept unauthorized access attempts, avert account takeovers, and reduce the risk of subsequent malevolent actions within the organization's systems. A true positive alert from this analytic would indicate an ongoing distributed password spraying campaign targeting the organization's Office 365 tenant. If such an attack is successful, it could lead to unauthorized access, especially to accounts with administrative privileges, resulting in data breaches, privilege escalation, persistent threats, and lateral movement within the organization's digital environment.

- **O365 Multiple Failed MFA Requests for User**
  - This analytic identifies potential "MFA fatigue" attacks targeting Office 365 users. Specifically, it detects scenarios where a user experiences more than nine Multi-Factor Authentication (MFA) prompts within a 10-minute timeframe. Attackers may exploit MFA fatigue by repeatedly triggering MFA requests, hoping that the user, out of frustration or oversight, will approve a malicious authentication attempt. The detection leverages O365 management activity logs, focusing on Azure Active Directory events. It looks for the UserLoginFailed operation combined with a Success ResultStatus and an ErrorNumber of 500121, which indicates MFA prompts. By monitoring these specific events and conditions,

the analytic captures and alerts on potential MFA fatigue scenarios. With MFA being a cornerstone of modern cybersecurity defenses, attackers are constantly seeking ways to bypass or exploit it. MFA fatigue is one such tactic, where attackers rely on user frustration or confusion caused by frequent MFA prompts. Detecting potential MFA fatigue scenarios allows security teams to proactively investigate and ensure that users aren't inadvertently granting access to malicious actors. If this detection flags a true positive, it suggests a potential attempt by an attacker to exploit MFA mechanisms to gain unauthorized access to an O365 account. Successful exploitation could lead to data breaches, unauthorized data access, or further compromise within the O365 environment. Immediate investigation and response would be crucial to safeguard the affected account and assess the full scope of the potential breach.

- **Windows Modify Registry AuthenticationLevelOverride**
  - The following analytic identifies a modification in the Windows registry related to authentication level settings. This registry is the configuration for authentication level settings within the Terminal Server Client settings in Windows. AuthenticationLevelOverride might be used to control or override the authentication level used by the Terminal Server Client for remote connections. DarkGate malware modifies this registry as part of its malicious installation in a targeted host for its remote desktop capabilities.

- **Windows Modify Registry DisableSecuritySettings**
  - The following analytic identifies a modification in the Windows registry to disable security settings of Terminal Services, altering or disabling security settings within Terminal Services. Terminal Services, now known as Remote Desktop Services (RDS) in more recent Windows versions, allows users to access applications, data, and even an entire desktop, remotely. DarkGate malware modifies this registry as part of its malicious installation in a targeted host for its remote desktop capabilities.

- **Windows Modify Registry ProxyEnable**
  - The following analytic identifies a modification in the Windows registry to enable proxy. This method has been exploited by various malware and adversaries to establish proxy communication on compromised hosts, facilitating connections to malicious Command and Control (C2) servers. Identifying this anomaly serves as a crucial indicator to unveil suspicious processes attempting to activate the proxy feature within the Windows operating system. Detecting such attempts becomes pivotal in flagging potential threats, especially those aiming to leverage proxy configurations for unauthorized communication with malicious entities.

- **Windows Credentials from Password Stores Creation**
  - The following analytic identifies a process execution of Windows OS cmdkey.exe tool. This tool is being abused or used by several post exploitation tools and malware, such as DarkGate malware, to create stored usernames, passwords, or credentials in the targeted Windows OS host. This information can be used by the attacker to gain privilege escalation and persistence in the targeted hosts for further attacks.

- **Windows Archive Collected Data via RAR**
  - The following analytic identifies a process to execute a RAR utilities to archive files. This method has been exploited by various threat actors, including red-teamers and malware like DarkGate, to gather and compress collected data on compromised hosts. Subsequently, these archives are transmitted to command-and-control servers as part of their data exfiltration techniques. These adversaries leverage RAR archiving to consolidate and compress collected data on compromised hosts. Once the data is compiled into these archives, it serves as a means for these entities to effectively exfiltrate sensitive information. This process involves transferring the archived data to command-and-control servers, facilitating the extraction and retrieval of critical information from compromised systems.

- **O365 Mail Permissioned Application Consent Granted by User**
  - The following analytic identifies instances where a user grants consent to an application that requests mail-related permissions within the Office 365 environment. This could involve

permissions to read, send, or manage mail settings. It leverages the O365 audit logs, specifically events related to application permissions and user consent actions. By filtering for mail-related permissions and user-granted consents, the analytic pinpoints potential security concerns. While many legitimate applications request mail permissions for valid reasons, malicious actors can exploit these permissions for data exfiltration, spear phishing, or other malicious activities. By monitoring for user-granted mail permissions, security teams can identify and review potentially risky consents, ensuring that only trusted applications have access to sensitive email data. If the detection is a true positive, it indicates that an application now has access to the user's mail data as permitted. In the hands of a malicious actor, this could lead to unauthorized data access, email forwarding, or even the sending of malicious emails from the compromised account. It's crucial to validate the legitimacy of the application and the context of the consent to prevent potential data breaches or further malicious activities.

- **Windows Modify Registry ProxyServer**
  - The following analytic identifies a modification in the Windows registry to setup proxy server. This method has been exploited by various malware and adversaries to establish proxy communication on compromised hosts, facilitating connections to malicious Command and Control (C2) servers. Identifying this anomaly serves as a crucial indicator to unveil suspicious processes attempting to activate the proxy feature within the Windows operating system. Detecting such attempts becomes pivotal in flagging potential threats, especially those aiming to leverage proxy configurations for unauthorized communication with malicious entities.

- **O365 Block User Consent for Risky Apps Disabled**
  - This analytic detects when the "risk-based step-up consent" security setting in Microsoft 365 is disabled. This setting, when enabled, prevents regular users from granting consent to potentially malicious OAuth applications, requiring an administrative "step-up" for consent instead. Disabling this feature could expose the organization to OAuth phishing threats. The detection operates by monitoring Azure Active Directory logs for events where the "Update authorization policy" operation is performed. It specifically looks for changes to the "AllowUserConsentForRiskyApps" setting, identifying instances where this setting is switched to "true," effectively disabling the risk-based step-up consent. Monitoring for changes to critical security settings like the "risk-based step-up consent" is vital for maintaining the integrity of an organization's security posture. Disabling this feature can make the environment more susceptible to OAuth phishing attacks, in which attackers trick users into granting permissions to malicious applications. Identifying when this setting is disabled can help blue teams to quickly respond, investigate, and potentially uncover targeted phishing campaigns against their users. If an attacker successfully disables the "risk-based step-up consent" and subsequently launches an OAuth phishing campaign, they could gain unauthorized access to user data and other sensitive information within the M365 environment. This could lead to data breaches, unauthorized access to emails, and potentially further compromise within the organization.

- **Windows Credentials from Password Stores Deletion**
  - The following analytic identifies a process execution of Windows OS cmdkey.exe tool. This tool is being abused or used by several post exploitation tools and malware, such as Darkgate malware, to delete stored usernames, passwords or credentials in the targeted Windows OS host. This information can be used by the attacker to gain privilege escalation and persistence in the targeted hosts for further attacks.

- **Windows Defender ASR Rule Disabled**
  - The following analytic identifies when a Windows Defender ASR rule has disabled events. ASR is a feature of Windows Defender Exploit Guard that prevents actions and apps that are typically used by exploit-seeking malware to infect machines. ASR rules are applied to processes and applications. When a process or application attempts to perform an action that is blocked by an ASR rule, an event is generated. This detection searches for ASR rule disabled events that are generated when an ASR rule is disabled.

- **Windows Defender ASR Audit Events**

    - This detection searches for Windows Defender ASR audit events. ASR is a feature of Windows Defender Exploit Guard that prevents actions and apps that are typically used by exploit-seeking malware to infect machines. ASR rules are applied to processes and applications. When a process or application attempts to perform an action that is blocked by an ASR rule, an event is generated. This detection searches for ASR audit events that are generated when a process or application attempts to perform an action that would be blocked by an ASR rule but is allowed to proceed for auditing purposes.

- **Windows Defender ASR Rules Stacking**
    - This hunting analytic targets a range of security events from Microsoft Defender, focusing on the Exploit Guard and Attack Surface Reduction (ASR) features. It monitors specific Event IDs. Event IDs 1121 and 1126 indicate active blocking of unauthorized operations or dangerous network connections, whereas Event IDs 1122 and 1125 represent audit logs for similar activities. Event ID 1129 shows user overrides on blocked operations. For ASR-related activities, Event IDs 1131 and 1133 signal blocked operations, while 1132 and 1134 are audit logs. Event ID 5007 alerts on configuration changes, possibly indicating security breaches. Additionally, the analytic utilizes a lookup to correlate ASR rule GUIDs with their descriptive names, enhancing understanding of the context behind these security alerts. This includes rules for blocking vulnerable drivers, restricting actions of Adobe Reader and Office applications and protecting against various malware and unauthorized system changes. This comprehensive approach aids in assessing policy enforcement and potential security risks.

- **Windows Modify Registry DontShowUI**
    - The following analytic identifies a modification in the Windows Error Reporting registry to DontShowUI. DarkGate malware modifies this registry as part of its malicious installation in a targeted host for its remote desktop capabilities. When this registry value is present and set to a specific configuration, it can influence the behavior of error reporting dialogs or prompts, suppressing them from being displayed to the user. For instance, setting DontShowUI to a value of 1 often indicates that the Windows Error Reporting UI prompts will be suppressed, meaning users won't see error reporting pop-ups when errors occur.

- **O365 Mailbox Read Access Granted to Application**
    - The following analytic identifies instances where the Mail.Read Graph API permissions are granted to an application registration within an Office 365 tenant. It leverages O365 audit logs, specifically events related to changes in application permissions within the AzureActiveDirectory workload. The Mail.Read permission allows applications to access and read all emails within a user's mailbox. Emails often contain sensitive or confidential information, and unauthorized access can lead to data breaches or leakage. Monitoring the assignment of this permission ensures that only legitimate applications have such access and that any inadvertent or malicious assignments are promptly identified. If an attacker successfully grants this permission to a malicious or compromised application, they can read all emails in the affected mailboxes. This can lead to data exfiltration, spear-phishing attacks, or further compromise based on the information gathered from the emails.

- **O365 File Permissioned Application Consent Granted by User**
    - This analytic identifies instances where a user in the Office 365 environment grants consent to an application that requests file permissions, specifically targeting OneDrive or SharePoint. Such permissions mean the application could potentially access, modify, or delete files stored within these services. The detection process leverages O365 audit logs, particularly focusing on events related to OAuth application consents. By examining these logs, the analytic is designed to capture and alert on any actions where users grant consent to applications requesting file-related permissions for OneDrive or SharePoint. The sensitivity of file permissions, especially in platforms as widely utilized as OneDrive and SharePoint, cannot be overstated. While many legitimate applications might require such permissions to operate, there's an inherent risk with malicious or overly permissive applications. Attackers could craft or exploit applications to gain file permissions, aiming to access, exfiltrate, or manipulate sensitive data housed in OneDrive or SharePoint. It's crucial for security operations centers to monitor these consents to ensure that only trustworthy applications gain access and that users aren't inadvertently granting

permissions to potentially harmful applications. If this detection flags a true positive, it indicates that an application has been granted permissions which could allow it to interact with OneDrive or SharePoint files in potentially malicious ways. Such actions could lead to data breaches, data loss, or unauthorized data manipulation. Immediate investigation would be required to validate the application's legitimacy, understand the nature of its requested permissions, and assess the potential risks associated with the access it's been granted.

- **Windows Defender ASR Block Events**
  - This detection searches for Windows Defender ASR block events. ASR is a feature of Windows Defender Exploit Guard that prevents actions and apps that are typically used by exploit-seeking malware to infect machines. ASR rules are applied to processes and applications. When a process or application attempts to perform an action that is blocked by an ASR rule, an event is generated. This detection searches for ASR block events that are generated when a process or application attempts to perform an action that is blocked by an ASR rule. Typically, these will be enabled in block most after auditing and tuning the ASR rules themselves. Set to TTP once tuned.

- **O365 ApplicationImpersonation Role Assigned**
  - The following analytic identifies the assignment of the ApplicationImpersonation role in Office 365, either to a user or an application. This analytic leverages the Office 365 Management Activity API, specifically monitoring for events related to role assignments and changes within the Azure Active Directory audit logs. The ApplicationImpersonation role allows a security principal to impersonate any user within the organization and perform actions on their behalf, such as accessing or modifying their mailbox. This role, if misused or granted inappropriately, can pose a significant security risk. Monitoring the assignment of this role is crucial as it can be an indicator of potential malicious activity or misconfigurations. If an attacker successfully assigns the ApplicationImpersonation role to a malicious user or application, they can gain the ability to impersonate any user within the organization. This can lead to unauthorized access to sensitive information, manipulation of mailbox data, and other malicious actions. The attacker can effectively masquerade as a legitimate user, making their actions harder to detect and potentially causing significant harm to the organization.

- **O365 Multiple Users Failing to Authenticate From IP**
  - This analytic identifies instances where multiple users (more than 10 unique accounts) have failed to authenticate from a single IP address within a short time span (five minutes). Such a pattern can be indicative of malicious activities, such as brute-force attacks or password spraying attempts. The detection leverages O365 audit logs, specifically focusing on Azure Active Directory login failures (AzureActiveDirectoryStsLogon). By aggregating these failures based on the source IP address and time, the analytic captures patterns where multiple unique user accounts have authentication failures from the same IP within a five-minute window. Multiple authentication failures from a single IP address targeting various accounts can be a strong indicator of an attacker trying to gain unauthorized access. It could represent a brute-force attack, password spraying, or other malicious login attempts. Identifying and responding to such patterns promptly is crucial to prevent potential account compromises and unauthorized access to organizational resources. If the detection is a true positive, it suggests that an external entity is actively trying to breach the security by targeting multiple user accounts. While the attempts have been unsuccessful (as indicated by the login failures), it's a clear sign of malicious intent. Immediate action is required to block or monitor the suspicious IP, investigate the nature of the attempts, and potentially notify affected users to take precautionary measures like password changes or enabling multi-factor authentication.

- **O365 High Number of Failed Authentications for User**
  - The following analytic identifies an O365 account that has experienced more than 20 failed authentication events within a span of five minutes. This could be indicative of an attacker attempting to brute force or guess the password for a user account. It leverages the O365 Unified Audit Logs, specifically the "UserLoginFailed" events. By monitoring the frequency and volume of these events for individual users, the analytic can flag accounts that exceed the set threshold of failed attempts within the defined timeframe. Multiple failed login

attempts in a short period can be a strong indicator of malicious activity. While there could be benign reasons, such as a user forgetting their password, the rapid succession of failed attempts is often a sign of an attacker trying to gain unauthorized access. By detecting and alerting on this behavior, the SOC can quickly investigate and take appropriate action, potentially stopping an attack in its early stages. Given that environments differ across organizations, security teams should consider customizing the threshold of this detection to better suit their specific needs and risk profile. If an attacker successfully guesses or brute-forces a user's password after numerous attempts, they can gain unauthorized access to the O365 environment. This unauthorized access could allow them to view sensitive emails, documents, and other data.

- **Windows Indicator Removal Via Rmdir**
  - The following analytic identifies a process executing rmdir via command line to delete files and directory tree. This technique has been observed in the actions of various malware strains, such as DarkGate, as they attempt to eliminate specific files or components during their cleanup operations within compromised hosts. Notably, this deletion method doesn't exclusively require elevated privileges and can be executed by regular users or network administrators, although it's not the typical approach used for file deletion.

- **Windows Parent PID Spoofing with Explorer**
  - The following analytic identifies a suspicious explorer.exe process that has "/root" process command line. The presence of this parameter is considered a significant indicator as it could indicate attempts at spoofing the parent process by a specific program or malware. By spoofing the parent process, the malicious entity aims to circumvent detection mechanisms and operate undetected within the system. This technique of manipulating the command-line parameter (/root) of explorer.exe is a form of masquerading utilized by certain malware or suspicious processes. The objective is to obscure the true nature of the activity by imitating a legitimate system process. By doing so, it attempts to evade scrutiny and evade detection by security measures.

- **O365 Mailbox Inbox Folder Shared with All Users**
  - The following analytic identifies instances where the inbox folder of a mailbox in Office 365 is shared with all users within the tenant. Sharing the inbox folder with all users is an unusual and risky configuration. Attackers have been known to exploit this setting to surreptitiously read a target user's emails from another account. Such unauthorized access can lead to data breaches, leakage of confidential information, or further compromise based on the information gathered from the emails. Monitoring for this configuration change ensures that inadvertent or malicious sharing is promptly identified and addressed. If an attacker successfully configures the inbox to be shared with all users, they can access and read all emails in the affected mailbox from any account within the tenant. This can lead to data exfiltration, spear-phishing attacks based on the information in the emails, or further malicious activities using sensitive information gathered from the mailbox.

- **O365 Tenant Wide Admin Consent Granted**
  - The following analytic identifies instances where admin consent is granted to an application within an Azure AD and Office 365 tenant. It leverages O365 audit logs, specifically events related to the admin consent action within the AzureActiveDirectory workload. The admin consent action allows applications to access data across the entire tenant, potentially encompassing a vast amount of organizational data. Given its broad scope and the sensitivity of some permissions that can only be granted via admin consent, it's crucial to monitor this action. Unauthorized or inadvertent granting of admin consent can lead to significant security risks, including data breaches, unauthorized data access, and potential compliance violations. If an attacker successfully tricks an administrator into granting admin consent to a malicious or compromised application, they can gain extensive and persistent access to organizational data. This can lead to data exfiltration, espionage, further malicious activities within the tenant, and potential breaches of compliance regulations.

- **O365 Application Registration Owner Added**
  - The following analytic identifies instances where a new owner is assigned to an application registration within an Azure AD and Office 365 tenant. It leverages O365 audit logs,

specifically events related to changes in owner assignments within the AzureActiveDirectory workload for application registrations. Assigning a new owner to an application registration can grant significant control over the application's configuration, permissions, and behavior. An unauthorized or inadvertent change in ownership can lead to misuse of the application, potentially affecting data access, user permissions, or the application's interactions within the tenant. Monitoring for such changes ensures that only legitimate and authorized personnel have control over application registrations. If an attacker successfully assigns themselves or a compromised account as an owner to an application registration, they can modify the application's settings, permissions, and behavior. This can lead to unauthorized data access, escalation of privileges, or the introduction of malicious behavior within the application's operations.

- **Web Remote ShellServlet Access**
  - The following analytic identifies an attempt to access the Remote ShellServlet on a web server. This servlet is used to execute commands on the server. This activity is often associated with web shells and other malicious activity. This activity was identified against a Confluence server related to CVE-2023-22518 and CVE-2023-22515. Activity prior to access to the shell servlet includes adding a plugin to Confluence. In addition, monitor for ShellServlet?act=3, ShellServlet or obfuscated variations including Sh3llServlet1.

- **Windows Masquerading Msdtc Process**
  - The following analytic identifies a suspicious msdtc.exe with specific command-line parameters, particularly -a or -b, which are regarded as potential indicators of the presence of the insidious PlugX malware. This malware is notorious for its covert operations and is frequently utilized by threat actors for unauthorized access, data exfiltration, and espionage. The analytic's focus on the -a or -b command-line parameters within msdtc.exe is rooted in the PlugX malware's sophisticated tactic of masquerading its activities. To elude detection, PlugX employs a technique where it injects a concealed, headless PlugX Dynamic Link Library (DLL) module into the legitimate msdtc.exe process. By leveraging these specific command-line parameters, the malware attempts to disguise its presence within a system's legitimate processes, thereby evading immediate suspicion.

- **O365 User Consent Blocked for Risky Application**
  - The following analytic identifies instances where Office 365 has blocked a user's attempt to grant consent to an application deemed risky or potentially malicious. This suggests that the application has exhibited behaviors or characteristics that are commonly associated with malicious intent or that it poses a security risk. This detection leverages the O365 audit logs, specifically focusing on events related to user consent actions and system-driven blocks. By filtering for blocked consent actions associated with applications, the analytic highlights instances where O365's built-in security measures have intervened. Applications that are flagged and blocked by O365 typically exhibit suspicious characteristics or behaviors. Monitoring for these blocked consent attempts helps security teams identify potential threats early on and can provide insights into users who might be targeted or susceptible to such risky applications. It's an essential layer of defense in ensuring that malicious or risky applications don't gain access to organizational data. If the detection is a true positive, it indicates that the built-in security measures of O365 successfully prevented a potentially harmful application from gaining access. However, the attempt itself suggests that either a user might be targeted or that there's a presence of malicious applications trying to infiltrate the organization. Immediate investigation is required to understand the context of the block and to take further preventive measures.

- **O365 Advanced Audit Disabled**
  - The following analytic identifies instances where the O365 advanced audit is disabled for a specific user within the Office 365 tenant. It leverages O365 audit logs, specifically events related to audit license changes or modifications within the AzureActiveDirectory workloads. The O365 advanced audit provides granular logging and insights into user and administrator activities, making it a crucial tool for security monitoring and incident response. Disabling this audit for a user can blind security teams to potential malicious or unauthorized activities related to that user's mailbox or account. Attackers may disable these audits to obscure their actions and reduce the chances of detection. If an attacker

successfully disables the O365 advanced audit for a user, they can operate within that user's mailbox or account with reduced risk of detection. This can lead to unauthorized data access, data exfiltration, account compromise, or other malicious activities without leaving a detailed audit trail.

- **O365 High Privilege Role Granted**
    - This analytic detects when high-privilege roles, specifically "Exchange Administrator," "SharePoint Administrator," or "Global Administrator," are granted within Office 365. By monitoring O365 audit logs for events where these administrative roles are assigned to any user or service account, the analytic provides insight into critical role changes. The assignment of these roles is of paramount importance to Security Operations Centers (SOCs) as they grant extensive permissions, allowing for broad access and control over critical organizational resources and data. An unexpected or unauthorized role assignment could indicate potential malicious activity, insider threats, or misconfigurations. If an attacker or unauthorized individual is granted one of these roles, the potential impact includes gaining significant control over O365 resources, accessing, modifying, or deleting critical data, making configuration changes, and potentially compromising the overall security and functionality of the O365 environment.

- **Windows Defender ASR Registry Modification**
    - This detection searches for Windows Defender ASR registry modification events. ASR is a feature of Windows Defender Exploit Guard that prevents actions and apps that are typically used by exploit-seeking malware to infect machines. ASR rules are applied to processes and applications. When a process or application attempts to perform an action that is blocked by an ASR rule, an event is generated. This detection searches for ASR registry modification events that are generated when a process or application attempts to modify a registry key that is blocked by an ASR rule. Typically, these will be enabled in block most after auditing and tuning the ASR rules themselves. Set to TTP once tuned.

- **O365 Service Principal New Client Credentials**
    - The following analytic identifies the addition of new credentials for service principals in addition to existing legitimate credentials within an Office 365 tenant. These credentials include both x509 certificates and passwords. It leverages O365 audit logs, specifically events related to credential modifications or additions within the AzureActiveDirectory workload for service principals. Service principals represent application identities in Office 365 / AzureAD, and their credentials allow applications to authenticate and access resources. Adding new credentials or modifying existing ones can be an indication of configuration changes, but it can also be a sign of malicious intent. If an attacker successfully adds or modifies credentials for a service principal, they can potentially use those credentials to authenticate as the application, gaining access to resources and data the application is permitted to access. This can lead to unauthorized data access, data exfiltration, or malicious operations performed under the guise of the application.

- **O365 User Consent Denied for OAuth Application**
    - The following analytic identifies instances where a user has actively denied consent to an OAuth application seeking permissions within the Office 365 environment. This suggests that the user either recognized something suspicious about the application or chose not to grant it the requested permissions for other reasons. This detection leverages the O365 audit logs, specifically focusing on events related to user consent actions. By filtering for denied consent actions associated with OAuth applications, the analytic captures instances where users have actively rejected permission requests. While user-denied consents can be routine, they can also be indicative of users spotting potentially suspicious or unfamiliar applications. By monitoring these denied consent attempts, security teams can gain insights into applications that might be perceived as risky or untrusted by users. It can also serve as a feedback loop for security awareness training, indicating that users are being cautious about granting permissions. If the detection is a true positive, it indicates that a user has actively prevented an OAuth application from gaining the permissions it requested. While this is a proactive security measure on the user's part, it's essential for security teams to review the context of the denial. Understanding why certain applications are being denied can help in refining application whitelisting policies and ensuring that no malicious

applications are attempting to gain access.

- **Commands to Disable Shadow Copy Service**
    - "Looking for commands used to disable Volume Shadow Copy Services. These have been used by ransomware operators, but it is also used in commercial backup software making it hard to distinguish. The commands will look like the following:
        - sc stop vss
        - sc.exe stop ""Acronis VSS Provider""
        - sc.exe stop BackupExecVSSProvider
        - sc config vss start=disabled
        - net stop ShadowProtectSvc
        - net stop vss
        - net.exe stop BackupExecVSSProvider
        - net.exe stop ""Acronis VSS Provider""
        - net stop VSNAPVSS"

[i] https://blogs.blackberry.com/en/2023/11/aeroblade-on-the-hunt-targeting-us-aerospace-industry

[ii] https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-347a, https://www.fortinet.com/blog/threat-research/teamcity-intrusion-saga-apt29-suspected-exploiting-cve-2023-42793

[iii] https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-339a

[iv] https://blog.talosintelligence.com/lazarus_new_rats_dlang_and_telegram/

[v] https://securityintelligence.com/x-force/itg05-ops-leverage-Israel-Hamas-conflict-lures-to-deliver-headlace-malware/, https://unit42.paloaltonetworks.com/russian-apt-fighting-ursa-exploits-cve-2023-233397/, https://www.microsoft.com/en-us/security/blog/2023/03/24/guidance-for-investigating-attacks-using-cve-2023-23397/

[vi] https://lab52.io/blog/mustang-pandas-plugx-new-variant-targetting-taiwanese-government-and-diplomats/