

Understanding
Vectra AI Detections

VECTRA®

August 2025

Table of Contents

| | |
|--|-----------|
| Network Detections | 6 |
| Command & Control | 7 |
| External Remote Access | 8 |
| ICMP Tunnel | 9 |
| Hidden Tunnel | 10 |
| Hidden DNS Tunnel | 11 |
| Hidden HTTP Tunnel | 12 |
| Hidden HTTPS Tunnel | 13 |
| Malware Update | 14 |
| Multi-home Fronted Tunnel | 15 |
| Peer-To-Peer | 16 |
| Stealth HTTP Post | 17 |
| Suspect Domain Activity | 18 |
| Suspicious HTTP | 19 |
| Suspicious Relay | 20 |
| TOR Activity | 21 |
| Threat Intelligence Match | 22 |
| Vectra Threat Intelligence Match | 23 |
| Botnet Activity | 24 |
| Brute-Force | 25 |
| Cryptocurrency Mining | 26 |
| Outbound DoS | 27 |
| Outbound Port Sweep | 28 |
| Reconnaissance | 29 |
| File Share Enumeration | 30 |
| Internal Darknet Scan | 31 |
| Kerberoasting: Weak Cipher Request | 32 |
| Kerberoasting: SPN Sweep | 33 |
| Kerberos Account Scan | 34 |
| Kerberos Brute-Sweep | 35 |
| RDP Recon | 36 |
| RPC Recon | 37 |
| RPC Targeted Recon | 38 |
| SMB Account Scan | 39 |
| Suspicious LDAP Query | 40 |
| Suspicious Port Scan | 41 |
| Suspicious Port Sweep | 42 |
| Lateral Movement | 43 |
| Automated Replication | 44 |
| Brute-Force | 45 |
| Credential Access via NTLM Relay | 46 |
| ICMP Tunnel: Client | 47 |
| ICMP Tunnel: Server | 48 |
| Kerberoasting: Targeted Weak Cipher Response | 49 |
| Privilege Anomaly: Unusual Account on Host | 50 |
| Privilege Anomaly: Unusual Host | 51 |
| Privilege Anomaly: Unusual Service | 52 |
| Privilege Anomaly: Unusual Service - Insider | 53 |
| Privilege Anomaly: Unusual Service from Host | 55 |
| Privilege Anomaly: Unusual Trio | 56 |
| Ransomware File Activity | 57 |
| SMB Brute-Force | 58 |
| Shell Knocker Client | 59 |
| Shell Knocker Server | 60 |
| SQL Injection Activity | 61 |
| Stage Loader | 62 |
| Suspicious Active Directory Operations | 63 |
| Suspicious Admin | 65 |
| Suspicious Remote Desktop | 67 |
| Suspicious Remote Execution | 68 |
| Threat Intelligence Match | 69 |

| | |
|---|------------|
| Exfiltration | 70 |
| Data Gathering | 71 |
| Data Smuggler | 72 |
| Hidden DNS Tunnel | 73 |
| Hidden HTTP Tunnel | 74 |
| Hidden HTTPS Tunnel | 75 |
| ICMP Tunnel | 76 |
| Smash and Grab | 77 |
| Threat Intelligence Match | 78 |
| Suspect Protocol Activity Detections | 79 |
| Info | 80 |
| Network Detection Profiles | 82 |
| Botnet | 83 |
| Cloud Services | 84 |
| External Adversary | 85 |
| Insider Threat: Admin | 86 |
| Insider Threat: User | 87 |
| IT Discovery | 88 |
| IT Services | 89 |
| Potentially Unwanted Program | 90 |
| Ransomware | 91 |
| Vulnerability Discovery | 92 |
| Worm | 93 |
| Observed Privilege Scores | 94 |
| Microsoft 365 Detections | 97 |
| Command & Control | 98 |
| M365 Power Automate HTTP Flow Creation | 99 |
| M365 Suspicious Power Automate Flow Creation | 100 |
| Reconnaissance | 101 |
| M365 Suspicious Compliance Search | 102 |
| M365 Unusual eDiscovery Search | 103 |
| M365 Suspect eDiscovery Usage | 104 |
| M365 Suspicious Copilot for M365 Access | 105 |
| M365 Copilot Sensitive Data Discovery | 106 |
| Lateral Movement | 107 |
| M365 Attacker Tool: Ruler | 108 |
| M365 Disabling of Security Tools | 109 |
| M365 DLL Hijacking Activity | 110 |
| M365 External Teams Access | 111 |
| M365 Internal Spearphishing | 112 |
| M365 Log Disabling Attempt | 113 |
| M365 Malware Stage: Upload | 114 |
| M365 Ransomware | 115 |
| M365 Risky Exchange Operation | 116 |
| M365 Phishing Simulation Configuration Change | 117 |
| M365 SecOps Mailbox Change | 118 |
| M365 Suspicious Teams Application | 119 |
| M365 Suspicious Mailbox Manipulation | 120 |
| M365 Suspicious Mailbox Rule Creation | 121 |
| Exfiltration | 122 |
| M365 eDiscovery Exfil | 123 |
| M365 Exfiltration Before Termination | 124 |
| M365 Suspicious Download Activity | 125 |
| M365 Suspicious Exchange Transport Rule | 126 |
| M365 Suspicious Mail Forwarding | 127 |
| M365 Suspect Power Automate Activity | 128 |
| M365 Suspicious Sharing Activity | 129 |

| | |
|--|------------|
| Info | 130 |
| Azure AD Detections | 132 |
| Command & Control | 133 |
| Azure AD Admin Account Creation | 134 |
| Azure AD Cross Tenant Access Change | 135 |
| Azure AD Domain Settings Modified | 136 |
| Azure AD Login From Suspicious Location | 137 |
| Azure AD MFA-Failed Suspicious Sign-On | 138 |
| Azure AD New Certification Authority Registered | 139 |
| Azure AD New Partner Added to Organization | 140 |
| Azure AD Redundant Access Creation | 141 |
| Azure AD Suspicious OAuth Application | 142 |
| Azure AD Suspicious Sign-on | 143 |
| Azure AD Suspicious Access from Cloud Provider | 144 |
| Azure AD Suspected Compromised Access | 145 |
| Azure AD TOR Activity | 146 |
| Lateral Movement | 147 |
| Azure AD Successful Brute-Force | 148 |
| Azure AD Successful Brute-Force – Failed Login | 149 |
| Azure AD Suspicious Device Registration | 150 |
| Azure AD Suspicious Factor Registration | 151 |
| Azure AD Change to Trusted IP Configuration | 152 |
| Azure AD MFA Disabled | 153 |
| Azure AD Newly Created Admin Account | 154 |
| Azure AD Privilege Operation Anomaly | 155 |
| Azure AD Unusual Scripting Engine Usage | 156 |
| Info | 157 |
| Azure Cloud Detections | 159 |
| Command & Control | 160 |
| Azure Suspicious Access from AWS Cloud | 161 |
| Azure Suspicious Access from GCP Cloud | 162 |
| Azure Suspicious Hybrid Machine Extension Installation | 163 |
| Azure Suspicious Hybrid Machine Run Command Execution | 164 |
| Azure Suspicious VM Extension Installation | 165 |
| Azure Suspicious VM Run Command Execution | 166 |
| Azure Suspicious VM Scale Set Extension Installation | 167 |
| Azure Suspicious VM Scale Set Run Command Execution | 168 |
| Azure Suspicious VM Automation Test | 169 |
| Azure Suspicious Hybrid Automation Test | 170 |
| Azure Suspicious Automation Staged | 171 |
| Azure Suspicious VM Automation Execution | 172 |
| Azure Suspicious Hybrid Automation Execution | 173 |
| Azure TOR Activity | 174 |
| Reconnaissance | 175 |
| Azure Suspect Credential Dump | 176 |
| Azure Suspect Key Vault Enumeration | 178 |
| Lateral Movement | 179 |
| Azure Anomalous App Service WebJob Activity | 180 |
| Azure Diagnostic Logging Disabled | 181 |
| Azure Managed Identity Anomaly | 182 |
| Azure Mass Resource Deletion | 183 |
| Azure Privilege Anomaly: Management Group Scope | 184 |
| Azure Privilege Anomaly: Root Scope | 185 |
| Azure Resource Group Admin Privilege Granting | 186 |
| Azure Resource Group Admin Role Unassigned | 187 |
| Azure Subscription Admin Privilege Granting | 188 |
| Azure Subscription Admin Role Unassigned | 189 |

| | |
|--|------------|
| Azure Suspect App Service Deployment Activity | 190 |
| Azure Suspect Key Vault Privilege Granting | 191 |
| Azure Suspicious App Service Creation or Modification | 192 |
| Azure Suspicious App Service Credential Download | 193 |
| Azure Suspicious App Service Deployment Configuration Download | 194 |
| Azure Suspicious Automation DSC Execution | 195 |
| Azure Suspicious Policy Assignment | 197 |
| Azure Suspicious Policy Creation or Modification | 198 |
| Azure Suspicious Policy Remediation Task | 199 |
| Azure Suspicious Serial Console Usage | 200 |
| Exfiltration | 201 |
| Azure Suspect Public Storage Account Change | 202 |
| Azure Suspicious Disk Download | 203 |
| Azure Suspicious Key Vault Extraction | 204 |
| Botnet | 205 |
| Azure Cryptomining | 206 |
| Info | 207 |
| Azure Public Storage Container Change | 208 |
| AWS Detections | 209 |
| Kingpin Technology | 210 |
| Command & Control | 212 |
| AWS Root Credential Usage | 213 |
| AWS Suspicious Credential Usage | 214 |
| AWS TOR Activity | 215 |
| Reconnaissance | 216 |
| AWS External Network Discovery | 217 |
| AWS Network Configuration Discovery | 218 |
| AWS Organization Discovery | 219 |
| AWS S3 Enumeration | 220 |
| AWS Suspect Credential Access from EC2 | 221 |
| AWS Suspect Credential Access from ECS | 222 |
| AWS Suspect Credential Access from SSM | 223 |
| AWS Suspect Discovery from EC2 Instance | 224 |
| AWS Suspect Escalation Reconnaissance | 225 |
| AWS Suspicious EC2 Enumeration | 226 |
| AWS User Permissions Enumeration | 227 |
| Lateral Movement | 228 |
| AWS Attack Tools | 229 |
| AWS Bedrock Logging Configuration Disabled | 230 |
| AWS Bedrock Novel Model Enabled | 231 |
| AWS ECR Hijacking | 232 |
| AWS Lambda Hijacking | 233 |
| AWS Logging Disabled | 234 |
| AWS Ransomware S3 Activity | 235 |
| AWS Root MFA Disabled | 236 |
| AWS S3 Objection Deletion | 237 |
| AWS Security Tools Disabled | 238 |
| AWS Suspect Admin Privilege Granting | 239 |
| AWS Suspect Console Pivot | 240 |
| AWS Suspect Login Profile Manipulation | 241 |
| AWS Suspect Organization Exit | 242 |
| AWS Suspect Privilege Escalation | 243 |
| AWS Suspect Region Activity | 244 |
| AWS Suspicious Bedrock Activity | 245 |
| AWS Suspicious S3 Batch Deletion | 246 |
| AWS Suspicious S3 Encryption | 247 |
| AWS User Hijacking | 248 |

| | |
|--------------------------------------|------------|
| Exfiltration | 249 |
| AWS Suspect External Access Granting | 250 |
| AWS Suspect Public AMI Change | 251 |
| AWS Suspect Public EBS Change | 252 |
| AWS Suspect Public RDS Change | 253 |
| AWS Suspect Public EC2 Change | 254 |
| AWS Suspect Public S3 Change | 255 |
| Botnet | 256 |
| AWS Cryptomining | 257 |
| AWS Suspect Traffic Mirror Creation | 258 |
| Info | 259 |

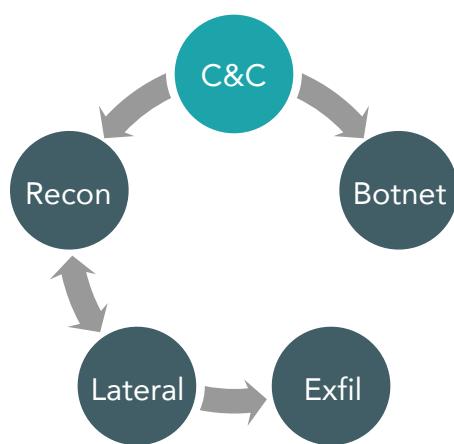
Network Detections



Category

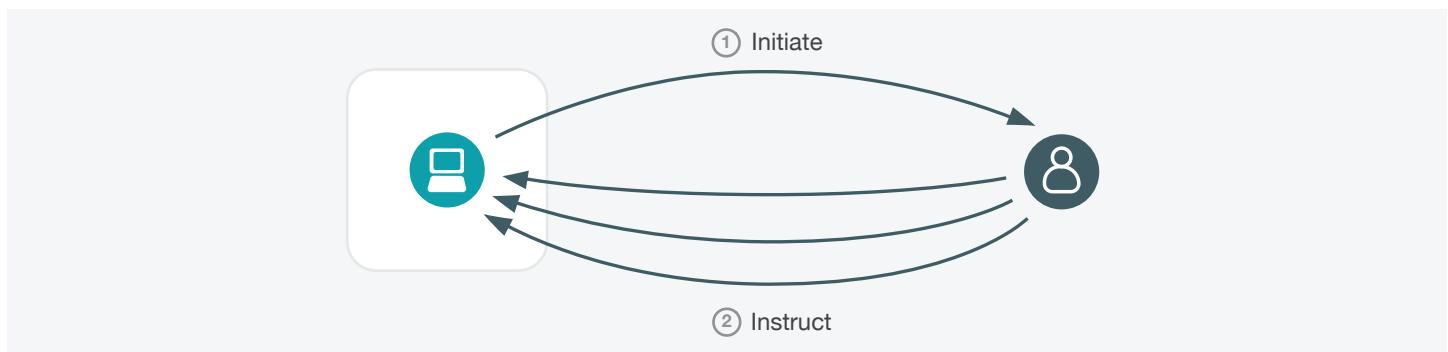
Command & Control

- A host or account appears to be under control of an external entity
- Most often, the control is automated as the host or account is part of a botnet or has adware or spyware installed
- The host or account may be manually controlled from the outside – this is the most threatening case and makes it highly likely that this is a targeted attack



External Remote Access

Command & Control



MITRE | ATT&CK®

[T1005 Data from Local System](#)

[T1115 Clipboard Data](#)

[T1071 Application Layer Protocol](#)

[T1125 Video Capture](#)

[T1090 Proxy](#)

[T1113 Screen Capture](#)

[T1010 Application Window Discovery](#)

[T1037 Boot or Logon Initialization Scripts](#)

[T1111 Two-Factor Authentication Interception](#)

[T1572 Protocol Tunneling](#)

[T1573 Encrypted Channel](#)

[T1048 Exfiltration Over Alternative Protocol](#)

[T1204 User Execution](#)

[T1056 Input Capture](#)

[T1001 Data Obfuscation](#)

[T1571 Non-Standard Port](#)

[T1059 Command and Scripting Interpreter](#)

[T1518 Software Discovery](#)

[T1176 Browser Extensions](#)

[T1123 Audio Capture](#)

[T1008 Fallback Channels](#)

[T1219 Remote Access Software](#)

[T1105 Ingress Tool Transfer](#)

[T1133 External Remote Services](#)

[T1095 Non-Application Layer Protocol](#)

[T1132 Data Encoding](#)

Triggers

- An internal host is connecting to an external server and the pattern looks reversed from normal client to server traffic; the client appears to be receiving instructions from the server and a human on the outside appears to be controlling the exchange

Possible Root Causes

- A host includes malware with remote access capability (e.g. Meterpreter, Poison Ivy) that connects to its C&C server and receives commands from a human operator
- A user has intentionally installed and is using remote desktop access software and is accessing the host from the outside (e.g. GotoMyPC, RDP)
- This behavior can also be exhibited through very active use of certain types of chat software that exposes similar human-driven behavior

Business Impact

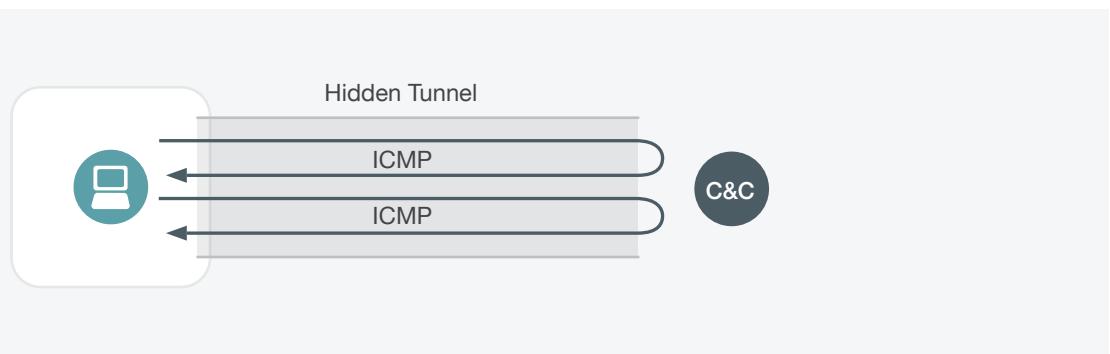
- Presence of malware with human-driven C&C is a property of targeted attacks
- Business risk associated with outside human control of an internal host is very high
- Provisioning of this style of remote access to internal hosts poses substantial risks as compromise of the service provides direct access into your network

Steps to Verify

- Look at the detection details and the PCAP to determine whether this may be traffic from chat software
- Check if a user has knowingly installed remote access software and decide whether the resulting risk is acceptable
- Scan the computer for known malware and potentially reimage it, noting that some remote access toolkits leave no trace on disk and reside entirely in memory

ICMP Tunnel

Command & Control



MITRE | ATT&CK®

[T1008 Fallback Channels](#)

[T1095 Non-Application Layer Protocol](#)

Triggers

- A host was observed using ICMP in ways inconsistent with standard implementation of the protocol.
- More precisely, a host's ICMP traffic was observed to contain datagrams which vary in size more frequently than typical ICMP traffic would.
- An attacker may be using the host to communicate with or transfer data to an external host.

Possible Root Causes

Malicious Detection

- An attacker is using ICMP as a staging and/or control channel.
- An attacker has established persistence & has chosen ICMP as a backup channel.

Benign Detection

- A network device like a vulnerability scanner is crafting nonstandard ICMP datagrams.

Business Impact

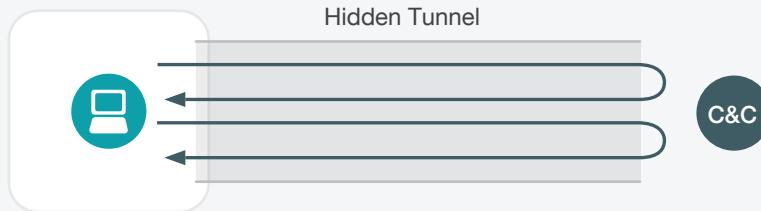
- The presence of an ICMP tunnel indicates the host was compromised & that an attacker has remote access to the machine.
- Recon, data exfiltration, lateral movement, privilege escalation, & establishing a tunnel over a more reliable protocol like HTTPS are all likely next steps.
- ICMP tunnels can be stealthy and are often used to evade sophisticated perimeter security controls.

Steps to Verify

- Check the destination IP & determine if the observed traffic arrives at a trusted endpoint.
- Investigate the host for malware, there may be code present which establishes a C2 channel with another host.

Hidden Tunnel

Command & Control



MITRE | ATT&CK®

[TA0011 Command and Control](#)

[T1571 Non-Standard Port](#)

[T1572 Protocol Tunneling](#)

[T1001 Data Obfuscation](#)

[T1132 Data Encoding](#)

Triggers

- An internal host is communicating with an external server using protocol tunneling to mask the contents of the communication channel

Possible Root Causes

- An attacker may utilize protocol tunneling to egress the network and establish a communication channel with a command-and-control server
- Software can use tunneling to simplify traversing corporate firewalls not requiring modification to existing rulesets

Business Impact

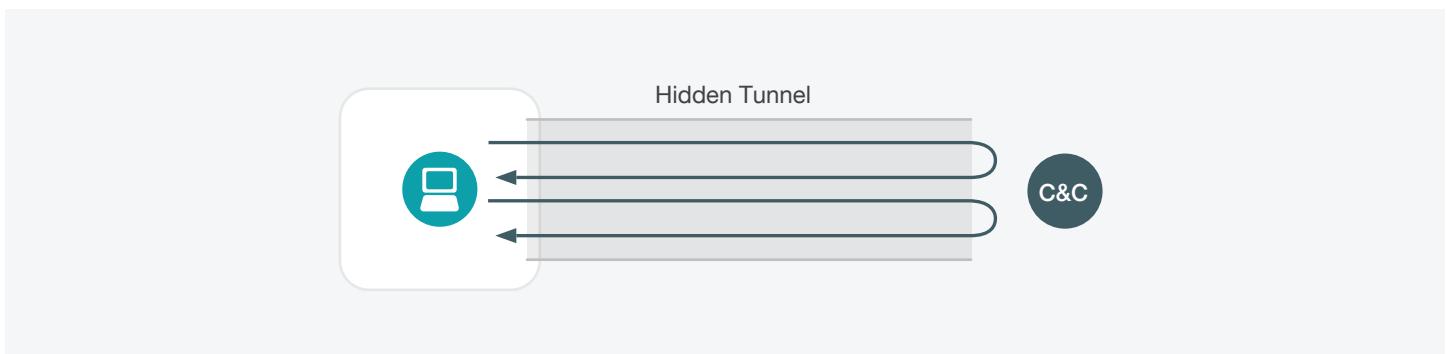
- The use of a hidden tunnel by some software may be benign, but it represents significant risk as the intention is to bypass security controls
- Hidden tunnels used as part of a targeted attack are meant to slip by your perimeter security controls and indicate a sophisticated attacker
- Hidden tunnels are rarely used by botnets, though more sophisticated bot herders with more ambitious goals may utilize them

Steps to Verify

- Check to see if the destination IP or domain of the tunnel is an entity you trust for your network
- Ask the user of the host whether they are using hidden tunnel software for any purpose
- Before removing the offending software via antivirus or reimaging, take a memory snapshot for future analysis of the incident
- If the behavior reappears shortly after a reimaging, this may be a hardware/BIOS tunnel

Hidden DNS Tunnel

Command & Control



MITRE | ATT&CK®

[T1005 Data from Local System](#)

[T1071 Application Layer Protocol](#)

[T1010 Application Window Discovery](#)

[T1037 Boot or Logon Initialization Scripts](#)

[T1572 Protocol Tunneling](#)

[T1573 Encrypted Channel](#)

[T1056 Input Capture](#)

[T1001 Data Obfuscation](#)

[T1059 Command and Scripting Interpreter](#)

[T1008 Fallback Channels](#)

[T1105 Ingress Tool Transfer](#)

[T1132 Data Encoding](#)

Triggers

- An internal host is communicating with an outside IP using DNS where another protocol is running over the top of the DNS sessions
- This represents a hidden tunnel involving multiple sessions over longer periods of time mimicking normal DNS traffic

Possible Root Causes

- A targeted attack may use hidden tunnels to hide communication with command and control servers
- A user is utilizing tunneling software to communicate with Internet services which might not otherwise be accessible
- Intentionally installed software is using a hidden tunnel to bypass expected firewall rules

Business Impact

- The use of a hidden tunnel by some software may be benign, but it represents significant risk as the intention is to bypass security controls
- Hidden tunnels used as part of a targeted attack are meant to slip by your perimeter security controls and indicate a sophisticated attacker
- Hidden tunnels are rarely used by botnets, though more sophisticated bot herders with more ambitious goals may utilize them

Steps to Verify

- Check to see if the destination domain of the tunnel is an entity you trust for your network
- Ask the user of the host whether they are using hidden tunnel software for any purpose
- Before removing the offending software via antivirus or reimaging, take a memory snapshot for future analysis of the incident
- If the behavior reappears shortly after a reimaging, this may be a hardware/BIOS tunnel

Hidden HTTP Tunnel

Command & Control



MITRE | ATT&CK®

[T1005 Data from Local System](#)

[T1115 Clipboard Data](#)

[T1071 Application Layer Protocol](#)

[T1185 Man in the Browser](#)

[T1125 Video Capture](#)

[T1113 Screen Capture](#)

[T1010 Application Window Discovery](#)

[T1037 Boot or Logon Initialization Scripts](#)

[T1111 Two-Factor Authentication Interception](#)

[T1572 Protocol Tunneling](#)

[T1204 User Execution](#)

[T1056 Input Capture](#)

[T1001 Data Obfuscation](#)

[T1571 Non-Standard Port](#)

[T1059 Command and Scripting Interpreter](#)

[T1518 Software Discovery](#)

[T1176 Browser Extensions](#)

[T1123 Audio Capture](#)

[T1008 Fallback Channels](#)

[T1105 Ingress Tool Transfer](#)

[T1132 Data Encoding](#)

Triggers

- An internal host is communicating with an outside IP using HTTP where another protocol is running over the top of the HTTP sessions
- This represents a hidden tunnel involving multiple sessions over longer periods of time mimicking normal Web traffic

Possible Root Causes

- A targeted attack may use hidden tunnels to hide communication with command and control servers
- A user is utilizing tunneling software to communicate with Internet services which might not otherwise be accessible
- Intentionally installed software is using a hidden tunnel to bypass expected firewall rules

Business Impact

- The use of a hidden tunnel by some software may be benign, but it represents significant risk as the intention is to bypass security controls
- Hidden tunnels used as part of a targeted attack are meant to slip by your perimeter security controls and indicate a sophisticated attacker
- Hidden tunnels are rarely used by botnets, though more sophisticated bot herders with more ambitious goals may utilize them

Steps to Verify

- Check to see if the destination IP or domain of the tunnel is an entity you trust for your network
- Ask the user of the host whether they are using hidden tunnel software for any purpose
- Before removing the offending software via antivirus or reimaging, take a memory snapshot for future analysis of the incident
- If the behavior reappears shortly after a reimaging, this may be a hardware/BIOS tunnel

Hidden HTTPS Tunnel

Command & Control



MITRE | ATT&CK®

[T1005 Data from Local System](#)

[T1115 Clipboard Data](#)

[T1071 Application Layer Protocol](#)

[T1185 Man in the Browser](#)

[T1125 Video Capture](#)

[T1113 Screen Capture](#)

[T1010 Application Window Discovery](#)

[T1037 Boot or Logon Initialization Scripts](#)

[T1111 Two-Factor Authentication Interception](#)

[T1572 Protocol Tunneling](#)

[T1573 Encrypted Channel](#)

[T1204 User Execution](#)

[T1056 Input Capture](#)

[T1001 Data Obfuscation](#)

[T1571 Non-Standard Port](#)

[T1059 Command and Scripting Interpreter](#)

[T1518 Software Discovery](#)

[T1176 Browser Extensions](#)

[T1123 Audio Capture](#)

[T1008 Fallback Channels](#)

[T1132 Data Encoding](#)

Triggers

- An internal host is communicating with an outside IP using HTTPS where another protocol is running over the top of the HTTPS sessions
- This represents a hidden tunnel involving one long session or multiple shorter sessions over a longer period of time mimicking normal encrypted Web traffic
- When it can be determined whether the tunneling software is console-based or driven via a graphical user interface, that indicator will be included in the detection

Possible Root Causes

- A targeted attack may use hidden tunnels to hide communication with command and control servers over SSL on port 443
- A user is utilizing tunneling software to communicate with Internet services which might not otherwise be accessible
- Intentionally installed software is using a hidden tunnel to bypass expected firewall rules

Business Impact

- The use of a hidden tunnel by some software may be benign, but it represents significant risk as the intention is to bypass security controls
- Hidden tunnels used as part of a targeted attack are meant to slip by your perimeter security controls and indicate a sophisticated attacker
- Hidden tunnels are rarely used by botnets, though more sophisticated bot herders with more ambitious goals may utilize them

Steps to Verify

- Check to see if the destination IP or domain of the tunnel is an entity you trust for your network
- Ask the user of the host whether they are using hidden tunnel software for any purpose
- Before removing the offending software via antivirus or reimaging, take a memory snapshot for future analysis of the incident
- If the behavior reappears shortly after a reimaging, this may be a hardware/BIOS tunnel

Malware Update

Command & Control



MITRE | ATT&CK®

T1105 Ingress Tool Transfer

Triggers

- An internal host is downloading and installing software from the Internet
- The downloads are over HTTP, appear to be machine- driven, and follow a suspicious pattern of checking for availability of files before downloading them

Possible Root Causes

- The initial exploit on this host may be loading malware to continue the attack
- Malware installed on the host may be updating itself to enhance its functionality
- Malware installed on the host may be updating itself to a new version of its software

Business Impact

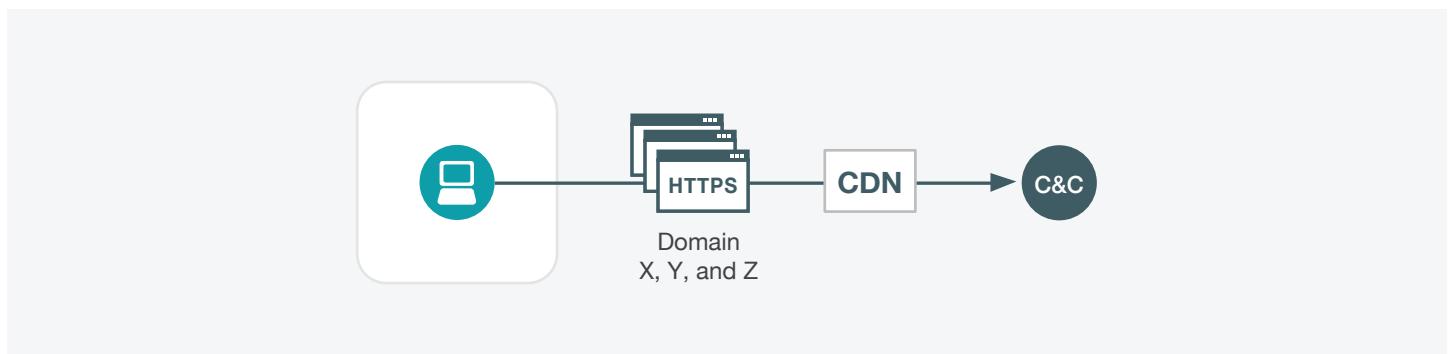
- An infected host can attack other organizations (e.g. spam, DoS, ad clicks) thus causing harm to your organization's reputation, potentially causing your IP addresses to be black listed and impacting the performance of business-critical applications
- If this is a targeted attack, it can spread further into your network and ultimately exfiltrate data from it
- The malware which infected the host can create nuisances and affect user productivity

Steps to Verify

- Look up the domain and IP address to which the communication is being sent via reputation services to see if this is known malware; such lookups are supported directly within the UI
- Search for the domain + "virus" via a search engine; this is effective for finding references to known adware or spyware
- Download the supplied PCAP and look at the HTTP payload being sent to see if any data is being leaked in clear text or whether the identity of the program is visible

Multi-home Fronted Tunnel

Command & Control



MITRE | ATT&CK®

[T1005 Data from Local System](#)

[T1115 Clipboard Data](#)

[T1071 Application Layer Protocol](#)

[T1125 Video Capture](#)

[T1113 Screen Capture](#)

[T1010 Application Window Discovery](#)

[T1037 Boot or Logon Initialization Scripts](#)

[T1111 Two-Factor Authentication Interception](#)

[T1572 Protocol Tunneling](#)

[T1573 Encrypted Channel](#)

[T1204 User Execution](#)

[T1056 Input Capture](#)

[T1001 Data Obfuscation](#)

[T1571 Non-Standard Port](#)

[T1059 Command and Scripting Interpreter](#)

[T1518 Software Discovery](#)

[T1176 Browser Extensions](#)

[T1123 Audio Capture](#)

[T1008 Fallback Channels](#)

[T1132 Data Encoding](#)

Triggers

- An internal host is communicating with an outside IP using HTTPS where another protocol is running over the top of the HTTPS sessions. The sessions appear to go to different domains but are all served by a single Content Delivery Network (CDN) and all utilize a JA3 hash which is only used by this host with this one CDN.
- This represents a hidden tunnel involving multiple shorter sessions over a longer period of time mimicking normal encrypted Web traffic

Possible Root Causes

- A targeted attack may use hidden tunnels to hide communication with command and control servers over TLS on port 443 and other ports
- Intentionally installed software is using a domain-fronted hidden tunnel utilizing multiple benign domains to bypass expected firewall rules

Business Impact

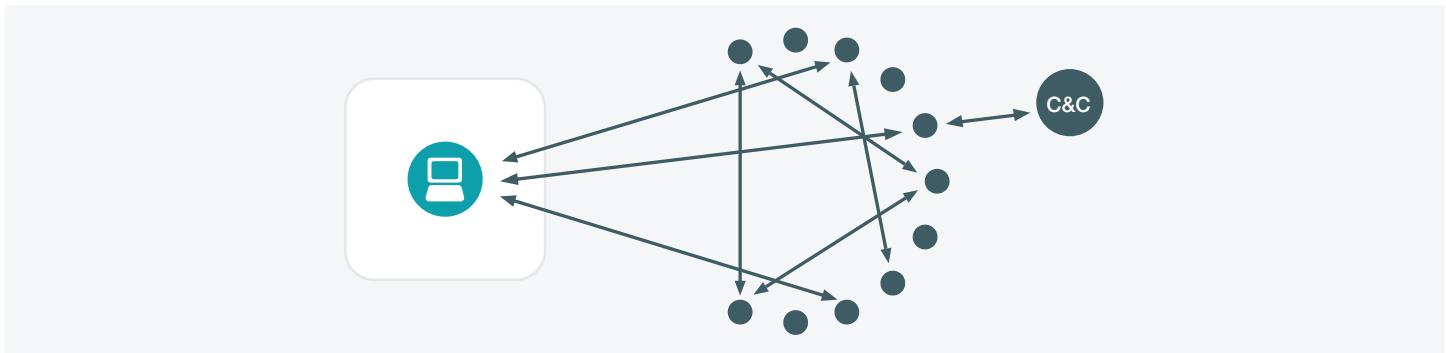
- The use of a hidden tunnel with multi-domain fronting is quite unusual, and it represents significant risk as the intention is to bypass security controls
- Hidden tunnels used as part of a targeted attack are meant to slip by your perimeter security controls and indicate a sophisticated attacker

Steps to Verify

- Ask the user of the host whether they are using hidden tunnel software for any purpose and if not, whether they intentionally connected to the list of domains in the detection (the JA3 hash in the detection may provide a clue to the software utilized)
- Before removing the offending software via antivirus or reimaging, take a memory snapshot for future analysis of the incident
- If the behavior reappears shortly after a reimaging, this may be a hardware/BIOS tunnel

Peer-To-Peer

Command & Control



MITRE | ATT&CK®

T1090 Proxy

Triggers

- An internal host is communicating with a set of external IP addresses with a pattern and low data rate common to peer-to-peer command and control

Possible Root Causes

- The internal host is infected with malware which is using peer-to-peer communication for its command and control; some botnets utilize this form of command and control as it is more resilient to attempts at disrupting or sink holing it
- Legitimate peer-to-peer software is running idle in the background without any data (e.g. BitTorrent) or voice (e.g. Skype) transfer activity and as such exhibits patterns similar to command and control traffic

Business Impact

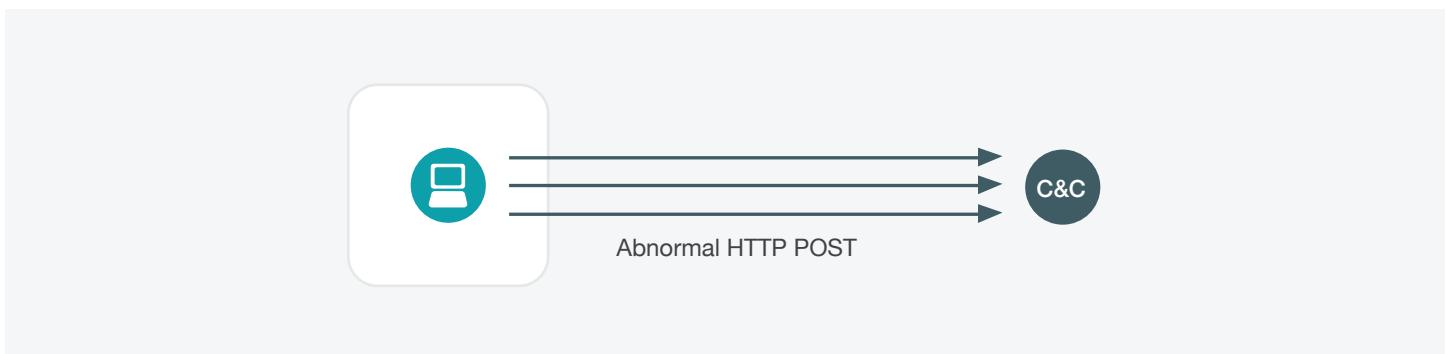
- An infected host can attack other organizations (e.g. spam, DoS, ad clicks) thus causing harm to your organization's reputation, potentially causing your IP addresses to be black listed and impacting the performance of business-critical applications
- The host can also be instructed to spread further into your network and ultimately exfiltrate data from it
- Software which infected the host can create nuisances and affect user productivity

Steps to Verify

- If the detection is generated as a result of a purposely installed peer-to-peer application, make sure the software complies with IT security policy
- If the detection cannot be attributed to such an application, the host is likely infected with a malware and should be fixed through the use of AV software or reimaged

Stealth HTTP Post

Command & Control



MITRE | ATT&CK®

T1071 Application Layer
Protocol

Triggers

- An internal host is sending data to an external system in multiple HTTP Post requests without being referred and without software identification
- These posts appear to be machine generated since they occur with a regular timing pattern

Possible Root Causes

- Adware, spyware or malware installed on an internal host is communicating back to its command and control server
- The communication may include some data leakage from the local host, which is particularly common with spyware

Business Impact

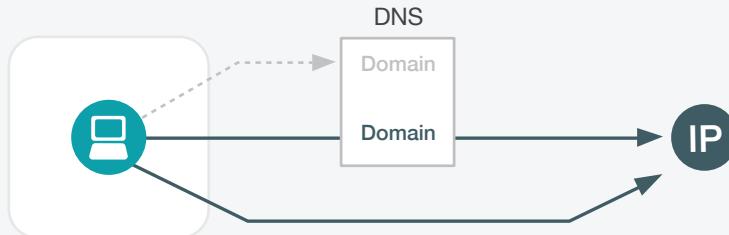
- An infected host can attack other organizations (e.g. spam, DoS, ad clicks) thus causing harm to your organization's reputation, potentially causing your IP addresses to be black listed and impacting the performance of business-critical applications
- The host can also be instructed to spread further into your network and ultimately exfiltrate data from it
- Software which infected the host can create nuisances and affect user productivity

Steps to Verify

- Look up the domain and IP address to which the communication is being sent via VirusTotal or other reputation services to see if this is known malware; such lookups are supported directly within the UI
- Search for the domain + "virus" via a search engine – this is effective for finding references to known adware or spyware
- Download the supplied PCAP and look at the HTTP payload being sent to see if any data is being leaked in clear text or whether the identity of the program is visible in the payload

Suspect Domain Activity

Command & Control



MITRE | ATT&CK®

T1568 Dynamic Resolution

Triggers

- An internal host is looking up suspicious external domains
- Suspicious activity may involve looking up machine-generated domain names or non-existent domain names in rapid succession

Possible Root Causes

- An infected host which is part of a botnet is using a domain generation algorithm (DGA) to locate its command & control servers
- An infected host or adware installed by the user is accessing newly generated domains to present ad impressions to the user
- An internal user visits newly registered domains with unusual names (e.g., letter sequences not normally found in domains)

Business Impact

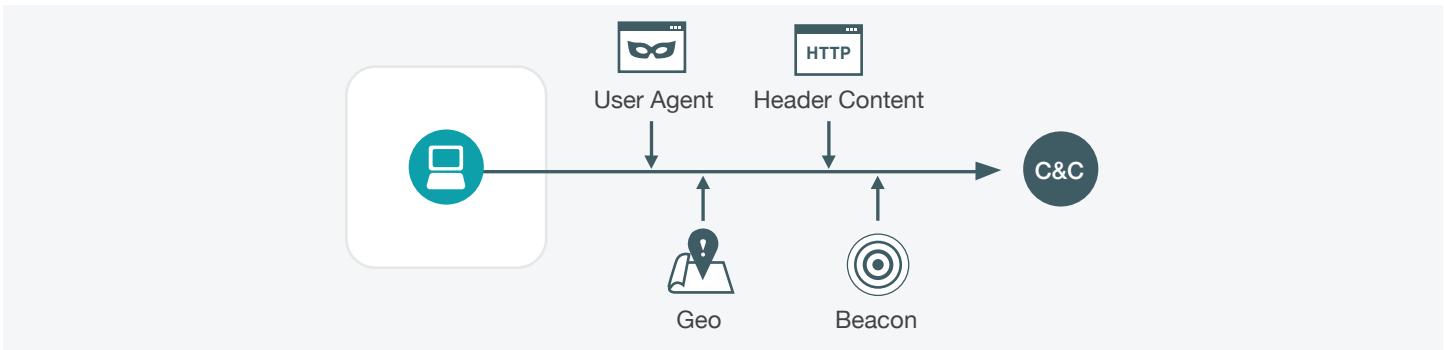
- An infected host can attack other organizations (e.g. spam, DoS, ad clicks) thus causing harm to your organization's reputation, potentially causing your IP addresses to be black listed and impacting the performance of business-critical applications
- The host can also be instructed to spread further into your network and ultimately exfiltrate data from it
- Software which infected the host can create nuisances and affect user productivity

Steps to Verify

- Do not go directly to the listed domain as it is likely to be malicious
- Look up the domain and IP address to which the communication is being sent via reputation services to see if this is known malware; such lookups are supported directly from the UI
- Inquire whether the user of the host would likely have gone to the listed domain
- Check to see if the host is also exhibiting other detected behaviors to understand the intent of the malware

Suspicious HTTP

Command & Control



MITRE | ATT&CK®

T1071 Application Layer Protocol

Triggers

- Software on an internal host is initiating one or more suspicious HTTP requests which form a pattern typically observed in command and control communications in recent malware samples
- The suspicious pattern may be the result of any combination of the following: (a) incorrect or malformed User-Agent, (b) absence or presence and order of a variety of HTTP headers, (c) presence and regularity of beaconing of the request and (d) connections to geographies which have a higher likelihood of hosting command and control servers

Possible Root Causes

- Malware installed on the host may be communicating back to its command and control server(s)
- Adware or spyware installed on the host may be communicating to its command and control server(s) or may be leaking data acquired on the host
- Software installed on the host is emitting HTTP requests that share two or more patterns with recent known malware samples: (a) malformed User-Agent, (b) unusual collection of HTTP headers, (c) communicating in an automated pattern and (d) communicating to out-of-the-ordinary geographies

Business Impact

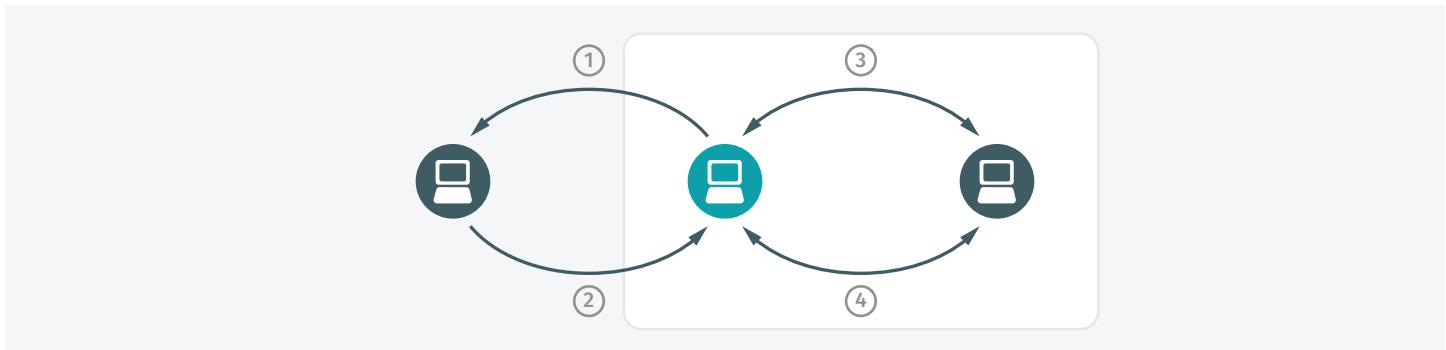
- An infected host can attack other organizations (e.g. spam, DoS, ad clicks) thus causing harm to your organization's reputation, potentially causing your IP addresses to be black listed and impacting the performance of business-critical applications
- The host can also be instructed to spread further into your network and ultimately exfiltrate data from it
- Software which infected the host can create nuisances and affect user productivity

Steps to Verify

- Look up the domain and IP address to which the communication is being sent via reputation services to see if this is known malware; such lookups are supported directly within the UI
- Search for the domain + "virus" via a search engine; this is effective for finding references to known adware or spyware
- Download the supplied PCAP and look at the HTTP payload being sent to see if any data is being leaked in clear text or whether the identity of the program is visible
- If there is no known reason why the user of the system would communicate to the geography in question, ask the end-user for a possible reason for the communication

Suspicious Relay

Command & Control



MITRE | ATT&CK®

[T1090 Proxy](#)

[T1104 Multi-Stage Channels](#)

Triggers

- This host appears to be acting as a relay for communication between an external system to another internal host—relays of this type involve a first (external) leg and a second (internal) leg
- This host also has another active command and control detection

Possible Root Causes

- A host is compromised and is being used to relay information to and from a host deeper inside the network
- An internal host is hosting some form of approved proxy (e.g. SOCKS) to allow other internal hosts to communicate with the Internet through it

Business Impact

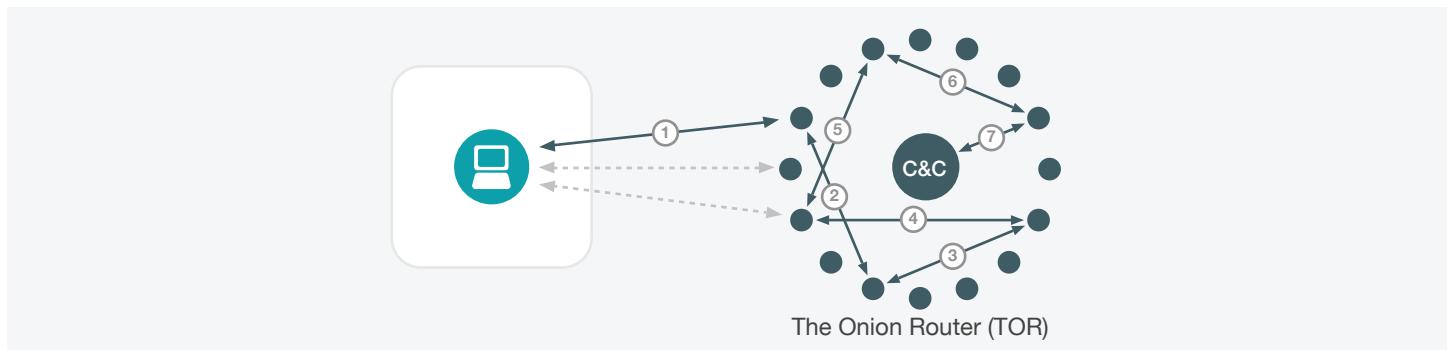
- An infected host which is enabling another internal host to hide its communication with the Internet by acting as a relay represents a high risk as this may allow a host which normally is not allowed to communicate with the outside to do so
- For hosts that have approved proxy software installed, ensure all the necessary security controls are in place to prevent unauthorized use

Steps to Verify

- Determine whether this host should be providing relay services to other internal hosts; if not, this is likely malicious behavior
- Look at the outside destination of the traffic and the payload of traffic, available in the PCAP, to determine what it is being sent and where it is going; this will help further calibrate the risk

TOR Activity

Command & Control



MITRE | ATT&CK®

T1090 Proxy

Triggers

- An internal host establishes connections with outside servers where protocol usage approximates communicating via The Onion Router (TOR)
- The algorithm inspects the protocol handshake of each session and triggers if characteristics of the session setup are similar to those observed in TOR connections

Possible Root Causes

- A targeted attack is utilizing TOR to hide communications with command and control servers or to exfiltrate your organization's data
- An infected host which is part of a botnet is utilizing TOR to communicate with its command and control servers or to leak small amounts of stolen data
- A user is utilizing a TOR-enabled program to anonymously communicate with servers available on the Internet or ones available only through TOR

Business Impact

- The use of TOR as part of a targeted attack is meant to slip by most standard perimeter defenses and indicates attacker sophistication
- The use of TOR as part of a botnet is relatively rare and would indicate a more sophisticated botnet
- The intentional use of TOR by employees may be allowed, but it does represent significant risk as the intention of TOR is to mask traffic source and destination

Steps to Verify

- Ask the user of the host whether they are using TOR for any purpose
- Check to see if any TOR-enabled software is installed on the host
- Check the TOR entry nodes listed in the detection against lists of known TOR entry nodes (e.g., search for "tor entry node list"), but note that these lists are seldom complete and shift over time

Threat Intelligence Match

Command & Control



Triggers

- An internal host is connecting to an external system and the connection has met criteria specified in one or more configured threat feeds

Possible Root Causes

- A host includes malware which is initiating the connection that triggered the detection
- A user on the host manually initiated the connection which triggered the detection

Business Impact

- Presence of command & control is a property of most attacks that originate from the outside
- The threat intel feed may have included additional context tied to the specific criteria that the connection met
- Business risk associated with outside control of an internal host is very high

Steps to Verify

- Refer to the information accompanying your threat feed as it may include verification and remediation instructions
- Determine which process on the internal host is sending the traffic which was flagged; in Windows systems, this can be done using a combination of netstat and tasklist commands
- Check if a user has knowingly installed remote access software and decide whether the resulting risk is acceptable
- Scan the computer for known malware and potentially reimage it, noting that some infections leave no trace on disk and reside entirely in memory

Vectra Threat Intelligence Match

Command & Control



MITRE | ATT&CK®

[T1008 Fallback Channels](#)

[T1041 Exfiltration Over C2 Channel](#)

[T1048 Exfiltration Over Alternative Protocol](#)

[T1059 Command and Scripting Interpreter](#)

[T1071 Application Layer Protocol](#)

[T1095 Non-Application Layer Protocol](#)

[T1105 Ingress Tool Transfer](#)

[T1132 Data Encoding](#)

[T1189 Drive-by Compromise](#)

[T1219 Remote Access Software](#)

[T1571 Non-Standard Port](#)

[T1573 Encrypted Channel](#)

Triggers

- After configuration of a [3rd party intelligence thread feed](#), an internal host has been observed either generating DNS activity or making direct connections associated with malicious external IPs or Domains identified by Vectra Threat Intelligence.

Possible Root Causes

- A host is communicating with a confirmed malicious IP or Domain that may be associated with staged malware, command and control, or client-side attacks.
- A user has been redirected to a site associated with phishing or credential compromise.
- A host is communicating with a benign service co-hosted on an IP or Domain with a poor or malicious reputation.

Business Impact

- Compromised assets or user credentials provide adversaries with the internal foothold necessary to begin to stage an attack.
- The identification of internal connections to known bad IP addresses or domains demonstrates positive risk to organizational assets and users and may indicate active attack progression.

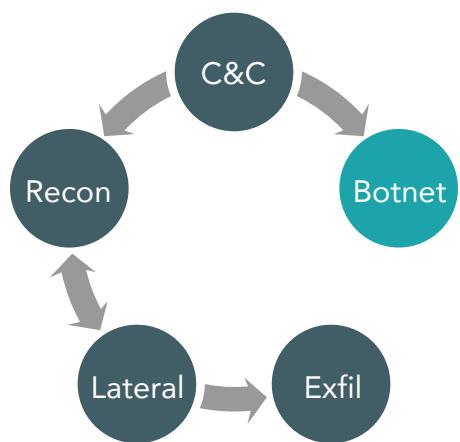
Steps to Verify

- Investigate the host and accounts associated for further indications of compromise.
- Using appropriate operational security and safeguards, verify the risk posed by this known bad IP or Domain by consulting external third party sources.
- Verify if supplemental preventative security controls protected the asset from full communication.
- In the case of phishing, verify with the user if credentials may have been compromised or take appropriate risk-based containment activities to include session revocation and password resets.
- Verify host integrity, the presence of new, unauthorized, or malicious software, and take appropriate incident handling or response activities.

Category

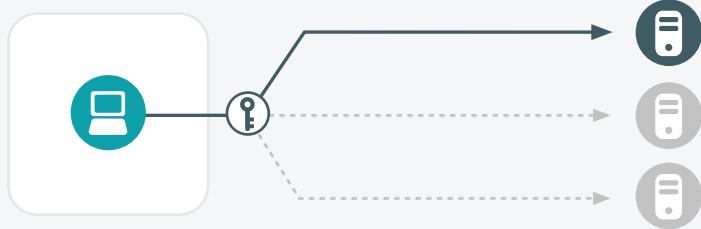
Botnet Activity

- A host is making money for its bot herder
- The ways in which an infected host can be used to produce value can range from mining bitcoins to sending spam emails to producing fake ad clicks
- The bot herder is utilizing the host computer, its network connection and, most of all, the unsullied reputation of the assigned IP to turn a profit



Brute-Force

Botnet Activity



MITRE | ATT&CK®

T1110 Brute Force

Triggers

- An internal host is making an unusually high number of login attempts, a behavior which is consistent with a brute-force password-guessing attack on one or more external servers
- Such attacks can be performed via a number of different protocols

Possible Root Causes

- An infected host is trying to guess passwords on one or more external systems; this is common botnet behavior where the host is instructed to breach internet-accessible systems that can be used as way points for command and control and data leakage
- A misconfigured internal host is constantly trying to connect to one or more external systems

Business Impact

- Botnet activity presents several risks to the organization: (1) it creates noise which may hide more serious issues; (2) there is a chance your organization's IP will end up on black lists; and (3) the compromised host can always be instructed to perform a direct attack on the organization
- Even if triggered due to a misconfiguration, the identified behavior is creating significant noise that may mask more serious issues and should be cleaned up

Steps to Verify

- If the internal host should not even be connecting to the external servers, this is likely malicious behavior
- Determine which process on the internal host is sending traffic to the external IP address(es) and ports; in Windows systems, this can be done using a combination of netstat and tasklist commands
- Verify that the process on the infected host should even be running and whether the process is configured correctly

Cryptocurrency Mining

Botnet Activity



MITRE | ATT&CK®

T1496 Resource Hijacking

Triggers

- An internal host is mining units of cryptocurrency of which Bitcoin, Litecoin, Ethereum, and Monero are some of the most common variants
- Cryptocurrency mining is a common way for botnet operators to make money
- Cryptocurrency mining may involve communication via HTTP or via the Stratum mining protocol

Possible Root Causes

- An infected host is mining cryptocurrency for its bot herder
- Some cryptocurrency mining can occur in the user's browser as a side effect of visiting compromised or low-reputation websites
- The user of the host on which the behavior has been detected has installed cryptocurrency mining software and is making money using your organization's systems, power, and network resources

Business Impact

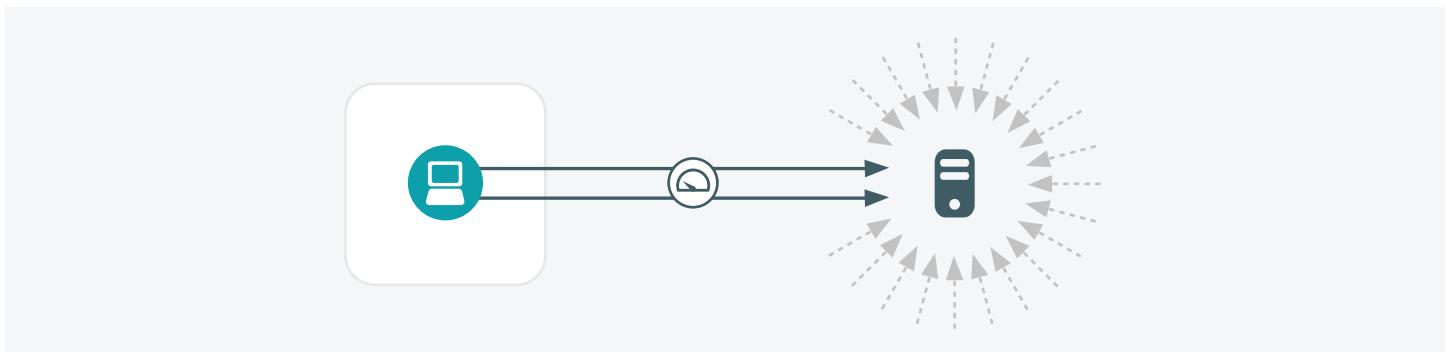
- Botnet activity presents several risks to the organization: (1) it creates noise which may hide more serious issues; (2) there is a chance your organization's IP will end up on black lists; and (3) the compromised host can always be instructed to perform a direct attack on the organization
- If the user of the host intentionally installed cryptocurrency mining software, the risk may be minimal, though such a user may also be prone to installing other "money making" software which may not prove to be as benign

Steps to Verify

- If the user intentionally installed cryptocurrency mining software, decide whether it should be removed
- If the user did not install cryptocurrency mining software, the host is likely infected and part of a botnet that performs "silent mining"
- Use anti-virus software or reimagine the host to remove the malware

Outbound DoS

Botnet Activity



MITRE | ATT&CK®

T1498 Network Denial of Service

Triggers

- An internal host appears to be taking part in a Denial- of-Service (DoS) campaign on an external IP address
- The form of DoS detection has two types: “SYN Flood” and “Slowloris”

Possible Root Causes

- The internal host is infected and has become part of a botnet and is being instructed by its bot herder to perform a DoS attack on an external system, which is a relatively common way for a botnet to make money
- An internal host is misconfigured and continually, in high volume, tries to connect to an external IP address

Business Impact

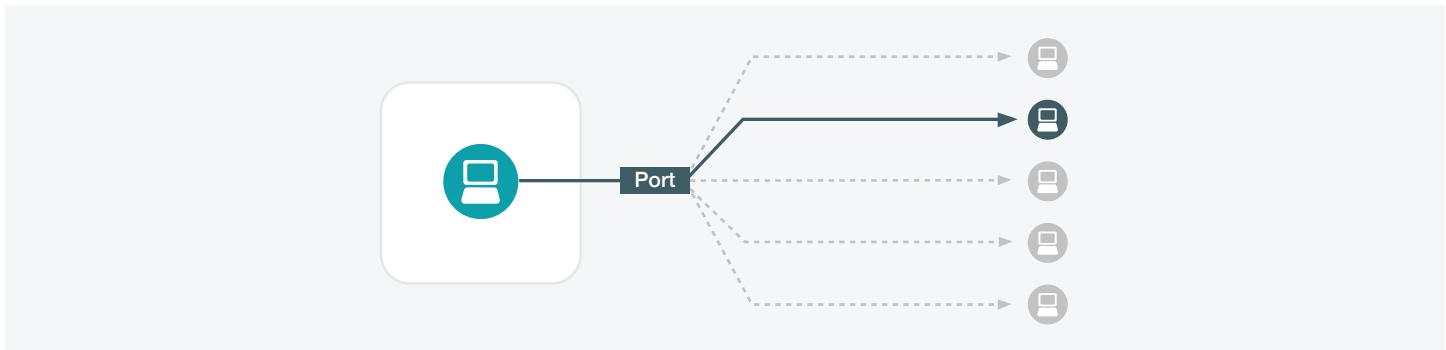
- Botnet activity presents several risks to the organization: (1) it creates noise which may hide more serious issues; (2) there is a chance your organization’s IP will end up on black lists; and (3) the compromised host can always be instructed to perform a direct attack on the organization
- The sheer volume of flood attacks may materially affect the amount of bandwidth available for legitimate functions which need to access the Internet

Steps to Verify

- Explore if there is a legitimate reason for the host to be connecting to the suspected victim of the attack
- Contact the user of the host to see whether they are trying to perform some unusual task which might trigger the DoS detection
- Check the host for presence of malware that is participating in a DoS attack

Outbound Port Sweep

Botnet Activity



MITRE | ATT&CK®

T1018 Remote System
Discovery

Triggers

- An internal host is generating many more unsuccessful attempts to connect to external services than successful ones

Possible Root Causes

- An internal host is part of a botnet and is being used by its bot herder to find other external services that could subsequently be attacked
- An internal host is misconfigured and is making many connection attempts to different IP addresses on the Internet

Business Impact

- Botnet activity presents several risks to the organization: (1) it creates noise which may hide more serious issues; (2) there is a chance your organization's IP will end up on black lists; and (3) the compromised host can always be instructed to perform a direct attack on the organization
- A misconfigured internal host may be using unnecessary bandwidth and slowing down both the host itself and other applications as a result of the traffic it is sending

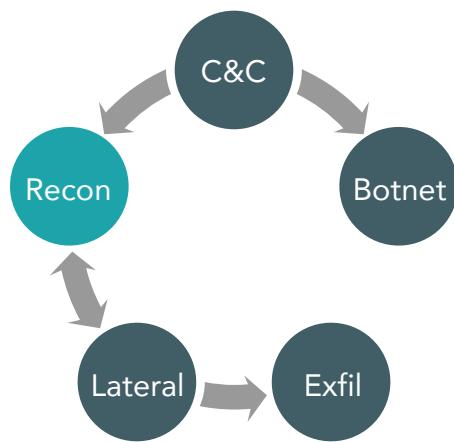
Steps to Verify

- Look at the pattern of IP addresses being scanned to determine the intent of the scan
- Verify whether there is misconfigured software on the host which is causing the scan
- If the behavior cannot be explained by user action or known software behavior, the host is likely infected and should be remediated

Category

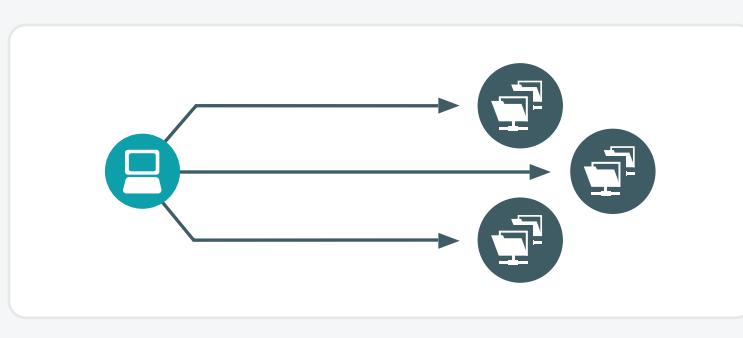
Reconnaissance

- A host or account is mapping out the inside of your network or cloud environment
- The activity may indicate that this is a targeted attack
- Detection types cover fast scans and slow scans
 - your vulnerability scanner will show up here as it performs much the same activity as an attacker



File Share Enumeration

Reconnaissance



MITRE | ATT&CK®

[T1039 Data from Network Shared Drive](#)

[T1119 Automated Collection](#)

[T1135 Network Share Discovery](#)

Triggers

- A host accesses a number of file shares significantly in excess of the number of file shares normally accessed in the network

Possible Root Causes

- An attacker is looking for data to exfiltrate or is looking for files which provide additional information necessary for achieving the goals of the attack
- The host is accessing a large number of file shares as an end user attempts to find a particular file or directory

Business Impact

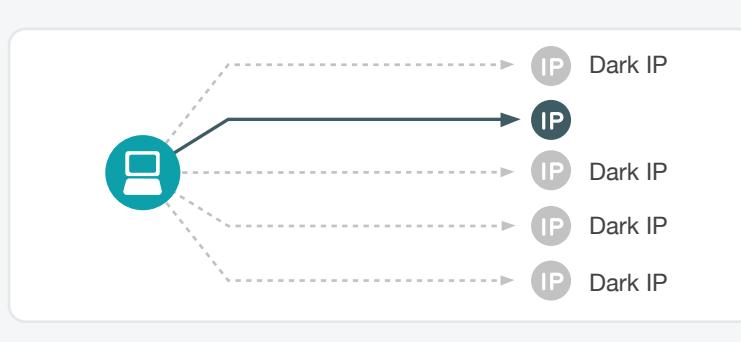
- An enumeration of the available file shares in a network is an effective way for an attacker to find data to exfiltrate or data that helps further the attack
- Reconnaissance within a network is a precursor to active attacks which ultimately exposes an organization to substantial risk of data acquisition and exfiltration
- This form of reconnaissance is often a lot less noticeable than a port sweep or a port scan so attackers feel they can use it with relatively little risk of detection

Steps to Verify

- Ask the user of the host whether they have any knowledge of accessing the listed file shares
- Check the file server logs to see what files were accessed on the shares
- If the file share access continues and remains unexplained, determine which process on the internal host is accessing the file shares; in Windows systems, this can be done using a combination of netstat and tasklist commands

Internal Darknet Scan

Reconnaissance



MITRE | ATT&CK®

[T1082 System Information Discovery](#)

[T1018 Remote System Discovery](#)

[T1072 Software Deployment Tools](#)

[T1046 Network Service Scanning](#)

[T1016 System Network Configuration Discovery](#)

Triggers

- An internal host has contacted a number of internal IPs that have not been active in the recent past
- Darknet detections cover longer periods than port scans and ignore contact to systems which do not respond to this host, but which are otherwise active

Possible Root Causes

- An infected internal system that is part of targeted attack is performing slow reconnaissance of your network by reaching out to different IP addresses in your network
- A vulnerability scanner or asset discovery system is mapping systems in your network
- A host has been moved to a new network and is unsuccessfully attempting to connect to many previously available services

Business Impact

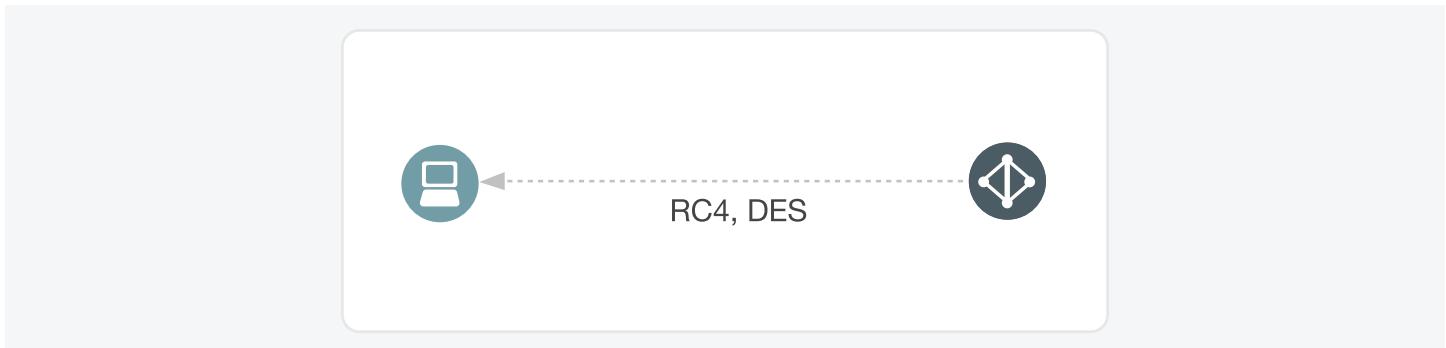
- Slow reconnaissance of your systems may represent the beginning of a targeted attack in your network
- Authorized reconnaissance by vulnerability scanners and asset discovery systems should be limited to a small number of hosts which can be whitelisted for this behavior

Steps to Verify

- Check to see if the detected host should be authorized for network scans
- Look at the pattern of IP addresses being scanned to determine the intent of the scan
- If the pattern appears random and distributed over time, determine which software on the host could be causing the connection requests

Kerberoasting: Weak Cipher Request

Reconnaissance



MITRE | ATT&CK®

T1558.003 Kerberoasting

Triggers

- A host that does not typically work with weak encryption types receives a service ticket that was signed using a weak cipher.

Possible Root Causes

- Malicious Detection: An attacker is requesting service tickets with weak encryption so that they may attempt to learn the service account's password.
- Benign Detection: Legacy systems may still require the use of weak encryption ciphers simply because they do not support newer, more secure ciphers.

Business Impact

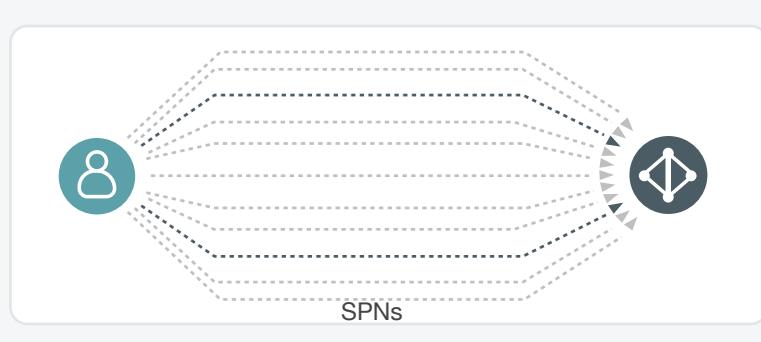
- Specific Risk: Kerberoasting may result in the discovery of a privileged account's password.
- Impact: Depending on the level of privilege a cracked account has (e.g. service account with domain admin), this could lead directly to a full domain compromise.

Steps to Verify

- Investigate the host, user, and service accounts involved when weak ciphers are returned to a host that doesn't typically request them.
- Conventionally, service accounts with a sufficiently complex password (cryptographically random, minimum 25 characters, rotates often) can be ignored, since these take long enough to crack that the cracked password has likely expired by the time its discovered.

Kerberoasting: SPN Sweep

Reconnaissance



MITRE | ATT&CK®

T1558.003 Kerberoasting

Triggers

- A host is observed requesting service tickets for a high volume of SPNs.

Possible Root Causes

- Malicious Detection: An attacker is performing recon in a domain to find favorable targets for offline password cracking.
- Benign Detection: Enterprise vulnerability scanners may also submit requests for a large volume of SPNs.

Business Impact

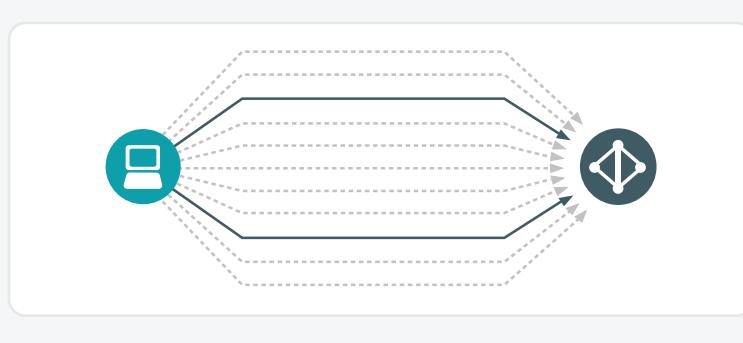
- Specific Risk: Kerberoasting may result in the discovery of a privileged account's password.
- Impact: Depending on the level of privilege a cracked account has (e.g. service account with domain admin), this could lead directly to a full domain compromise.

Steps to Verify

- Investigate the host making requests for high volume of SPNs, this behavior is not typical for general users and should only be conducted by authorized hosts.

Kerberos Account Scan

Reconnaissance



MITRE | ATT&CK®

T1087 Account Discovery

Triggers

- A Kerberos client attempts a suspicious amount of authentication requests using a large number of user accounts with many of them failing as a result of non-existent accounts

Possible Root Causes

- The internal Kerberos client is part of targeted attack which aims to spread horizontally within the network by first discovering the existence of user accounts and then stealing the account's credentials or Kerberos tickets
- A client is initiating a large number of authentication attempts with many of them failing

Business Impact

- An account scan to a Kerberos or Active Directory server is an effective way for an attacker to determine what accounts are available inside an organization's network
- Reconnaissance within a network is a precursor to active attacks which ultimately exposes an organization to substantial risk of data acquisition and exfiltration
- This form of reconnaissance is often a lot less noticeable than a port sweep or a port scan so attackers feel they can use it with relatively little risk of detection

Steps to Verify

- Examine the Kerberos or Active Directory server logs for a more detailed view of activity by this host
- Inquire whether the host should be utilizing the user accounts listed in the detection
- Verify that the host on which authentication is attempted is not a shared resource as this could generate a sufficient variety of authentications to resemble an account scan

Kerberos Brute-Sweep

Reconnaissance



MITRE | ATT&CK®

T1110 Brute Force

Triggers

- A host attempts a suspicious amount of authentication requests using a large number of user accounts with some of them failing because the accounts don't exist and others failing because the password is incorrect

Possible Root Causes

- The host is part of targeted attack which aims to spread horizontally within the network by first discovering the existence of user accounts and simultaneously attempting to login to them using credentials from a common set of passwords
- The host may be a portal (a shared resource) and the authentication requests are being performed on behalf of other systems inside or outside the organization

Business Impact

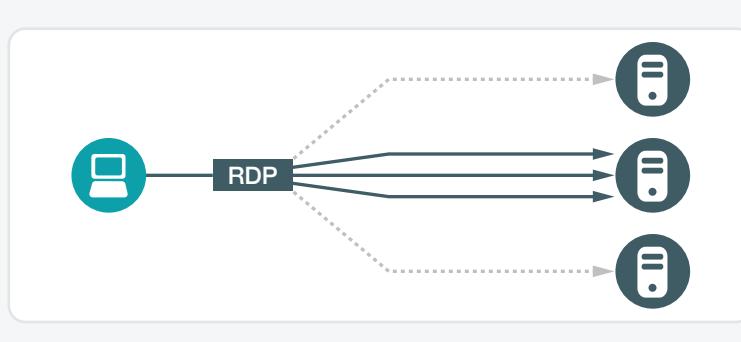
- An account brute sweep to a Kerberos or AD server is an effective way for an attacker to determine what accounts are available inside an organization's network and to simultaneously try to guess the accounts' passwords
- Reconnaissance within a network is a precursor to active attacks which ultimately exposes an organization to substantial risk of data acquisition and exfiltration
- This form of reconnaissance is often a lot less noticeable than a port sweep, a port scan, or even the widespread use of RPCs to many hosts, so attackers feel they can use it with relatively little risk of detection

Steps to Verify

- Examine the Kerberos or Active Directory server logs for a more detailed view of activity by this host
- Inquire whether the host should be utilizing the user accounts listed in the detection
- Verify that the host on which authentication is attempted is not a shared resource as this could generate a sufficient variety of authentications to resemble an account brute sweep

RDP Recon

Reconnaissance



MITRE | ATT&CK®

T1033 System Owner/User Discovery

T1018 Remote System Discovery

Triggers

- A host is making multiple RDP connection attempts with most of the connections failing to complete
- The connection attempts can target one or more RDP servers
- Even when a single RDP server is targeted, multiple accounts may still be involved in the encrypted part of the RDP connection setup

Possible Root Causes

- An attacker is trying to determine the existence of accounts in order to progress to the next step in the attack
- The attacker is working through a list of accounts with well-known default passwords in an attempt to find a working account/password combination
- This host is a jump server and several users are unsuccessfully attempting to RDP to other servers from it

Business Impact

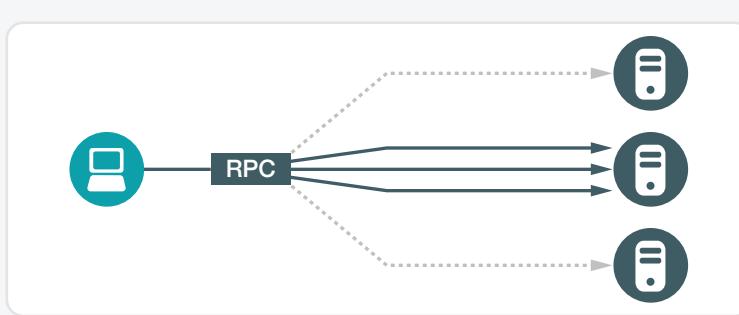
- A scan via RDP is an effective way for an attacker to determine what accounts are available inside an organization's network and which RDP servers accept logins via the accounts
- If one of the targets has not been normally accessed via RDP, the nature of the target server will provide additional guidance regarding the potential business impact
- Reconnaissance within a network is a precursor to active attacks which ultimately exposes an organization to substantial risk of data acquisition and exfiltration
- This form of reconnaissance is often a lot less noticeable than a port sweep or a port scan so attackers feel they can use it with relatively little risk of detection

Steps to Verify

- Inquire whether the target of the RDP connection attempts should even be setup to accept RDP connections
- Inquire whether this host should be initiating the number of RDP connections to the targets listed in the detection
- If this host is a jump server, retrieve the logs of the jump server to see what upstream connections are the originators of the large number of failed RDP connections

RPC Recon

Reconnaissance



MITRE | ATT&CK®

[T1082 System Information Discovery](#)

[T1201 Password Policy Discovery](#)

[T1087 Account Discovery](#)

[T1124 System Time Discovery](#)

[T1049 System Network Connections Discovery](#)

[T1007 System Service Discovery](#)

[T1057 Process Discovery](#)

[T1069 Permission Groups Discovery](#)

[T1033 System Owner/User Discovery](#)

[T1135 Network Share Discovery](#)

Triggers

- This host is making RPC calls to a large number of other hosts
- The number of hosts being contacted far exceeds the number of hosts normally contacted as observed on this network

Possible Root Causes

- An attacker is active inside the network and is mining information from individual hosts in order to build a better map of assets in the network
- The information mined can include what accounts have recently logged into which hosts and can be used in deciding where to steal privileged account credentials
- An admin is completing authorized system management activity
- Endpoint management software installed on a central server is performing periodic system management activity
- Specialized hardware, including IoT, is utilizing RPC for peer discovery and identification

Business Impact

- A scan of neighboring hosts' information is an effective way for an attacker to complete a detailed map of what happens where inside the target organization's network
- Reconnaissance within a network is a precursor to active attacks which ultimately exposes an organization to substantial risk of data acquisition and exfiltration
- This form of reconnaissance is often a lot less noticeable than a port sweep or a port scan so attackers feel they can use it with relatively little risk of detection

Steps to Verify

- Examine the local logs on the host making the RPC queries for a more detailed view of activity by this host
- Inquire whether the host should be contacting the hosts listed in the detection
- If the behavior continues and remains unexplained, determine which process on the internal host is establishing the connections over which the RPC requests are made; in Windows systems, this can be done using a combination of netstat and tasklist commands

RPC Targeted Recon

Reconnaissance



MITRE | ATT&CK®

[T1007 System Service Discovery](#)

[T1082 System Information Discovery](#)

[T1124 System Time Discovery](#)

[T1077 Windows Admin Shares](#)

[T1049 System Network Connections Discovery](#)

[T1057 Process Discovery](#)

[T1069 Permission Groups Discovery](#)

[T1087 Account Discovery](#)

[T1135 Network Share Discovery](#)

[T1201 Password Policy Discovery](#)

Triggers

- This host is making one or more RPC function calls indicative of information gathering to one or more other hosts
- The RPC function calls related to information gathering being made differ from ones normally made by this host or received by the target host.

Possible Root Causes

- An attacker is active inside the network and is mining information from individual hosts in order to better understand the usefulness of the target host to furthering the attack
- The information mined may include recently logged on accounts, running services, available network shares, or password hashes
- An admin is completing authorized system management activity
- Endpoint management software installed on a central server is performing periodic system management activity
- Specialized hardware, including IoT, is utilizing RPC for peer discovery and identification

Business Impact

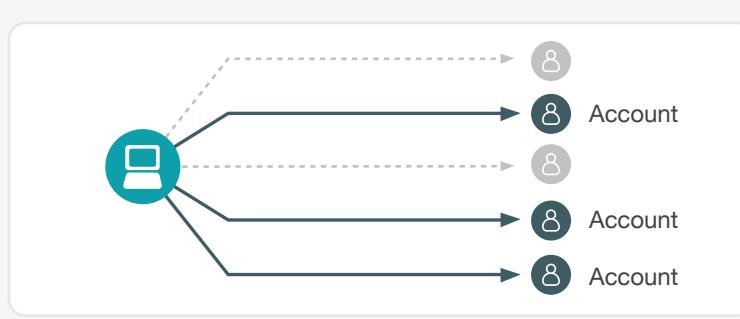
- Retrieval of a key host's information is an effective way for an attacker to further a "low-and-slow" attack on an organization's network
- Reconnaissance within a network is a precursor to active attacks which ultimately exposes an organization to substantial risk of data acquisition and exfiltration
- This form of reconnaissance is often a lot less noticeable than a port sweep, a port scan, or even the widespread use of RPCs to many hosts, so attackers feel they can use it with relatively little risk of detection

Steps to Verify

- Examine the local logs on the host making the RPC queries for a more detailed view of activity by this host
- Inquire whether the host should be contacting the hosts listed in the detection
- If the behavior continues and remains unexplained, determine which process on the internal host is establishing the connections over which the RPC requests are made; in Windows systems, this can be done using a combination of netstat and tasklist commands

SMB Account Scan

Reconnaissance



MITRE | ATT&CK®

T1087 Account Discovery

Triggers

- A host rapidly makes use of multiple accounts via the SMB protocol which can be used for file sharing, RPC and other activity

Possible Root Causes

- An attacker is trying to determine the existence of accounts in order to progress to the next step in the attack
- The attacker is working through a list of accounts with well-known default passwords in an attempt to find a working account/password combination
- This host provides services through a portal and many users are using the portal by logging in and requesting services which require an SMB connection to fulfill

Business Impact

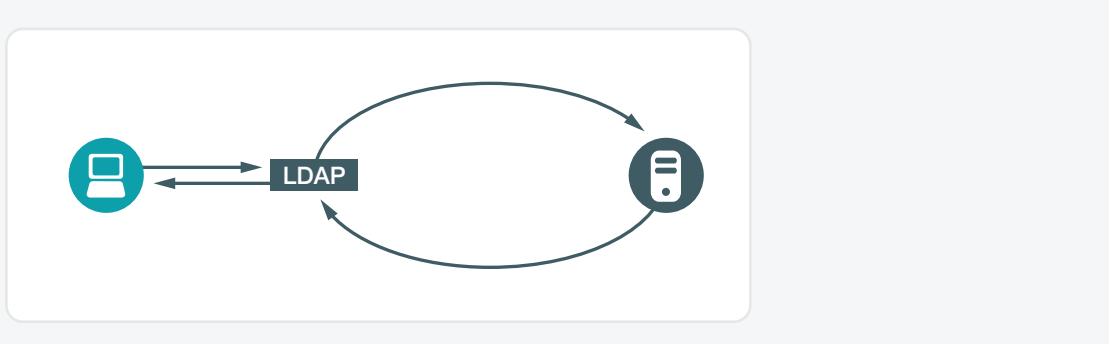
- An account scan is an effective way for an attacker to determine what accounts are available inside an organization's network
- Reconnaissance within a network is a precursor to active attacks which ultimately exposes an organization to substantial risk of data acquisition and exfiltration
- This form of reconnaissance is often a lot less noticeable than a port sweep or a port scan so attackers feel they can use it with relatively little risk of detection

Steps to Verify

- If logs of user session activity are available, examine the logs for a more detailed view of activity by this host
- Inquire whether the host should be utilizing the user accounts listed in the detection
- Verify that the host from which authentication is attempted is not a shared resource as this could generate a sufficient variety of account usage to resemble an account scan

Suspicious LDAP Query

Reconnaissance



MITRE | ATT&CK®

[T1087 Account Discovery](#)

[T1018 Remote System Discovery](#)

[T1482 Domain Trust Discovery](#)

Triggers

- This host is querying Active Directory using the LDAP protocol in a manner that appears like reconnaissance behavior
- The LDAP queries are either unusually broad in scope or are specifically targeting accounts and groups that have names which imply administrative privilege

Possible Root Causes

- An attacker is active inside the network and is mining information from one or more Active Directory servers in order to build a better map of assets in the network
- An admin is retrieving information from AD in order to complete a certain task or create a report
- An auditing application installed on this host is retrieving information from AD as part of its core functionality

Business Impact

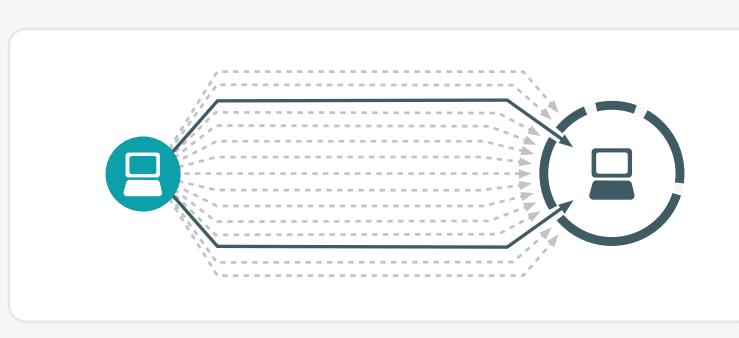
- A scan of information in an Active Directory server is an effective way for an attacker to determine what accounts are privileged inside an organization's network and what the names of servers and infrastructure components are
- Reconnaissance within a network is a precursor to active attacks which ultimately exposes an organization to substantial risk of data acquisition and exfiltration
- This form of reconnaissance is often a lot less noticeable than a port sweep or a port scan so attackers feel they can use it with relatively little risk of detection

Steps to Verify

- Examine the Kerberos or Active Directory server logs for a more detailed view of activity by this host
- Inquire whether the host should be making the queries listed in the detection
- If the LDAP queries continue and remain unexplained, determine which process on the internal host is making the queries; in Windows systems, this can be done using a combination of netstat and tasklist commands

Suspicious Port Scan

Reconnaissance



MITRE | ATT&CK®

[T1082 System Information Discovery](#)

[T1018 Remote System Discovery](#)

[T1072 Third Party Software](#)

[T1046 Network Service Scanning](#)

[T1016 System Network Configuration Discovery](#)

Triggers

- An internal host has attempted contact with many ports on a small number of internal IP addresses

Possible Root Causes

- An infected internal system that is part of a targeted attack is trying to locate any services which may be active on a small number of hosts by attempting connections on different ports on one or more IP addresses
- An IT-run vulnerability scanner or asset discovery system is mapping out system services on a host
- The detected host is communicating with another host using a peer-to-peer protocol and the traffic configuration on the switch is only supplying one direction of the traffic to the Vectra sensor

Business Impact

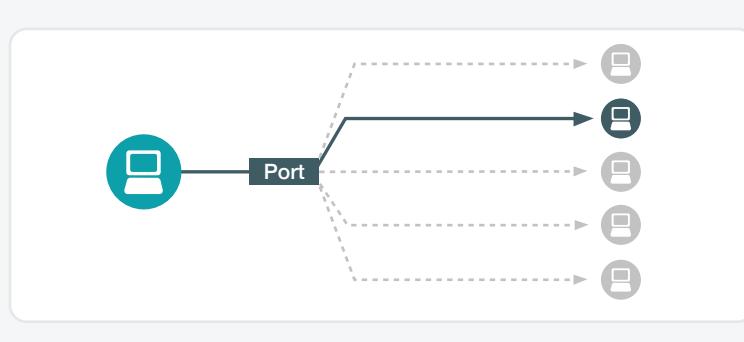
- Reconnaissance of individual systems may represent the beginning of a targeted attack in your network
- If the system being scanned is an important or critical asset, any unauthorized scan should be treated with utmost suspicion
- Authorized reconnaissance by vulnerability scanners and asset discovery systems should be limited to a small number of hosts which can be whitelisted for this behavior using triage filters

Steps to Verify

- Check to see if the detected host is authorized to perform port scans on the target hosts
- Look at the pattern of ports being scanned to try to determine what the detected host may be searching for
- If the pattern appears random and distributed over time, it is likely some form of reconnaissance and should be dealt with before the attack progresses further

Suspicious Port Sweep

Reconnaissance



MITRE | ATT&CK®

[T1082 System Information Discovery](#)

[T1018 Remote System Discovery](#)

[T1072 Third Party Software](#)

[T1046 Network Service Scanning](#)

[T1016 System Network Configuration Discovery](#)

Triggers

- An internal host has attempted contact with a large number of internal IP addresses on a small number of ports

Possible Root Causes

- An infected internal system that is part of a targeted attack is contacting a large number of internal IP addresses on a small number of ports to find systems which are running particular software that may be vulnerable to an attack
- An IT-run vulnerability scanner or asset discovery system is mapping out system services in your network
- A host with an unusual discovery mechanism is looking for a service on its local subnet
- Alarm equipment or IP cameras are performing large-scale scans due to misconfiguration or firmware bugs

Business Impact

- Reconnaissance of your systems may represent the beginning of a targeted attack in your network
- Authorized reconnaissance by vulnerability scanners and asset discovery systems should be limited to a small number of hosts which can be whitelisted for this behavior using triage filters

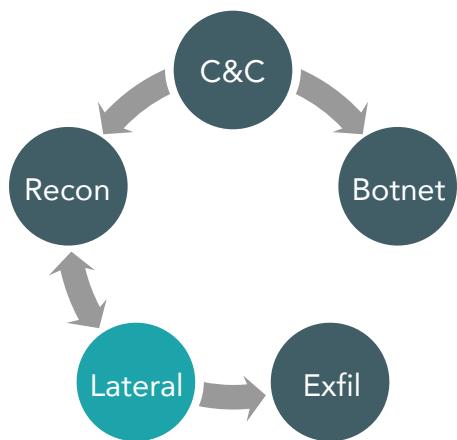
Steps to Verify

- Check to see if the detected host is authorized to perform port sweeps
- Look at the pattern of ports being scanned to determine the intent of the scan
- If the pattern appears random and distributed over time, it is likely some form of reconnaissance and should be dealt with before the attack progresses further

Category

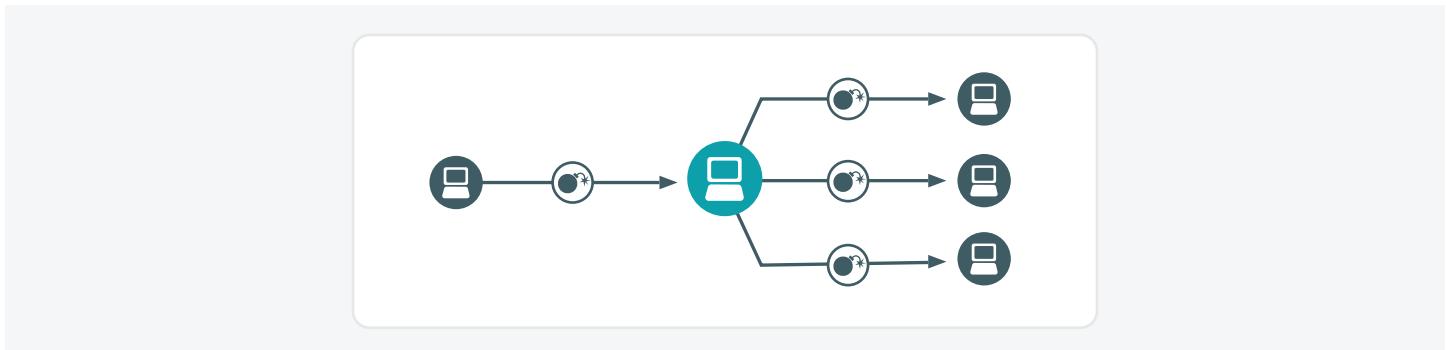
Lateral Movement

- Covers scenarios of lateral action meant to further a targeted attack
- This can involve attempts to steal account credentials or to steal data from another resource
- It can also involve compromising another host or account to make the attacker's foothold more durable or to get closer to target data



Automated Replication

Lateral Movement



MITRE | ATT&CK®

T1072 Software Deployment Tools

T1210 Exploitation of Remote Services

Triggers

- An internal host is sending very similar payloads to several internal targets
- This may be the result of an infected host sending one or more exploits to other hosts in an attempt to infect them

Possible Root Causes

- An infected host which is part of a botnet is trying to expand the botnet's footprint by infecting additional hosts
- An infected host which is taking part in a targeted attack is trying to spread laterally in an effort to get closer to data it wants to exfiltrate
- An agent on the host is utilizing unusual techniques to discover an available service

Business Impact

- Internal spreading of botnet-related malware often is repeated by the next infected host, thus mimicking a computer worm and rapidly infecting all possible hosts
- A wide scale spread of botnet-related malware will incur significant remediation costs
- Lateral spread which is part of a targeted attack makes the attack more resilient and gets it closer to your crown jewels

Steps to Verify

- Look at the protocol and port listed in the detection to determine what network service is being exploited
- Determine if there's any reason for this host to be communicating these services on the listed targets
- Try to ascertain what software on this host would emit the traffic being seen
- Examine the packet capture file to see if this appears to be a network discovery attempt

Brute-Force

Lateral Movement



MITRE | ATT&CK®

T1110 Brute Force

Triggers

- An internal host is making many login attempts on an internal system, behavior which is consistent with a brute-force password attack
- Such attacks can be performed via different protocols (e.g. RDP, VNC, SSH) and may also be a Heartbleed attack (e.g. memory scraping)

Possible Root Causes

- An infected host or a malicious insider in control of the host is trying to guess passwords on another internal system
- A misconfigured host is constantly trying to connect to one or more other internal systems

Business Impact

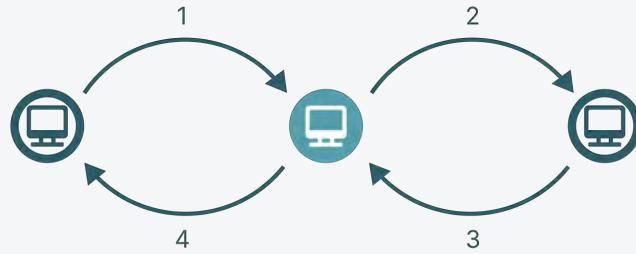
- Successful harvesting of account credentials (usernames and password) of other accounts, particularly more privileged accounts, is a classic progression of a targeted attack
- Even if triggered due to a misconfiguration, the identified misconfiguration is creating significant stress on the target system and should be cleaned up

Steps to Verify

- Determine whether the internal host in question should be connecting to the target host; if not, this is likely malicious behavior
- Determine which process on the internal host is sending traffic to the internal IP address(es) and ports; in Windows systems, this can be done using a combination of netstat and tasklist commands
- Verify that the process should be running on the infected host and whether the process is configured correctly

Credential Access via NTLM Relay

Lateral Movement



MITRE | ATT&CK®

T1557 - Adversary-in-the-Middle: LLMNR/NBT-NS Poisoning and SMB Relay

Triggers

- An internal host was observed sending NTLM authentication packets received from another internal source and 'relaying' them to the intended destination.
- The relay host may be attempting to intercept credential data to use them for nefarious purposes.

Possible Root Causes

- An infected host is being used to intercept and relay NTLM authentication traffic.

Business Impact

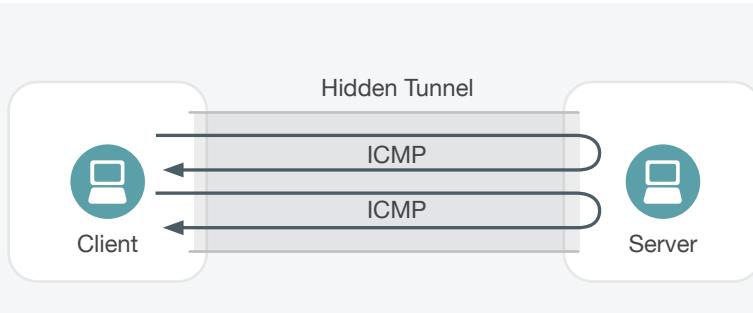
- Compromised credentials can be used by an attacker to conduct further unauthorized access
- Stolen credentials can also be used to access sensitive information and assets
- Attackers can further their foothold into victim networks using accounts that have high privileges

Steps to Verify

- Identify whether the relay host intercepting NTLM has been compromised
- Determine whether the relay host is operating in an authorized way
- Reset any account passwords that may have been compromised once the relay host has been remediated

ICMP Tunnel: Client

Lateral Movement



MITRE | ATT&CK®

[T1008 Fallback Channels](#)

[T1095 Non-Application Layer Protocol](#)

Triggers

- A host was observed using ICMP in ways inconsistent with standard implementation of the protocol.
- More precisely, a host's ICMP traffic was observed to contain datagrams which vary in size more frequently than typical ICMP traffic would.
- An attacker may be using this host to communicate with or transfer data to an internal host.

Possible Root Causes

Malicious Detection

- An attacker is using ICMP as a staging and/or control channel.
- An attacker has established persistence & has chosen ICMP as a backup channel.

Benign Detection

- A network device like a vulnerability scanner is crafting nonstandard ICMP datagrams.

Business Impact

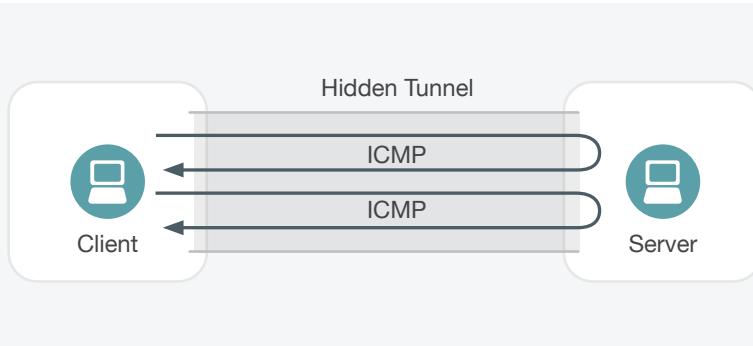
- The presence of an ICMP tunnel indicates the host was compromised & that an attacker has remote access to the machine.
- Recon, data exfiltration, lateral movement, privilege escalation, & establishing a tunnel over a more reliable protocol like HTTPS are all likely next steps.
- ICMP tunnels can be stealthy and are often used to evade sophisticated perimeter security controls.

Steps to Verify

- Check the destination IP & determine if the observed traffic arrives at a trusted endpoint.
- Investigate the host for malware, there may be code present which establishes a C2 channel with another host.

ICMP Tunnel: Server

Lateral Movement



MITRE | ATT&CK®

[T1008 Fallback Channels](#)

[T1095 Non-Application Layer Protocol](#)

Triggers

- A host was observed using ICMP in ways inconsistent with standard implementation of the protocol.
- More precisely, this host's ICMP traffic was observed to contain datagrams which vary in size more frequently than typical ICMP traffic would.
- An attacker may be using this host as a server to communicate with or transfer data to internal clients.

Possible Root Causes

Malicious Detection

- An attacker is using ICMP as a staging and/or control channel.
- An attacker has established persistence & has chosen ICMP as a backup channel.

Benign Detection

- A network device like a vulnerability scanner is crafting nonstandard ICMP datagrams.

Business Impact

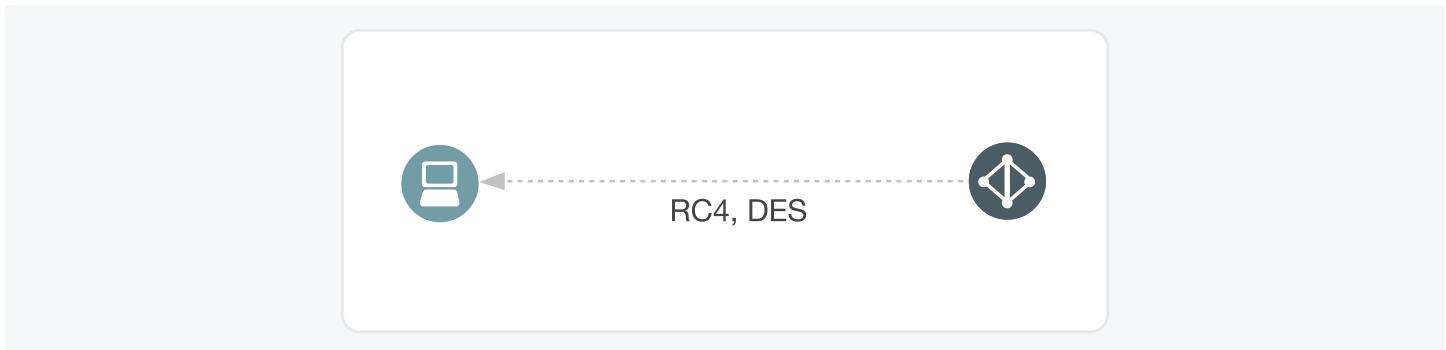
- The presence of an ICMP tunnel indicates the host was compromised & that an attacker has remote access to the machine.
- Recon, data exfiltration, lateral movement, privilege escalation, & establishing a tunnel over a more reliable protocol like HTTPS are all likely next steps.
- ICMP tunnels can be stealthy and are often used to evade sophisticated perimeter security controls.

Steps to Verify

- Check the destination IP & determine if the observed traffic arrives at a trusted endpoint.
- Investigate the host for malware, there may be code present which establishes a C2 channel with another host.

Kerberoasting: Targeted Weak Cipher Response

Lateral Movement



MITRE | ATT&CK®

T1558.003 Kerberoasting

Triggers

- A user has requested a Kerberos service ticket from the ticket granting service which has responded with a weak cipher. This service has either never responded with a weak cipher or the occurrence of this behavior is rare.
- This behavior could indicate the presence of an active attacker who is aiming to escalate privileges on the network by exploiting a weak cipher response and subsequently cracking the NTLM, AES, or RC4 hash. The objective is to obtain the password associated with the targeted service account.

Possible Root Causes

Malicious Detection

- A single ticket cipher downgrade attack may be leveraged in a targeted attack to escalate privileges to progress further into the network and ultimately achieve the goal set out by the malicious actor.
- Observing a weak cipher response from the request of a Ticket Granting Service (TGS) ticket where the service being requested has either never responded with a weak cipher or rarely responds with a weak cipher could be indicative of a malicious actor attempting to elevate privileges to progress further in carrying out their attack.

Benign Detection

- Legacy systems, which lack support for modern ciphers, may still rely on weak ciphers for their operations. In such cases, it becomes crucial to confirm all necessary patches are applied and to prioritize the security of accounts associated with these services. This involves enforcing strong and lengthy passwords for these accounts, while also adhering to regular password update schedules. Additionally, it is essential to assign minimal privileges to these accounts, ensuring they possess only the necessary permissions required for proper functioning within the organization.

Business Impact

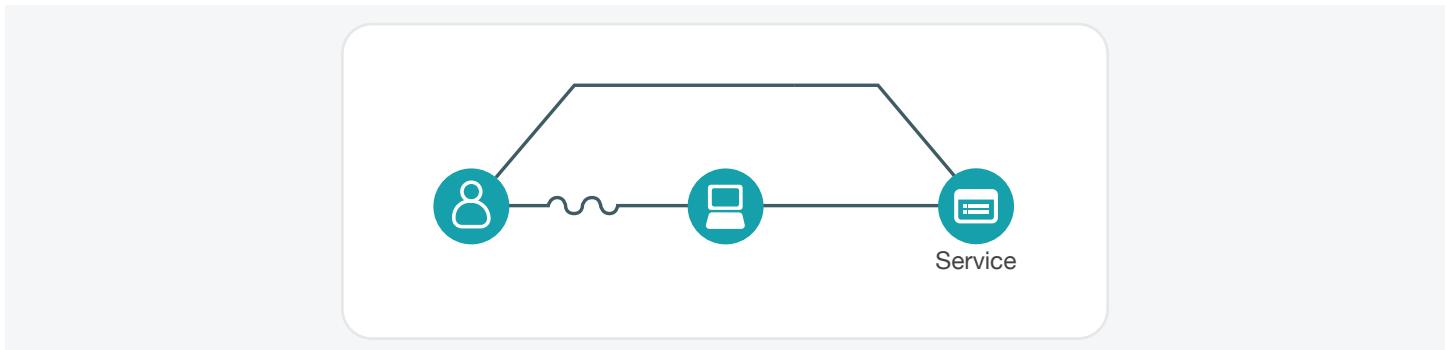
- Specific Risk: A cipher downgrade could lead to exposing credentials of a privileged account.
- Impact: Depending on the level of privilege a targeted account has (e.g. service account with domain admin), this could lead directly to a full domain compromise.

Steps to Verify

- Investigate the host, user, and service accounts involved when weak ciphers are returned to a host that doesn't typically request them.
- Conventionally, service accounts with a sufficiently complex password (cryptographically random, minimum 25 characters, rotates often) can be ignored, since these take long enough to crack that the cracked password has likely expired by the time its discovered.

Privilege Anomaly: Unusual Account on Host

Lateral Movement



MITRE | ATT&CK®

[T1078 Valid Accounts](#)

[T1098 Account](#)

[Manipulation](#)

[T1552 Unsecured Credentials](#)

[T1555 Credentials from Password Stores](#)

[T1040 Network Sniffing](#)

[T1033 System Owner/User Discovery](#)

[T1212 Exploitation for Credential Access](#)

[T1484 Group Policy Modification](#)

[T1556 Modify Authentication Process](#)

[T1558 Steal or Forge Kerberos Tickets](#)

[T1550 Use Alternate Authentication Material](#)

[T1539 Steal Web Session Cookie](#)

[T1003 OS Credential Dumping](#)

[T1136 Create Account](#)

Triggers

- A privileged account is used to access a privileged service, but is doing so from a host which the account has not been observed on but where the host (using other accounts) has been seen accessing the service

Possible Root Causes

- The privileged account has been compromised and is being used to access a privileged service normal for the account, but from a host that the account is typically not used from; additionally, the host used for the access is itself a normal place from which to connect to the privileged server, just not with this account
- A privileged employee has borrowed another privileged user's machine (either due to their primary laptop crashing or because they are away from their desk) to perform what is otherwise normal work for the account

Business Impact

- Lateral movement within a network involving privileged accounts, hosts, or services exposes an organization to substantial risk of data acquisition and exfiltration
- Unexplained unusual patterns of use of privileged accounts, hosts, and services are involved in almost all major breaches
- Attacks carried out by rogue insiders will often exhibit unusual patterns of use as well
- The accounts and hosts used and the services accessed provide a possible perspective on the potential business impact

Privilege Anomaly: Unusual Host

Lateral Movement



MITRE | ATT&CK®

[T1078 Valid Accounts](#)

[T1098 Account Manipulation](#)

[T1552 Unsecured Credentials](#)

[T1555 Credentials from Password Stores](#)

[T1040 Network Sniffing](#)

[T1033 System Owner/User Discovery](#)

[T1212 Exploitation for Credential Access](#)

[T1484 Group Policy Modification](#)

[T1556 Modify Authentication Process](#)

[T1558 Steal or Forge Kerberos Tickets](#)

[T1550 Use Alternate Authentication Material](#)

[T1539 Steal Web Session Cookie](#)

[T1003 OS Credential Dumping](#)

[T1136 Create Account](#)

Triggers

- An account is used to access a service from a host which the account is not usually on and from which the service is not usually accessed and at least the service (and likely the account) has a high privilege score OR the privilege score of the host is suspiciously low in comparison to the privilege levels of the account and service

Possible Root Causes

- The account is under the control of an attacker and is being used from an unusual host to connect to one or more services which are normal for the account but abnormal from the host
- An employee or contractor with approved access to the network who pretty consistently works from a particular set of hosts has been assigned a new host or has temporarily decided to work from another host

Business Impact

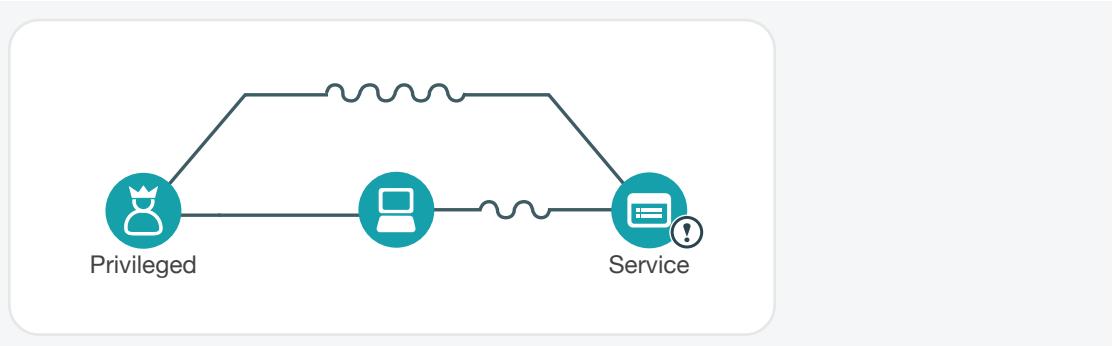
- Lateral movement within a network involving privileged accounts, hosts or services exposes an organization to substantial risk of data acquisition and exfiltration
- Unexplained unusual patterns of use of privileged accounts, hosts and services are involved in almost all major breaches
- Attacks carried out by rogue insiders will often exhibit unusual patterns of use as well
- The accounts and hosts used and the services accessed provide a possible perspective on the potential business impact

Steps to Verify

- Examine the Kerberos or Active Directory server logs for a more detailed view of activity by this account across all hosts
- Carefully inquire into whether the owner of the host in question should be using the specified accounts to access the listed services
- Verify that the host from which authentication is attempted is not a shared resource as this could mean that the attacker is using it as a pivot point

Privilege Anomaly: Unusual Service

Lateral Movement



MITRE | ATT&CK®

[T1078 Valid Accounts](#)

[T1098 Account Manipulation](#)

[T1552 Unsecured Credentials](#)

[T1555 Credentials from Password Stores](#)

[T1040 Network Sniffing](#)

[T1033 System Owner/User Discovery](#)

[T1212 Exploitation for Credential Access](#)

[T1484 Group Policy Modification](#)

[T1556 Modify Authentication Process](#)

[T1558 Steal or Forge Kerberos Tickets](#)

[T1550 Use Alternate Authentication Material](#)

[T1539 Steal Web Session Cookie](#)

[T1003 OS Credential Dumping](#)

[T1136 Create Account](#)

Triggers

- An account which is typically used from this host is accessing a service which the account has not been observed accessing from any host and at least two entities (account and service) have high privilege scores

Possible Root Causes

- The host is under the control of an attacker and the account on the host is being used to connect to one or more services which are abnormal for the account and may or may not be abnormal for the host
- An employee or contractor with approved access to the network has been assigned a new project or job which involve new privileged services which are quite abnormal given their prior role

Business Impact

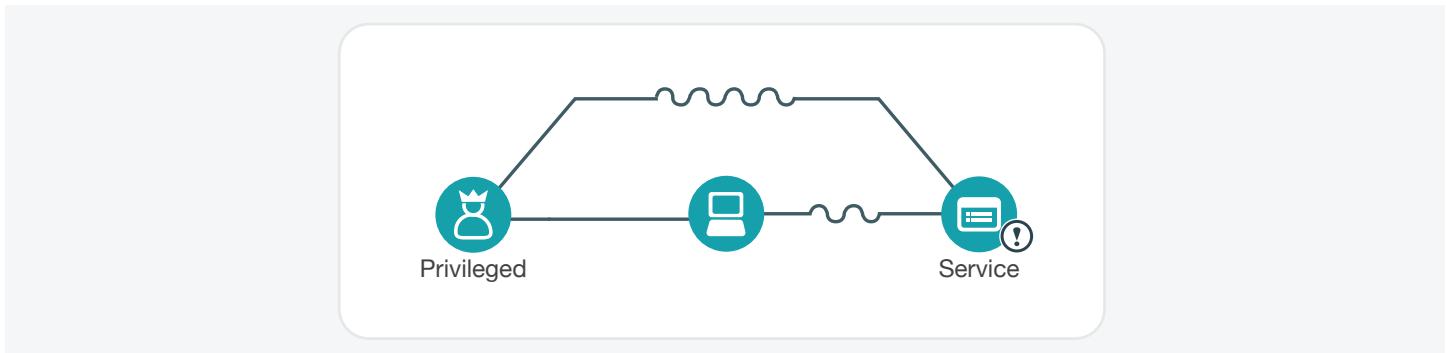
- Lateral movement within a network involving privileged accounts, hosts or services exposes an organization to substantial risk of data acquisition and exfiltration
- Unexplained unusual patterns of use of privileged accounts, hosts and services are involved in almost all major breaches
- Attacks carried out by rogue insiders will often exhibit unusual patterns of use as well
- The accounts and hosts used and the services accessed provide a possible perspective on the potential business impact

Steps to Verify

- Examine the Kerberos or Active Directory server logs for a more detailed view of activity by this host and account since if the host is compromised, the account must be considered to be compromised as well
- Carefully inquire into whether the owner of the host in question should be using the specified accounts to access the listed services
- Verify that the host from which authentication is attempted is not a shared resource as this could mean that the attacker is using it as a pivot point

Privilege Anomaly: Unusual Service - Insider

Lateral Movement



MITRE | ATT&CK®

[T1078 Valid Accounts](#)

[T1098 Account](#)

[Manipulation](#)

[T1552 Unsecured Credentials](#)

[T1555 Credentials from Password Stores](#)

[T1040 Network Sniffing](#)

[T1033 System Owner/User Discovery](#)

[T1212 Exploitation for Credential Access](#)

[T1484 Group Policy Modification](#)

[T1556 Modify Authentication Process](#)

[T1558 Steal or Forge Kerberos Tickets](#)

[T1550 Use Alternate Authentication Material](#)

[T1539 Steal Web Session Cookie](#)

[T1003 OS Credential Dumping](#)

[T1136 Create Account](#)

Triggers

- An account with a low privilege score is used from a host that has a low privilege score to access a service which has a substantially higher privilege score

Possible Root Causes

- The host is under the control of an attacker and the account on the host is being used to connect to one or more higher privileged services
- The account is under the control of an attacker and is being used from multiple hosts to connect to one or more higher privileged services
- A new admin has been hired and as the account used by the admin is new and the machine assigned to the admin is new, both have low privilege scores; when the admin then begins to perform legitimate work, detections are triggered until the privilege scores of the admin's account and host are raised based on observed activity
- A new service is being rolled out and it was initially only used by higher privileged admin accounts (and thus considered to be a high privilege service) but then released for use by a broader set of lower privileged accounts
- A rarely used service is generally accessed by higher privileged accounts, but is technically also available to lower privileged accounts is accessed by one such low privileged accounts

Business Impact

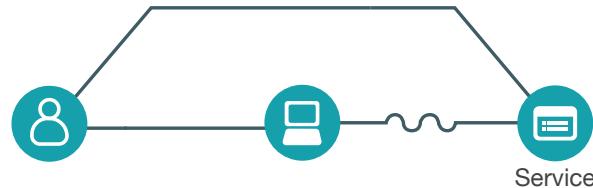
- Lateral movement within a network involving privileged accounts, hosts or services exposes an organization to substantial risk of data acquisition and exfiltration
- Unexplained unusual patterns of use of privileged accounts, hosts and services are involved in almost all major breaches
- Attacks carried out by rogue insiders will often exhibit unusual patterns of use as well
- The accounts and hosts used and the services accessed provide a possible perspective on the potential business impact

Steps to Verify

- Examine the Kerberos or Active Directory server logs for a more detailed view of activity by this host and account since if the host is compromised, the account must be considered to be compromised as well
- Carefully inquire into whether the owner of the host in question should be using the specified accounts to access the listed services
- Verify that the host from which authentication is attempted is not a shared resource as this could mean that the attacker is using it as a pivot point

Privilege Anomaly: Unusual Service from Host

Lateral Movement



MITRE | ATT&CK®

[T1078 Valid Accounts](#)

[T1098 Account](#)

[Manipulation](#)

[T1552 Unsecured Credentials](#)

[T1555 Credentials from Password Stores](#)

[T1040 Network Sniffing](#)

[T1033 System Owner/User Discovery](#)

[T1212 Exploitation for Credential Access](#)

[T1484 Group Policy Modification](#)

[T1556 Modify Authentication Process](#)

[T1558 Steal or Forge Kerberos Tickets](#)

[T1550 Use Alternate Authentication Material](#)

[T1539 Steal Web Session Cookie](#)

[T1003 OS Credential Dumping](#)

[T1136 Create Account](#)

Triggers

- A privileged account is used to access a privileged service, and is doing so from a host which the account has been observed on but where the host has not been seen accessing the service

Possible Root Causes

- The privileged account has been compromised and is being used to access a privileged service normal for the account, but from a host that the service is typically not accessed from; additionally, the host used for the access is itself a normal place for this account, but not a place from which this service is accessed by any account
- A privileged employee has decided to use their backup/secondary machine (either due to their primary laptop crashing or because they are away from their desk) to perform what is otherwise normal work for the account

Business Impact

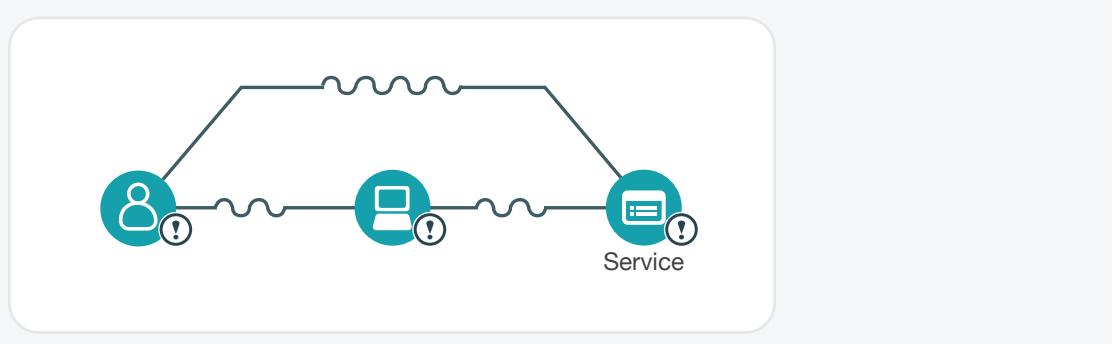
- Lateral movement within a network involving privileged accounts, hosts, or services exposes an organization to substantial risk of data acquisition and exfiltration
- Unexplained unusual patterns of use of privileged accounts, hosts, and services are involved in almost all major breaches
- Attacks carried out by rogue insiders will often exhibit unusual patterns of use as well
- The accounts and hosts used and the services accessed provide a possible perspective on the potential business impact

Steps to Verify

- Examine the Kerberos or Active Directory server logs for a more detailed view of activity by this account since if it has been compromised, all hosts the account has been on must be considered to be compromised as well
- Verify that the host in question is a secondary machine owned by the account owner
- Verify that the host from which authentication is attempted is not a shared resource as this could mean that the attacker is using it as a pivot point

Privilege Anomaly: Unusual Trio

Lateral Movement



MITRE | ATT&CK®

[T1078 Valid Accounts](#)

[T1098 Account Manipulation](#)

[T1552 Unsecured Credentials](#)

[T1555 Credentials from Password Stores](#)

[T1040 Network Sniffing](#)

[T1033 System Owner/User Discovery](#)

[T1212 Exploitation for Credential Access](#)

[T1484 Group Policy Modification](#)

[T1556 Modify Authentication Process](#)

[T1558 Steal or Forge Kerberos Tickets](#)

[T1550 Use Alternate Authentication Material](#)

[T1539 Steal Web Session Cookie](#)

[T1003 OS Credential Dumping](#)

[T1136 Create Account](#)

Triggers

- An account is used from a host to request access to a service where none of the pairings (account-host, account-service and host-service) are consistent with prior observed behavior and at least the service is considered privileged

Possible Root Causes

- The account or host (or both) are under the control of an attacker and are being used to in a manner which is abnormal for all three entities (account, host and service) involved
- An employee or contractor with approved access to the network is attacking the organization by using their account on an unusual host or someone else's account on their host to access a service which neither the account nor the host usually connects to

Business Impact

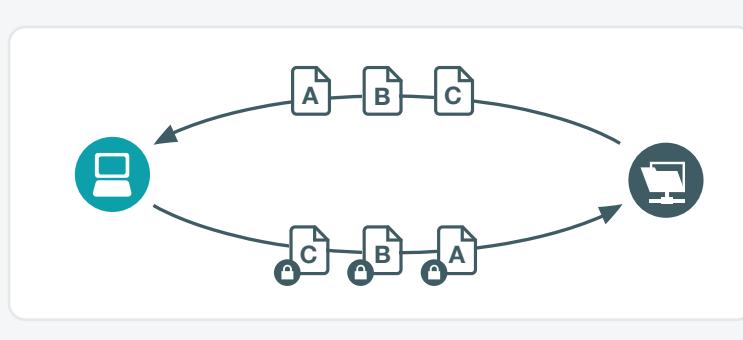
- Lateral movement within a network involving privileged accounts, hosts or services exposes an organization to substantial risk of data acquisition and exfiltration
- Unexplained unusual patterns of use of privileged accounts, hosts and services are involved in almost all major breaches
- Attacks carried out by rogue insiders will often exhibit unusual patterns of use as well
- The accounts and hosts used and the services accessed provide a possible perspective on the potential business impact

Steps to Verify

- Examine the Kerberos or Active Directory server logs for a more detailed view of activity by this account since, if it has been compromised, all hosts the account has been on must be considered to be compromised as well
- Carefully inquire into whether the owner of the host in question would expect the account listed in the detection to be used on this host
- Verify that the host from which authentication is attempted is not a shared resource as this could mean that the attacker is using it as a pivot point

Ransomware File Activity

Lateral Movement



MITRE | ATT&CK®

T1486 Data Encrypted for Impact

Triggers

- An internal host is connected to one or more file servers via the SMB protocol and is rapidly reading files and writing files of roughly the same size and with roughly the same file name
- This pattern is highly correlated with how ransomware interacts with file servers

Possible Root Causes

- The internal host is infected with a variant of ransomware
- A benign application on the host is rapidly reading files from and writing files to a networked file share
- A user is compiling a large set of source files located on a file share, causing a pattern of reading and writing files that exhibits a similar pattern

Business Impact

- Ransomware encrypts files and transmits the encryption key to the attacker
- The attacker then attempts to extract a ransom (typically payable in an untraceable cyber currency) from the organization in return for a promise to release the encryption key which allows the files to be recovered
- Even if your organization is willing to pay the ransom, there is no guarantee that the encryption key will be provided by the attacker
- Absent the encryption key, files will have to be restored from a backup and any changes since the last backup will be lost

Steps to Verify

- Examine the sample files referenced in the detection and see if the original files are missing and the files that have replaced them carry strange but similar file names or file extensions
- Check the directory in which the files reside for ransom notes with instructions on how to pay the ransom and retrieve the encryption key

SMB Brute-Force

Lateral Movement



MITRE | ATT&CK®

T1110 Brute Force

Triggers

- An internal host is utilizing the SMB protocol to make many login attempts using the same account(s), behavior which is consistent with a brute-force password attack
- Many, though not necessarily all, of these authentications are observed to fail

Possible Root Causes

- An infected host or a malicious insider in control of the host is trying to guess passwords for an account on another internal system
- A misconfigured host is constantly trying to connect to one or more other internal systems using an incorrect password or trying to log into an account which no longer exists or is locked out

Business Impact

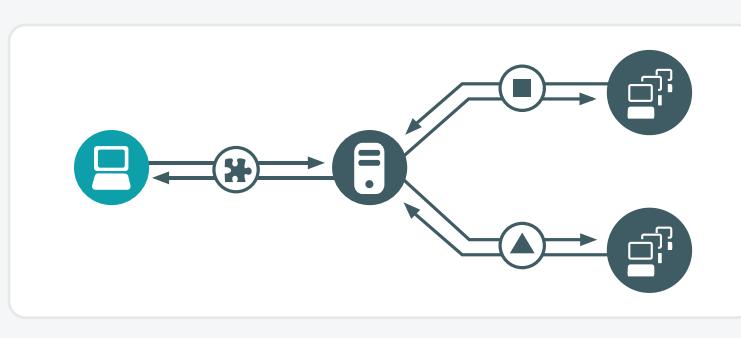
- Successful harvesting of account credentials (usernames and passwords) of other accounts, particularly more privileged accounts, is a classic progression of a targeted attack
- Even if triggered due to a misconfiguration, the identified behavior is creating significant stress on the target system and should be cleaned up

Steps to Verify

- Determine whether the internal host in question should be connecting to the target host using the indicated account(s); if not, this is likely malicious behavior
- Determine which process on the internal host is initiating the SMB requests; in Windows systems, this can be done using a combination of netstat and tasklist commands
- Verify that the process should be running on the internal host and whether the process is configured correctly

Shell Knocker Client

Lateral Movement



MITRE | ATT&CK®

T1205 Traffic Signaling

Triggers

- The host is communicating in an unusual manner with an internal server on a port that has previously shown a stable pattern for requests and responses
- The request sent to the internal server and the response received from it don't conform to any of the previously observed patterns

Possible Root Causes

- The server has been compromised and the port has been hijacked to enable communication to the compromised part of the system without requiring a new port to be utilized for the communication
- The client or the server has been recently upgraded and the pattern of use on the server port has changed
- This client has an unusual configuration in that it communicates with the port on the server in a manner unlike all the other observed communication on that port

Business Impact

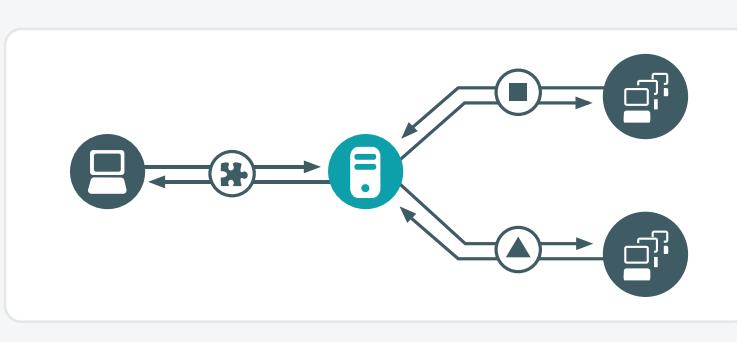
- Port hijacking is a technique attackers use to enable communication to a compromised server without raising alarms which may go off when a new port is used on an existing server
- Compromised servers are often more valuable than compromised laptops as they remain on the network at all times and are often located in the data center where most of an organization's important data resides

Steps to Verify

- See if the pattern of the flagged request and response represent acceptable deviations from the normal patterns or are significant departures such as binary data in an otherwise character-based protocol
- Inquire whether the software which emitted the request on this host has recently been updated as this may cause detections for a short period of time after the update
- Inquire whether the software on the server which responded to the request has recently been updated as this may cause detections for a short period of time after the update
- If the changed pattern remains unexplained, boot the client and server using a known good image on a USB device, then mount the local drive and scan it for signs of compromise

Shell Knocker Server

Lateral Movement



MITRE | ATT&CK®

T1205 Traffic Signaling

Triggers

- The server is communicating in an unusual manner with an internal client on a port that has previously shown a stable pattern for requests and responses
- The request received by the server and the response sent by it don't conform to any of the previously observed patterns

Possible Root Causes

- The server has been compromised and the port has been hijacked to enable communication to the compromised part of the system without requiring a new port to be utilized for the communication
- The client or the server has been recently upgraded and the pattern of use on the server port has changed
- The client which triggered the detection has an unusual configuration in that it communicates with the port on this server in a manner unlike all the other observed communication on the port

Business Impact

- Port hijacking is a technique attackers use to enable communication to a compromised server without raising alarms which may go off when a new port is used on an existing server
- Compromised servers are often more valuable than compromised laptops as they remain on the network at all times and are often located in the data center where most of an organization's important data resides

Steps to Verify

- See if the pattern of the flagged request and response represent acceptable deviations from the normal patterns or are significant departures such as binary data in an otherwise character-based protocol
- Inquire whether the software which emitted the request on the client has recently been updated as this may cause detections for a short period of time after the update
- Inquire whether the software on this server which responded to the request has recently been updated as this may cause detections for a short period of time after the update
- This type of backdoor is most likely to be in a kernel module, so produce a list of all installed kernel modules and verify against list of good known kernel modules
- If the changed pattern remains unexplained, boot the client and server using a known good image on a USB device, then mount the local drive and scan it for signs of compromise

SQL Injection Activity

Lateral Movement



MITRE | ATT&CK®

T1190 Exploit Public Facing Application

Triggers

- An internal host sends requests to a Web server and embeds SQL fragments into HTTP Post data or the URL to gain access to the backend database; the requests appear machine-generated due to the large volume and rate of arrival

Possible Root Causes

- An infected system that is part of targeted attack is looking for vulnerabilities in an internal Web app through which to access the database integrated into it, or is harvesting information for later exfiltration
- An IT-operated vulnerability scanner is scanning for Web app vulnerabilities
- A software application on the host uses the unsafe practice of passing SQL statements in HTTP Post data or in a URL

Business Impact

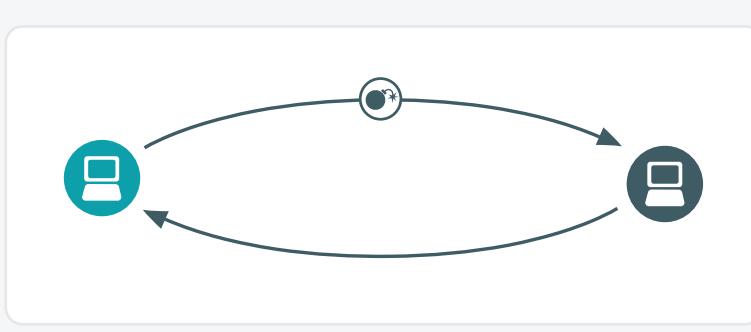
- Probing and potentially exploiting an internal Web application's vulnerabilities can be a prelude to a targeted attack getting access to data and then exfiltrating it
- Application software that passes SQL statements in HTTP Post data or as part of a URL may be vulnerable to attackers as they can send very different input than the application writer expects

Steps to Verify

- Verify systems identified as the source of SQL injection attacks should be communicating directly with SQL servers; download the PCAP to see the entire HTTP Post data or the URL to determine if its behaving as expected
- If this pattern is coming from neither an IT-run vulnerability scanner nor from software that by design sends SQL statements in requests, check for presence of malware on the host

Stage Loader

Lateral Movement



MITRE | ATT&CK®

T1210 Exploitation of Remote Services

T1570 Lateral Tool Transfer

Triggers

- The detection results from the observation of two closed sessions where an internal host is attacking another internal host by uploading a payload which causes the destination host to connect back to the initial host to download additional stages of software

Possible Root Causes

- The initial host is transmitting an exploit to a destination host which runs a stage loader and connects back to the initial host to load the rest of the malware necessary for the attacker to make progress toward their goal
- Bidirectional transaction-based protocols where commands or requests are issued over one port/protocol and data is returned shortly thereafter over another port/protocol can also trigger the detection—common protocols which behave in this manner include the WinRM 2.0 Framework (used for Windows remote management), PostgreSQL, and SNPP (Simple Network Paging Protocol)

Business Impact

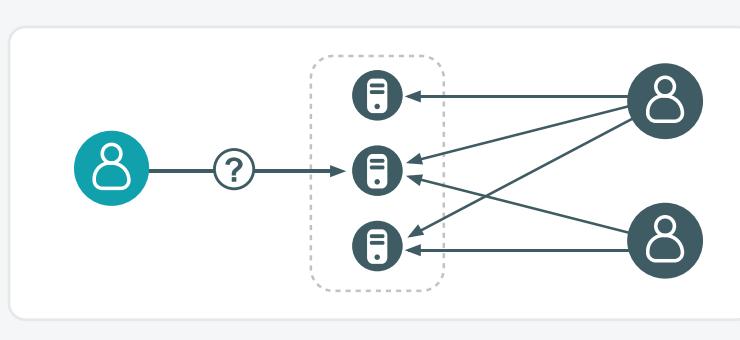
- Lateral movement within a network expands an attacker's footprint and exposes an organization to substantial risk of data acquisition and exfiltration
- Lateral movement through exploits or leveraging stolen credentials is involved in almost all high-profile breaches
- The destination host which is attacked provides a possible perspective on the potential business impact

Steps to Verify

- Determine whether there is any reason for the two hosts involved in a stage loading sequence to be communicating with each other
- Check to see whether any connections between the initial and destination host (in either direction) persist after the stage loading sequence
- Run all available endpoint checks on both the initial and the destination host to check for unwanted malware, but realize that fileless malware will typically escape detection

Suspicious Active Directory Operations

Lateral Movement



MITRE | ATT&CK®

TA0006 Credential Access
T1207 Rogue Domain Control

Triggers

- Either a new or non-domain controller host successfully triggered an anomalous Active Directory replication request against a legitimate domain controller. This functionality is normally limited to usage by domain controllers and limited high-privilege service accounts.

Possible Root Causes

Malicious Detection

- Provided the malicious actor has the required permissions and connectivity to a domain controller, they can leverage the DRS RPC protocol to successfully execute the following attacks:
- DCSync:** A malicious actor mimics a domain controller and targets a legitimate domain controller to invoke a Replication request (GetNCChanges) of the targeted AD Database containing hashed passwords.
- DCShadow:** A malicious actor creates a rogue domain controller by targeting a legitimate domain controller to add itself to a group of hosts permitted to receive these requests (domain controllers). The attacker will then force replication, dumping the Active Directory database and hashed password to the rogue domain controller. The attacker then typically removes itself from the list of hosts permitted to receive the requests.

Benign Detection

- A new domain controller has been deployed and hasn't had enough history to be identified as a domain controller.

Business Impact

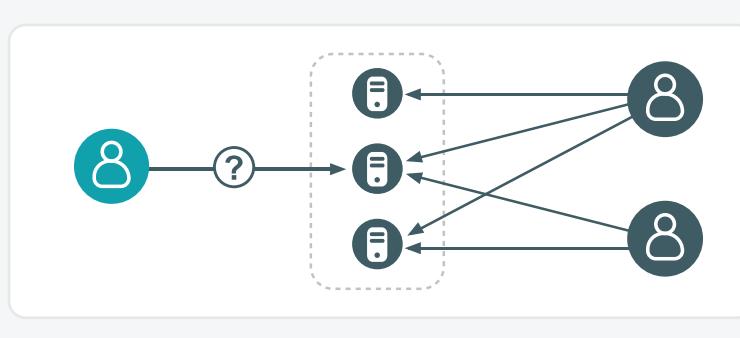
- Specific Risk:** Successful execution of either attack results in access to both usernames and hashed passwords of the targeted Active Directory infrastructure. An attacker can then perform offline attacks against the hashed passwords to escalate access.
- Impact:** These attacks likely result in a full domain compromise due to malicious actor having access to privileged account hashed passwords which will either be cracked or used to authenticate (NTLM) to other services/hosts.

Steps to Verify

- Investigate the host involved in the alert, verify if the host is a true domain controller through either an internal CMDB or Active Query of Domain Controller hosts on your environment.
 - Either the addition or removal of a domain controller on an environment is a rare event in comparison to other events within the environment and more specially within the RPC metadata stream.
 - Usage of requests like GetNCChanges, ReplicaAdd, or UpdateRefs are explicit are specific to only domain controllers.
 - If this host is a domain controller you should add it to the Domain Controllers Group, and apply a triage filter to exclude this host from generating a detection.
- Based on your environments configuration the replication requests should occur on a timely interval (default 15 minutes). In normal usage, you should see subsequent replication events.
- In malicious cases, these events will typically occur once, as there is no requirement for another replication of the database.
- Review logs for indications of either privileged accounts with the following:
 - Privileged accounts using old/odd authentication types such as NTLM to new hosts and services.
 - Privileged accounts invoking actions across multiple hosts on network within the RPC metadata stream

Suspicious Admin

Lateral Movement



MITRE | ATT&CK®

[T1003 OS Credential Dumping](#)

[T1078 Valid Accounts](#)

[T1212 Exploitation For Credential Access](#)

[T1552 Unsecure Credentials](#)

[T1555 Credentials From Password Stores](#)

[T1021 Remote Services](#)

[T1563 Remote Service Session Hijacking](#)

Triggers

- The host is using protocols correlated with administrative activity (RDP, SSH, IPMI, iDRAC, etc.) in ways which are considered suspicious

Possible Root Causes

- The host has begun using an administrative protocol to connect to a system for which one or more other hosts have already been observed to be regular administrators using the same protocol
- Administrative connections via a particular administrative protocol to a system which has no known regular administrators using that protocol will not result in a detection
- Administrative connections to a system which has a known regular administrator host using the chosen protocol will also not result in a detection if there is significant overlap in administrative connections to other systems between this host and the other known administrator host
- If such an administrative connection recurs over a period of several days, it is considered normal and no longer will trigger a detection
- The detection may be benign when it involves a host assigned to a new employee authorized to administrate the target systems or when the role of the employee has undergone a significant change

Business Impact

- Administrative protocols are a primary tool for attackers to move laterally inside a network in which they have already established a toehold
- Given that administrative connections are typically used in conjunction with administrative credentials, the attacker may have almost unconstrained access to systems and data that are the organization's key assets
- Unexpected and unexplained administrative connections represent a huge potential risk in the lifecycle of a major breach

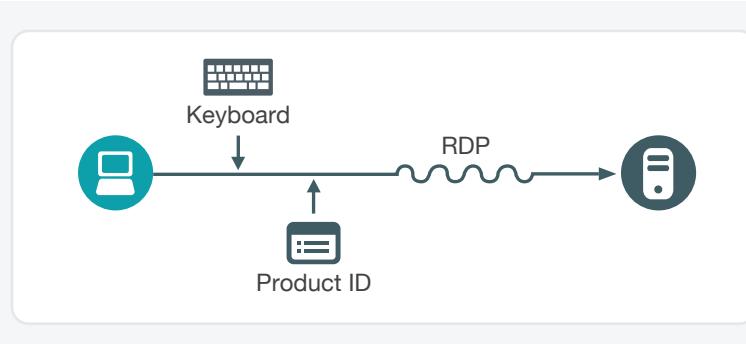
Steps to Verify

- Verify whether the host belongs to an employee whose job function requires administrative access to other systems
- Verify whether the employee who has been assigned the host should be using the particular administrative protocol to administer the identified system

- Inquire whether the owner of the host actually initiated the administrative connection in order to determine whether the host has been compromised
- Check the logs on the administered target for the creation of new accounts, the launch of abnormal processes and the modification of registry key
- If employee associated with this host was not the originator of the admin session, reset all domain and local admin credentials belonging to the employee across the local machine and the network
- If the credentials of the employee whose machine was compromised had domain administrative privileges, the secret key of the domain controller may have been compromised and may need to be reset – search for “krbtgt account password change” to find instructions on how to do this
- Verify that the host from which the administrative connection was originated is a jump system as this may mean that the originator of the administrative connection is an upstream host which connected to the jump system

Suspicious Remote Desktop

Lateral Movement



MITRE | ATT&CK®

[T1003 OS Credential Dumping](#)

[T1078 Valid Accounts](#)

[T1212 Exploitation For Credential Access](#)

[T1552 Unsecure Credentials](#)

[T1555 Credentials From Password Stores](#)

[T1021 Remote Services](#)

Triggers

- A host connects to an internal RDP server with a keyboard layout or a product ID different than the one usually seen in conjunction with the specified RDP client token
- A host connects to an internal RDP server with a keyboard layout that is unusual for that RDP server
- A host connects to an internal RDP server with a keyboard layout that is different from those usually seen on the network

Possible Root Causes

- An external foreign attacker who has taken over control of an internal host is using it with unusual keyboard layouts to connect to RDP servers to move laterally in the network
- An external attacker who has taken over control of an internal host has brought along their own RDP stack and is using it to connect to internal RDP servers to move laterally in the network
- An employee has switched to their native keyboard layout while accessing an RDP server
- An employee has installed a new RDP client with a new product ID and is accessing an RDP server

Business Impact

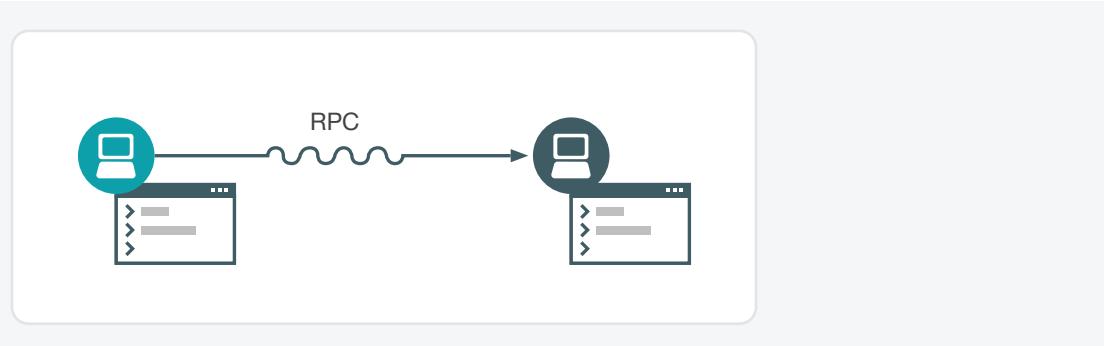
- Along with SSH, RDP is one of the most useful lateral movement protocols for attackers as it allows remote control of the target as well as the copying of files across the connection
- This type of control and data acquisition may happen well in advance of actual exfiltration attempts and represents a great chance to head off attacks before any substantial damage occurs

Steps to Verify

- For keyboard layout anomalies, inquire whether the user of the internal host is fluent in the language of the flagged keyboard layout
- For an RDP product id anomaly, inquire whether IT has installed new RDP client software or ask the user of the host whether they have done so

Suspicious Remote Execution

Lateral Movement



MITRE | ATT&CK®

[T1569 System Services](#)

[T1021 Remote Services](#)

[T1047 Windows Management Instrumentation](#)

[T1053 Scheduled Task/Job](#)

[T1078 Valid Accounts](#)

[T1570 Lateral Tool Transfer](#)

[T1571 Non-Standard Port](#)

[T1572 Protocol Tunneling](#)

Triggers

- An internal host is utilizing the SMB or DCE RPC protocol to make one or more suspicious RPC requests and referencing functions related to remote execution of code
- The combination of source host, destination host, user account and RPC UUID has not previously been observed

Possible Root Causes

- An infected host, a malicious insider or a red team participant who is in control of the host is trying to spread laterally by executing code on systems to which it has connected
- Newly installed software or software that is infrequently used is legitimately making use of remote execution RPCs; this behavior is relatively common for system management software

Business Impact

- Lateral movement via remote execution is a key element of many different attacks and the SMB channel allows both for the copying of executables and the use of RPCs to execute them
- Even systems which are permitted to perform remote execution should be monitored because those systems are the most valuable for an attacker to compromise

Steps to Verify

- Determine whether the internal host in question should be using remote execution RPCs
- Determine whether the user account flagged in the detection is one with administrative privileges and whether that administrator logged into the host which triggered the detection
- Determine whether the user account flagged in the detection is a service account associated with a specific product and whether that product should be running on the host which triggered the detection
- Determine which process on the internal host is initiating the SMB requests that includes the RPC request; in Windows systems, this can be done using a combination of netstat and tasklist commands
- Verify that the process should be running on the internal host and whether the process is configured correctly

Threat Intelligence Match

Lateral Movement



Triggers

- After configuration of a [3rd party intelligence thread feed](#), an internal host has initiated communications with another internal host and the connection has met criteria specified in one or more configured threat feeds

Possible Root Causes

- A host includes malware which is initiating the connection that triggered the detection
- A user on the host manually initiated the connection which triggered the detection

Business Impact

- The internal connection may be used by the originating host to compromise the target host or to maintain communication with a previously compromised host
- If the connection is to a target host which contains important data, this may represent an attempt to acquire data for later exfiltration
- The threat intel feed may have included additional context tied to the specific criteria that the connection met
- Lateral movement and data acquisition are present in almost all large-scale breaches

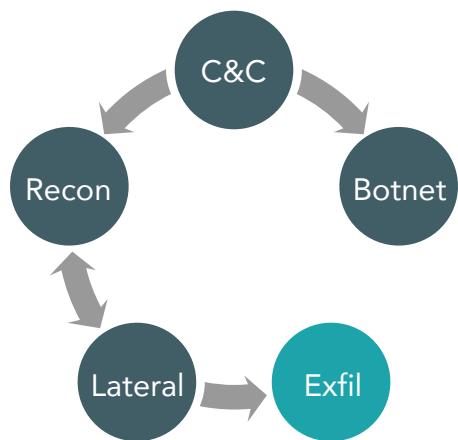
Steps to Verify

- Refer to the information accompanying your threat feed as it may include verification and remediation instructions
- Determine which process on the internal host is sending the traffic which was flagged; in Windows systems, this can be done using a combination of netstat and tasklist commands
- Check if a user has knowingly installed remote access software and decide whether the resulting risk is acceptable
- Scan the computer for known malware and potentially reimage it, noting that some infections leave no trace on disk and reside entirely in memory

Category

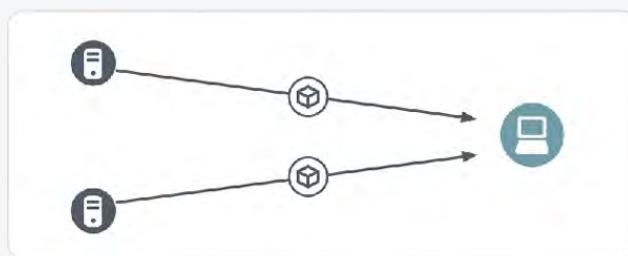
Exfiltration

- Covers scenarios where data is being sent outside or collected in a way meant to hide the data transfer
- While data is constantly being sent out of the network or cloud environment, it usually does not involve the use of techniques meant to hide the transfer
- The host or account transmitting the data, where it is transmitting the data, the amount of data and the technique used to send it all provide indicators of exfiltration



Data Gathering

Exfiltration



MITRE | ATT&CK®

[T1213 Data From Information Repositories](#)

[T1074 Data Staged](#)

[T1119 Automated Collection Alternative Protocol](#)

[T1039 Data from Network Shared Drive](#)

Triggers

- Pre-exfiltration behaviors have been observed on a host that has received abnormally high amounts of data from one or more hosts within a short period of time.

Possible Root Causes

- An attacker has pivoted to a host to use for dumping/staging data prior to exfiltrating, likely taking advantage of the trusted nature of this host to bypass security controls and evade detection.
- A malicious insider is collecting data they intend to steal from a position of trust.
- A user has joined a new team, changed organizational roles, or otherwise been given reason to significantly depart from their typical data access and retrieval activities.
- An application has been observed on an unusual or infrequent backup or update cycle.

Business Impact

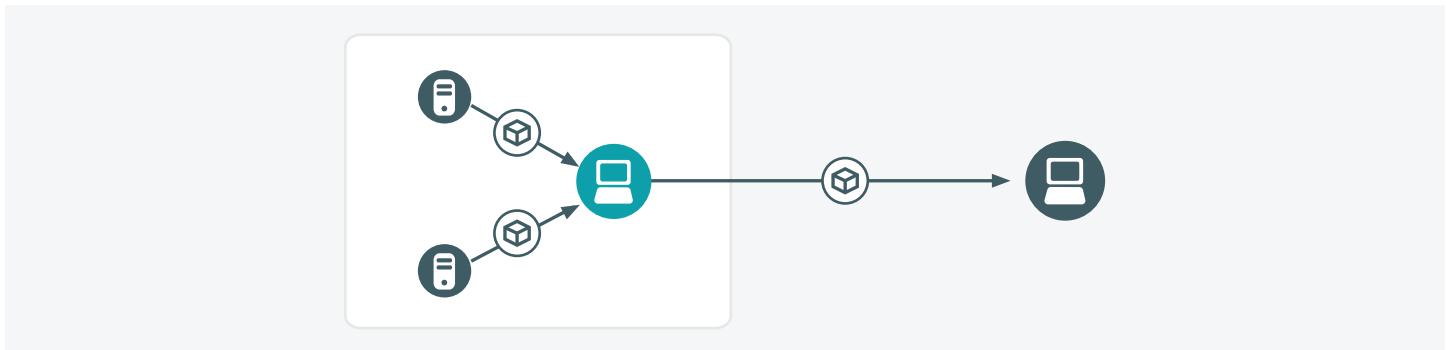
- Failure to identify and respond to pre-exfiltration activities in an organization increases the likelihood of data loss.
- When successful, data exfiltration places an organization at the risk of the loss of intellectual property, financial data, or other regulated or sensitive data sources.

Steps to Verify

- Verify if the data gathered supports valid and authorized business activities.
- Investigate the host and associated accounts for other signs of compromise.

Data Smuggler

Exfiltration



MITRE | ATT&CK®

[T1041 Exfiltration Over C2 Channel](#)

[T1213 Data From Information Repositories](#)

[T1560 Archive Collected](#)

[Data](#)

[T1074 Data Staged](#)

[T1048 Exfiltration Over Alternative Protocol](#)

[T1020 Automated Exfiltration](#)

[T1030 Data Transfer Size Limits](#)

[T1567 Exfiltration Over Web Service](#)

Triggers

- An internal host is acquiring a large amount of data from one or more internal servers and is subsequently sending a significant amount of data to an external system

Possible Root Causes

- A host infected with malware as part of a targeted attack or a malicious insider may be acquiring and exfiltrating company data
- While acquiring and transmitting a large quantity of data to the outside within a short period of time may be pure coincidence, the outbound data transfer is significant enough to warrant further examination

Business Impact

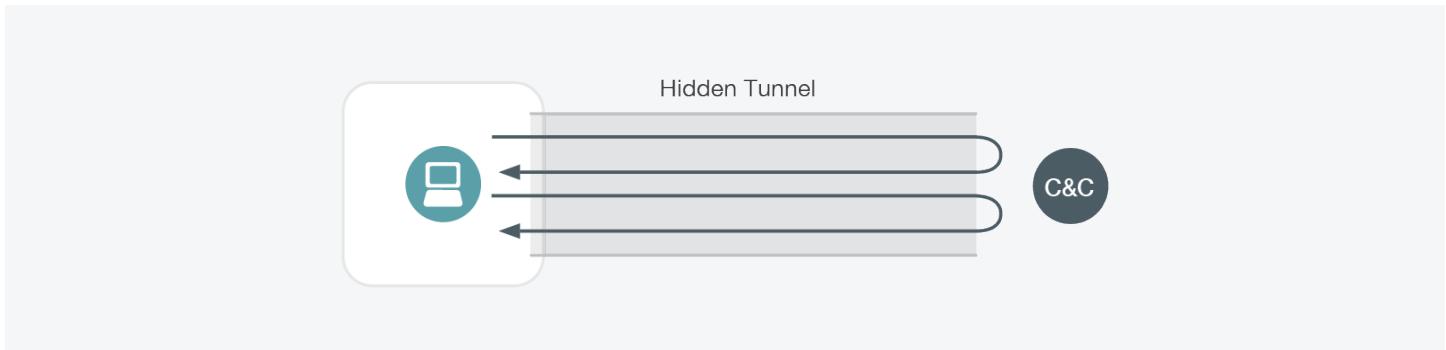
- The detection signals possible exfiltration of company data
- The internal servers from which the data was retrieved provides some indication of the data which was acquired; if those servers contain valuable information and the external service to which data was uploaded is not an IT- sanctioned service, the potential business risk is high

Steps to Verify

- Decide whether this may be a malicious insider or an infected host
- If the signs point to an infected host, contact the user to inquire if they initiated the uploading behavior in question
- For potential malicious insiders, perform a complete analysis of recent behavior
- Look up the external system IP addresses and domain names on sites that maintain reputation lists as this may provide a clear indication that the internal host is infected; such lookups are supported directly within the UI

Hidden DNS Tunnel

Exfiltration



MITRE | ATT&CK®

[T1005 Data from Local System](#)

[T1115 Clipboard Data](#)

[T1071 Application Layer Protocol](#)

[T1125 Video Capture](#)

[T1113 Screen Capture](#)

[T1572 Protocol Tunneling](#)

[T1123 Audio Capture](#)

[T1041 Exfiltration Over C2 Channel](#)

Triggers

- An internal host is communicating with an outside IP using DNS where another protocol is running over the top of the DNS sessions
- This represents a hidden tunnel involving multiple sessions over longer periods of time mimicking normal DNS traffic

Possible Root Causes

- A targeted attack may use hidden tunnels to hide exfiltration activity
- A user is utilizing tunneling software to communicate with Internet services which might not otherwise be accessible
- Intentionally installed software is using a hidden tunnel to bypass expected firewall rules

Business Impact

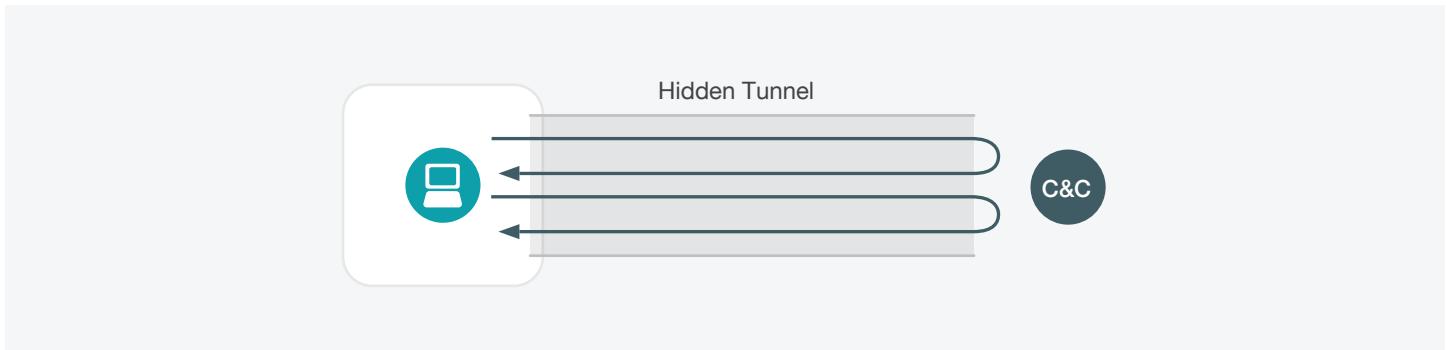
- The use of a hidden tunnel by some software may be benign, but it represents significant risk as the intention is to bypass security controls
- Hidden tunnels used as part of a targeted attack are meant to slip by your perimeter security controls and indicate a sophisticated attacker
- Hidden tunnels are rarely used by botnets, though more sophisticated bot herders with more ambitious goals may utilize them

Steps to Verify

- Check to see if the destination domain of the tunnel is an entity you trust for your network
- Ask the user of the host whether they are using hidden tunnel software for any purpose
- Before removing the offending software via antivirus or reimaging, take a memory snapshot for future analysis of the incident
- If the behavior reappears shortly after a reimaging, this may be a hardware/BIOS tunnel

Hidden HTTP Tunnel

Exfiltration



MITRE | ATT&CK®

[T1005 Data from Local System](#)

[T1115 Clipboard Data](#)

[T1071 Application Layer Protocol](#)

[T1125 Video Capture](#)

[T1113 Screen Capture](#)

[T1572 Protocol Tunneling](#)

[T1123 Audio Capture](#)

[T1041 Exfiltration Over C2 Channel](#)

Triggers

- An internal host is communicating with an outside IP using HTTP where another protocol is running over the top of the HTTP sessions
- This represents a hidden tunnel involving multiple sessions over longer periods of time mimicking normal Web traffic

Possible Root Causes

- A targeted attack may use hidden tunnels to hide exfiltration activity
- A user is utilizing tunneling software to communicate with Internet services which might not otherwise be accessible
- Intentionally installed software is using a hidden tunnel to bypass expected firewall rules

Business Impact

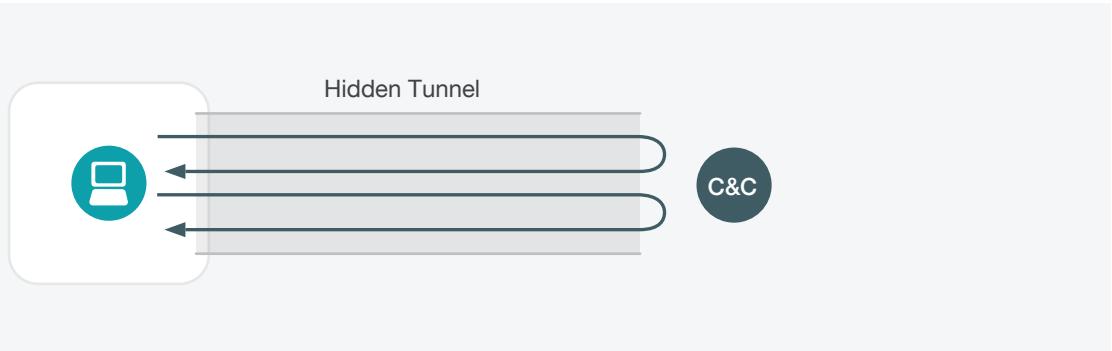
- The use of a hidden tunnel by some software may be benign, but it represents significant risk as the intention is to bypass security controls
- Hidden tunnels used as part of a targeted attack are meant to slip by your perimeter security controls and indicate a sophisticated attacker
- Hidden tunnels are rarely used by botnets, though more sophisticated bot herders with more ambitious goals may utilize them

Steps to Verify

- Check to see if the destination IP address or domain of the tunnel is an entity you trust for your network
- Ask the user of the host whether they are using hidden tunnel software for any purpose
- Before removing the offending software via antivirus or reimaging, take a memory snapshot for future analysis of the incident
- If the behavior reappears shortly after a reimaging, this may be a hardware/BIOS tunnel

Hidden HTTPS Tunnel

Exfiltration



MITRE | ATT&CK®

[T1005 Data from Local System](#)

[T1115 Clipboard Data](#)

[T1071 Application Layer Protocol](#)

[T1125 Video Capture](#)

[T1113 Screen Capture](#)

[T1572 Protocol Tunneling](#)

[T1123 Audio Capture](#)

[T1041 Exfiltration Over C2 Channel](#)

Triggers

- An internal host is communicating with an outside IP using HTTPS where another protocol is running over the top of the HTTPS sessions
- This represents a hidden tunnel involving one long session or multiple shorter sessions over a longer period of time mimicking normal encrypted Web traffic
- When it can be determined whether the tunneling software is console-based or driven via a graphical user interface, that indicator will be included in the detection

Possible Root Causes

- A targeted attack may use hidden tunnels over SSL on port 443 to hide exfiltration activity
- A user is utilizing tunneling software to communicate with Internet services which might not otherwise be accessible
- Intentionally installed software is using a hidden tunnel to bypass expected firewall rules

Business Impact

- The use of a hidden tunnel by some software may be benign, but it represents significant risk as the intention is to bypass security controls
- Hidden tunnels used as part of a targeted attack are meant to slip by your perimeter security controls and indicate a sophisticated attacker
- Hidden tunnels are rarely used by botnets, though more sophisticated bot herders with more ambitious goals may utilize them

Steps to Verify

- Check to see if the destination IP or domain of the tunnel is an entity you trust for your network
- Ask the user of the host whether they are using hidden tunnel software for any purpose
- Before removing the offending software via antivirus or reimaging, take a memory snapshot for future analysis of the incident
- If the behavior reappears shortly after a reimaging, this may be a hardware/BIOS tunnel

ICMP Tunnel

Exfiltration



MITRE | ATT&CK®

T1008 Fallback Channels

T1095 Non-Application Layer Protocol

Triggers

- A host was observed sending larger than expected ICMP payloads to an external address which may indicate data exfiltration to an external destination.

Possible Root Causes

Malicious Detection

- An attacker is using ICMP to exfiltrate data from the environment.

Benign Detection

- A networked device is using ICMP to assess connectivity in a nonstandard way.

Business Impact

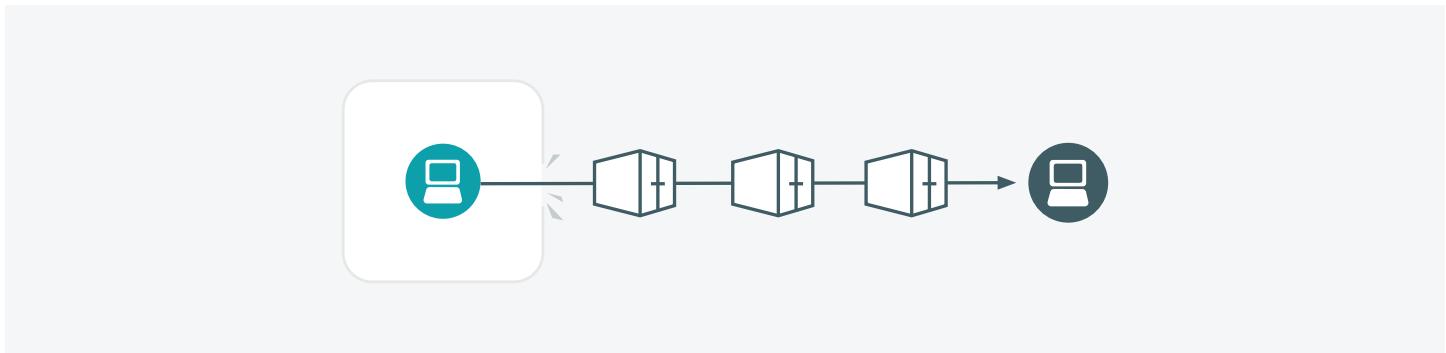
- The presence of an ICMP tunnel indicates the host was compromised & that an attacker has remote access to the machine.
- Recon, data exfiltration, lateral movement, privilege escalation, & establishing a tunnel over a more reliable protocol like HTTPS are all likely next steps.
- ICMP tunnels can be stealthy and are often used to evade sophisticated perimeter security controls.

Steps to Verify

- Check the destination IP & determine if the observed traffic arrives at a trusted endpoint.
- Investigate the host for malware, there may be code present which establishes a C2 channel with another host.

Smash and Grab

Exfiltration



MITRE | ATT&CK®

[T1041 Exfiltration Over C2 Channel](#)

[T1213 Data From Information Repositories](#)

[T1560 Archive Collected Data](#)

[T1029 Scheduled Transfer](#)

[T1119 Automated Collection](#)

[T1048 Exfiltration Over Alternative Protocol](#)

[T1020 Automated Exfiltration](#)

[T1030 Data Transfer Size Limits](#)

[T1567 Exfiltration Over Web Service](#)

Triggers

- A host transmits unusually large volumes of data to destinations which are not considered normal for this network

Possible Root Causes

- An attacker is rapidly exfiltrating large volumes of data from your network
- The host is sending large volumes of data to destinations that have not been previously used for large data transfers

Business Impact

- The detection signals possible exfiltration of company data
- The host from which the data was sent, the destination to which the data was sent and the volume of data transmitted may provide some clues to what data was transmitted
- If the external service to which data was uploaded is not an IT-sanctioned service, the potential business risk is high

Steps to Verify

- Check to see if the destination IP or domain to which data was moved is an entity you trust for your network
- Ask the user of the host whether they have any knowledge of the data transfer
- If the data transfer is unexplained and your endpoint security solution logs such things, determine what software on the host was responsible for the data transfer

Threat Intelligence Match

Exfiltration



Triggers

- After configuration of a [3rd party intelligence thread feed](#), an internal host is connecting to an external system and the connection has met criteria specified in one or more configured threat feeds

Possible Root Causes

- A host includes malware which is initiating the connection that triggered the detection
- A user on the host manually initiated the connection which triggered the detection

Business Impact

- The detection signals exfiltration of company data
- The host from which the data was sent, the destination to which the data was sent and the volume of data transmitted may provide some clues to what data was transmitted
- The threat intel feed may have included additional context tied to the specific criteria that the connection met
- If the external service to which data was uploaded is not an IT-sanctioned service, the potential business risk is high

Steps to Verify

- Refer to the information accompanying your threat feed as it may include verification and remediation instructions
- Determine which process on the internal host is sending the traffic which was flagged; in Windows systems, this can be done using a combination of netstat and tasklist commands
- Check if a user has knowingly installed remote access software and decide whether the resulting risk is acceptable
- Scan the computer for known malware and potentially reimage it, noting that some infections leave no trace on disk and reside entirely in memory

Category

Suspect Protocol Activity Detections

Overview

- Suspect Protocol Activity (SPA) detections are a feature of Vectra NDR and enhance the overall effectiveness of a Vectra NDR deployment. Vectra's existing AI models provide an incredible breadth of coverage for both known and unknown attack vectors because they examine attacker behavior. These AI models can take months to over a year to research, build, test, and deploy. SPA detections are signature based and can be developed in much shorter timeframes. Some attacks do not need an AI model to find them and can be efficiently found with Vectra-developed signatures. SPA detections compliment AI detections, providing comprehensive coverage of the environment. For a more information, please see [Suspect Protocol Detections \(Feature Overview\)](#).

Detections

- Vectra SPA detections cover all areas of compromise, including Command & Control, Exfil, Reconnaissance, and Lateral Movement. For details on specific detections, please see [Suspect Protocol Activity Detections \(Explained\)](#).

Details

- There is no additional licensing cost incurred for SPA detections. SPA detections are not enabled by default, instead the feature can be enabled by executing a command at the CLI of your Vectra Brain. SPA signatures are automatically provided to your Vectra deployment from Vectra's cloud. At this time, there is no ability to support air gap deployments for SPA detections.

Category

Info

- Reports on new and novel events without directly impacting scoring
- New and novel events occur normally in most network and cloud environments and in most cases are not directly linked to threats
- Awareness of new and novel events support better situational awareness and provide additional context when observed with kill chain alerts

New Host

- Reports on the first time a host was seen on the internal network.

New Host Role

- Reports when a host began operating with a particular infrastructure role.

Novel MAC Vendor

- Reports when a host appears with an unusual MAC vendor for the network.

Novel Admin Protocol Usage

- Reports when a host uses an administrative protocol (e.g., SSH) for the first time.

Novel External Destination Port

- Reports when a host is seen making an outbound connection on a destination port that is rare for the environment and lasted longer than 5 minutes.

Novel Access to SMB Admin Share

- Reports when a host is seen connecting to another host's SMB admin share and it is unusual for this host to connect to other systems in this way.

Single Weak Kerberos Cipher Response

- This behavior is likely the result of a Legacy system that lacks support for modern ciphers and is using weak ciphers. Vectra recommends applying all necessary patches and prioritizing the security of associated accounts.

Vectra Indicator Match

- A host was seen with network artifacts that are sometimes associated with attacker infrastructure. These events should be reviewed in the context of other threat detections.

Network Detection Profiles



Botnet

Detection Profile

General Behavioral Profile

- Programmatic discovery and asset monetization techniques
- External, persistent Command and Control behaviors

Possible Root Causes

- A host has been infected and is participating in a botnet
- SaaS enabled asset discovery services have been observed

Business Impact

- Investigations of entities matching this profile should be prioritized in alignment with malware remediation procedures and urgency
- Failure to take timely steps to respond to entities that match this profile may allow crypto-mining activities to persist, or open the door to more aggressive attacks from the compromised host over time

About Detection Profiles

Vectra AI supports security analyst investigative workflows by classifying the behavioral profile of an entity based on the active detections it has exhibited – the assignment of these profiles are useful for quickly wrapping context around the types of real world profiles that exhibit similar behaviors to the one under investigation.

When determining a behavioral profile for a host, only active detections are considered – this means that if detections on a host are triaged as marked-as-custom, marked-as-fixed, or become inactive, that host's behavioral profile may change.

Cloud Services

Detection Profile

General Behavioral Profile

- Complex, active external Command and Control and/or Data Exfiltration Services
- NOT PRESENT: Lateral movement focused behaviors

Possible Root Causes

- Entities are leveraging unauthorized cloud services

Business Impact

- Investigations of entities matching this profile may generally be prioritized in alignment with addressing the presence of unauthorized IT Services, or with risks associated with data exfiltration and data loss

About Detection Profiles

Vectra AI supports security analyst investigative workflows by classifying the behavioral profile of an entity based on the active detections it has exhibited – the assignment of these profiles are useful for quickly wrapping context around the types of real world profiles that exhibit similar behaviors to the one under investigation.

When determining a behavioral profile for a host, only active detections are considered – this means that if detections on a host are triaged as marked-as-custom, marked-as-fixed, or become inactive, that host's behavioral profile may change.

External Adversary

Detection Profile

General Behavioral Profile

- Technically sophisticated, objective-oriented activities
- Advanced discovery and lateral movement techniques
- External, persistent Command and Control and/or Data Exfiltration

Possible Root Causes

- Advanced Persistent Threat
- Full scope Red Team / Penetration Test

Business Impact

- Investigation of entities matching this profile should be considered urgent
- Failure to take timely steps to respond to entities that match this profile may increase the risk of a breach

About Detection Profiles

Vectra AI supports security analyst investigative workflows by classifying the behavioral profile of an entity based on the active detections it has exhibited – the assignment of these profiles are useful for quickly wrapping context around the types of real world profiles that exhibit similar behaviors to the one under investigation.

When determining a behavioral profile for a host, only active detections are considered – this means that if detections on a host are triaged as marked-as-custom, marked-as-fixed, or become inactive, that host's behavioral profile may change.

Insider Threat: Admin

Detection Profile

General Behavioral Profile

- Technically sophisticated, objective-oriented activities
- Advanced discovery and lateral movement techniques
- NOT PRESENT: External Command and Control and/or Data Exfiltration

Possible Root Causes

- Technically sophisticated insider threat with local network access
- Emerging External Adversary with out-of-band communication
- An Admin has begun performing authorized activities that were previously unknown to the system

Business Impact

- Investigations of entities matching this profile should be prioritized above less critical severity tasks
- Failure to take timely steps to respond to entities that match this profile may increase the risk of unauthorized or malicious activities

About Detection Profiles

Vectra AI supports security analyst investigative workflows by classifying the behavioral profile of an entity based on the active detections it has exhibited – the assignment of these profiles are useful for quickly wrapping context around the types of real world profiles that exhibit similar behaviors to the one under investigation.

When determining a behavioral profile for a host, only active detections are considered – this means that if detections on a host are triaged as marked-as-custom, marked-as-fixed, or become inactive, that host's behavioral profile may change.

Insider Threat: User

Detection Profile

General Behavioral Profile

- Low-sophistication, human-based network reconnaissance and objectives
- Data Exfiltration

Possible Root Causes

- A user is collecting and exfiltrating data outside of the organized authorized storage
- A user has been granted additional roles and privileges not previously known, or is moving data to previously unauthorized cloud storage locations

Business Impact

- Investigations of entities matching this profile should be prioritized in alignment with organizational tolerance to data loss
- Failure to take timely steps to respond to entities that match this profile may allow for the loss of intellectual property, competitive advantage, legally protected, or regulated data

About Detection Profiles

Vectra AI supports security analyst investigative workflows by classifying the behavioral profile of an entity based on the active detections it has exhibited – the assignment of these profiles are useful for quickly wrapping context around the types of real world profiles that exhibit similar behaviors to the one under investigation.

When determining a behavioral profile for a host, only active detections are considered – this means that if detections on a host are triaged as marked-as-custom, marked-as-fixed, or become inactive, that host's behavioral profile may change.

IT Discovery

Detection Profile

General Behavioral Profile

- Low-sophistication discovery and reconnaissance techniques

Possible Root Causes

- Asset Management or Change Management Infrastructure
- IP Address Management (IPAM) Infrastructure

Business Impact

- Investigations of entities matching this profile may generally be prioritized after more urgent activities are complete
- Failure to take timely steps to investigate may allow the perpetuation of unauthorized IT Discovery Services

About Detection Profiles

Vectra AI supports security analyst investigative workflows by classifying the behavioral profile of an entity based on the active detections it has exhibited – the assignment of these profiles are useful for quickly wrapping context around the types of real world profiles that exhibit similar behaviors to the one under investigation.

When determining a behavioral profile for a host, only active detections are considered – this means that if detections on a host are triaged as marked-as-custom, marked-as-fixed, or become inactive, that host's behavioral profile may change.

IT Services

Detection Profile

General Behavioral Profile

- Low-sophistication reconnaissance and discovery
- Lateral machine-to-machine communication
- Simple external data exfiltration services

Possible Root Causes

- IT Services are exhibiting machine-to-machine communication patterns

Business Impact

- Investigations of entities matching this profile may generally be prioritized after more urgent activities are complete
- Failure to take timely steps to investigate may allow the perpetuation of unauthorized IT Services

About Detection Profiles

Vectra AI supports security analyst investigative workflows by classifying the behavioral profile of an entity based on the active detections it has exhibited – the assignment of these profiles are useful for quickly wrapping context around the types of real world profiles that exhibit similar behaviors to the one under investigation.

When determining a behavioral profile for a host, only active detections are considered – this means that if detections on a host are triaged as marked-as-custom, marked-as-fixed, or become inactive, that host's behavioral profile may change.

Potentially Unwanted Program

Detection Profile

General Behavioral Profile

- External, persistent Command and Control behaviors
- Programmatic Discovery behaviors
- NOTE PRESENT: Asset monetization techniques

Possible Root Causes

- Adware or Potentially Unwanted Programs (PUP) are active.
- SaaS enabled asset discovery services have been observed

Business Impact

- Investigations of entities matching this profile may generally be prioritized in alignment with addressing the presence of unauthorized IT Services, or Unwanted or Unauthorized Software, or Policy and Acceptable Use violations.

About Detection Profiles

Vectra AI supports security analyst investigative workflows by classifying the behavioral profile of an entity based on the active detections it has exhibited – the assignment of these profiles are useful for quickly wrapping context around the types of real world profiles that exhibit similar behaviors to the one under investigation.

When determining a behavioral profile for a host, only active detections are considered – this means that if detections on a host are triaged as marked-as-custom, marked-as-fixed, or become inactive, that host's behavioral profile may change.

Ransomware

Detection Profile

General Behavioral Profile

- Behavioral patterns associated with ransomware

Possible Root Causes

- Malicious ransomware activity
- Technical services exhibiting behaviors similar to ransomware

Business Impact

- Investigation of entities matching this profile should be considered urgent
- Failure to take timely steps to respond to entities that match this profile may increase risk of loss of data and system availability

About Detection Profiles

Vectra AI supports security analyst investigative workflows by classifying the behavioral profile of an entity based on the active detections it has exhibited – the assignment of these profiles are useful for quickly wrapping context around the types of real world profiles that exhibit similar behaviors to the one under investigation.

When determining a behavioral profile for a host, only active detections are considered – this means that if detections on a host are triaged as marked-as-custom, marked-as-fixed, or become inactive, that host's behavioral profile may change.

Vulnerability Discovery

Detection Profile

General Behavioral Profile

- Discovery, Reconnaissance, Lateral movement, and/or Exploitation
- NOT PRESENT: External, persistent Command and Control and/or Data Exfiltration

Possible Root Causes

- An adversary that has yet to exhibit the full range of malicious behaviors, or a limited scope penetration testing activity
- Vulnerability discovery and management infrastructure behaviors observed

Business Impact

- Investigations of entities matching this profile should be prioritized in alignment with procedures associated with unauthorized vulnerability discovery or limited scope penetration testing
- Failure to take timely steps to investigate may allow additional dwell time for an adversary with unobserved, persistent command and control or allow the presence of unauthorized, rogue vulnerability discovery infrastructure

About Detection Profiles

Vectra AI supports security analyst investigative workflows by classifying the behavioral profile of an entity based on the active detections it has exhibited – the assignment of these profiles are useful for quickly wrapping context around the types of real world profiles that exhibit similar behaviors to the one under investigation.

When determining a behavioral profile for a host, only active detections are considered – this means that if detections on a host are triaged as marked-as-custom, marked-as-fixed, or become inactive, that host's behavioral profile may change.

Worm

Detection Profile

General Behavioral Profile

- Wide-but-shallow network recon, searching for specific services.
- Lateral machine-to-machine communication.
- NOT PRESENT: Deep, thorough network enumeration of many services on individual targets.

Possible Root Causes

- Malicious software is actively performing worm-like spreading behaviors across network
- Authorized IT software is leveraging risky, rare machine-to-machine discovery and update functionality

Business Impact

- Investigations of entities matching this profile may generally be prioritized in alignment with addressing the presence of destructive malware, ransomware, and worms.

About Detection Profiles

Vectra AI supports security analyst investigative workflows by classifying the behavioral profile of an entity based on the active detections it has exhibited – the assignment of these profiles are useful for quickly wrapping context around the types of real world profiles that exhibit similar behaviors to the one under investigation.

When determining a behavioral profile for a host, only active detections are considered – this means that if detections on a host are triaged as marked-as-custom, marked-as-fixed, or become inactive, that host's behavioral profile may change.

Observed Privilege Scores



The Vectra AI platform displays observed privilege scores for accounts, hosts, and services in some host and detection pages. The concept of “observed privilege” is distinct from that of “granted privilege”. A user may be given an account that has been granted a lot of privilege, but if the user only makes very modest use of that privilege, the observed privilege of the account will be low. Vectra AI focuses on observed privilege as it provides a clearer basis for implementing effective detection strategies related to advanced attackers’ use of stolen credentials.

All observed privilege scores, regardless of the object (account, host, or service) to which they refer, are expressed on the same scale. Each privilege score consists of two components: a numerical score from 1 to 10 (ranging from low to high privilege) and a label (low, medium, or high). Scores of 1 and 2 are labeled “low”, scores of 3 to 7 are labeled “medium”, and scores of 8-10 are labeled “high”. Vectra AI detection algorithms that are part of the Privileged Access Analytics (PAA) feature make extensive use of these privilege scores.

Account Scores

Observed privilege scores for accounts derive from the number of services an account connects to, either exclusively or in partnership with a small number of other accounts. An account that connects to 200 services, each of which is used by only a small number of other accounts, will score high. An account which connects to 5 services, each of which is used by a large number of other accounts, will score low.

Using this approach, service accounts tend to score high as they usually connect to many services that only the service account can access. Privileged users (aka admins) are typically given a normal account (to be used for normal non-privileged activity such as getting onto WiFi, requesting vacations, etc.) and a privileged account (to be used only for activities which require privileges). The first of these accounts will typically have a low score, the second a high score.

Service Scores

Let’s begin by defining what a “service” is. Given that PAA is constructed on Kerberos traffic and Active Directory data, a service is any distinct place (server) to which a system (client) can connect to request a service. Using this definition, RDP is not a service, but RDP to a particular system (e.g. RDP to serverA) *is* a service. Given such a methodology, it’s easy to see how a network can contain many services.

Observed privilege scores for services derive from the scores of the accounts that are used to connect to the service. Thus, if a service is only accessed via accounts that predominantly have high privilege scores, the service will also have a high privilege score. This can, for instance, happen when a small number of privileged accounts belonging to admins are used from each admin’s laptop to exclusively connect to a particular service. Another example is when a service account for a backup server connects to an agent running on 1,000 laptops. In both instances, the accounts used are high privilege—in the latter example, there is only single account in use. Conversely, a vacation request portal used by everyone in an organization (each logging in with their user accounts) will rate low on the privilege scale. And a service used exclusively by a low privilege account will also have a low privilege score.

Host Scores

Observed privilege scores for hosts derive from the scores of the accounts that are seen on the host. If a particular host has only high-privileged accounts on it, the host will have a high privilege. A jump system from which only privileged users initiate connections to downstream servers is an example of a high-privileged host. A laptop on which a privileged user uses both their normal account and their privileged account will score quite high (though not as high as the jump system described above).

Privilege scores for hosts often indicate how interesting the hosts would be to an attacker. If an attacker compromises a high privilege host, they can harvest the credentials of one or more high-privilege accounts on that host. In a scenario where an attacker wants to move laterally through the use of stolen credentials, this is exactly their goal. After all, stealing credentials which have little or no privilege won't get the attacker closer to their goal.

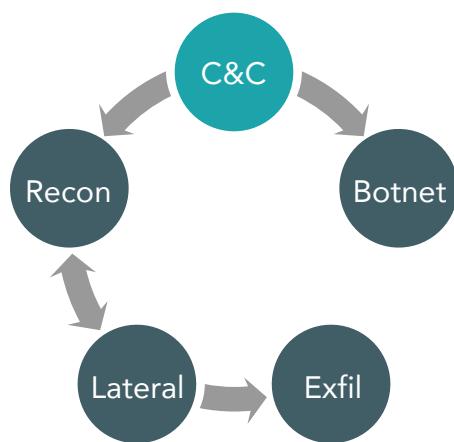
Microsoft 365 Detections



Category

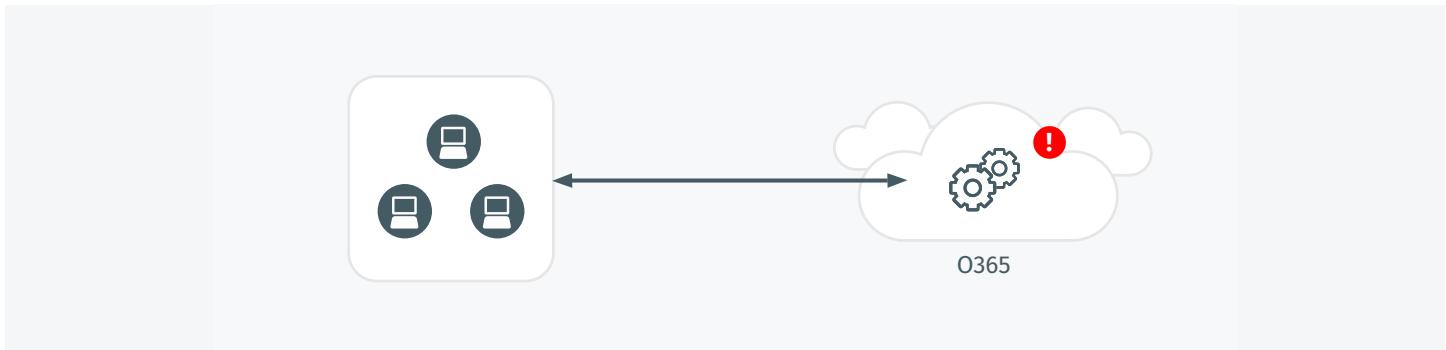
Command & Control

- A host or account appears to be under control of an external entity
- Most often, the control is automated as the host or account is part of a botnet or has adware or spyware installed
- The host or account may be manually controlled from the outside – this is the most threatening case and makes it highly likely that this is a targeted attack



M365 Power Automate HTTP Flow Creation

Command & Control



MITRE | ATT&CK®

[T1041 Exfiltration Over C2 Channel](#)

[T1008 Fallback Channels](#)

[T1105 Ingress Tool Transfer](#)

[T1059 Command and Scripting Interpreter](#)

[T1020 Automated Exfiltration](#)

Triggers

- An account has configured an internal resource for remote interaction through the use of a Power Automate HTTP Connector.

Possible Root Causes

- An attacker is leveraging Power Automate HTTP connectors to extend malicious access into internal resources.
- In rare cases, a Power Automate HTTP connector is used to enable legitimate external connectors which trigger approved internal actions.

Business Impact

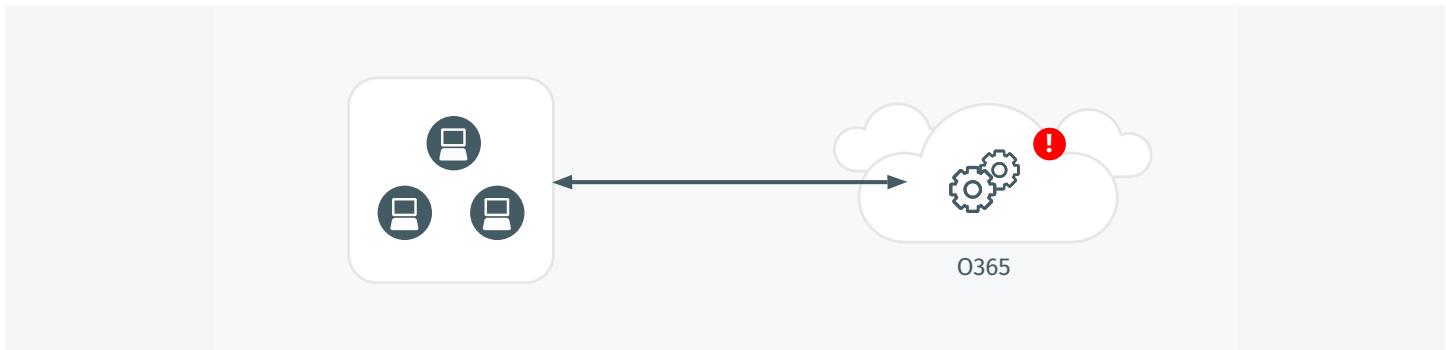
- Adversaries using this technique may gain malicious access to a wide range of internal resources including forms, pages, files, and emails.
- Use of this technique allows an adversary to bypass login and MFA requirements once the Power Automate flow is installed.

Steps to Verify

- Given the risk and relative rarity associated with Power Automate HTTP connectors, the legitimacy of associated flows should be investigated.

M365 Suspicious Power Automate Flow Creation

Command & Control



MITRE | ATT&CK®

[T1041 Exfiltration Over C2 Channel](#)

[T1008 Fallback Channels](#)

[T1105 Ingress Tool Transfer](#)

[T1059 Command and Scripting Interpreter](#)

[T1020 Automated Exfiltration](#)

Triggers

- Power Automate Flow creation has been observed by a user not typically associated with this activity.

Possible Root Causes

- An adversary has leveraged Power Automate as a persistence mechanism inside the environment.
- One of a small set of users who are authorized to perform Power Automate Flow creation has been observed doing so.

Business Impact

- Adversaries using this technique may gain malicious access to a wide range of internal resources including forms, pages, files, and emails.
- Use of this technique may enable persistence or lateral movement, or may be used to establish a means for subsequent data exfiltration.

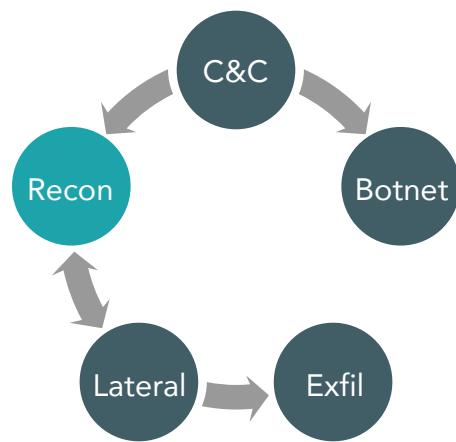
Steps to Verify

- Power Automate activities from unauthorized users should be immediately investigated
- Users authorized for Power Automate activities should be explicitly triaged in this system to avoid future detections.

Category

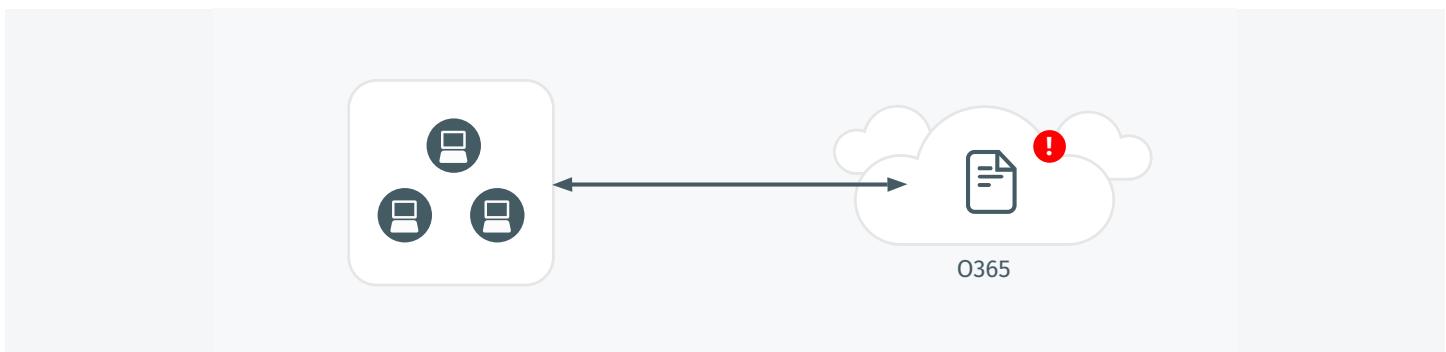
Reconnaissance

- A host or account is mapping out the inside of your network or cloud environment
- The activity may indicate that this is a targeted attack
- Detection types cover fast scans and slow scans
 - your vulnerability scanner will show up here as it performs much the same activity as an attacker



M365 Suspicious Compliance Search

Reconnaissance



MITRE | ATT&CK®

[T1119 Automated Collection](#)

[T1213 Data from Information Repositories](#)

[T1083 File and Directory Discovery](#)

Triggers

- The Exchange compliance search functionality was observed being used by an account that does not normally use this functionality.

Possible Root Causes

- Attackers may use compliance searches to search across Exchange mailboxes for sensitive data to collect and exfiltrate.
- Some internal users may use compliance searches to support legitimate business operations like legal and HR for litigation, audit, and compliance purposes.

Business Impact

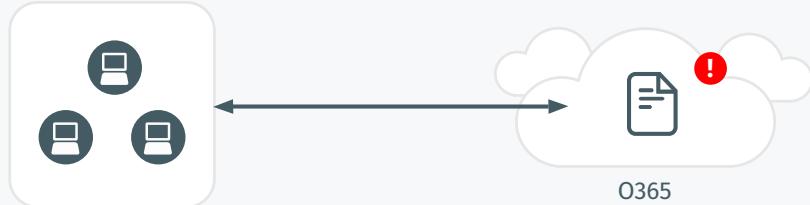
- Compliance search capabilities provide an enticing target for adversaries to abuse and may result in the loss of sensitive information up to and including passwords, encryption keys, and even financial data or intellectual property.

Steps to Verify

- Review the account in question to ensure they should be issuing compliance searches within the environment.
- Review the search being done to determine if the data being sought may be particularly interesting to attackers.
- Contact the user to ensure the searches are being done in compliance with company policy.

M365 Unusual eDiscovery Search

Reconnaissance



MITRE | ATT&CK®

[T1119 Automated Collection](#)

[T1213 Data from Information Repositories](#)

[T1083 File and Directory Discovery](#)

Triggers

- A user is creating or updating an eDiscovery search.

Possible Root Causes

- An adversary has gained access to eDiscovery capabilities and is using that access to perform reconnaissance across the environment.
- One of a small set of users authorized to perform eDiscovery has been observed doing so.

Business Impact

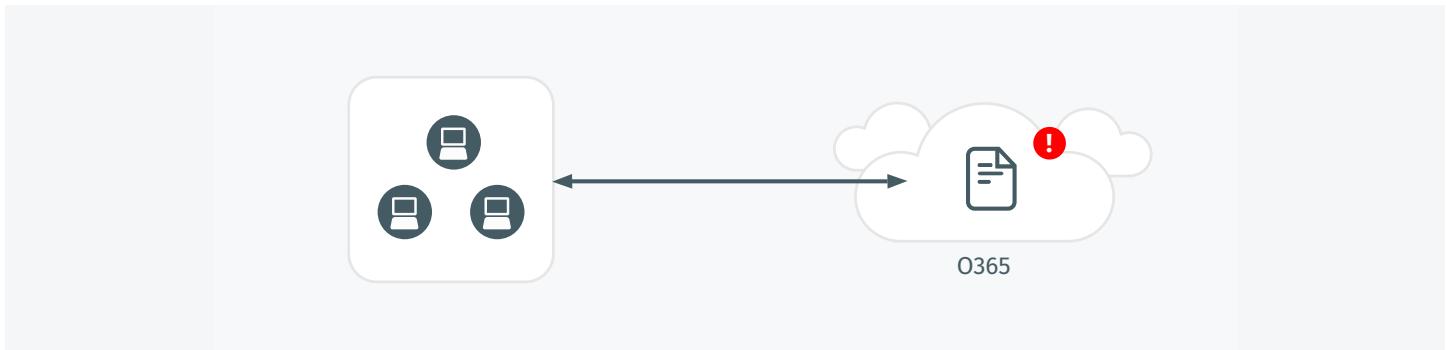
- eDiscovery capabilities provide an enticing target for adversaries to abuse and may result in the loss of sensitive information up to and including passwords, encryption keys, and even financial data or intellectual property.
- eDiscovery capabilities may include data traditionally inaccessible through other means but preserved as part of a litigation hold.

Steps to Verify

- eDiscovery search from unauthorized users should be immediately investigated.
- Users authorized for eDiscovery should be explicitly triaged in this system to avoid future detections.

M365 Suspect eDiscovery Usage

Reconnaissance



MITRE | ATT&CK®

[T1119 Automated Collection](#)

[T1213 Data from Information Repositories](#)

[T1083 File and Directory Discovery](#)

[T1562 Impair Defenses](#)

Triggers

- Behaviors commonly associated with covering up a potentially malicious eDiscovery search have been observed.

Possible Root Causes

- An attacker has compromised the eDiscovery system, is using it to actively collect and exfiltrate data, and is hiding their tracks.
- A legitimate user has abused the eDiscovery system to gain information and has deleted the search quickly to go unnoticed.
- An improperly created eDiscovery Search has been flagged for removal based on deviation from enterprise policies on accepted eDiscovery usage.
- An authorized test of the eDiscovery system has been observed and clean up actions from that test have been flagged as suspicious.

Business Impact

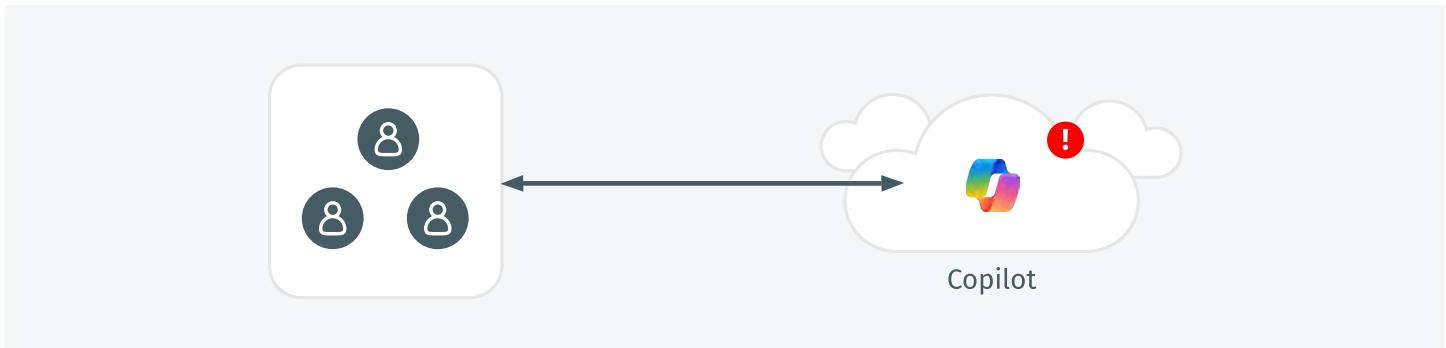
- eDiscovery search capabilities provide an enticing target for adversaries to abuse and may result in the loss of sensitive information up to and including passwords, encryption keys, and even financial data or intellectual property.
- Abuse of eDiscovery search could result in sensitive data exfiltration as well as advancing an attack deeper into the organization.

Steps to Verify

- Review the account in question to ensure they should be issuing compliance searches within the environment.
- Review any remaining and undeleted artifacts associated the search being done to determine if the data being sought may be particularly interesting to attackers.
- Contact the user to ensure the searches are being done in compliance with company policy.

M365 Suspicious Copilot for M365 Access

Reconnaissance



MITRE | ATT&CK®

T1213 Data from Information Repositories

T1114 Email Collection

Triggers

- A Copilot for M365 session was initiated by a user originating from a location that is unusual for the user and/or environment within the context of this functionality.

Possible Root Causes

- An attacker may be using the Copilot for M365 functionality to simplify their ability to discover knowledge documented within your environment that can help them enable their next steps within your environment (i.e. IT policies and procedures, documented static passwords/accounts, etc.).
- An attacker may be using the Copilot for M365 functionality to simplify the discovery and extraction of sensitive information from e-mails stored within the M365 environment.
- A legitimate user has accessed this functionality from a location that is not typical for your environment, but is using the functionality for benign/approved use cases.

Business Impact

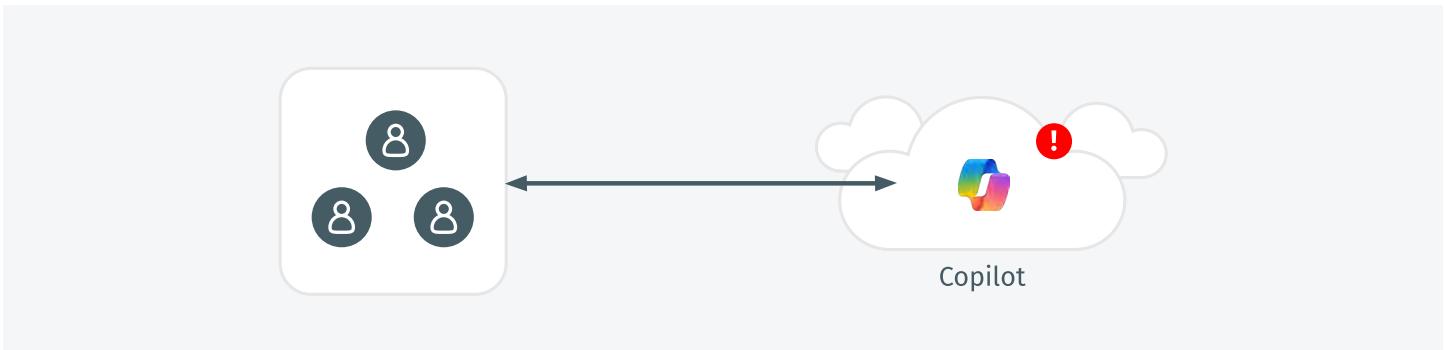
- An attacker utilizing Copilot for M365 can simplify the process of mining important knowledge about your organization and hide files that were accessed to support gaining that knowledge. This is because Copilot for M365 does not always log the files accessed to provide a response.

Steps to Verify

- This detection is most interesting when it is accompanied by other detections indicating this account may be compromised.
- Review whether the unusual location aligns with what is expected for this user.
- Consult the available logs to determine if the activity prior to the registration is as expected.
- If warranted, reach out to the account owner to confirm they accessed this functionality in this way.

M365 Copilot Sensitive Data Discovery

Reconnaissance



MITRE | ATT&CK®

[T1213 Data from Information Repositories](#)

[T1530 Data From Cloud Storage](#)

Triggers

- A Copilot for M365 session was leveraged by an identity to access file(s) that may contain sensitive information.

Possible Root Causes

- An attacker may be using a compromised account in the environment to get Copilot to reveal sensitive information.
- A legitimate user interacting with Copilot inadvertently reveals sensitive information.

Business Impact

- Malicious activities leveraging Copilot may result in leaks of sensitive or proprietary information.
- An attacker may use this method to explore and acquire backdoor access to the environment.
- Attackers may also tamper with files that were accessed using Copilot for M365 to exploit gaps in logging of file modifications.

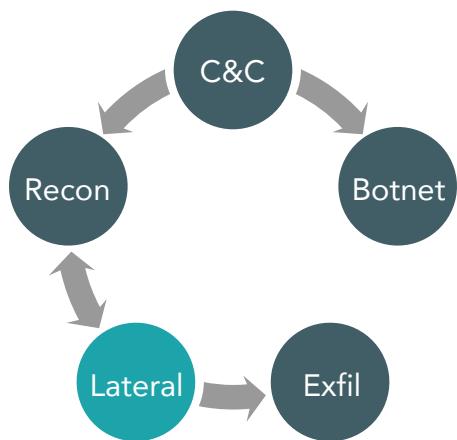
Steps to Verify

- Leverage Copilot logging activity to determine which files were accessed by the identity. Review them for sensitive content.
- If warranted, use the eDiscovery functionality to review Copilot prompts and responses.
- If review indicates suspicious behavior, reach out to the account owner to confirm if their interaction with Copilot was legitimate.

Category

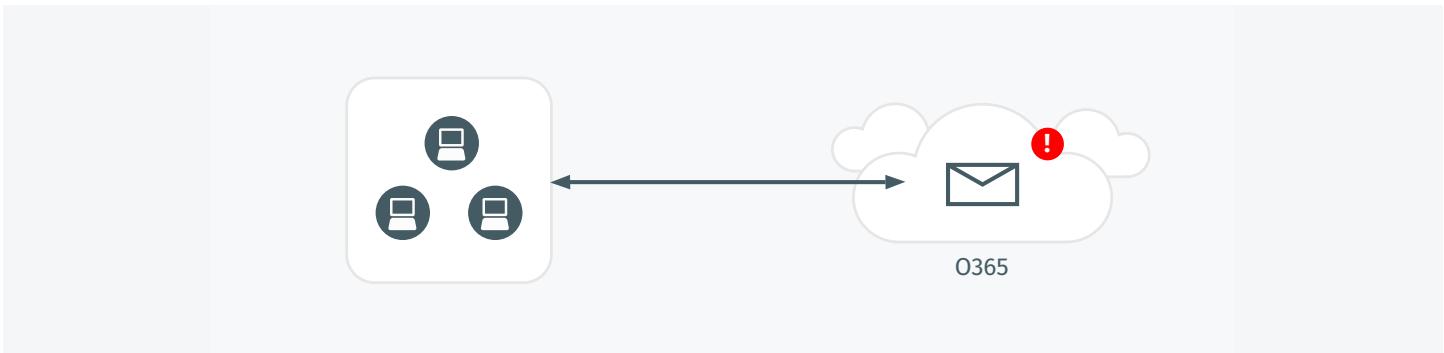
Lateral Movement

- Covers scenarios of lateral action meant to further a targeted attack
- This can involve attempts to steal account credentials or to steal data from another machine
- It can also involve compromising another host or account to make the attacker's foothold more durable or to get closer to target data



M365 Attacker Tool: Ruler

Lateral Movement



MITRE | ATT&CK®

[T1114 Email Collection](#)

[T1137 Office Application Startup](#)

Triggers

- The Ruler attack tool has been observed.

Possible Root Causes

- An adversary has used compromised account credentials in conjunction with the Ruler attack tool to enable malicious code or command execution.
- As this is a known attacker tool, there are no non-malicious use cases.

Business Impact

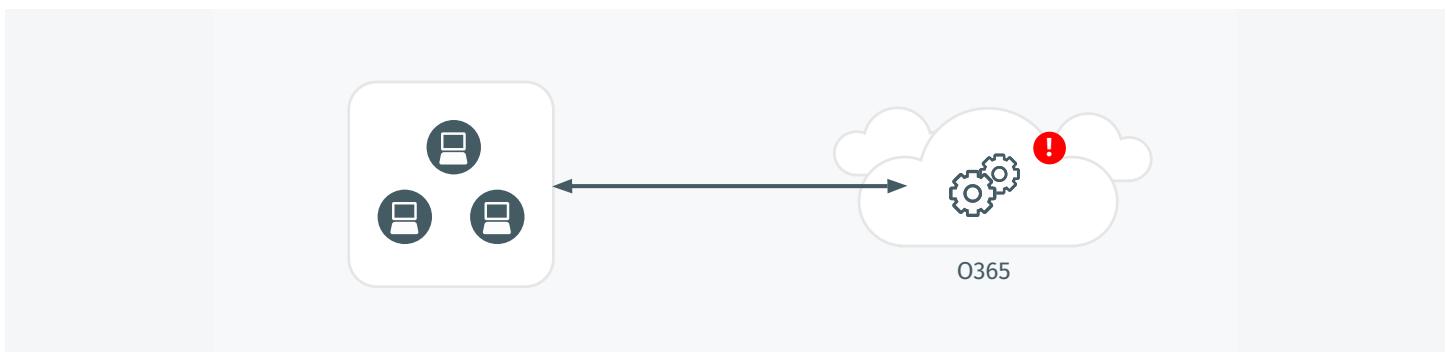
- Use of this tool may allow an adversary to install malware or execute commands on the endpoint running the exchange client associated with this compromised account. Malware or arbitrary command execution may be used for a variety of malicious activities, such as additional credential compromise, data collection and exfiltration, or to further attack progression.

Steps to Verify

- Investigate the compromised account for additional malicious actions and respond according to findings.

M365 Disabling of Security Tools

Lateral Movement



MITRE | ATT&CK®

T1562 Impair Defenses

Triggers

- Activities which weaken or disable Office 365 protective security features and tools.

Possible Root Causes

- Attackers will attempt to disable or downgrade Office 365 security mechanisms to blind defenders or to enable further malicious activities without the risk of detection.
- In some cases, administrators may disable security mechanisms while troubleshooting problems.

Business Impact

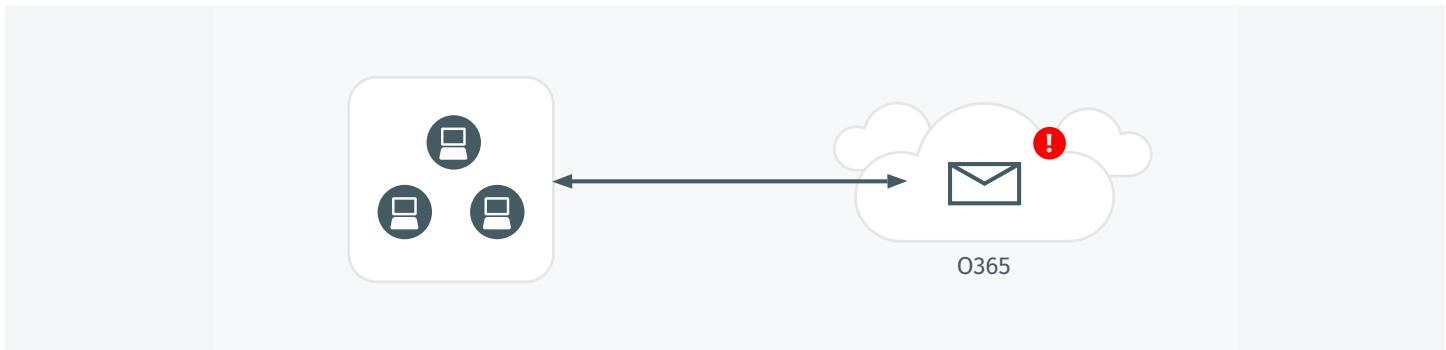
- Attackers who have successfully degraded, disabled, or bypassed security controls can more easily progress towards their objectives.
- Degraded or disabled security controls increase the potential impact of both present and future attacks against the organization.

Steps to Verify

- Review if this configuration is expected and appropriate in light of any available compensating controls.
- If this is a temporary configuration for troubleshooting purposes, confirm it has been reenabled once that troubleshooting is complete.

M365 DLL Hijacking Activity

Lateral Movement



MITRE | ATT&CK®

T1574 Hijack Execution Flow

Triggers

- An account that may not download DLLs typically has been observed downloading a DLL file under conditions that highlight the risk of DLL hijacking, such as both a non-DLL and DLL file being downloaded from the same directory in a short time frame.

Possible Root Causes

- An attacker has abused the way applications search for DLLs by placing a malicious DLL file into a shared directory with the intention of compromising any endpoint that loads the malicious DLL file rather than the intended application DLL file.
- In some cases, developers collaborating from a cloud hosted repository could intentionally download and access DLLs this way.

Business Impact

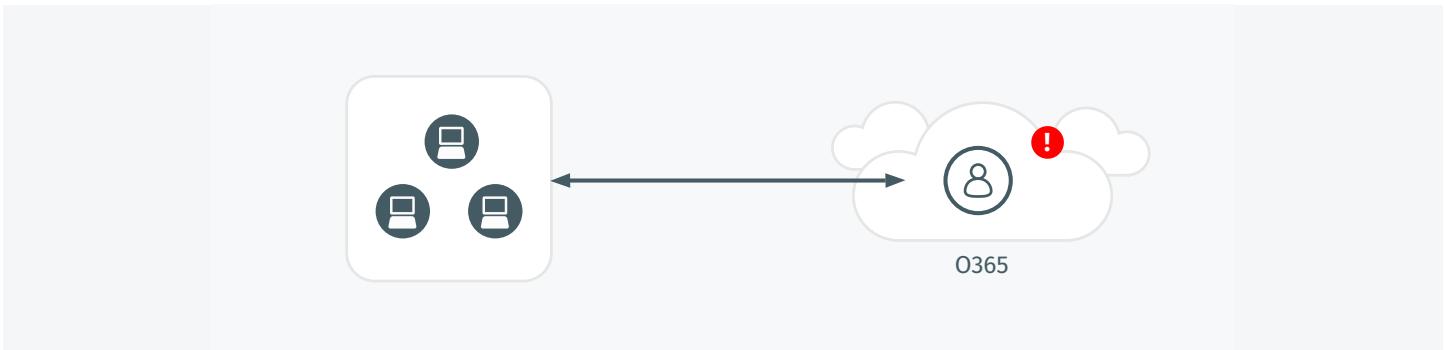
- DLL Hijacking may result in the complete compromise of a targeted system, and associated accounts and data.
- Endpoints compromised through DLL Hijacking give an attacker an additional foothold in the environment and an opportunity for additional lateral movement, increasing the risk of impact to enterprise systems, users, and data.

Steps to Verify

- Investigate the user associated with this action, and verify if this user would be downloading DLL files as part of their expected workflows.
- Investigate presence of additional files accessed as part of this detection, and assess if this is indicative of an authorized remote application, used for legitimate business purposes.

M365 External Teams Access

Lateral Movement



MITRE | ATT&CK®

T1213 Data from Information Repositories

Triggers

- A new team member has been added to a team in O365 Teams consisting of an external account from a domain rarely associated with O365 Teams access.

Possible Root Causes

- An adversary has added an external account under their control as a new member of a team by abusing an existing O365 Teams account.
- Sometimes legitimate external users (such as partners, contractors, lawyers, auditors, etc.) are added to an O365 Team as part of an authorized activity.

Business Impact

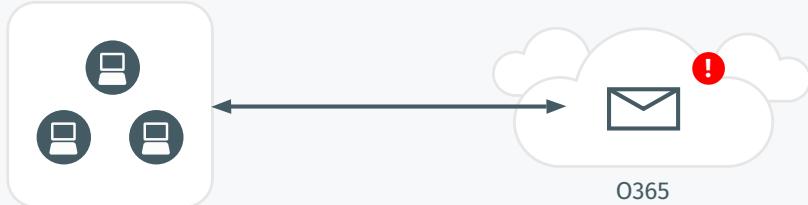
- This type of access enables an attacker to perform additional discovery or collection activities by exposing sensitive business information which may include shared files, meeting content, or chat transcripts.
- The impact of such access may include information necessary to enable further attack progression or facilitate the loss of proprietary information or intellectual property, and regulated data.
- In some cases, access to the team's communication fabric and conversation history can enable successful blackmail or extortion against enterprise personnel.

Steps to Verify

- Validate that the account added is an authorized member of the O365 Team.

M365 Internal Spearphishing

Lateral Movement



MITRE | ATT&CK®

T1534 Internal Spearphishing

Triggers

- A user was observed sending multiple emails to internal recipients which were flagged by O365 reputation scanning as likely phishing emails.

Possible Root Causes

- An attacker has compromised a single account and is abusing its access and implicit trust within an organization to attack additional accounts via spearphishing emails.
- Benign emails have been flagged as suspicious based on their content or attachments, which are most frequently associated with invoices sent to distribution lists.

Business Impact

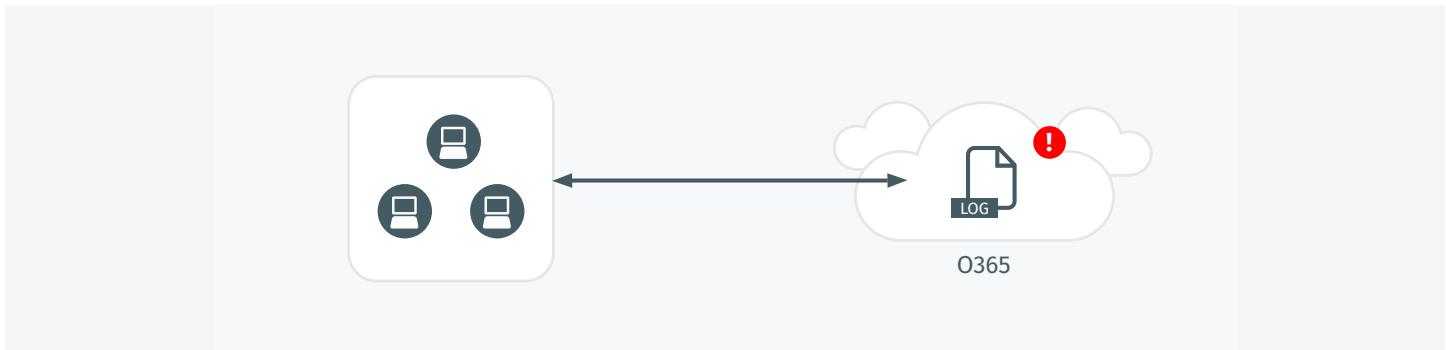
- Spearphishing is one of the predominant ways attackers gain and expand access to credentials within an environment and is particularly effective when utilizing the implicit trust of an internal sender.
- Successful internal spearphishing campaigns result in broad access to a large range of resources within the environment, resulting in a significant increase in overall impact of a compromised account incident within an organization.

Steps to Verify

- Review the details and contents of the email to validate it is malicious.
- Review additional detections and events by the source user which may indicate their account has been compromised.
- Validate the source user is aware of and sent the email that was flagged.

M365 Log Disabling Attempt

Lateral Movement



MITRE | ATT&CK®

T1562 Impair Defenses

Triggers

- An attempt has been made to disable important Office 365 logs that enhance security.

Possible Root Causes

- Attackers will seek to disable logging to blind detection mechanisms and cover their tracks.
- Logging may be temporarily turned off by an admin while changing configuration or troubleshooting a problem.

Business Impact

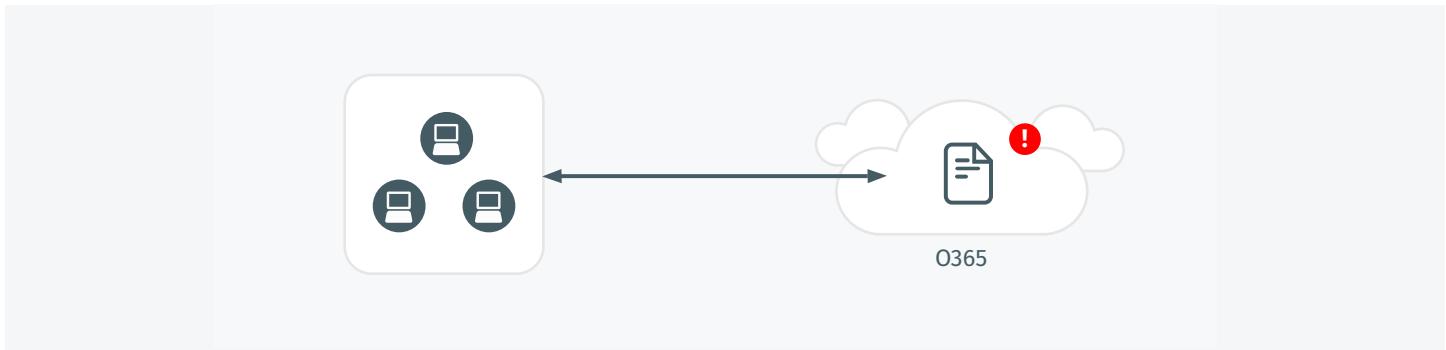
- An attacker who has disabled logging may progress parts of an attack without being detected, and without producing an auditable record to aid in forensics.
- Disabling logging degrades a critical component of an organization's security architecture.
- Many audit and compliance requirements can only be met through the collection of activity logs.

Steps to Verify

- Review whether this logging configuration is expected and appropriate.
- If this is a temporary configuration for troubleshooting purposes, confirm it has been reenabled once that troubleshooting is complete.

M365 Malware Stage: Upload

Lateral Movement



MITRE | ATT&CK®

T1203 Exploitation for Client Execution

Triggers

- Files which were subsequently flagged as malware were uploaded into the environment by this account.

Possible Root Causes

- Attackers will stage malicious files in preparation for an attempt to infect other users from a trusted file repository.
- On rare occasions, benign files may be classified as malicious.

Business Impact

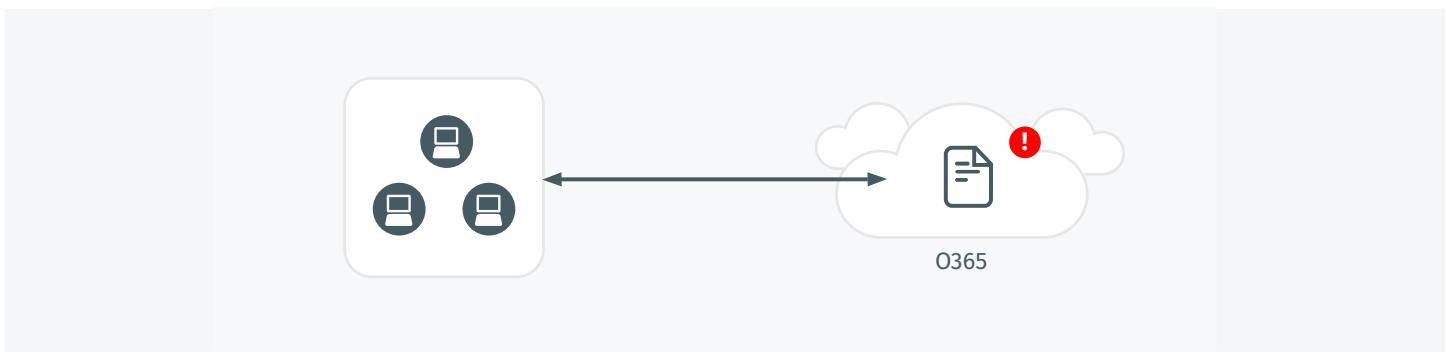
- An attacker who has disabled logging may progress parts of an attack without being detected, and without producing an auditable record to aid in forensics.
- Disabling logging degrades a critical component of an organization's security architecture.
- Many audit and compliance requirements can only be met through the collection of activity logs.

Steps to Verify

- Review whether this logging configuration is expected and appropriate.
- If this is a temporary configuration for troubleshooting purposes, confirm it has been reenabled once that troubleshooting is complete.

M365 Ransomware

Lateral Movement



MITRE | ATT&CK®

T1486 Data Encrypted for Impact

Triggers

- A series of file modifications typically associated with ransomware.

Possible Root Causes

- An account is being used to access an organization's cloud storage and encrypt and rewrite files.
- In some cases, automated jobs or services that perform widespread file renaming may trigger this detection.

Business Impact

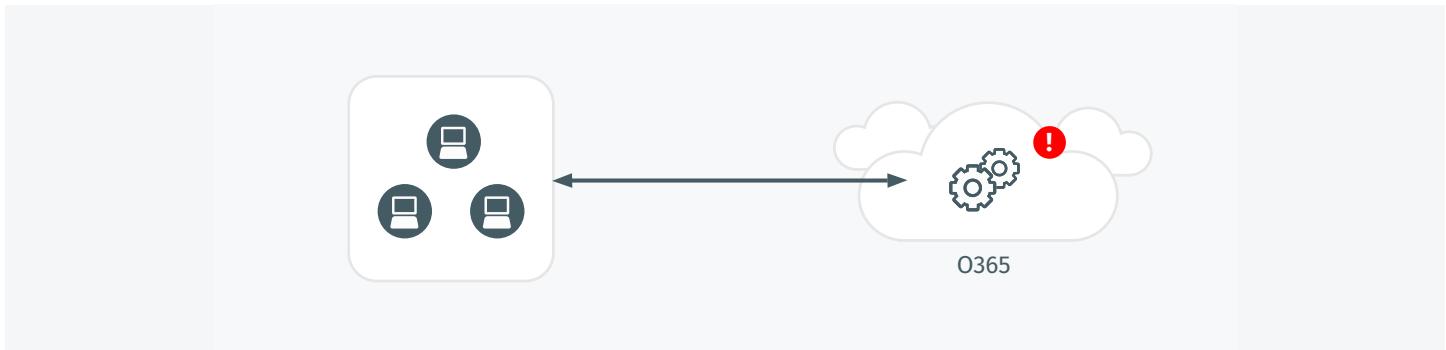
- Ransomware attacks directly impact access to the organization's data and are popular among attackers due to the possibility of a quick transition from attack to monetization.
- After files have been encrypted, the attacker will ask the organization to pay a ransom in return for a promise to provide the encryption key which would allow the files to be decrypted.
- Even if an organization is willing to pay the ransom, there is no guarantee that the encryption key will be provided by the attacker or that the decryption process will work.
- Absent the encryption key, an organization must rely on restoration of files from backups.

Steps to Verify

- Review the integrity of the affected files and determine whether they appear encrypted.

M365 Risky Exchange Operation

Lateral Movement



MITRE | ATT&CK®

T1484 Group Policy Modification

T1098 Account Manipulation

Triggers

- High risk Exchange operations which range from allowing the exfiltration of data, the creation of backdoor rules, execution of VBS scripts, or forwarding and collecting sensitive information.

Possible Root Causes

- An attacker is manipulating Exchange to gain access to a specific set of data or to enable continued attack progression.
- In some cases, these operations may be authorized activities for a small set of highly privileged users who perform them so infrequently that they are outside what the detection model considers normal.
- Authorized configurations in cases of a permanent employee separation or temporary leave of absence may involve activities that would otherwise compromise mailbox integrity.

Business Impact

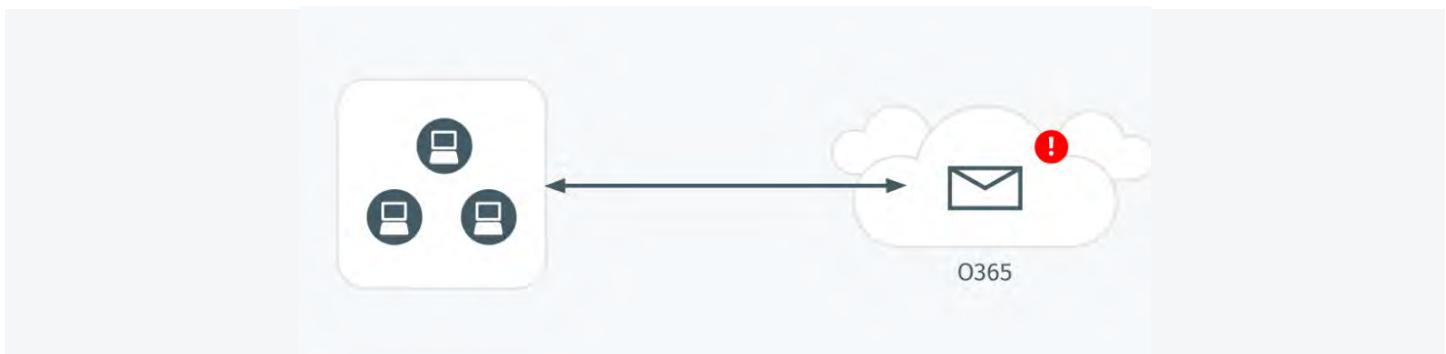
- Sensitive data and content may be contained within Exchange which may be useful or desirable to an adversary.
- Compromising Exchange may allow an attacker to continue their attack progression.

Steps to Verify

- Verify whether these changes to the configurations are intentional and have been made with appropriate compensating safeguards.

M365 Phishing Simulation Configuration Change

Lateral Movement



MITRE | ATT&CK®

[T1534 Internal Spearphishing](#)

[T1562 Impair Defenses](#)

Triggers

- A new phishing simulation provider is defined in a Microsoft 365 environment.

Possible Root Causes

- A malicious entity is designated to be a phishing simulation provider.
- A phishing simulation business is onboarded into the M365 environment in order to provide security training for the employees.

Business Impact

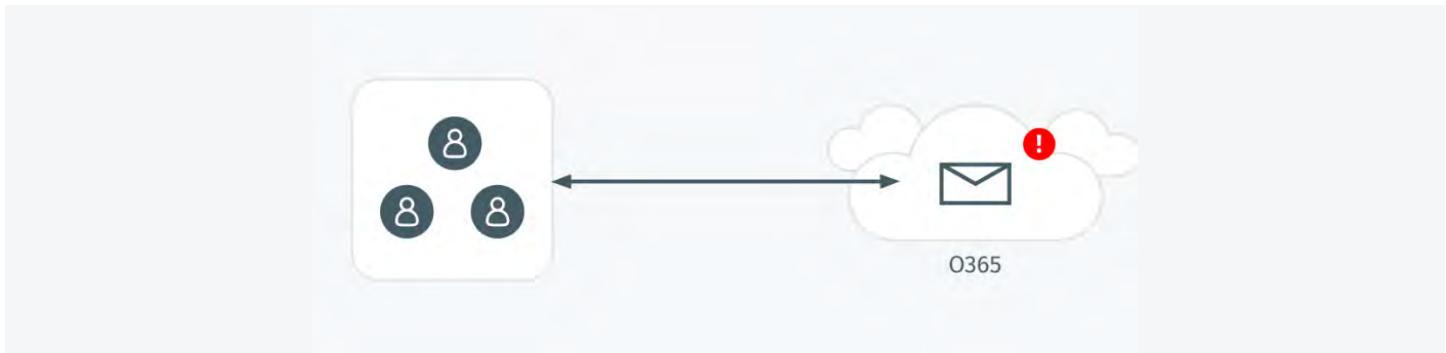
- A phishing simulation provider is exempt from all e-mail phishing checks in Microsoft 365. If a malicious provider is defined, any malicious e-mail sent by it will not be stopped by security filters.

Steps to Verify

- Verify if there was a legitimate phishing provider onboarded into the M365 environment.

M365 SecOps Mailbox Change

Lateral Movement



MITRE | ATT&CK®

[T1534 Internal Spearphishing](#)

[T1562 Impair Defenses](#)

Triggers

- One or more SecOps accounts are set up in Microsoft 365.

Possible Root Causes

- An attacker designates a victim account as a SecOps account in Microsoft 365.
- An account is designated as a SecOps account following a legitimate business requirement.

Business Impact

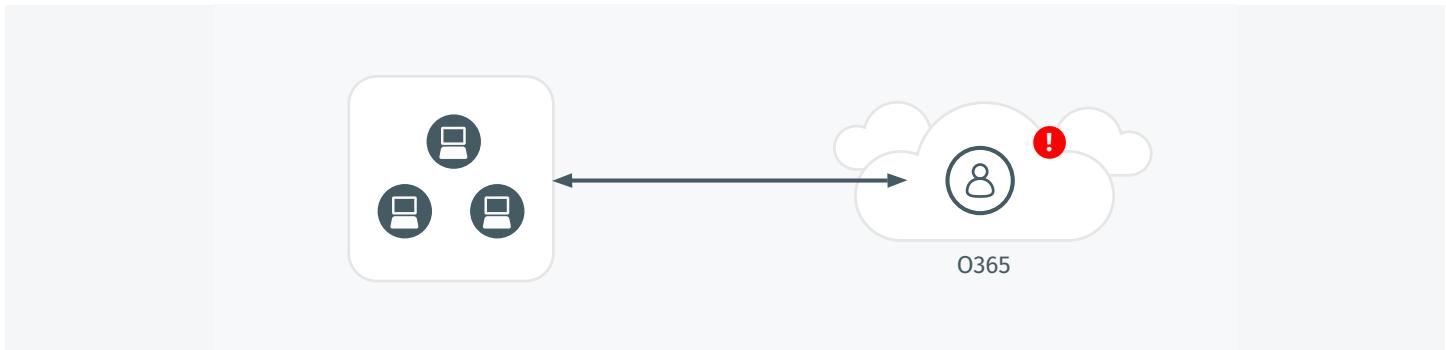
- SecOps accounts are exempt from spam and malware filtering. By defining a victim's account as a SecOps account, attackers can send the victim malicious e-mails without any obstruction from Microsoft 365.

Steps to Verify

- Determine if there was a legitimate business reason to designate a user's account to be a SecOps account.

M365 Suspicious Teams Application

Lateral Movement



MITRE | ATT&CK®

[T1550 Use Alternate Authentication Material](#)

[T1528 Steal Application Access Token](#)

Triggers

- A rarely used, third-party Microsoft Teams integrated application has been granted excessive or risky permissions that may enable malicious activities to be taken on behalf of the authorizing user.

Possible Root Causes

- An attacker is trying to trick the user into authorizing a third-party app that will allow the attacker to execute malicious actions.
- In some cases rare, legitimate applications do require a set of permissions that are authorized despite the risk they present.

Business Impact

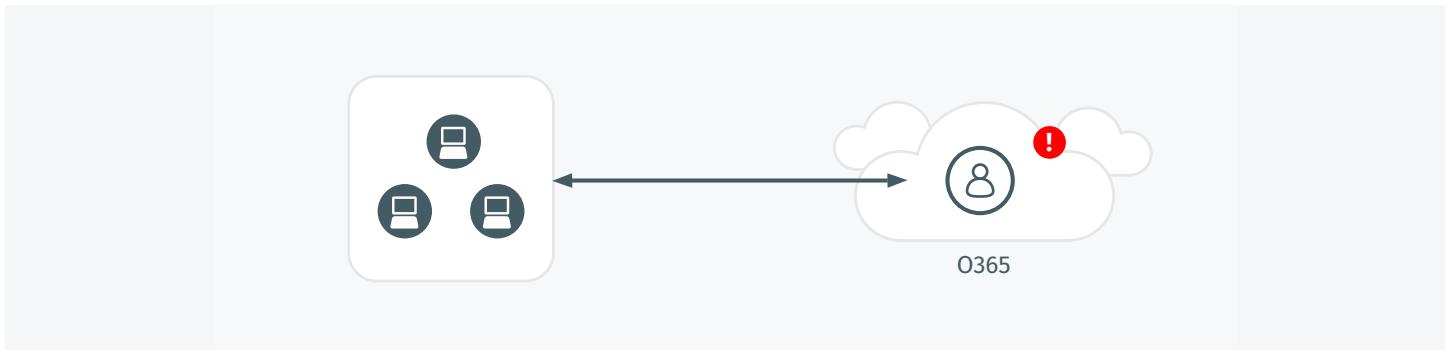
- Malicious third-party apps can be used to undermine existing security controls, such as multi-factor authentication (MFA), and enable malicious action on behalf of the authorizing user, increasing risk to enterprise system and data and increasing the likelihood of further attack progression.
- A suspicious teams application could result in outcomes ranging from the compromise of an individual account or host, to broader compromise of a full teams channel.
- Malicious apps may enable a foothold into the environment as a means of maintaining persistent access.
- Malicious apps could allow the collection of sensitive information or act as a mechanism to support data exfiltration.

Steps to Verify

- Verify that the application in question is authorized for the associated user.
- Validate that the required permission set is appropriate for the authorized business process associated with this application.
- Investigate for additional malicious indicators associated with this application or user.

M365 Suspicious Mailbox Manipulation

Lateral Movement



MITRE | ATT&CK®

T1098 Account Manipulation

Triggers

- Access has been granted to more resources than a user has had historically and has occurred outside of learned administrator behaviors.

Possible Root Causes

- An attacker has escalated the account's Exchange access rights to enable business email compromise or the collection of additional information to aid in the next step of the attack.
- Employee life-cycle activities such as permanent separation or temporary leaves of absence may legitimately require mailbox modifications which could trigger this detection.
- Some service-specific mailboxes are intentionally granted these permissions.

Business Impact

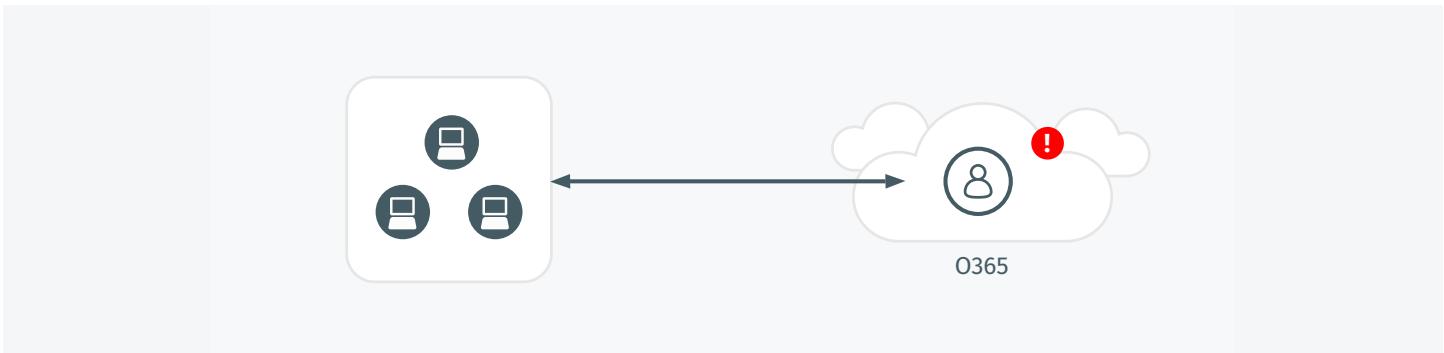
- Sensitive data and content may be contained within Exchange which may be useful or desirable to an adversary.
- Data may leak from a user's mailbox by being transmitted to unauthorized entities.

Steps to Verify

- Validate that the permissions granted are appropriate to the entity in question.

M365 Suspicious Mailbox Rule Creation

Lateral Movement



MITRE | ATT&CK®

T1564 Hide Artifacts

Triggers

- An account was observed creating suspicious mailbox rules in Exchange that allow an attacker to manipulate, hide, or delete incoming emails.

Possible Root Causes

- An attacker with control of an account created mailbox rules that hide or manipulate emails to either evade notice by the mailbox owner or impact business processes.
- A user created a benign but broad or abnormal inbox rule as part of normal business email management.

Business Impact

- Instances of malicious mailbox rules may indicate an adversary has control of an internal mailbox and can access the user's email data and send emails internally and externally on behalf of the user.
- A successful attack can result in immediate data theft or reputation loss from the compromised account.
- A successful attack can result in additional business impact through targeted phishing from the internal account, as they are often trusted and subsequent to less strict security controls relative to external accounts.

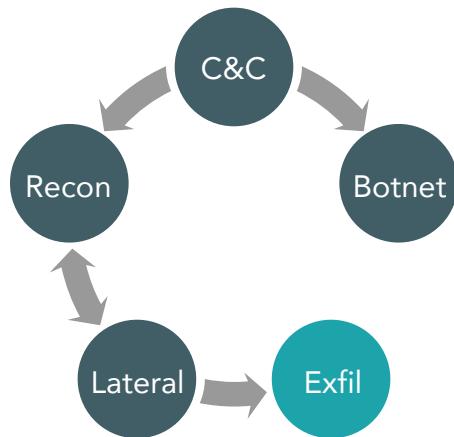
Steps to Verify

- Investigate the account that performed the action for other indications of malicious activity
- If review indicates possible malicious actions, revert configuration and disable credentials associated with this alert, then perform a comprehensive investigation.

Category

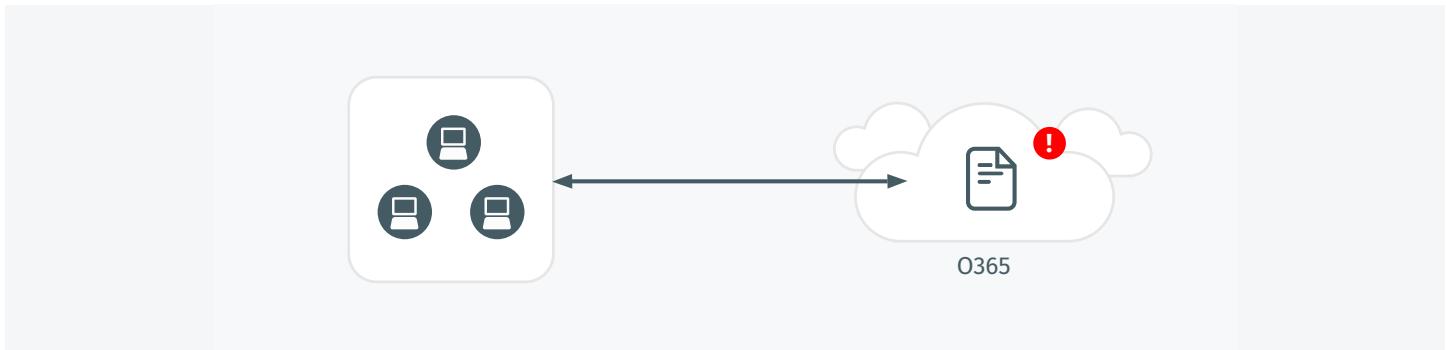
Exfiltration

- Covers scenarios where data is being sent outside or collected in a way meant to hide the data transfer
- While data is constantly being sent out of the network or cloud environment, it usually does not involve the use of techniques meant to hide the transfer
- The host or account transmitting the data, where it is transmitting the data, the amount of data and the technique used to send it all provide indicators of exfiltration



M365 eDiscovery Exfil

Exfiltration



MITRE | ATT&CK®

T1048 Exfiltration Over Alternative Protocol

Triggers

- A user is previewing or downloading the results of an eDiscovery activity.

Possible Root Causes

- An adversary has gained access to eDiscovery capabilities and is using that access to collect or exfiltrate data.
- One of a small set of users authorized to perform eDiscovery has been observed doing so.

Business Impact

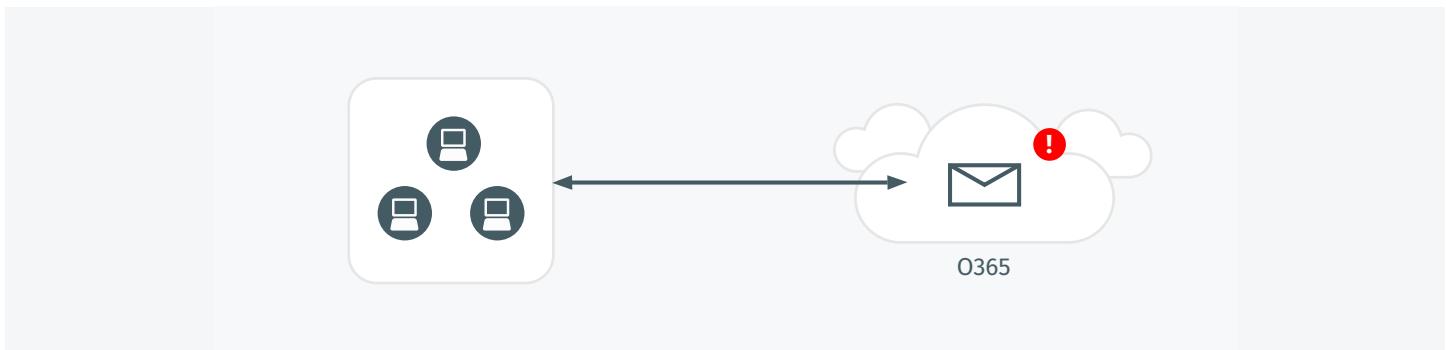
- eDiscovery capabilities provide an enticing target for adversaries to abuse and may result in the loss of sensitive information up to and including passwords, encryption keys, and even financial data or intellectual property.
- eDiscovery capabilities may include data traditionally inaccessible through other means but preserved as part of a litigation hold.

Steps to Verify

- eDiscovery activities from unauthorized users should be immediately investigated.
- Users authorized for eDiscovery should be explicitly triaged in this system to avoid future detections.

M365 Exfiltration Before Termination

Exfiltration



MITRE | ATT&CK®

T1213 Data from Information
Repositories

Triggers

- The risk of insider threat has been observed by an account downloading or exfiltrating files prior to that account being deleted or disabled.

Possible Root Causes

- A user with foreknowledge of separation or reassignment has intentionally acquired or stolen organizational data prior to departure with the intent to retain access to information or data for which they will no longer be authorized access.
- In some cases, suspicious data acquisition by a user prior to a separation or reassignment event may be part of an authorized activity.

Business Impact

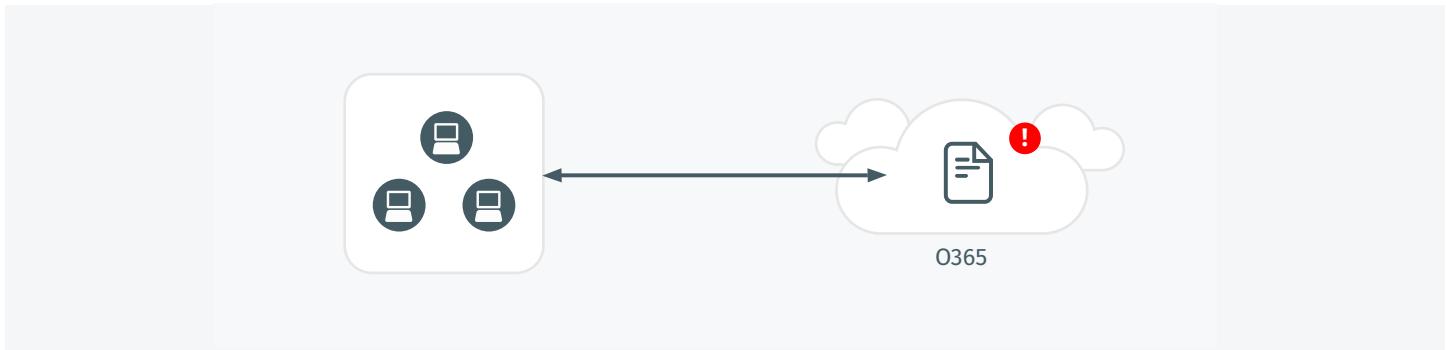
- Insider threat places an organization at risk of loss of sensitive information such as intellectual property, financial data, or other data associated with legal and compliance protections.
- The successful exfiltration of data by an insider may lead to regulatory fines or penalties, loss of competitive advantages, or other outcomes detrimental to business and organizational success.

Steps to Verify

- Investigate the reason this account was disabled or deleted, and if maintaining access to these files continues to be authorized.
- Investigate if the files associated with this detection include sensitive information.

M365 Suspicious Download Activity

Lateral Movement



MITRE | ATT&CK®

T1567 Exfiltration Over Web Service

Triggers

- An account was seen downloading an unusual number of objects compared to the user's past behavior or the behavior of other O365 users.

Possible Root Causes

- An attacker may be using SharePoint / OneDrive download functions to exfiltrate data.
- Users downloading an unusually large number of files as they start new projects, back up data or access multiple files to support their job function.

Business Impact

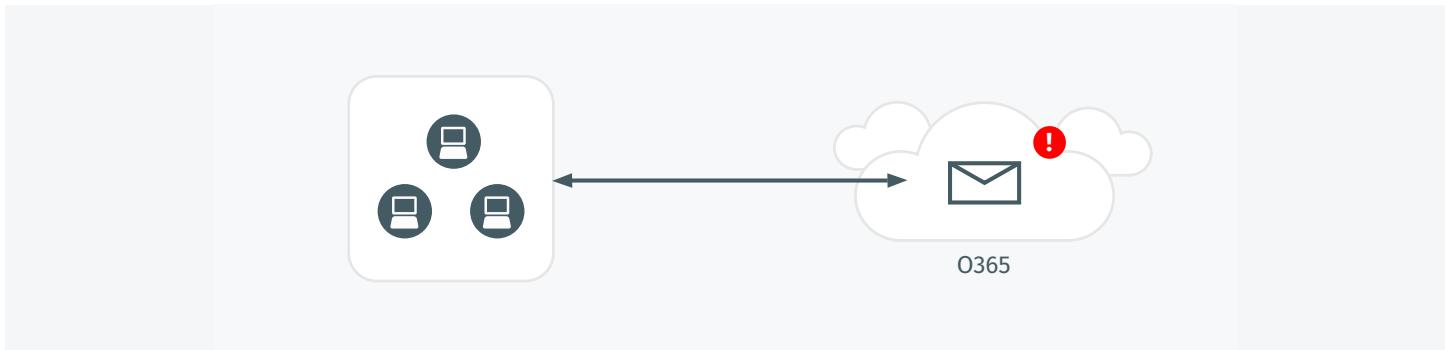
- Ability to exfiltrate a significant number of sensitive files from the enterprise is often the last stage of the security compromise.
- Exfiltration of sensitive business data may lead to loss of control of company secrets and intellectual property.

Steps to Verify

- Review the details and contents of the files to assess risk, and validate these are authorized downloads.
- Review additional detections and events by the source user which may indicate their account has been compromised.

M365 Suspicious Exchange Transport Rule

Exfiltration



MITRE | ATT&CK®

T1114 Email Collection

Triggers

- A new Exchange transport rule has been created with a potentially risky action that may provide email collection, exfiltration, or deletion capabilities (BlindCopyTo, CopyTo, Delete).

Possible Root Causes

- An attacker has gained Exchange administrator access with the capabilities of forwarding sensitive emails prior to their arrival in a user's inbox to an attacker controlled email address (internal or external).
- An attacker may be preparing to delete important emails prior to their arrival in a user's inbox to prevent important alerts or notifications from occurring.
- A legitimate transport rule was added to support business requirements or prevent dangerous emails from reaching user inboxes.

Business Impact

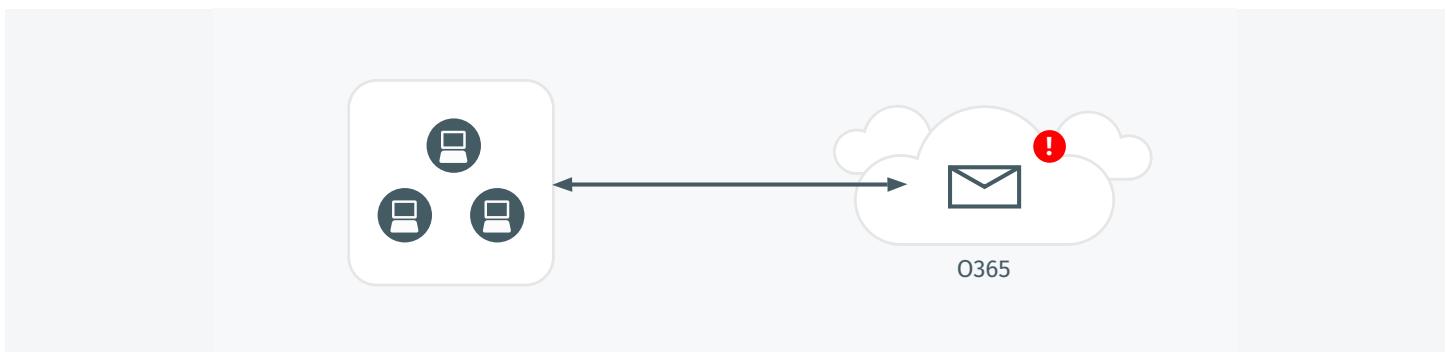
- Because email services are critical to so many enterprise activities, attackers prioritize access both as a means of progressing an attack as well as a mechanism for data exfiltration.
- Forwarded emails may expose sensitive data.
- Deleted emails may mask security alerts or important emails alerting an organization to a breach.
- The combination of forwarded and deleted emails may allow an external party to impersonate internal users to further their goals.

Steps to Verify

- Validate the new transport rule serves a business purpose, does not create a risk of data exposure, and has been implemented according to proper change control processes.

M365 Suspicious Mail Forwarding

Exfiltration



MITRE | ATT&CK®

T1114 Email Collection

Triggers

- Mail forwarding which may be used as a collection or exfiltration channel for an adversary has been observed.

Possible Root Causes

- An external attacker has established persistent access to contents of a specific mailbox without the need to otherwise maintain any kind of persistence through installing software.
- Employee life-cycle activities such as a permanent separation or a temporary leave of absence may legitimately require mailbox modifications which could trigger this detection.
- Emails belonging to executives may be forwarded to their associated administrative assistants.
- Emails for service accounts may be forwarded to the staff members who manage those services.

Business Impact

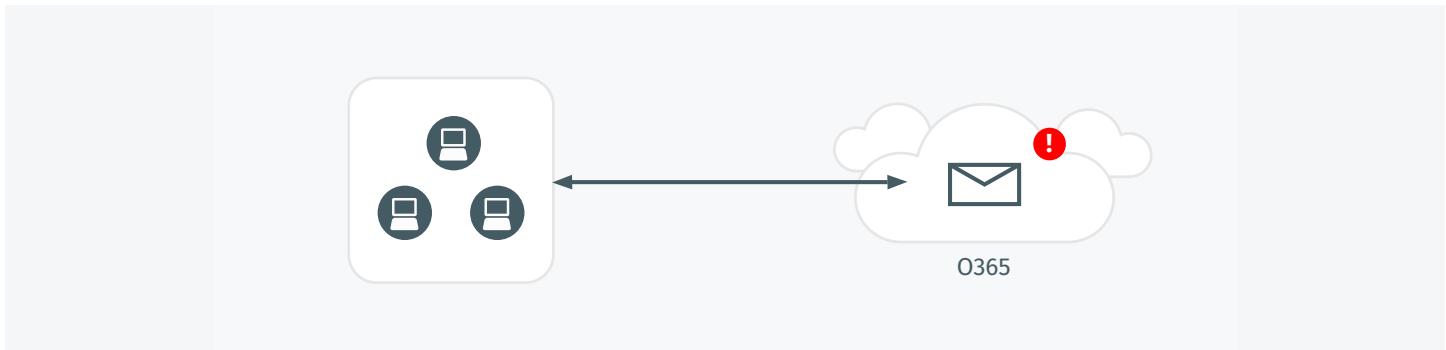
- Attackers who have gained persistence through the email systems may passively collect and exfiltrate data.
- Sensitive business information often resides in email systems and may be leaked through e-mail theft.

Steps to Verify

- Verify if sensitive data has been unintentionally forwarded using this feature.

M365 Suspect Power Automate Activity

Exfiltration



MITRE | ATT&CK®

[T1041 Exfiltration Over C2 Channel](#)

[T1008 Fallback Channels](#)

[T1059 Command and Scripting Interpreter](#)

[T1020 Automated Exfiltration](#)

Triggers

- Abnormal Power Automate activity was observed from a user in the environment.
- A user leveraged a Power Automate flow connector that was unusual for either the user or the environment.
- A user modified another user existing flow in a suspect manner.

Possible Root Causes

- An attacker may be creating automated tasks within the environment to secretly exfil, manipulate data for impact, or create network control channels.
- A normal user is attempting to subvert normal IT policies by leveraging native Microsoft infrastructure without authorization.
- One of a small set of users who are authorized to leverage Power Automate flow was observed doing so.

Business Impact

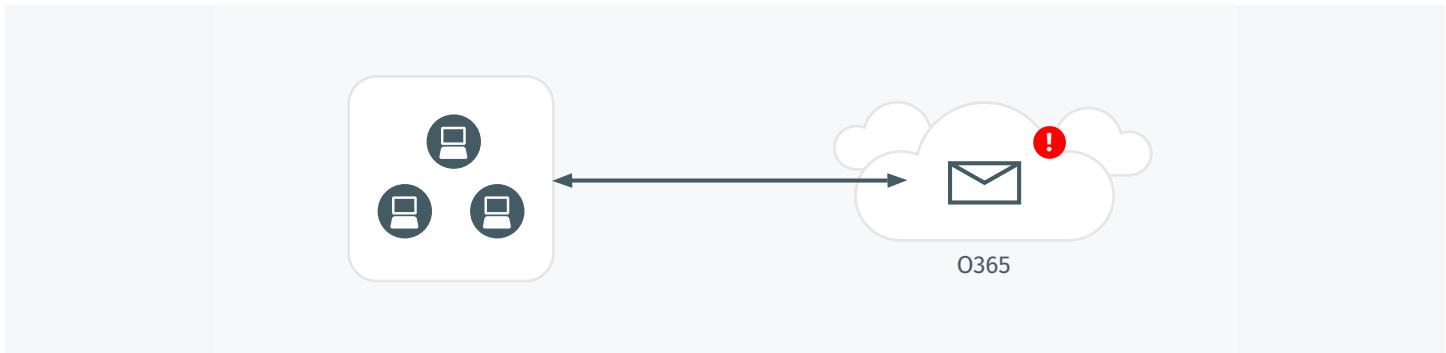
- Power Automate, Microsoft's native and on-by-default O365 automation tool, can be leveraged by attackers to interact directly with internal data and infrastructure to facilitate data exfil or attack automation.

Steps to Verify

- Power Automate activities involving unauthorized connectors should be investigated immediately.
- Users modifying other user's Power Automate flows should have explicit permission to do so.
- Users authorized for Power Automate activities should be explicitly triaged to avoid future detections.

M365 Suspicious Sharing Activity

Exfiltration



MITRE | ATT&CK®

T1213 Data from Information Repositories

Triggers

- An account was seen sharing files and/or folders at a volume that is higher than is normal for both the environment and for the account.

Possible Root Causes

- Attackers may use SharePoint/OneDrive sharing functions to exfiltrate data and enable ongoing access to data over extended periods of time.
- Use of sharing enables attackers to maintain access to data after a compromised account is remediated
- Users who rarely share files may periodically share more files than most other users in the environment as part of their job function.

Business Impact

- While some level of sharing may be normal for an environment or user, those users who emerge as sharing unusual amounts of data should be reviewed to validate the sharing is legitimate and does not pose a risk.
- Sharing of a large volume or breadth of files or folders exposes the organization to an increased risk of data theft or loss.

Steps to Verify

- Review the data being shared to determine if the information should be exposed to external parties.
- Review the sharing permissions to ensure the least possible data is exposed.
- Validate with the user that the sharing was intended and follows organizational policies on data sharing with external parties.

Category

Info

- Reports on new and novel events without directly impacting scoring
- New and novel events occur normally in most network and cloud environments and in most cases are not directly linked to threats
- Awareness of new and novel events support better situational awareness and provide additional context when observed with kill chain alerts

M365 Link Sent By External Teams User

- An external Teams user was observed sending a URL or attachments to internal users.
External users are not directly part of the organization (for example, contractors).
- This could be indicative of spearphishing/whaling campaigns or an attempt to infect target accounts with malware.

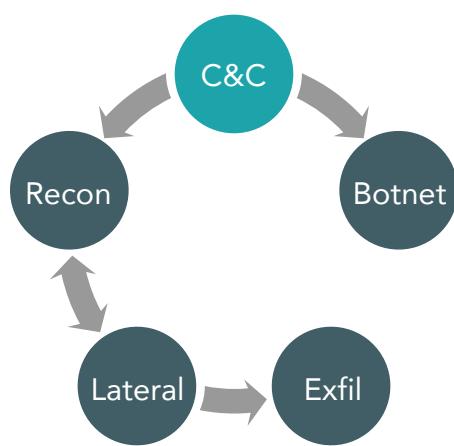
Azure AD Detections



Category

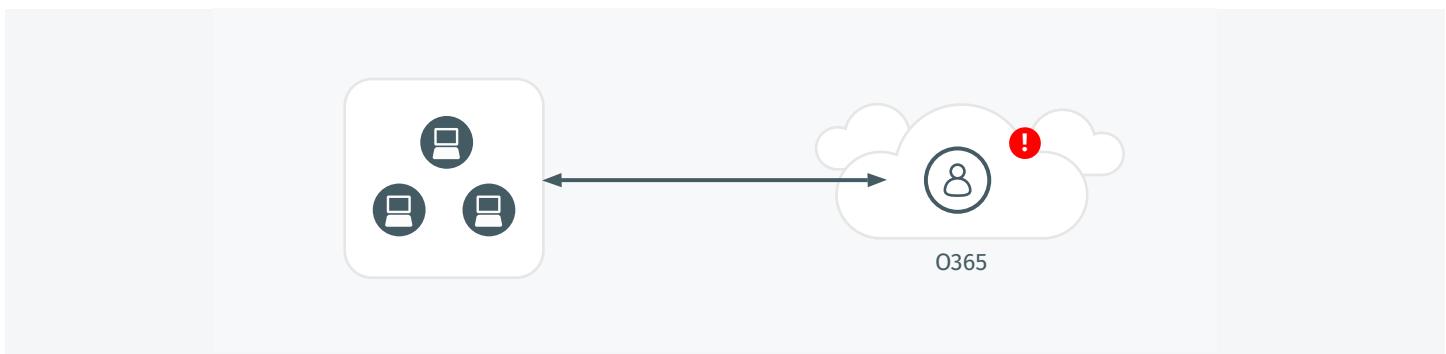
Command & Control

- A host or account appears to be under control of an external entity
- Most often, the control is automated as the host or account is part of a botnet or has adware or spyware installed
- The host or account may be manually controlled from the outside – this is the most threatening case and makes it highly likely that this is a targeted attack



Azure AD Admin Account Creation

Command & Control



MITRE | ATT&CK®

[T1528 Steal Application Access Token](#)

[T1550 Use Alternate Authentication Material](#)

Triggers

- An account has been created with administrative privileges (TenantAdmins, PrivilegedRoleAdmins, ApplicationAdministrators) that provide broad access to the environment.

Possible Root Causes

- An attacker that has gained administrative rights has added additional administrative accounts to the environment as a back-up access method if their existing access is disabled or otherwise removed at a future date.
- Existing legitimate administrators may add additional administrative users unintentionally or via social engineering.
- A new, legitimate, administrative account was added.

Business Impact

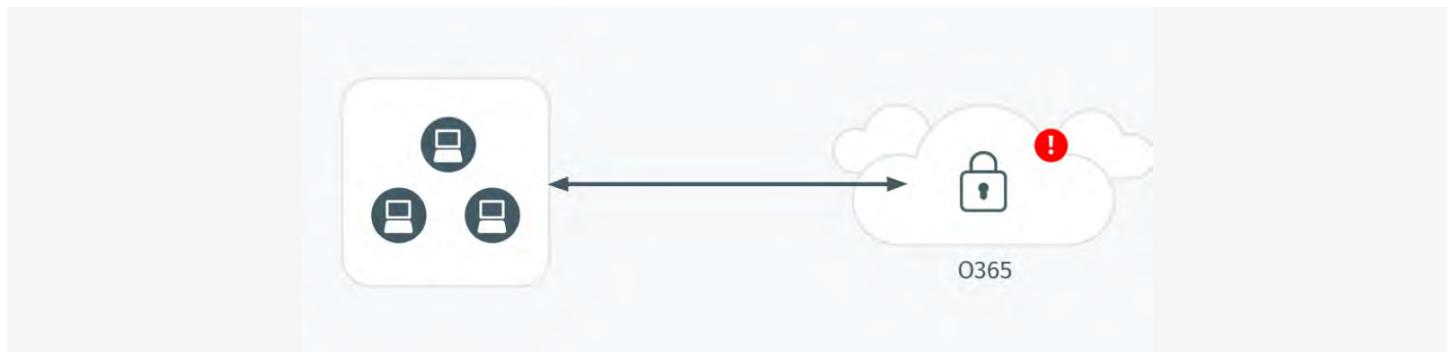
- Unauthorized administrative users have complete control within the environment, creating significant on-going risk to a broad range of resources.
- Attackers with access to the identified administrative rights will be able to operate unfettered within the environment.
- Attackers using multiple administrative accounts improve their resilience to an incident response and are able to silo operations to prevent the detection of a single compromised admin account from affecting access and actions undertaken from other compromised admin accounts.

Steps to Verify

- Validate the administrative account was created according to organizational change control policies and that the access granted is appropriate and necessary.

Azure AD Cross Tenant Access Change

Command & Control



MITRE | ATT&CK®

T1484.002 Trust Modification

Triggers

- A change was made to the cross-tenant access settings in the environment.

Possible Root Causes

- An attacker is modifying cross-tenant access settings in order to give themselves backdoor access to the environment.
- Cross-tenant access settings were modified following a legitimate business decision.

Business Impact

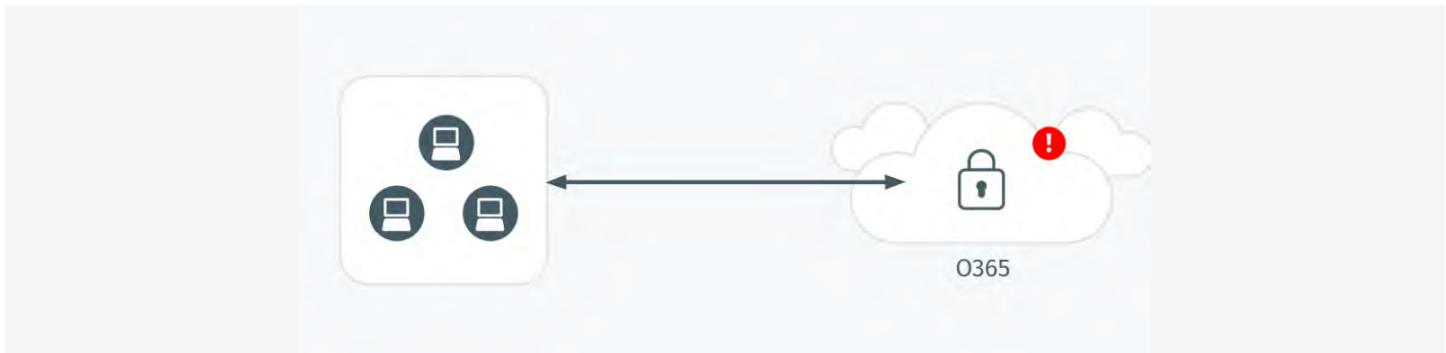
- Cross-tenant access settings manage how an organization collaborates with other organizations. They control the level of access users in external organizations have to your resources and let you trust MFA and device claims of other organizations. These settings can be configured by an attacker to get a backdoor access to the target environment from a malicious organization that they control.

Steps to Verify

- Verify that there was a legitimate business need to modify cross-tenant access settings in the environment.

Azure AD Domain Settings Modified

Command & Control



MITRE | ATT&CK®

T1484.002 Trust Modification

Triggers

- A new domain was added to the Entra ID environment.

Possible Root Causes

- An attacker is adding a malicious domain to the Entra ID environment.
- A new domain is added to the environment following a legitimate business request.

Business Impact

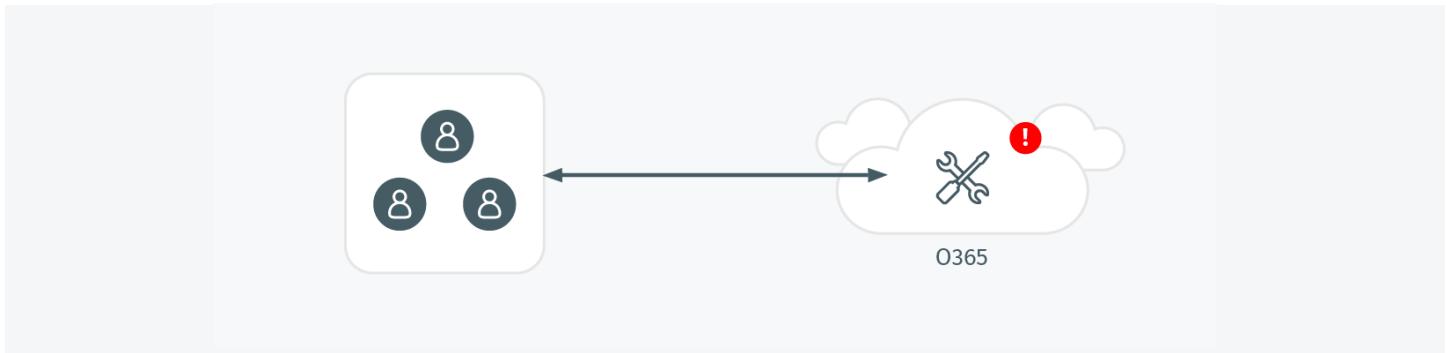
- Adding domains to an Entra ID environment may allow configuration of malicious federation. These domains can be used to create backdoors that an attacker can use to access to the environment.

Steps to Verify

- Verify that there was a legitimate business decision to add a domain to the environment.

Azure AD Login From Suspicious Location

Command & Control



MITRE | ATT&CK®

T1078 Valid Accounts

Triggers

- A successful login was observed to an account from a country that is unusual for this tenant.

Possible Root Causes

- An attacker may sign into the account they have compromised from their true location, or from a random proxy system that does not take into account the valid user's normal expected location.
- A user may be traveling to a new country on business or on vacation, and is signing into their account from there.

Business Impact

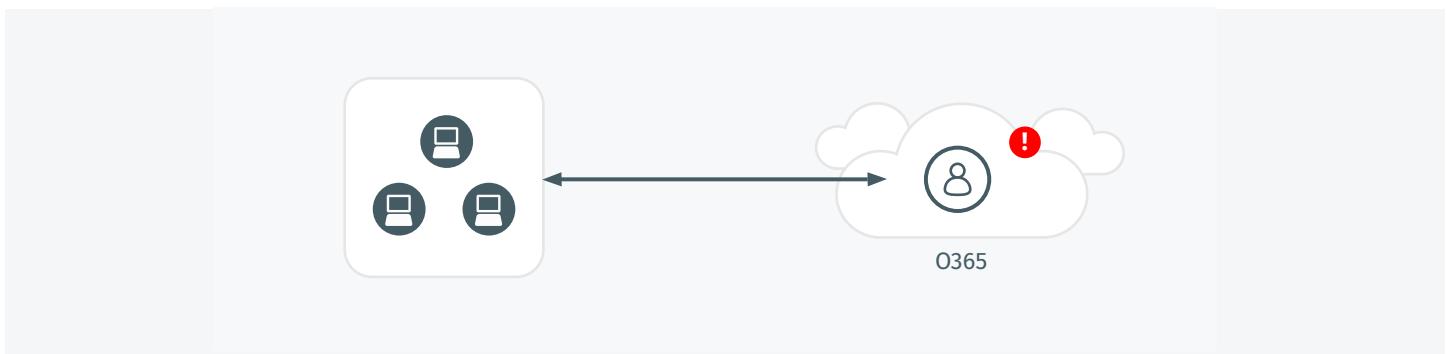
- Adversaries frequently bypass security controls through the malicious, unauthorized use of valid credentials.
- The compromise of a valid account may lead to the loss of confidentiality and integrity of any data and services that account may access, and it may be used in service of additional lateral movement or attacks against other internal users.

Steps to Verify

- Validate whether the user in question is expected to sign in from this location (e.g. as part of a business trip).

Azure AD MFA-Failed Suspicious Sign-On

Command & Control



MITRE | ATT&CK®

T1078 Valid Accounts

Triggers

- A login attempt occurred to an account where both conditional access policies were not met and where sign-on attributes (such as location, device, etc.) that are unusual for the account.

Possible Root Causes

- An adversary has stolen a valid account and is attempting to use it as part of an attack but had not yet succeeded in circumventing MFA or other conditional access policies.
- A user has moved and performed a full refresh of their devices and failed to pass MFA or other conditional access policies.

Business Impact

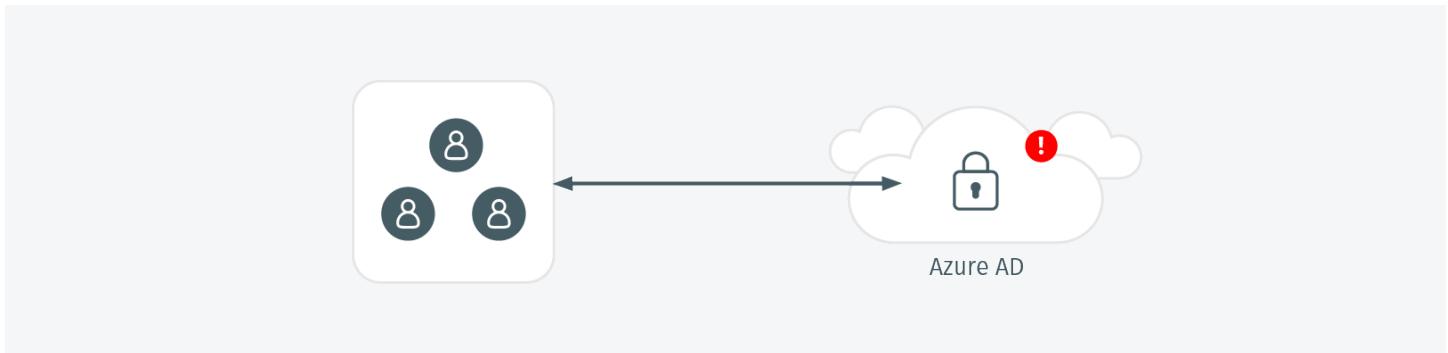
- Adversaries will continue to attempt to bypass security controls until successful unless directly stopped.
- The compromise of a valid account may lead to the loss of confidentiality and integrity of any data and services that the account may access, and it may be used in service of additional lateral movement or attacks against other internal users.

Steps to Verify

- Investigate irregularities associated with this user's login events for indications of a successful compromise.
- Validate whether these attempts were performed by the account's proper owner.

Azure AD New Certification Authority Registered

Command & Control



MITRE | ATT&CK®

[T1199 Trusted Relationship](#)

[T1484.002 Trust Modification](#)

[T1649 Steal or Forge Authentication Certificates](#)

Triggers

- An administrator has registered a new certification authority (CA) in Azure tenant.

Possible Root Causes

- An attacker is abusing a compromised administrator account in order to create a new certification authority in the tenant. Such authority can then serve as a backdoor into the tenant.
- The CA is created as part of legitimate business requirements in order to provide certificate-based authentication (CBA).

Business Impact

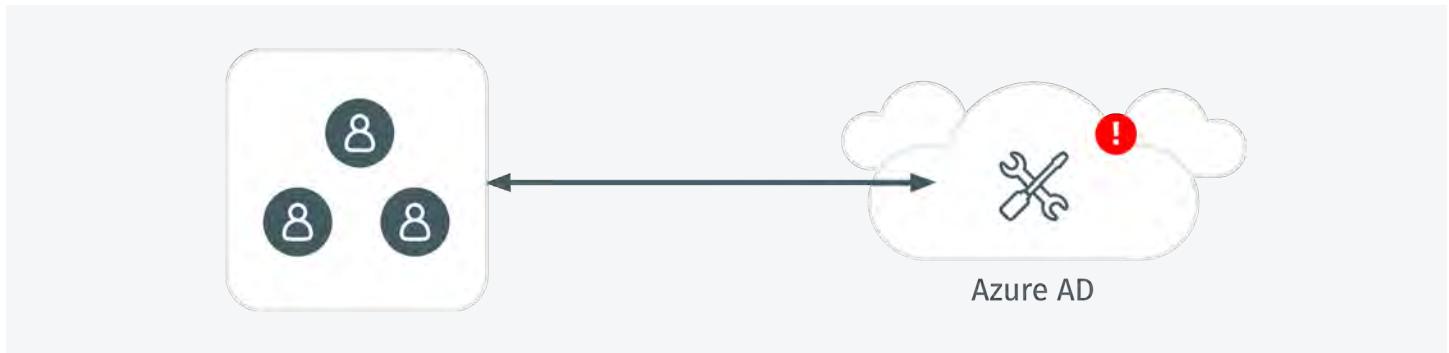
- A malicious CA can be used to impersonate and authenticate any user account without requiring knowledge of the user's credentials.

Steps to Verify

- Investigate whether there was a legitimate business requirement to create a CA.

Azure AD New Partner Added to Organization

Command & Control



MITRE | ATT&CK®

T1199 Trusted Relationship

Triggers

- A new Partner entity is added to the environment with the ability to manage, configure, and support Azure services on the organization's behalf.

Possible Root Causes

- Partner is added to the organization maliciously, by exploiting an internal admin account.
- Partner is added to the organization with legitimate intent and following a valid request by the management.

Business Impact

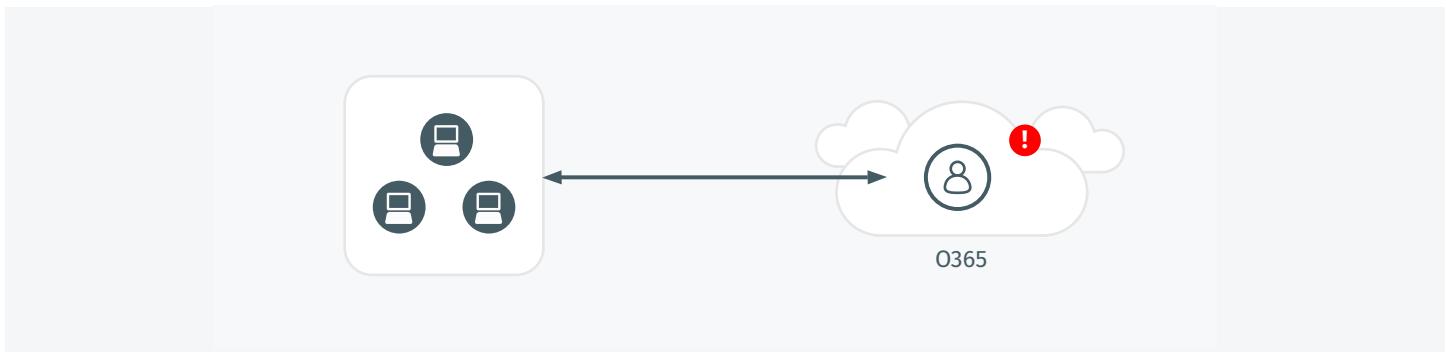
- Adding a partner entity to the organization gives that partner the ability to fully manage the environment. These privileges can be abused for complete environment compromise.

Steps to Verify

- Determine if the request to add a partner to the organization was a legitimate action. If not, delete the partner entity and investigate further.

Azure AD Redundant Access Creation

Command & Control



MITRE | ATT&CK®

T1098 Account Manipulation

Triggers

- A service principal, application, or user has been provisioned membership into the 'Privileged Role Administrator' AzureAD role.

Possible Root Causes

- An adversary has provisioned access into a sensitive role to create redundant access into the network.
- In some cases, administrators performing deployment testing will grant permissions associated with this role to the app or related service principal.

Business Impact

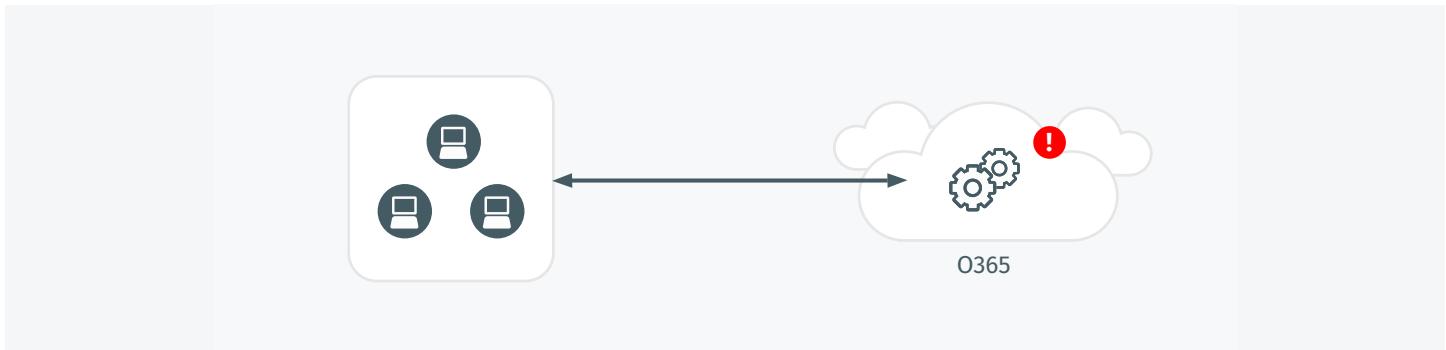
- Adversaries will create redundant access mechanisms so that they are able to continue to maintain persistence despite their primary access method being discovered and remediated.
- Redundant access allows malicious activities to continue well beyond initial discovery and response phases, increasing risks to enterprise services or data.

Steps to Verify

- Validate that this activity is not associated with authorized administrative testing activities.

Azure AD Suspicious OAuth Application

Command & Control



MITRE | ATT&CK®

T1550 Use Alternate Authentication Material

T1528 Steal Application Access Token

Triggers

- A third-party cloud application has requested excessive or risky access, which may allow malicious activities to be performed on behalf of the grantor of the permission.

Possible Root Causes

- An attacker is trying to trick the user into delegating permissions to them which will enable further malicious activities.
- A new legitimate 3rd party application is installed in the organization which requires elevated permissions from users.

Business Impact

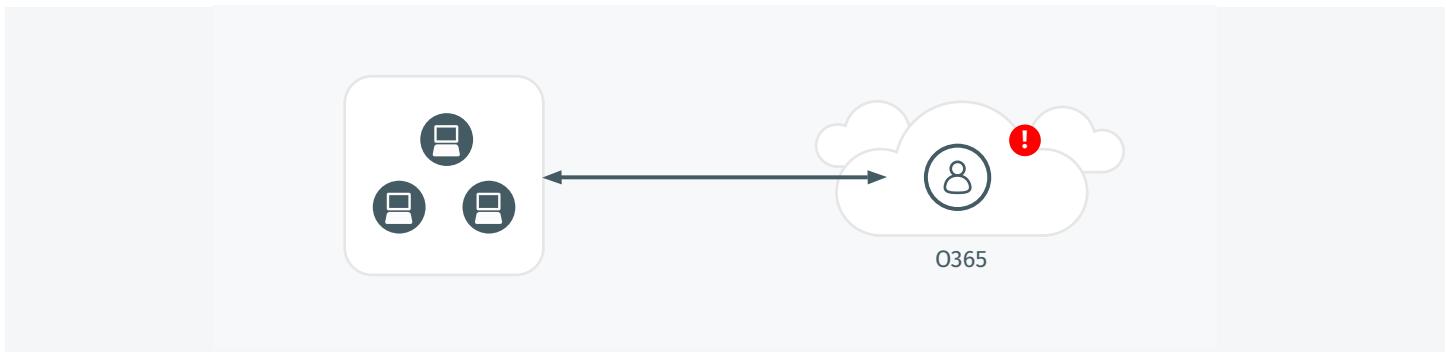
- Malicious applications are able to perform actions with delegated permissions without a user's knowledge and may be difficult to detect.
- Depending on the delegated privileges involved, the impact may range from single account takeover to full subscription compromise.

Steps to Verify

- Validate that this is an authorized application which has been vetted for risk by the security team.

Azure AD Suspicious Sign-on

Command & Control



MITRE | ATT&CK®

T1078 Valid Accounts

Triggers

- A successful login has occurred to an account with sign-on attributes (such as location, device, etc.) that are unusual for the account.

Possible Root Causes

- An adversary has stolen a valid account and is using it as part of an attack.
- A user has moved and performed a full refresh of their devices and performed login activities across these devices with new sign-on attributes.

Business Impact

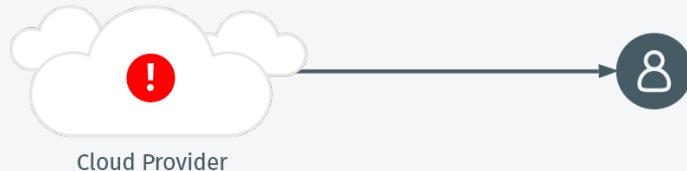
- Adversaries frequently bypass security controls through the malicious, unauthorized use of valid credentials.
- The compromise of a valid account may lead to the loss of confidentiality and integrity of any data and services that account may access, and it may be used in service of additional lateral movement or attacks against other internal users.

Steps to Verify

- Investigate irregularities associated with these login events for indications of compromise.
- Validate the login activities have been performed in accordance with organizational MFA policies, enforcing re-login with MFA if required.

Azure AD Suspicious Access from Cloud Provider

Command & Control



MITRE | ATT&CK®

T1078 Valid Accounts

T1090 Proxy

Triggers

- An account has been accessed successfully from a public cloud IP, e.g. AWS, Azure, GCP which is unusual for this account.
- Vectra's AI will continuously learn whether a cloud provider and region are normal for a given user based on their history.

Possible Root Causes

- An attacker has successfully logged into an account using a public cloud IP. The attacker has used a public IP to hide their true location and appear from a normal geolocation and IP space.
- A user or a user connected software has successfully logged into an account from a public cloud IP provider and region for the first time. The user may be actively using the cloud provider for a legitimate business function, or a cloud host software associated with the account has connected.

Business Impact

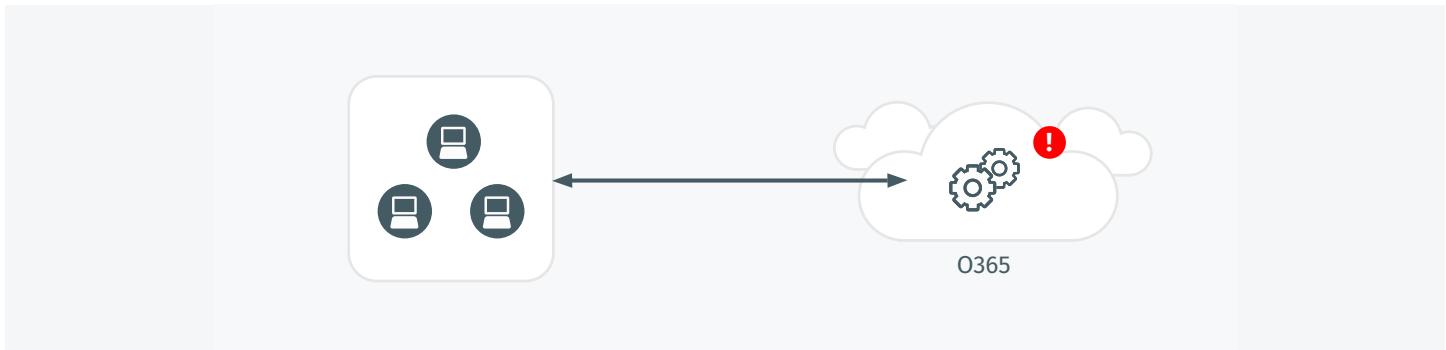
- An attacker who has access to an internal account can access any connected applications and progress their attack.

Steps to Verify

- Review whether the account owner would have a reason to be accessing their account from a public cloud.
- Consult the available logs to determine if there has been any attack progression.
- Reach out to the account owner to confirm that they were behind the sign-in event

Azure AD Suspected Compromised Access

Command & Control



MITRE | ATT&CK®

T1078 Valid Accounts

Triggers

- A successful login has occurred to an account with many characteristics that are both unusual for the account and highly correlated with account compromise.

Possible Root Causes

- An adversary has stolen a valid account and is using it as part of an attack.
- A user has shifted multiple aspects of their normal sign-on behavior which match multiple behaviors associated with malicious account takeovers.

Business Impact

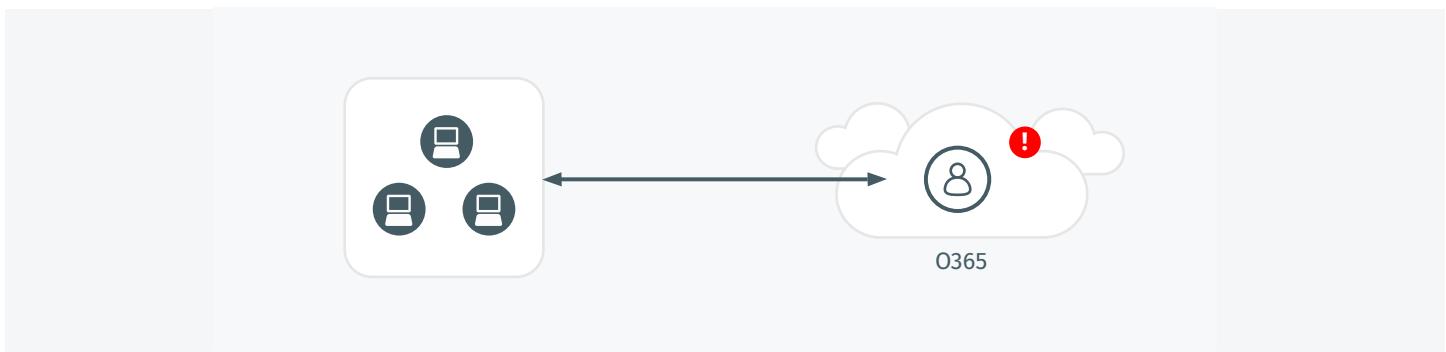
- Adversaries frequently bypass security controls through the malicious, unauthorized use of valid credentials.
- The compromise of a valid account may lead to the loss of confidentiality and integrity of any data and services that the account may access, and it may be used in service of additional lateral movement or attacks against other internal users.

Steps to Verify

- Investigate irregularities associated with these login events for indications of compromise.
- Validate the login activities have been performed in accordance with organizational MFA policies, enforcing re-login with MFA if required.

Azure AD TOR Activity

Command & Control



MITRE | ATT&CK®

T1090 Proxy

Triggers

- A user was observed accessing the environment from a known anonymized (TOR) exit node, post authentication.

Possible Root Causes

- An attacker is using an anonymizing proxy like TOR to obfuscate details of their source connection or make investigation more difficult by using multiple source IP addresses.
- A user may be intentionally using TOR to circumvent restrictions preventing access to the resources in question, such as those applied by the country they are accessing from.

Business Impact

- Attackers identified under this detection are actively operating within the environment while maintaining some level of operational security by obfuscating their source details.
- Attackers operating using TOR will reduce the ability of teams to connect identified attacker behavior with other behaviors not yet identified since it enables the attacker to regularly change the source detail of their connections while undertaking operations within the environment.

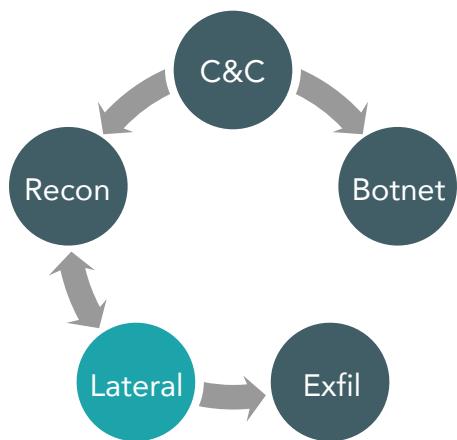
Steps to Verify

- Review the actions being undertaken by the user during and just before the identified activity to determine resources accessed and potential risk posed by that access.
- Review security policy to determine if use of TOR is allowed.
- Discuss with user to determine if use of TOR is known and legitimate.
- If review determines there is a high risk to data or the environment, disable the account and perform a comprehensive investigation.

Category

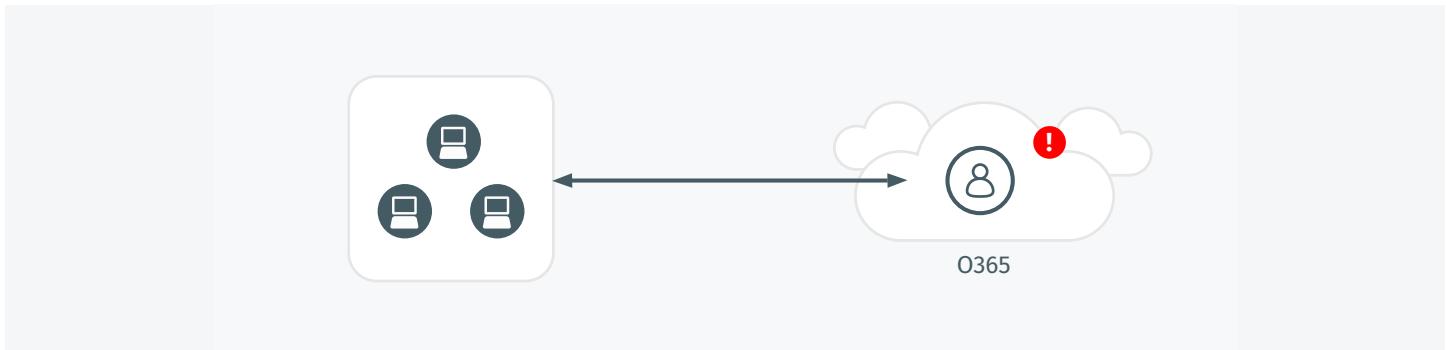
Lateral Movement

- Covers scenarios of lateral action meant to further a targeted attack
- This can involve attempts to steal account credentials or to steal data from another machine
- It can also involve compromising another host or account to make the attacker's foothold more durable or to get closer to target data



Azure AD Successful Brute-Force

Lateral Movement



MITRE | ATT&CK®

T1110 Brute Force

Triggers

- A successful login with suspicious IP Address or User-Agent after frequent failed login attempts.

Possible Root Causes

- Adoption of weak or reused credentials is common among users and attackers exploit this behavior by repeatedly attempting to login to discovered accounts using leaked or common passwords.
- Legitimate users who repeatedly mistype their password may trigger this detection
- Automated systems or services may attempt to continuously login with incorrect credentials.

Business Impact

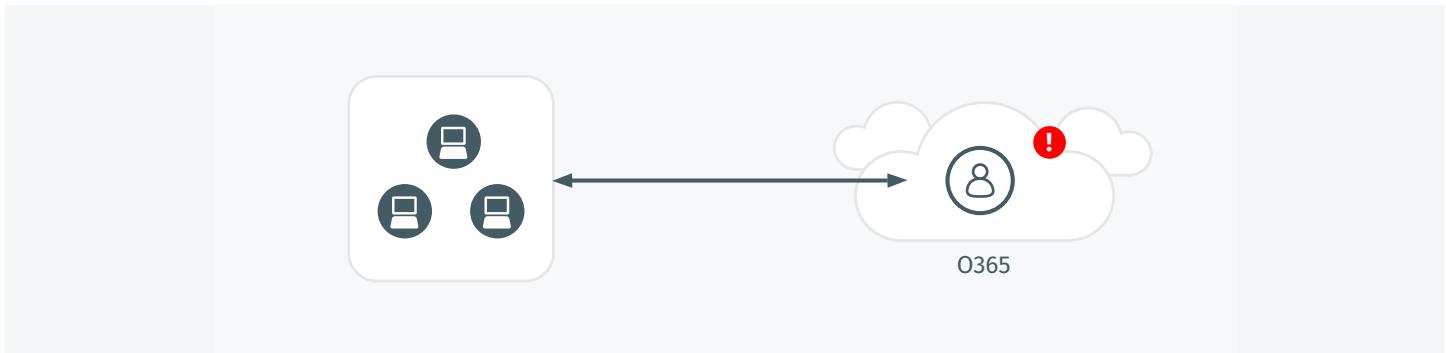
- Accounts compromised through brute-force attacks provide attackers a foothold in the enterprise.
- Attackers who have taken over administrative, executive, or high-value accounts put the enterprise at considerable risk.

Steps to Verify

- Brute-force attacks that end with a successful login should immediately be investigated for abnormal or threatening behavior.

Azure AD Successful Brute-Force – Failed Login

Lateral Movement



MITRE | ATT&CK®

T1110 Brute Force

Triggers

- Credentials were successfully validated after frequent login attempts but sign in was not accomplished.
- Password brute force tools are known to use techniques that are able to validate credentials without performing a successful authentication. The intent is to remain undetected.

Possible Root Causes

- Adoption of weak or reused credentials is common among users and attackers exploit this behavior by repeatedly attempting to login to discovered accounts using leaked or common passwords.
- Legitimate users who repeatedly mistype their password may trigger this detection
- Automated systems or services may attempt to continuously login with incorrect credentials.

Business Impact

- Accounts compromised through brute-force attacks provide attackers a foothold in the enterprise.
- Attackers who have taken over administrative, executive, or high-value accounts put the enterprise at considerable risk.

Steps to Verify

- This detection indicates that credentials were validated but authentication was not possible due to a reason indicated by an error code.
- Below is the list of error-codes and associated meanings that fall in this category:
 - 50055: The password has expired.
 - 50057: The account is disabled.
 - 50158: The user is redirected to another authentication provider.
 - 50074: MFA is required.
 - 50076: MFA is required because a new location is observed.
 - 53003: Blocked by Conditional Access Policy.
 - 50053: The account is temporarily locked.
- The account should be immediately investigated for abnormal or threatening behavior

Azure AD Suspicious Device Registration

Lateral Movement



MITRE | ATT&CK®

T1098 Account Manipulation

Triggers

- A new device has suspiciously been registered to an account which may provide an attacker with persistent access to your tenant.

Possible Root Causes

- An attacker may have compromised an account and registered a new device in the environment to maintain continued persistence. By registering a new device in the tenant, the attacker's ongoing access may be extended beyond the method of initial compromise.
- A legitimate user might have registered a new personal or official work device under unexpected circumstances.

Business Impact

- An attacker who controls tenant-registered devices could bypass policies related to login requirements and access, enabling persistent access to cloud and potentially network data.

Steps to Verify

- Review whether the location of the registration and device type aligns with the user's expected activity.
- Consult the available logs to determine if the activity prior to the registration is as expected.
- Reach out to the account owner to confirm that they registered the device.

Azure AD Suspicious Factor Registration

Lateral Movement



MITRE | ATT&CK®

[T1556 Modify Authentication Process](#)

Triggers

- A suspicious new authentication factor has been registered to an account which may provide an attacker with persistent access to your tenant.

Possible Root Causes

- An attacker may have compromised an account and registered a new authentication method (such as an MFA method) in the environment to maintain continuous access. By registering a new authentication method, the attacker's ongoing access may be extended beyond the method of initial compromise.
- A legitimate user may have added a new authentication method under circumstances that were unexpected for the environment.
- Note: The specific authentication method itself may be typical for the environment.

Business Impact

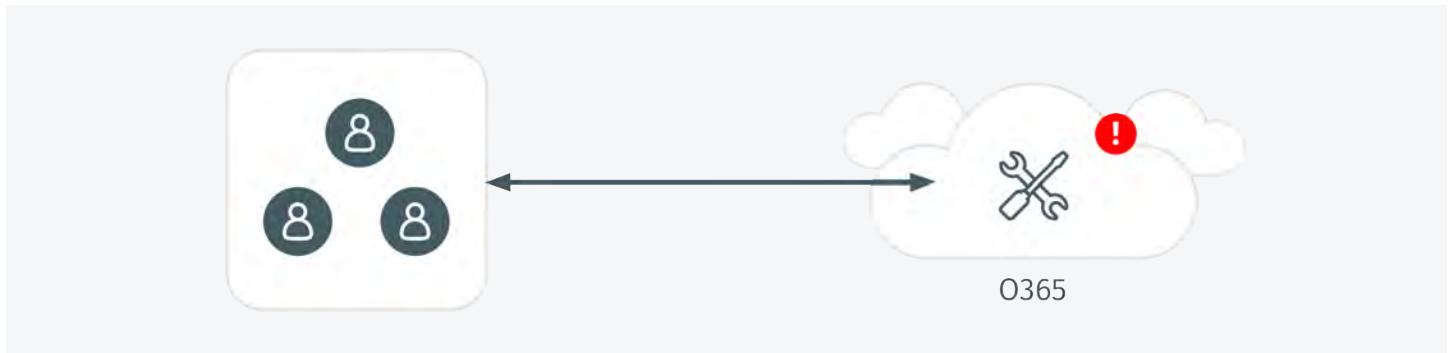
- An attacker who registers an authentication factor could bypass policies related to login requirements and access, enabling persistent access to cloud and potentially network data.

Steps to Verify

- Review whether the location of the registration aligns with the user's expected activity.
- Consult the available logs to determine if the activity prior to the registration is as expected.
- Reach out to the account owner to confirm that they registered the factor.

Azure AD Change to Trusted IP Configuration

Lateral Movement



MITRE | ATT&CK®

T1562 Impair Defenses

Triggers

- A change to a trusted IP configuration in Azure was observed in either the AzureAD Known Networks configuration or the configuration for trusted networks for multi-factor authentication.

Possible Root Causes

- Attackers may add networks to the trusted networks ranges to allow them to bypass security controls under conditional access policies or to bypass MFA requirements.
- System administrators may add trusted networks to allow trusted environments to have different security policies applied to them.

Business Impact

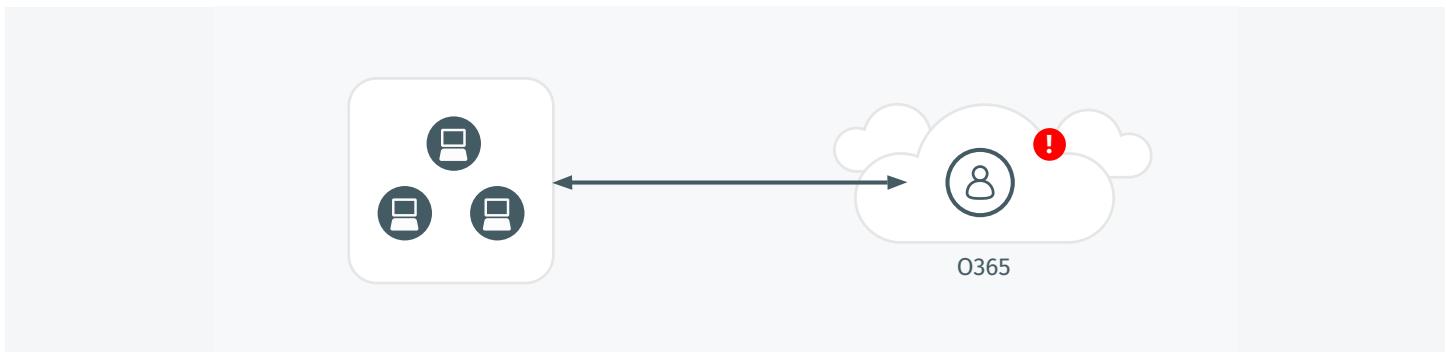
- Modifications to the trusted network configuration may introduce risks by allowing particular IP addresses/ranges to bypass critical security controls.
- Trade-offs in favor of usability over security can be achieved through the configuration of trusted IPs, but when abused or misconfigured can increase risk to an organization by disabling expected security controls.

Steps to Verify

- Investigate the IP addresses to determine if they should be trusted by the organization.
- Contact the owner of the account that made the change to verify it was done legitimately.

Azure AD MFA Disabled

Lateral Movement



MITRE | ATT&CK®

T1562 Impair Defenses

Triggers

- An account was observed disabling Multi-Factor Authentication (MFA) for another account.

Possible Root Causes

- An attacker is disabling MFA on an account to bypass this security control as a means of maintaining or acquiring additional access to the environment.
- Administrators may disable MFA for accounts used by automated processes or to temporarily enable users to access an environment after losing their second factor device.

Business Impact

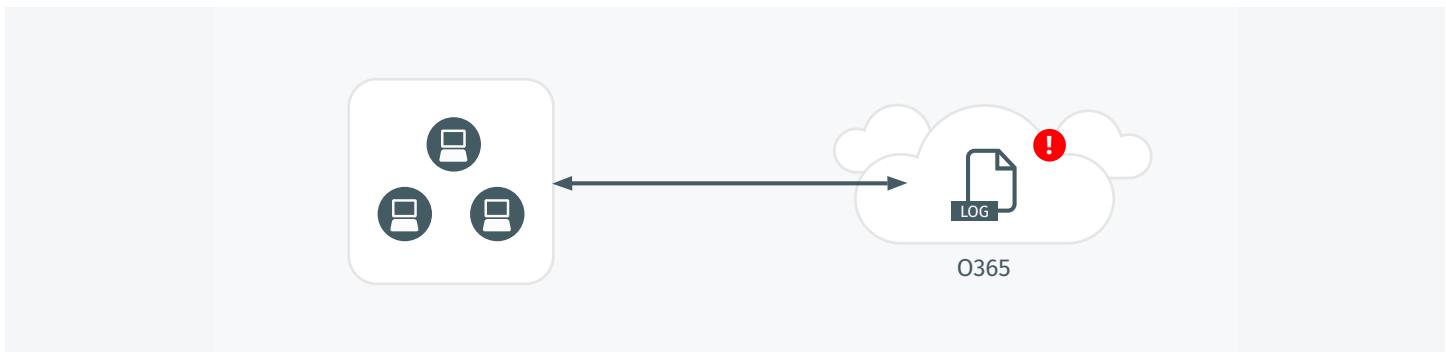
- MFA is a critical security control that if bypassed may be indicative of an active threat in the environment or increase risk of the account becoming compromised in the future.
- Compromised accounts provide attackers with access to critical systems and data which may be stolen, modified, or deleted.

Steps to Verify

- Review the account and internal policy to determine if MFA should be enabled for this account.
- Verify the action of disabling MFA on this account was intentional and followed internal security policies and change control processes.

Azure AD Newly Created Admin Account

Lateral Movement



MITRE | ATT&CK®

T1136 Create Account

Triggers

- An account has been created with administrative privileges (TenantAdmins, PrivilegedRoleAdmins, ApplicationAdministrators) that provide broad access to the environment.

Possible Root Causes

- An attacker that has gained administrative rights has added additional administrative accounts to the environment as a back-up access method if their existing access is disabled or otherwise removed at a future date.
- Existing legitimate administrators may add additional administrative users unintentionally or via social engineering.
- A new, legitimate, administrative account was added.

Business Impact

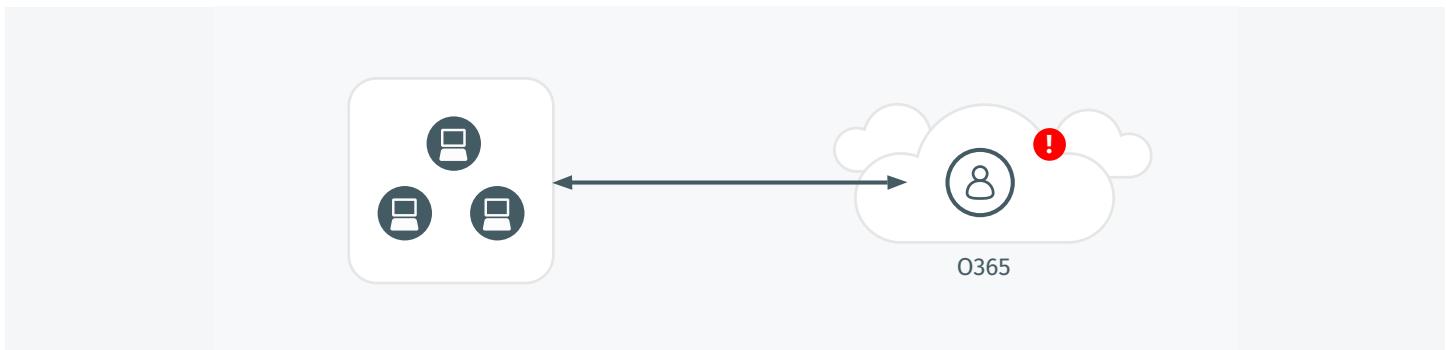
- Unauthorized administrative users have complete control within the environment, creating significant on-going risk to a broad range of resources.
- Attackers with access to the identified administrative rights will be able to operate unfettered within the environment.
- Attackers using multiple administrative accounts improve their resilience to an incident response and are able to silo operations to prevent the detection of a single compromised admin account from affecting access and actions undertaken from other compromised admin accounts.

Steps to Verify

- Validate the administrative account was created according to organizational change control policies and that the access granted is appropriate and necessary.

Azure AD Privilege Operation Anomaly

Lateral Movement



MITRE | ATT&CK®

T1078 Valid Accounts

Triggers

- Abnormal Azure AD operations that may be associated with privilege escalation or account takeover.

Possible Root Causes

- Attackers may be escalating privileges and performing admin-level operations after regular account takeover.
- A user whose learned activity baseline has been lost as a result of a prolonged leave of absence or a change in job function has returned to their regular job.
- A user's role may have evolved as part of a special project or assignment and the user is performing Azure AD activities previously outside of their learned baseline.

Business Impact

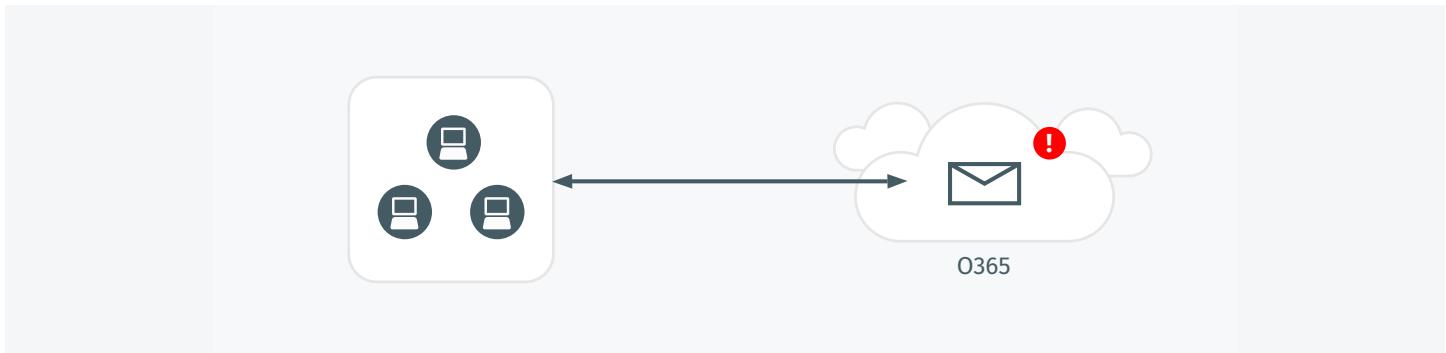
- Users substantially deviating from their learned baseline in ways that correspond to threats associated with privilege escalation or account takeover often indicate an adversary foothold.
- Account takeover and privilege escalation can lead to sensitive information leakage, ransomware attacks, and other abuses.

Steps to Verify

- Investigate both the target and result of these operations to understand the potential impact.

Azure AD Unusual Scripting Engine Usage

Lateral Movement



MITRE | ATT&CK®

T1059 Command and Scripting Interpreter

Triggers

- An account has executed O365 operations with either tools, scripting engines or command line interfaces which could be maliciously used by attackers.

Possible Root Causes

- An attacker is \"living off the land\" through the misuse of authorized tools (curl, AutoHotKey32, etc.) to extend their attack.
- An attacker has used a scripting engine (Powershell, Python, and others) to build and execute attack tools.
- When attacker is not careful, the default User Agent strings are submitted by these tools, indicating that the operation is not done interactively by a legitimate human user.
- Automation tools and scripts are sometimes used by power users and IT personnel to access O365.

Business Impact

- Automated tools increase attack speed and volume while reducing human error, and attackers that successfully leverage them have an opportunity to move faster and in some cases with a lower chance of detection.
- Use of automation tools is a \"force multiplier\" that increases chances of successful breaches and data exfiltration, significantly increasing risks to the enterprise.

Steps to Verify

- Investigate O365 operation in context of the user, verify if this user would reasonably conduct these types of operations.
- Investigate tooling or scripting engine to validate if this is an appropriate and approved tool for a user of this type.

Category

Info

- Reports on new and novel events without directly impacting scoring
- New and novel events occur normally in most network and cloud environments and in most cases are not directly linked to threats
- Awareness of new and novel events support better situational awareness and provide additional context when observed with kill chain alerts

Azure AD Login Attempt to Disabled Account

- A login attempt for an account that has been explicitly disabled was observed within the environment, potentially indicating environment probing or attempted access by former employees.

Azure AD Brute Force Attempt

- Reports when one or more external IPs are seen attempting to brute-force into an account without any successful attempts.

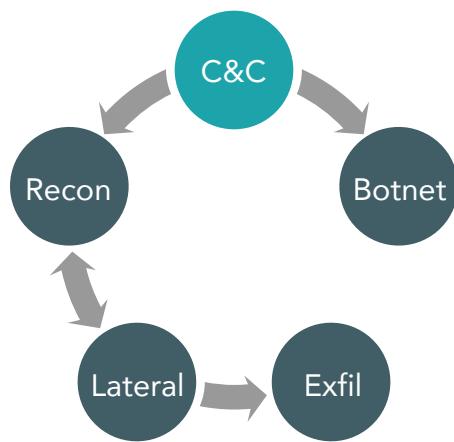
Azure Cloud Detections



Category

Command & Control

- An attacker is controlling an Azure account to orchestrate their attack against Azure infrastructure and services
- Attackers will access credential through various methods
- Attackers will control accounts in order to then perform reconnaissance and lateral movement in the name of achieving an object like data exfiltration or resource impact
- Actions are associated with [Command and Control](#), [Initial Access](#), [Reconnaissance](#) MITRE Tactics



Azure Suspicious Access from AWS Cloud

Command & Control



MITRE | ATT&CK®

[T1078 Valid Accounts](#)

[T1548 Abuse Elevation Control Mechanism](#)

Triggers

- An account has been accessed successfully from an AWS public cloud IP which is unusual for this account.
- Vectra's AI will continuously learn whether a cloud provider and region are normal for a given user based on their history.

Possible Root Causes

- An attacker has successfully logged into an account using an AWS public cloud IP. The attacker has used a public IP to hide their true location and appear from a normal geolocation and IP space.
- A user or a user connected software has successfully logged into an account from an AWS public cloud IP provider and region for the first time. The user may be actively using the cloud provider for a legitimate business function, or a cloud host software associated with the account has connected.

Business Impact

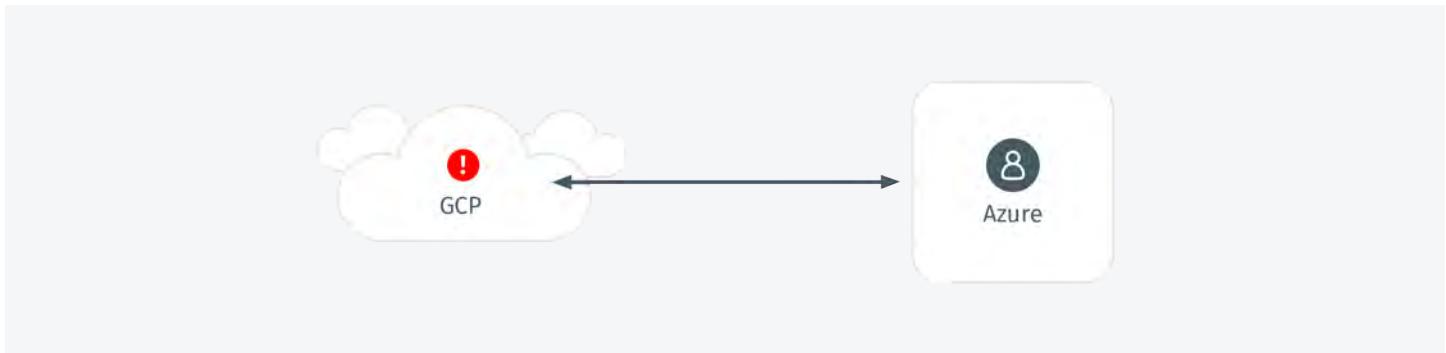
- An attacker who has access to an internal account can access connected applications and progress their attack.

Steps to Verify

- Review whether the account owner would have a reason to be accessing their account from the AWS public cloud.
- Consult the available logs to determine if there has been any attack progression.
- Reach out to the account owner to confirm that they were behind the activity observed.

Azure Suspicious Access from GCP Cloud

Command & Control



MITRE | ATT&CK®

[T1078 Valid Accounts](#)

[T1548 Abuse Elevation Control Mechanism](#)

Triggers

- An account has been accessed successfully from a GCP public cloud IP which is unusual for this account.
- Vectra's AI will continuously learn whether a cloud provider and region are normal for a given user based on their history.

Possible Root Causes

- An attacker has successfully logged into an account using a GCP public cloud IP. The attacker has used a public IP to hide their true location and appear from a normal geolocation and IP space.
- A user or a user connected software has successfully logged into an account from a GCP public cloud IP provider and region for the first time. The user may be actively using the cloud provider for a legitimate business function, or a cloud host software associated with the account has connected.

Business Impact

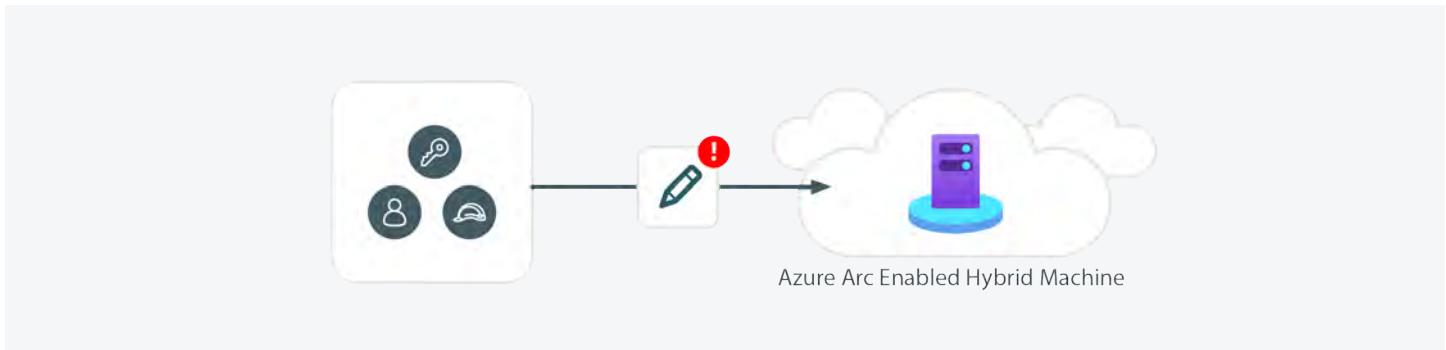
- An attacker who has access to an internal account can access connected applications and progress their attack.

Steps to Verify

- Review whether the account owner would have a reason to be accessing their account from the GCP public cloud.
- Consult the available logs to determine if there has been any attack progression.
- Reach out to the account owner to confirm that they were behind the activity observed.

Azure Suspicious Hybrid Machine Extension Installation

Command & Control



MITRE | ATT&CK®

T1651 Cloud Administration Command

T1543 Create or Modify System Process

Triggers

- A Microsoft Entra Identity is using a highly permissive role within a Resource Group or Subscription scope to install and execute a hybrid machine extension without explicit consent or an audit trail. This is unusual for the identity.

Possible Root Causes

- Unauthorized Access: An attacker has gained access to a Microsoft Entra Identity and is using it to install a hybrid machine extension without permission.
- Misconfigured security settings: A Microsoft Entra Identity is granted excessive permissions, allowing them to install hybrid machine extensions.
- Exploitation of a vulnerability in Azure's hybrid machine extension deployment process.
- Human error during hybrid machine management or maintenance.
- A legitimate identity that does not usually interact with hybrid machine extensions installing an extension to execute their responsibilities.

Business Impact

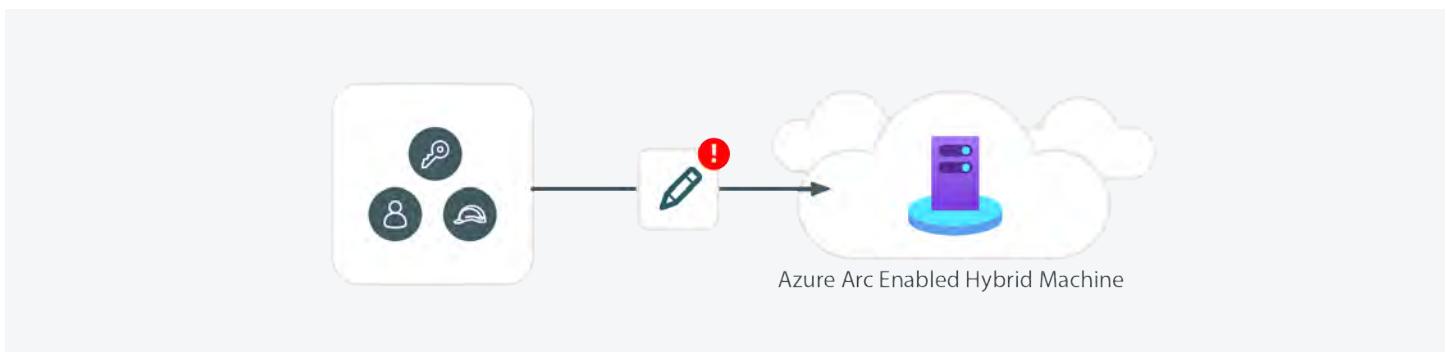
- Data loss or corruption due to unauthorized access.
- Denial-of-Service (DoS) attacks or resource exhaustion due to malware execution.
- System compromise or ransomware attacks.
- Unplanned changes, service disruptions or downtime for critical services.
- Compliance and regulatory issues due to inadequate security controls.

Steps to Verify

- Investigate the subscription or resource group scope where the hybrid machine extension was installed.
- Review the Azure Activity Logs for suspicious activity around the time of the hybrid machine extension was installed.
- Investigate the user or service principal responsible for the deployment.
- Validate the permissions associated with the hybrid machine extension.
- Analyze network traffic and system logs for other potential indicators of compromise (IOCs).
- Conduct a thorough security audit and risk assessment to identify vulnerabilities and implement remediation actions.

Azure Suspicious Hybrid Machine Run Command Execution

Command & Control



MITRE | ATT&CK®

T1651 Cloud Administration Command

T1059 Command and Scripting

Triggers

- The Microsoft Entra Identity executing the command does not usually perform Run Commands on Azure hybrid machines, indicating potential malicious activity.
- The execution of the Run Command is occurring outside of regular business hours.
- Multiple instances of suspicious Run Command executions are detected within a short timeframe or Run Commands are being executed anomalously across multiple hybrid machines.

Possible Root Causes

- Unauthorized Access: An attacker has gained access to a Microsoft Entra Identity and is using it to execute Run Commands on hybrid machines without permission.
- Misconfigured Permissions: A Microsoft Entra Identity is granted excessive permissions, allowing them to execute Run Commands on hybrid machines they shouldn't have access to.
- Accidental or unintentional execution of ARC commands due to user error.
- A legitimate Microsoft Entra Identity is running a script to configure a new hybrid machine for production use, which involves executing Run Commands with elevated privileges.
- A legitimate application is using Azure's Run Command feature to automate the deployment of updates and patches to a hybrid machine or fleet of hybrid machines.

Business Impact

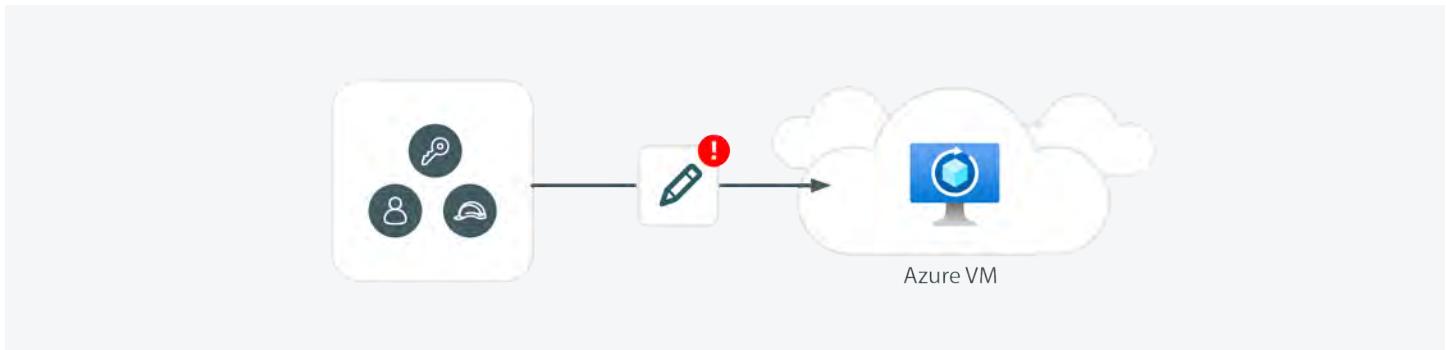
- Data loss or corruption due to unauthorized deletions. Security breaches due to exploitation of vulnerabilities or abuse of privileges.
- Compliance issues due to non-compliance with security policies
- Downtime and revenue loss due to malicious activity

Steps to Verify

- Investigate the Microsoft Entra Identity executing the Run Command for any signs of unauthorized access or excessive privileges.
- Investigate the user's or service principal's permissions and access levels within Azure.
- Examine the parameters used in the Run Command to identify potential security risks such as deletion of sensitive data. Verify if any additional security alerts or notifications were triggered around the time of the suspicious event. Check for any unusual or unexplained changes to Azure resources or services.
- Consult with Azure administrators, security teams, and relevant stakeholders to determine the cause and scope of the incident.

Azure Suspicious VM Extension Installation

Command & Control



MITRE | ATT&CK®

T1651 Cloud Administration Command

T1543 Create or Modify System Process

Triggers

- A Microsoft Entra Identity is using a highly permissive role within a Resource Group or Subscription scope to install and execute a VM extension without explicit consent or an audit trail. This is unusual for the identity.

Possible Root Causes

- Unauthorized Access: An attacker has gained access to a Microsoft Entra Identity and is using it to install a VM extension without permission.
- Misconfigured security settings: A Microsoft Entra Identity is granted excessive permissions, allowing them to install VM extensions.
- Exploitation of a vulnerability in Azure's VM extension deployment process.
- Human error during VM management or maintenance.
- A legitimate identity that does not usually interact with VM extensions installing an extension to execute their responsibilities.

Business Impact

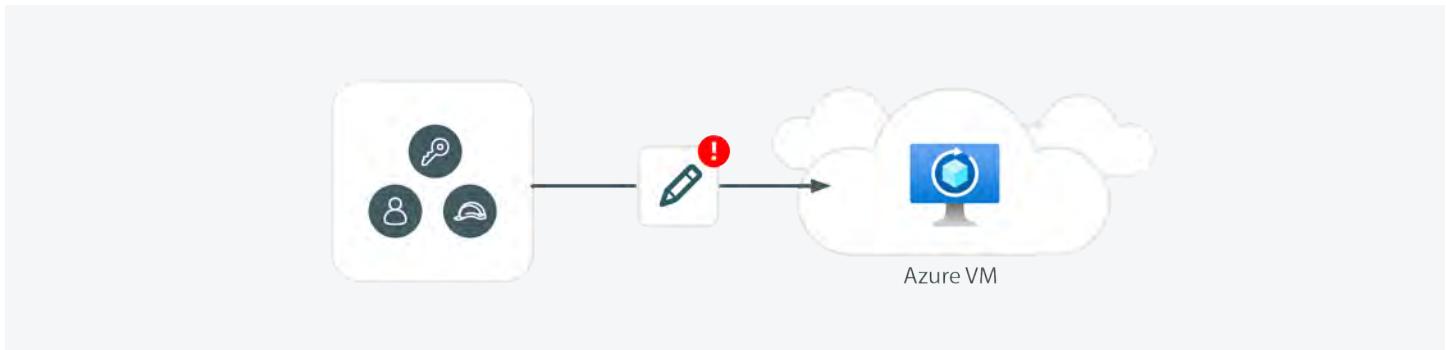
- Data loss or corruption due to unauthorized access.
- System compromise or ransomware attacks.
- Unplanned changes, service disruptions or downtime for critical services.
- Compliance and regulatory issues due to inadequate security controls.

Steps to Verify

- Review Azure activity logs for other suspicious extension deployments.
- Investigate the user or service principal responsible for the deployment.
- Validate the VM extension configuration and permissions.
- Analyze network traffic and system logs for other potential indicators of compromise (IOCs).
- Conduct a thorough security audit and risk assessment to identify vulnerabilities and implement remediation actions.

Azure Suspicious VM Run Command Execution

Command & Control



MITRE | ATT&CK®

T1651 Cloud Administration Command

T1059 Command and Scripting

Triggers

- The Microsoft Entra Identity executing the command does not usually perform Run Commands on Azure VMs.
- The execution of the Run Command is occurring outside of regular business hours.
- Multiple instances of suspicious Run Command executions are detected within a short timeframe.
- Run Commands are being executed anomalously across multiple VMs.

Possible Root Causes

- Unauthorized Access: An attacker has gained access to a Microsoft Entra Identity and is using it to execute Run Commands on VMs without permission.
- Misconfigured Permissions: A Microsoft Entra Identity is granted excessive permissions, allowing them to execute Run Commands on VMs they shouldn't have access to.
- Malware or Exploitation tool: An attacker is using malware or tools to abuse in the VM's operating system to gain elevated privileges and execute Run Commands.
- A legitimate Microsoft Entra Identity is running a script to configure a new VM for production use, which involves executing Run Commands with elevated privileges.
- A legitimate application is using Azure's Run Command feature to automate the deployment of updates and patches to a VM or fleet of VMs.

Business Impact

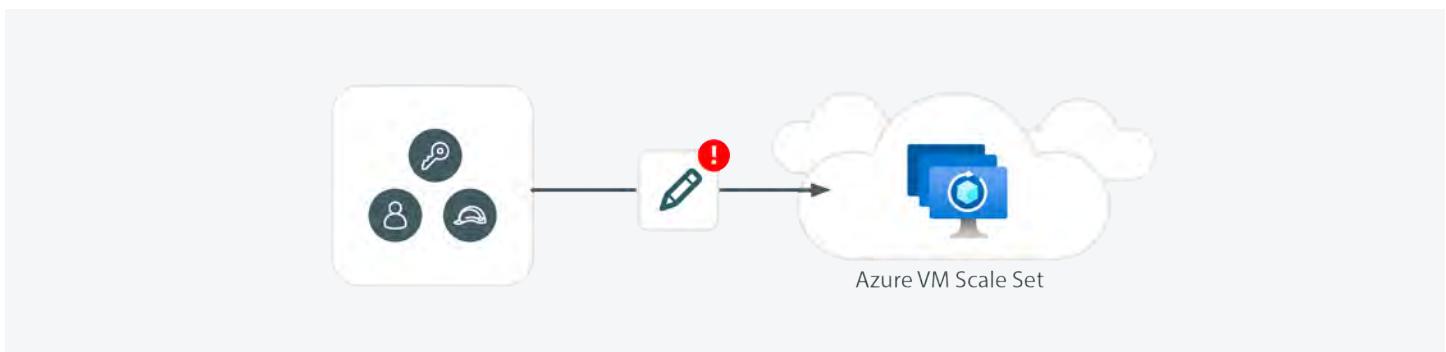
- Data loss or corruption due to unauthorized deletions.
- Security breaches due to exploitation of vulnerabilities or abuse of privileges.
- Compliance issues due to non-compliance with security policies.
- Downtime and revenue loss due to malicious activity.

Steps to Verify

- Review the actions being undertaken by the user after the identified activity and investigate the Microsoft Entra Identity: Review the Microsoft Entra Identity executing the Run Command for any signs of unauthorized access or excessive privileges.
- Analyze Command Parameters: Examine the parameters used in the Run Command to identify potential security risks such as deletion of sensitive data.
- Check Execution Time and Frequency: Verify if the suspicious Run Commands are occurring within regular business hours or if there are multiple instances within a short timeframe.
- Review VM Configuration and Permissions: Ensure that the VM's configuration and permissions align with organizational security policies.

Azure Suspicious VM Scale Set Extension Installation

Command & Control



MITRE | ATT&CK®

T1651 Cloud Administration Command

T1543 Create or Modify System Process

Triggers

- A Microsoft Entra Identity is using a highly permissive role within a Resource Group or Subscription scope to install and execute an extension within a VM Scale Set (VMSS) without explicit consent or an audit trail. This is unusual for the identity.

Possible Root Causes

- Unauthorized Access: An attacker has gained access to a Microsoft Entra Identity and is using it to install extensions on VM Scale Sets without permission.
- Misconfigured security settings: A Microsoft Entra Identity is granted excessive permissions, allowing them to install extensions on VM Scale Sets.
- Malware or Vulnerability Exploitation: An attacker uses malware or exploits a vulnerability in the VM's operating system to gain elevated privileges and install high-risk extensions.
- Human error during VM Scale Set management or maintenance.
- A legitimate identity that does not usually interact with VM Scale Set extensions installing an extension to execute their responsibilities.

Business Impact

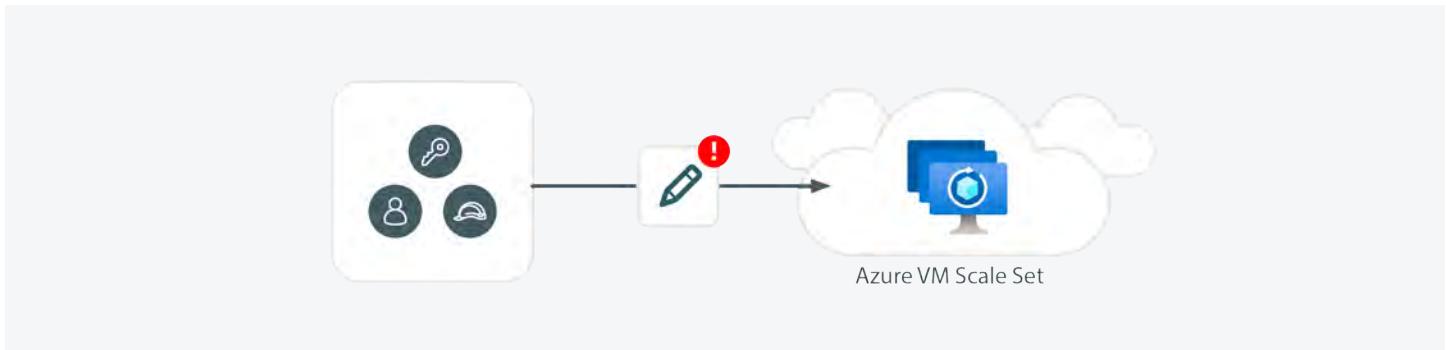
- Unauthorized Access: An attacker has gained access to a Microsoft Entra Identity and is using it to install extensions on VM Scale Sets without permission.
- Misconfigured security settings: A Microsoft Entra Identity is granted excessive permissions, allowing them to install extensions on VM Scale Sets.
- Malware or Vulnerability Exploitation: An attacker uses malware or exploits a vulnerability in the VM's operating system to gain elevated privileges and install high-risk extensions.
- Human error during VM Scale Set management or maintenance.
- A legitimate identity that does not usually interact with VM Scale Set extensions installing an extension to execute their responsibilities.

Steps to Verify

- Investigate the user or service principal responsible for the deployment. Review the identity for any signs of unauthorized access or excessive privileges.
- Analyze Extension Parameters: Examine the parameters used in the extension installation to identify potential security risks such as deletion sensitive data.
- Check Execution Time and Frequency: Review the Azure Activity logs to verify if the suspicious extension installations are occurring within regular business hours or if there are multiple instances within a short timeframe.
- Review VMSS Configuration and Permissions: Ensure that the VMSS's configuration and permissions align with organizational security policies.
- Conduct a thorough security audit and risk assessment to identify vulnerabilities and implement remediation actions.

Azure Suspicious VM Scale Set Run Command Execution

Command & Control



MITRE | ATT&CK®

[T1651 Cloud Administration Command](#)

[T1059 Command and Scripting](#)

Triggers

- A Microsoft Entra Identity is using a highly permissive role, such as an Azure Virtual Machine Contributor or Owner role, to execute Run Commands within a Resource Group that contains a VM Scale Set (VMSS). This is unusual for the identity.
- The execution of the Run Command is occurring outside of regular business hours.
- Multiple instances of suspicious Run Command executions are detected within a short timeframe.

Possible Root Causes

- Unauthorized Access: An attacker has gained access to a Microsoft Entra Identity and is using it to execute Run Commands on VM Scale Sets without permission.
- Misconfigured Permissions or role assignments: A Microsoft Entra Identity is granted excessive permissions, allowing them to execute Run Commands on VM Scale Sets they shouldn't have access to.
- A DevOps engineer is using Run Commands to automate a deployment or troubleshoot processes on multiple VMs in the scale set, but due to a misconfigured script or permissions issue, the commands appear suspicious.
- A legitimate automated backup or maintenance task is running on the VM scale set.
- A system update or patching process being executed causing changes in file system permissions or access patterns.
- A legitimate administrator is using Run Commands to troubleshoot an issue, but their credentials are not recognized by the security system.

Business Impact

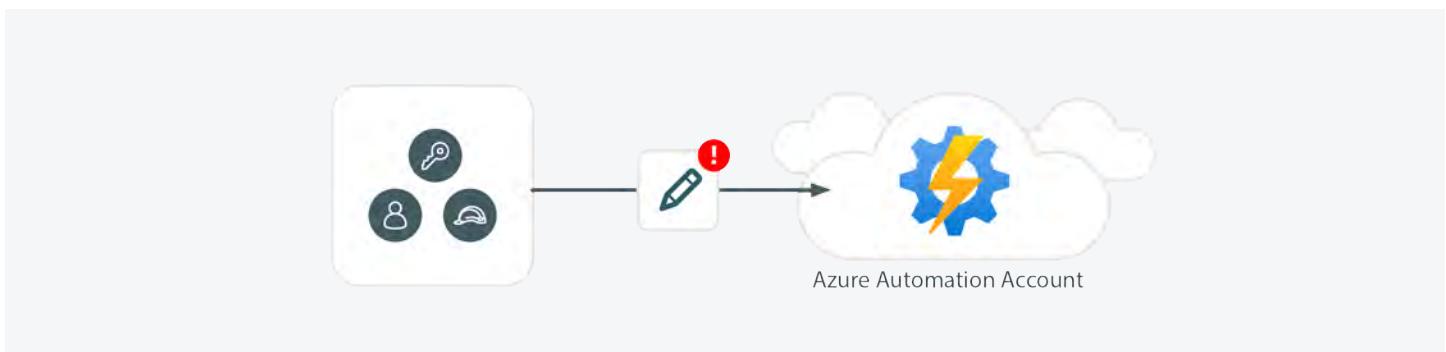
- Data loss or corruption due to unauthorized access.
- Security breaches due to exploitation of vulnerabilities or abuse of privileges.
- Unplanned changes or downtime for critical services. Exploitation of security vulnerabilities.

Steps to Verify

- Review the Resource Group and VMSS permissions, ensuring only authorized users or service principals have access. Investigate the identity of the user or service principal that executed the Run Command.
- Check the command's parameters and actions taken, comparing them with expected or approved activities. Analyze Azure Audit Logs to determine the timing and source of the suspicious activity.
- Confirm whether the VMSS is part of a larger infrastructure as code (IaC) configuration or automated deployment process.

Azure Suspicious VM Automation Test

Command & Control



MITRE | ATT&CK®

T1098 Account Manipulation

T1071 Enterprise

Triggers

- Unusual modifications were made to a Runbook in an Azure Automation Account.
- Unusual test job was run on an Azure Automation Account Runbook for an Azure VM before it is staged for execution.
- A VM Runbook was updated or tested as part of legitimate business workflows.

Possible Root Causes

- A compromised service principal account is attempting unauthorized modifications.
- An insider threat with elevated privileges is making malicious changes to the Runbook.
- Automated deployment scripts are updating a Runbook without proper authorization.
- An authorized user making changes to a Runbook as part of a legitimate development process.
- Misconfigured permissions or access levels within Azure, allowing unauthorized users to modify the Automation Account.

Business Impact

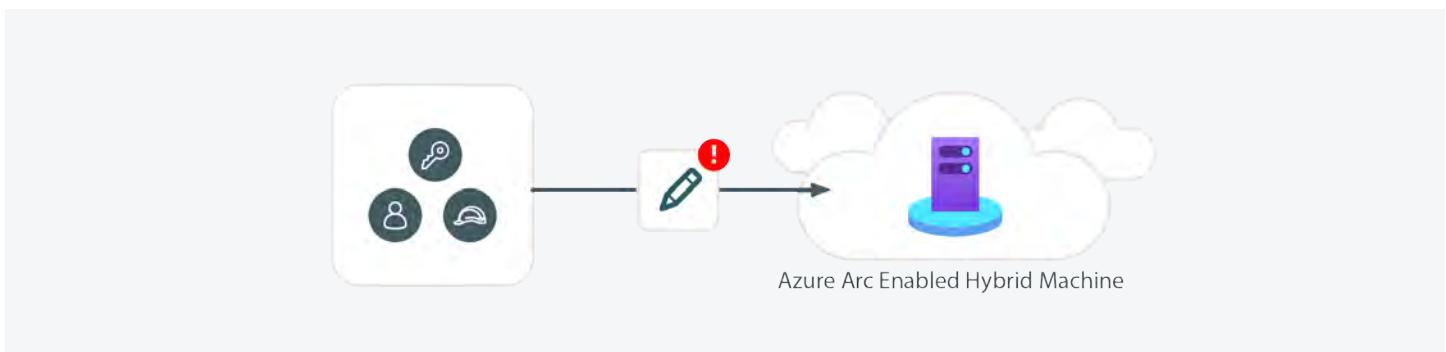
- Exposure of sensitive data through unauthorized access or data leaks.
- Security vulnerabilities exploited due to misconfigured Runbooks.
- Unplanned changes to business logic or workflows.
- Potential data breaches, unauthorized access to resources, disruption of critical business services, and reputational damage.

Steps to Verify

- Review the Azure Activity Logs for the suspicious event, focusing on the user/service principal and the modified Runbook.
- Inspect the Runbook code for signs of malicious activity. To view the Runbook:
 - Navigate to the 'Automation Accounts' service in Azure
 - Identify the Automation Account associated to the Runbook
 - The Runbooks can be found under the 'Process Automation' tab for the selected Automation Account
- Investigate the user's or service principal's permissions and access levels within Azure Automation Account.
- Verify if other security alerts or notifications were triggered around the time of the suspicious event.
- Check the Azure Virtual Machine for signs of malicious activity, such as unexpected changes to configuration or data.

Azure Suspicious Hybrid Automation Test

Command & Control



MITRE | ATT&CK®

T1098 Account Manipulation

T1071 Enterprise

Triggers

- Unusual modifications were made to a hybrid Runbook in an Azure Automation Account.
- Unusual test job was run on an Azure Automation Account Runbook for an Azure hybrid machine before it is staged for execution.
- A hybrid machine Runbook was updated or tested as part of legitimate business workflows.

Possible Root Causes

- A compromised principal account is attempting unauthorized modifications to the Automation Account.
- An authorized user making changes to the hybrid Runbook as part of a legitimate development process.
- Automated deployment scripts not previously used are updating the hybrid Runbook code or configuration.
- A developer creating/modifying a hybrid Runbook for testing purposes.
- Misconfigured permissions or access levels within Azure, allowing unauthorized users to modify the Automation Account.

Business Impact

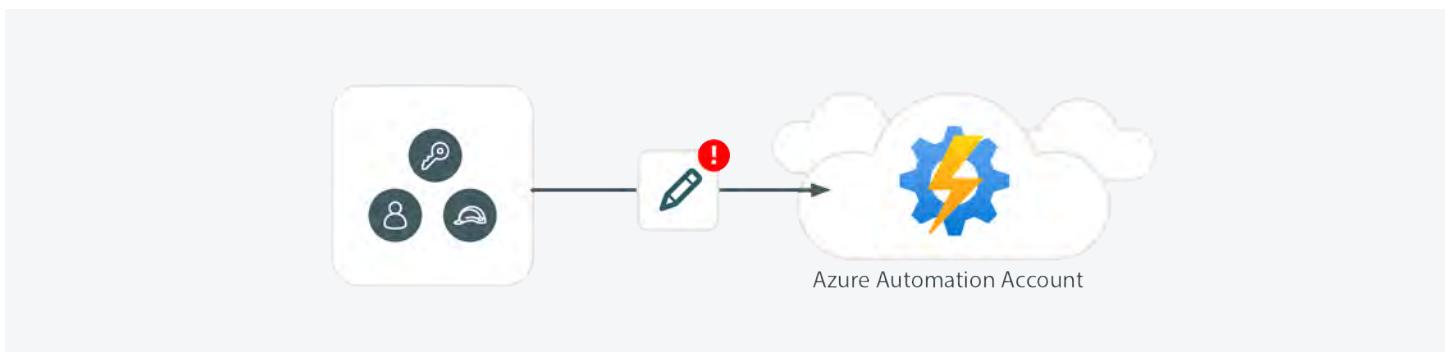
- Exposure of sensitive data through unauthorized access or data leaks.
- Security vulnerabilities exploited due to misconfigured or malicious hybrid Runbooks.
- Unplanned changes to business logic or workflows.
- Potential data breaches, unauthorized access to resources, disruption of critical business services, and reputational damage.

Steps to Verify

- Review the Azure Activity Logs for the suspicious event, focusing on the user/service principal and the modified hybrid Runbook.
- Inspect the Runbook code for signs of malicious activity, such as code injection or data exfiltration. To view the Runbook:
 - Navigate to the 'Automation Accounts' service in Azure
 - Identify the Automation Account associated to the Runbook
 - The Runbooks can be found under the 'Process Automation' tab for the selected Automation Account
- Investigate the user's or service principal's permissions and access levels within Azure Automation. Verify if other security alerts or notifications were triggered around the time of the suspicious event. Consult with Azure administrators, security teams, and relevant stakeholders to determine the cause and scope of the incident.

Azure Suspicious Automation Staged

Command & Control



MITRE | ATT&CK®

T1098 Account Manipulation

T1071 Enterprise

T1563 Remote Service Session Hijacking

Triggers

- Unusual changes to Runbook publishing permissions.
- Unexpected publication (staging) of new or modified Runbooks, especially if they contain unusual logic or parameters.
- Runbooks being published from an unexpected or unauthorized user or service principal.
- Runbooks are being published (staged) for legitimate business use-cases.

Possible Root Causes

- A compromised account is attempting unauthorized modifications.
- A developer is staging a new Runbook as part of a legitimate development process but with unusual parameters or access levels.
- Automated deployment scripts not previously used are publishing Runbooks with malicious intent.
- Insider threat: an authorized user intentionally publishes malicious Runbooks.

Business Impact

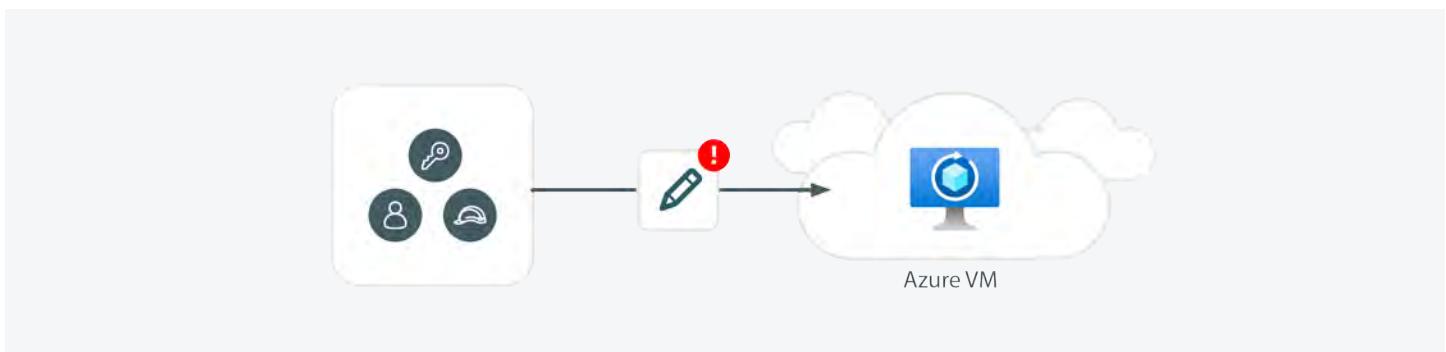
- Exposure of sensitive data through unauthorized access or data leaks.
- Security vulnerabilities exploited due to misconfigured Runbooks or storage accounts.
- Unplanned changes to business logic or workflows, potentially disrupting critical services.
- Potential data breaches, unauthorized access to resources, disruption of critical business services, and reputational damage.

Steps to Verify

- Review the Azure Activity Logs for the suspicious event, focusing on the user/service principal and the published runbook.
- Investigate the user's or service principal's permissions and access levels within Azure Automation.
- Verify if other security alerts or notifications were triggered around the time of the suspicious event.
- Inspect the Runbook code for signs of malicious activity, such as code injection or data exfiltration. To view the Runbook:
 - Navigate to the 'Automation Accounts' service in Azure
 - Identify the Automation Account associated to the Runbook
 - The Runbooks can be found under the 'Process Automation' tab for the selected Automation Account
- Consult with Azure administrators, security teams, and relevant stakeholders to determine the cause and scope of the incident.

Azure Suspicious VM Automation Execution

Command & Control



MITRE | ATT&CK®

T1563 Remote Service Session Hijacking

Triggers

- Unusual execution of a VM Azure Runbook by an unexpected or unauthorized user/service principal.
- Unusual changes to Runbook execution permissions.
- Runbooks are being executed (deployed) for legitimate business use-cases.

Possible Root Causes

- A compromised principal account is attempting to execute malicious Runbooks.
- An administrator or developer has inadvertently created or modified a Runbook with unusual logic.
- Automated deployment scripts are updating Runbooks without proper authorization.
- A legitimate business process involves running Runbooks, but the execution frequency or parameters have been altered.

Business Impact

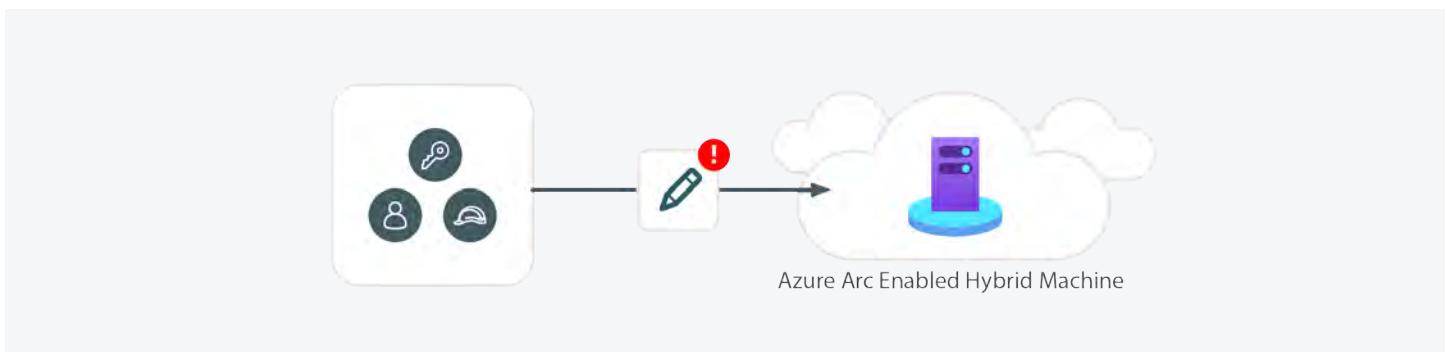
- Exposure of sensitive data due to unauthorized access or data leaks.
- Security vulnerabilities exploited through misconfigured VMs or Runbooks.
- Disruption of critical business services and reputational damage due to unplanned system downtime.
- Unauthorized changes to business logic or workflows, leading to financial losses or compliance issues.

Steps to Verify

- Review Azure Activity Logs for the suspicious event, focusing on the user/service principal and the executed Runbook.
- Investigate the user's or service principal's permissions and access levels within Azure.
- Verify if other security alerts or notifications were triggered around the time of the suspicious event.
- Inspect the Runbook code for signs of malicious activity. To view the Runbook:
 - Navigate to the 'Automation Accounts' service in Azure
 - Identify the Automation Account associated to the Runbook
 - The Runbooks can be found under the 'Process Automation' tab for the selected Automation Account
- Consult with Azure administrators, security teams, and relevant stakeholders to determine the cause and scope of the incident.
- If review indicates malicious actions, isolate the virtual machine for further investigation and remediation.

Azure Suspicious Hybrid Automation Execution

Command & Control



MITRE | ATT&CK®

T1563 Remote Service Session Hijacking

Triggers

- Unusual execution of an Azure Automation Runbook on a hybrid machine by an unexpected or unauthorized user/service principal.
- Unusual changes to Runbook execution permissions.
- Runbooks are being executed (deployed) on hybrid machines for legitimate business use-cases.

Possible Root Causes

- A compromised principal account is attempting to execute malicious Runbooks.
- An administrator or developer has inadvertently created or modified a Runbook with unusual logic.
- Automated deployment scripts are updating Runbooks without proper authorization.
- A legitimate business process involves running Runbooks, but the execution frequency or parameters have been altered.

Business Impact

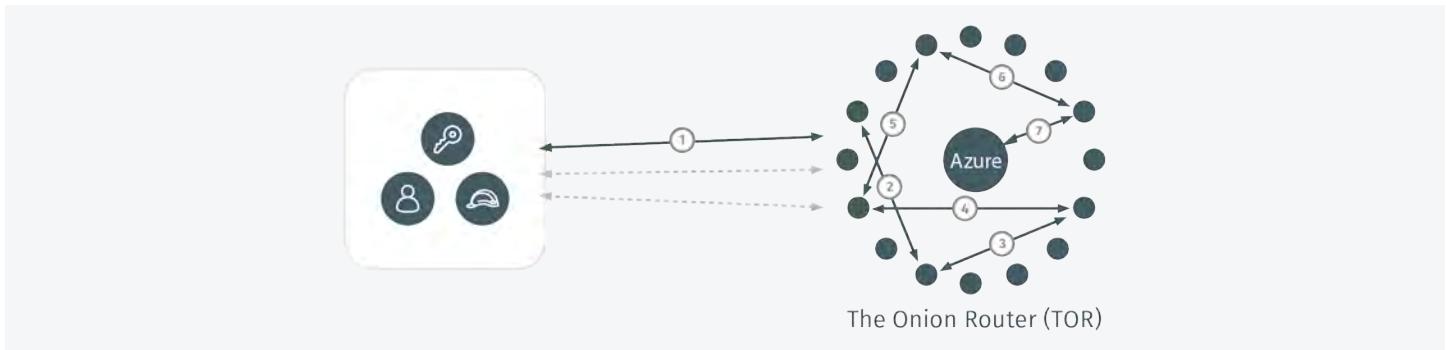
- Exposure of sensitive data due to unauthorized access or data leaks.
- Security vulnerabilities exploited through misconfigured VMs or runbooks or hybrid environments.
- Disruption of critical business services and reputational damage due to unplanned system downtime.
- Unauthorized changes to business logic or workflows, leading to financial losses or compliance issues.

Steps to Verify

- Review Azure Activity Logs for the suspicious event, focusing on the user/service principal and the executed hybrid Runbook.
- Investigate the user's or service principal's permissions and access levels within Azure.
- Verify if other security alerts or notifications were triggered around the time of the suspicious event.
- Inspect the Runbook code for signs of malicious activity. To view the Runbook:
 - Navigate to the 'Automation Accounts' service in Azure
 - Identify the Automation Account associated to the Runbook
 - The Runbooks can be found under the 'Process Automation' tab for the selected Automation Account
- Consult with Azure administrators, security teams, and relevant stakeholders to determine the cause and scope of the incident.
- If review indicates malicious actions, isolate the hybrid machine group for further investigation or member workers and remediate.

Azure TOR Activity

Command & Control



MITRE | ATT&CK®

T1090 Proxy

Triggers

- Control plane activity was observed in patterns consistent with TOR usage.

Possible Root Causes

- Unauthorized Access: An attacker uses TOR to access an Azure identity and execute Azure operations without permission.
- Misconfigured Permissions: A user or service account is granted excessive permissions, allowing them to use TOR for malicious activities.
- Malware or Vulnerability Exploitation: Malware or a vulnerability in the VM's operating system allows an attacker to access and utilize TOR from within the Azure environment.

Business Impact

- Potential data breaches due to unauthorized access to sensitive information stored on Azure resources.
- Compliance risks related to security regulations, such as GDPR, HIPAA, or PCI-DSS, if sensitive data is compromised.
- Downtime and revenue loss due to the disruption of critical business services hosted on Azure.

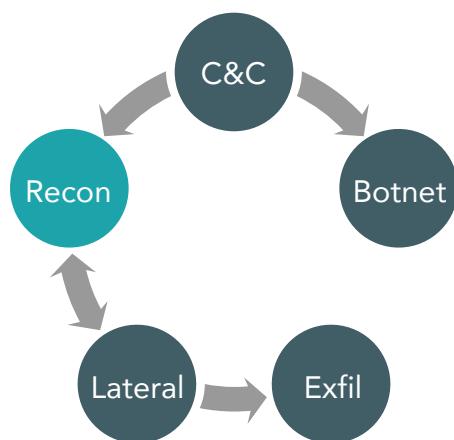
Steps to Verify

- Analyze Network Traffic: Use network monitoring tools and Azure logs to investigate the source IP addresses connected to the TOR nodes and verify their legitimacy.
- Inspect Security Group Rules: Review NSG rules to ensure they are configured correctly and not inadvertently allowing TOR traffic.
- Check for Malware or Vulnerabilities: Perform reviews on affected Azure resources to identify potential entry points used by attackers to access the TOR network.
- Review User Access and Permissions: Investigate user accounts with excessive permissions or unusual activity patterns, which may be linked to the observed TOR usage.

Category

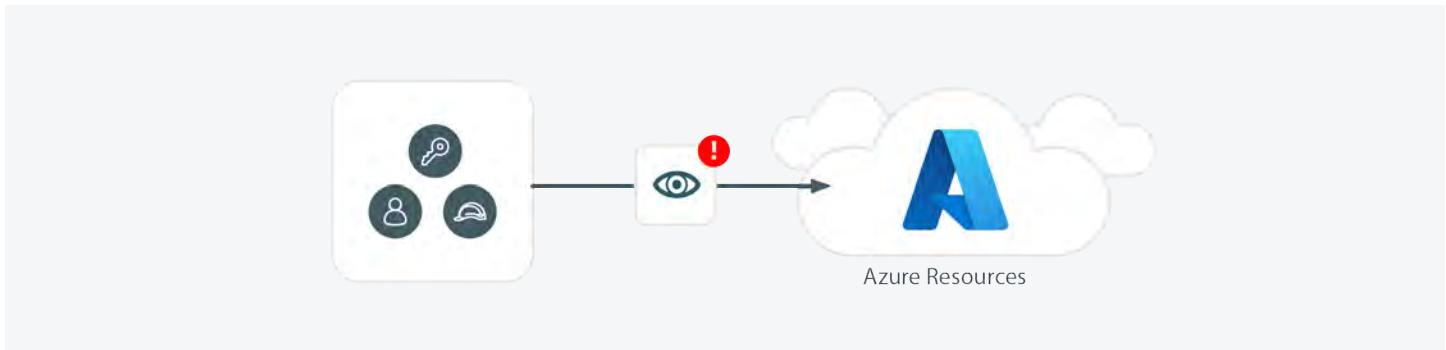
Reconnaissance

- An attacker is surveying and learning about Azure infrastructure
- Attackers with control of an account will look to identify paths to their final objectives
- Attackers will probe Azure services in-order to find credentials to gain additional access
- Attackers will identify services in order to collect, exfiltrate, or impact data
- Actions are associated with the [Discovery](#) MITRE Tactic.



Azure Suspect Credential Dump

Reconnaissance



MITRE | ATT&CK®

[T1555 Credentials from Password Stores](#)

[T1552 Unsecured Credentials](#)

Triggers

- Unusually high number of export actions were performed targeting credentials stored within Azure services, such as Azure Storage.
- Credentials were accessed or exported from multiple Azure services in a short duration, indicating potential credential harvesting.
- Credentials were accessed by accounts or service principals that have not previously accessed these services or credentials.

Possible Root Causes

- A threat actor has obtained access to an account or service principal and is attempting to enumerate or exfiltrate credentials.
- Automated scripts or scheduled jobs may be exporting credentials for backup, disaster recovery, or compliance purposes.
- Credential export actions driven by key rotation policies or regulatory requirements. Internal teams may conduct security tests or compliance audits that involve bulk access to credentials across services.

Business Impact

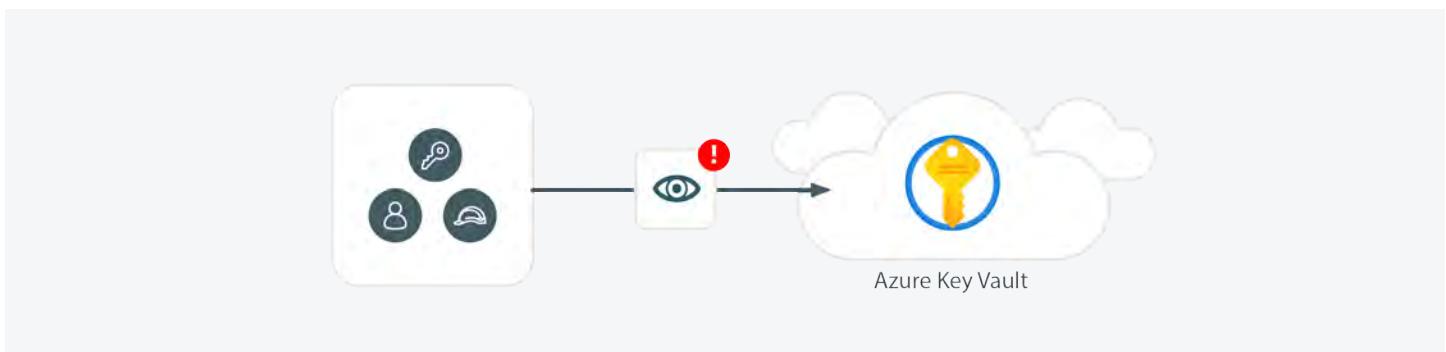
- Exposed Azure credentials may enable attackers to access critical resources, disrupt operations, and escalate privileges within the environment.
- Exposed credentials could allow attackers to steal sensitive data, such as customer information or intellectual property.
- Attackers may use compromised credentials to deploy costly resources, like VMs, for cryptomining.
- Unauthorized access could disrupt critical applications by modifying or deleting essential resources.
- Credential exposure could lead to privilege escalation, granting attackers full control over cloud environments.

Steps to Verify

- Analyze Check the activity logs for unusual credential export activity, focusing on the frequency, timing, and signs of automated exports or bulk actions inconsistent with normal operations.
- Verify if the account or service principal involved shows deviations from its typical access patterns, especially across Azure services in question.
- Inspect the permissions of the account or principal to ensure they were not recently elevated.
- Actions if confirmed suspicious:
 - Disable or rotate affected credentials to prevent further misuse.
 - Restrict access from unusual IP ranges or locations involved in the activity.
 - Perform a deeper investigation of affected resources to assess the extent of access and identify any further compromised assets.
- Notify relevant stakeholders and document the event in the incident management system to ensure appropriate follow-up.

Azure Suspect Key Vault Enumeration

Reconnaissance



MITRE | ATT&CK®

T1595 Active Scanning

Triggers

- Unusually high number of list or metadata access requests for secrets, certificates, or keys within a short time period.
- Enumeration attempts by accounts or service principals that have not previously accessed Key Vault resources.
- Multiple requests within a short duration, suggesting automated or scripted enumeration.

Possible Root Causes

- A compromised account or service principal is attempting to gather information on available Key Vault resources.
- Legitimate security or compliance scans are being conducted by internal teams, but with high volume or atypical timing.
- Automated scripts or tools are being used by internal or external actors to list or enumerate Key Vault items for discovery.

Business Impact

- Enumerating Key Vault resources may enable attackers to identify high-value secrets or keys, increasing the risk of targeted attacks.
- Results of enumerating key vaults can inform lateral movement or privilege escalation within the cloud environment.

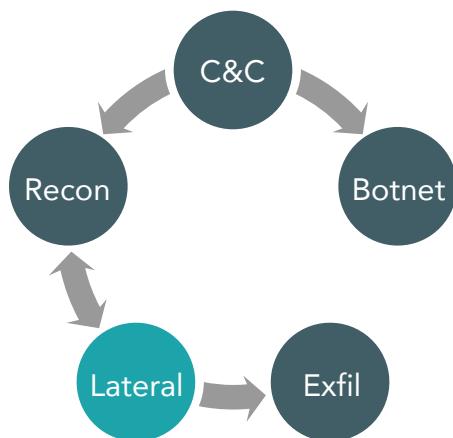
Steps to Verify

- Check Key Vault logs for an unusually high frequency of list or metadata access requests focused on secrets, certificates, or keys
- Verify if the accessing account or service principal shows unusual behavior, such as first-time access or access outside regular hours.
- Confirm that Key Vault permissions were not recently modified.
- Actions if confirmed suspicious:
 - Limit or revoke permissions for any accounts or service principals involved in the suspicious enumeration activity.
 - Apply temporary restrictions to IPs or regions associated with the unusual access patterns.
 - Conduct a review of Key Vault permissions and access policies to ensure they are properly restricted and monitored.
 - Notify the security team and document the event in the incident management system for follow-up and potential escalation.

Category

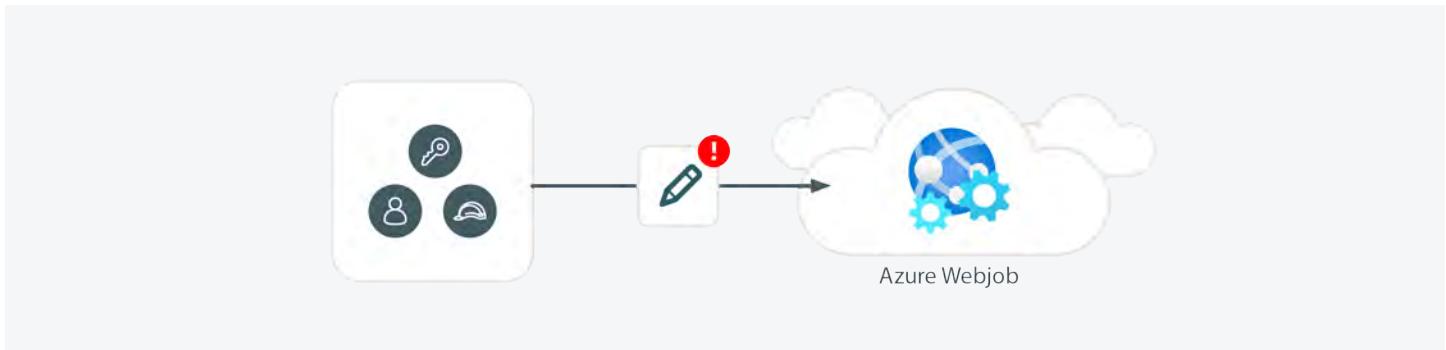
Lateral Movement

- An attacker in the Azure environment is spreading and taking actions that ensure continuous undetected access
- Attackers after gaining access to the credentials and discovering the environment will propagate and solidify their access
- Attackers will take actions to modify services and identifies to ensure continued access
- Attackers will leverage credentials within their defined permissions but for non-intended purposes to further the attacker's objective
- Actions are associated with [Execution](#), [Persistence](#), [Privilege Escalation](#), [Defense Evasion](#), [Lateral Movement](#), [Credential Access](#) MITRE Tactics



Azure Anomalous App Service WebJob Activity

Lateral Movement



MITRE | ATT&CK®

[T1546 Event Triggered Execution](#)

[T1574 Hijack Execution Flow](#)

Triggers

- Unusual creation or modification Azure App Service WebJob configurations.
- Frequent or repetitive execution of WebJobs with unknown or suspicious triggers.
- Changes to existing WebJobs that could indicate a malicious intent.

Possible Root Causes

- A compromised identity is attempting deployment of a malicious application.
- Unauthorized user or service principal creating or modifying WebJobs with elevated permissions.
- Misconfigured or compromised Azure resources leading to unexpected behavior.
- Inadequate monitoring or security practices allowing suspicious activity to go unnoticed.
- A developer creating a new WebJob or modifying an existing one as part of a legitimate development processes.

Business Impact

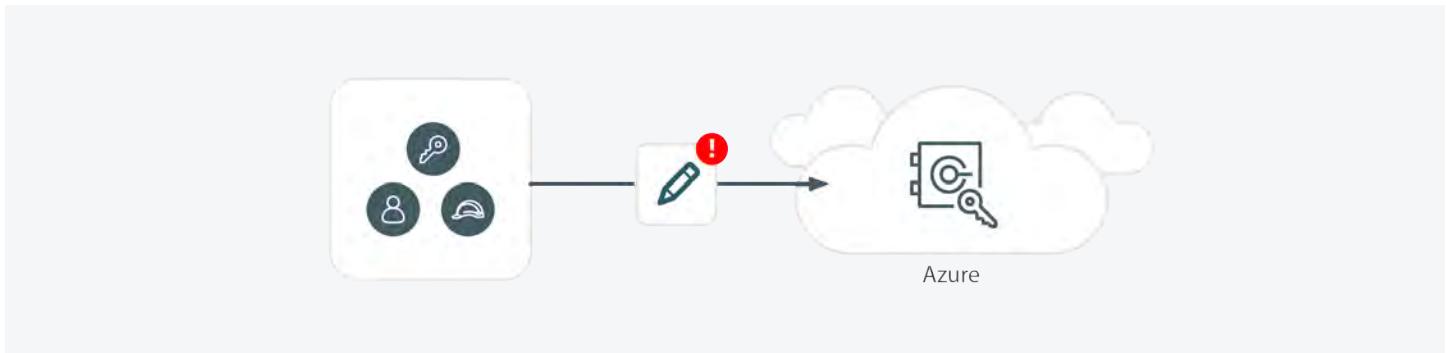
- Data breaches and unauthorized access to sensitive information.
- Denial-of-Service (DoS) attacks or resource exhaustion due to malicious code execution.
- Compliance and regulatory risks due to security vulnerabilities.
- Unintended changes to production data or systems due to unknown or untested WebJobs.

Steps to Verify

- Investigate the Azure resource group or subscription where suspicious WebJob activity was detected.
- Review the Azure Activity Logs for any suspicious activity related to WebJobs, such as creation, modification, or execution.
- Investigate the WebJob configurations and code for any signs of malicious intent or unauthorized access.
- Verify user and service principal permissions to ensure they align with expected roles and responsibilities.

Azure Diagnostic Logging Disabled

Lateral Movement



MITRE | ATT&CK®

[T1070 Indicator Removal on Host](#)

[T1578 Modify Cloud Compute Infrastructure](#)

[T1562 Impair Defenses](#)

Triggers

- Disable or delete Diagnostic logging for an Azure resource or subscription.

Possible Root Causes

- An attacker has deleted Diagnostic logging to hide their tracks.
- An administrator has disabled Diagnostic logging as part of normal changes to the environment

Business Impact

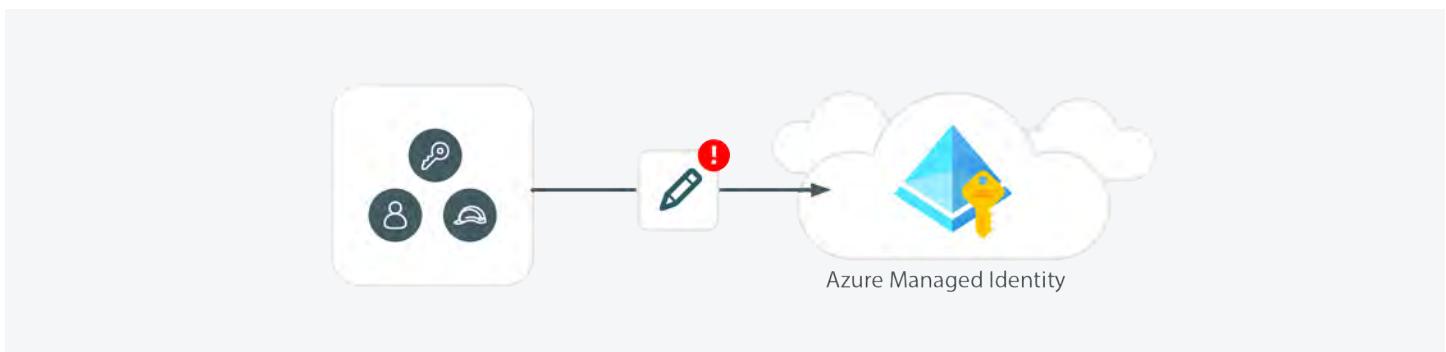
- Inability to detect future attacks, investigate future or historical attacks, or audit activity within the environment.
- Increased risk of activity that may negatively impact the business going unnoticed.

Steps to Verify

- Review the actions being undertaken by the user after the identified activity and potential risk posed by that access where logging continues to be enabled.
- Review security policy to determine if the removal of Diagnostic logging capabilities is allowed.
- Discuss with the user to determine if the activity is known and legitimate.
- If the review determines there is a high risk to data or the environment, disable the credentials and perform a comprehensive investigation.

Azure Managed Identity Anomaly

Lateral Movement



MITRE | ATT&CK®

T1563 Remote Service Session Hijacking

Triggers

- Successful login attempts using / usage of Azure Managed Identity credentials from unusual IP addresses or locations.
- Successful usage of Azure Managed Identity to access unusual subscription, resource type or services.
- Successful usage of Azure Managed Identity to perform unusual operations.

Possible Root Causes

- Successful login attempts using / usage of Azure Managed Identity credentials from unusual IP addresses or locations.
- Successful usage of Azure Managed Identity to access unusual subscription, resource type or services.
- Successful usage of Azure Managed Identity to perform unusual operations.

Business Impact

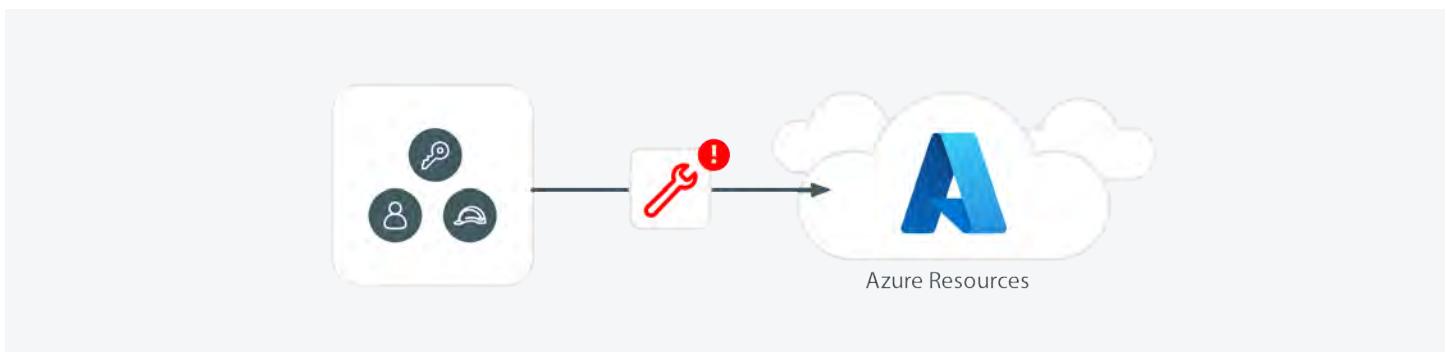
- Successful login attempts using / usage of Azure Managed Identity credentials from unusual IP addresses or locations.
- Successful usage of Azure Managed Identity to access unusual subscription, resource type or services.
- Successful usage of Azure Managed Identity to perform unusual operations.

Steps to Verify

- Review Azure Activity Logs for the suspicious event, focusing on the user/service principal and the created or modified Managed Identity.
- Investigate the user's or service principal's permissions and access levels within Azure.
- Check the Managed Identity configuration and settings for unusual changes.
- Consult with Azure administrators, security teams, and relevant stakeholders to determine the cause and scope of the incident.
- Analyze login attempts and successful authentication events associated with the compromised managed identity.
- Perform a thorough audit of Azure subscriptions and resources accessed by the compromised managed identity.

Azure Mass Resource Deletion

Lateral Movement



MITRE | ATT&CK®

[T1485 Data Destruction](#)

[T1561 Disk Wipe](#)

[T1562 Impair Defenses](#)

Triggers

- Unusual pattern of deletion associated with various Azure cloud resources.
- Deletion of resources as part of legitimate downsizing or other supporting business use-cases.

Possible Root Causes

- An attacker is attempting to interrupt operations by destroying core resources in the Azure environment.
- An attacker is attempting to evade defenders by removing Azure resources which facilitate blue team response.
- An automated process within the customer environment is deploying and removing a significant volume of resources.

Business Impact

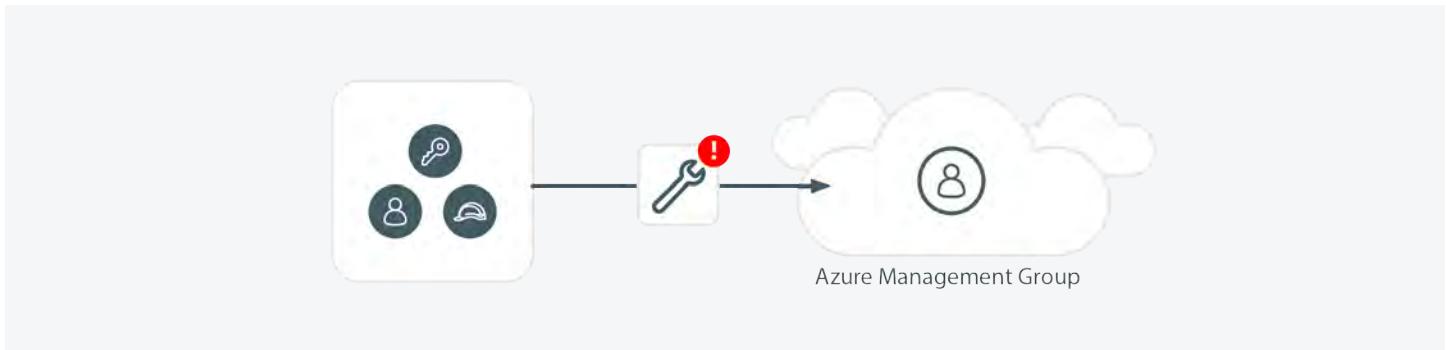
- Mass deletion of resources results significant risk of impact to normal operations.
- By default, many Azure resources cannot be recovered after they are deleted, increasing the level of effort required to restore normal operations.

Steps to Verify

- Mass deletion of resources results significant risk of impact to normal operations.
- By default, many Azure resources cannot be recovered after they are deleted, increasing the level of effort required to restore normal operations.

Azure Privilege Anomaly: Management Group Scope

Lateral Movement



MITRE | ATT&CK®

T1548 Abuse Elevation Control

T1098 Account Manipulation

Triggers

- Unusual use of a highly permissive role observed within a Management Group

Possible Root Causes

- An attacker is leveraging elevated permissions to gain additional or persistent access to the environment.
- An administrator is using a highly permissive role to enable them complete access to the environment.

Business Impact

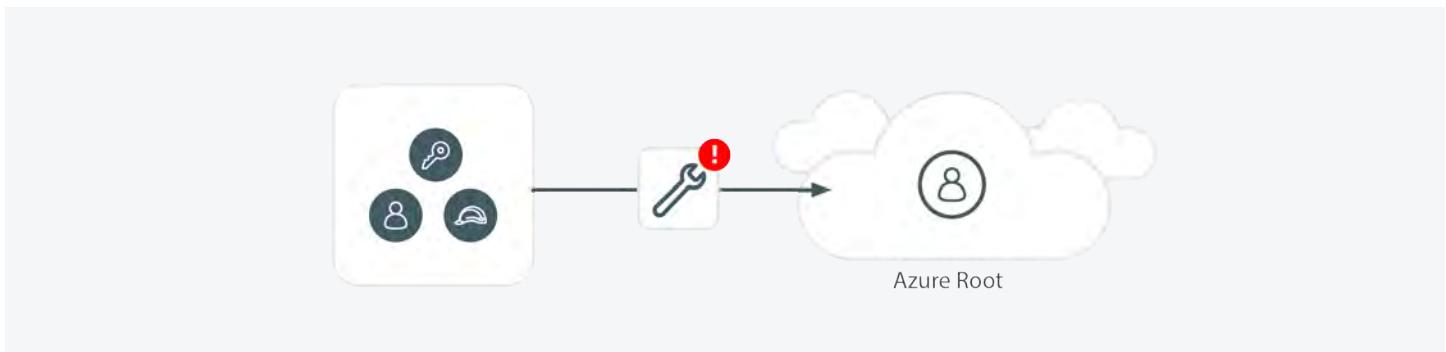
- Lateral movement may indicate that an adversary has established a foothold in the environment and is progressing towards their objective, increasing the risk of material impact

Steps to Verify

- Investigate the Principal which performed the actions for other signs of malicious activity.
- Review whether this account should have the assigned level of privilege for their normal duties.
- If review indicates possible malicious actions or high-risk configurations:
 - Revert any configuration changes.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and scope of impacted resources.

Azure Privilege Anomaly: Root Scope

Lateral Movement



MITRE | ATT&CK®

T1548 Abuse Elevation Control

T1098 Account Manipulation

Triggers

- Unusual use of a highly permissive role observed within a Root Scope.

Possible Root Causes

- An attacker is leveraging elevated permissions to gain additional or persistent access to the environment.
- An administrator is using a highly permissive role to enable them complete access to the environment.

Business Impact

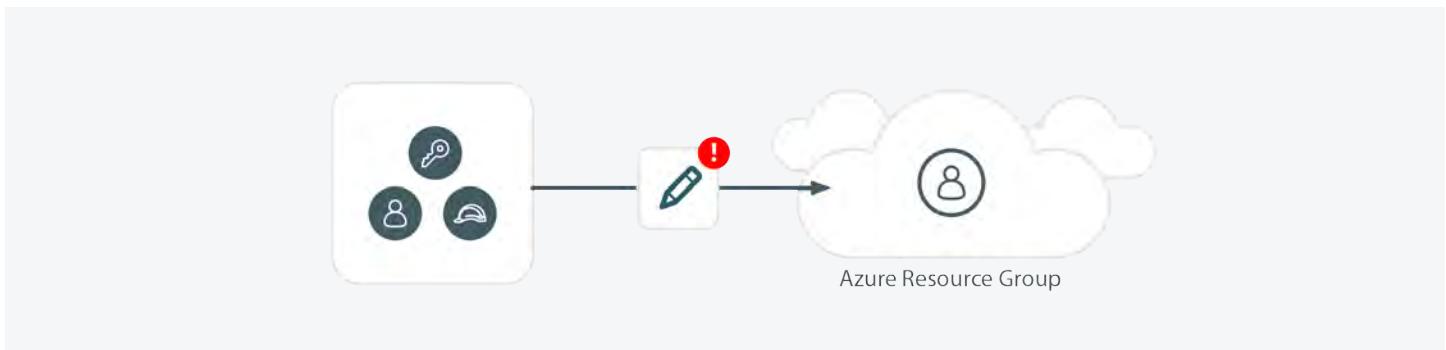
- Lateral movement may indicate that an adversary has established a foothold in the environment and is progressing towards their objective, increasing the risk of material impact.

Steps to Verify

- Investigate the Principal which performed the actions for other signs of malicious activity.
- Review whether this account should have the assigned level of privilege for their normal duties.
- If review indicates possible malicious actions or high-risk configurations:
 - Revert any configuration changes.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and scope of impacted resources.

Azure Resource Group Admin Privilege Granting

Lateral Movement



MITRE | ATT&CK®

T1548 Abuse Elevation Control

T1098 Account Manipulation

Triggers

- Assignment of a highly permissive role to an entity at a Resource Group level.

Possible Root Causes

- An attacker is changing the permissions of an entity to enable them to leverage those permissions to gain additional or persistent access to the environment.
- An administrator has been granted a highly permissive role to enable them complete access to the environment.

Business Impact

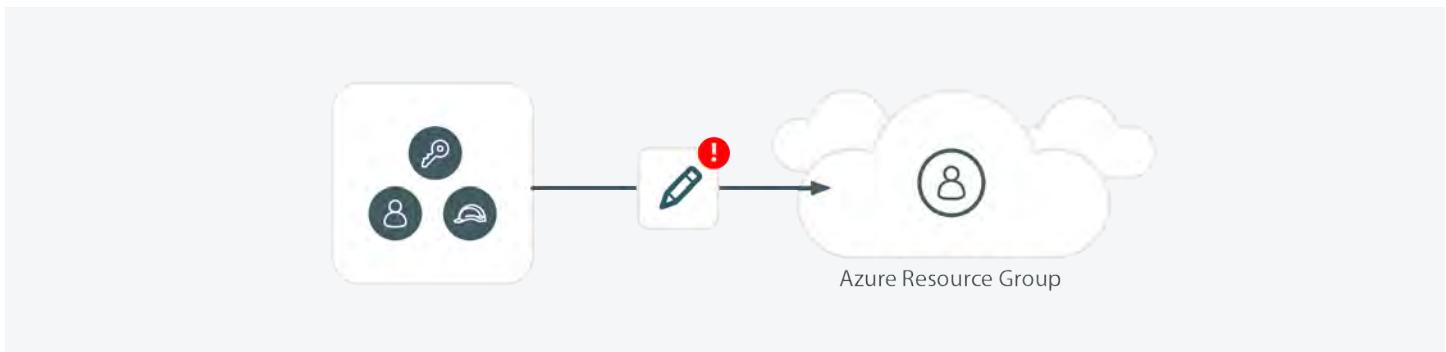
- Lateral movement may indicate that an adversary has established a foothold in the environment and is progressing towards their objective, increasing the risk of material impact.

Steps to Verify

- Investigate the Principal which performed the actions for other signs of malicious activity.
- Review whether this account should have the assigned level of privilege for their normal duties.
- If review indicates possible malicious actions or high-risk configurations:
 - Revert any configuration changes.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and scope of impacted resources.

Azure Resource Group Admin Role Unassigned

Lateral Movement



MITRE | ATT&CK®

T1531 Account Access Removal

Triggers

- Removal of a highly permissive role assigned to an entity at the Resource Group level.

Possible Root Causes

- An attacker is attempting to isolate access to a Resource Group by removing a legitimate administrator. This is done in order to impair defenses and disrupt logging visibility.
- An administrator is performing authorized changes to permissions.

Business Impact

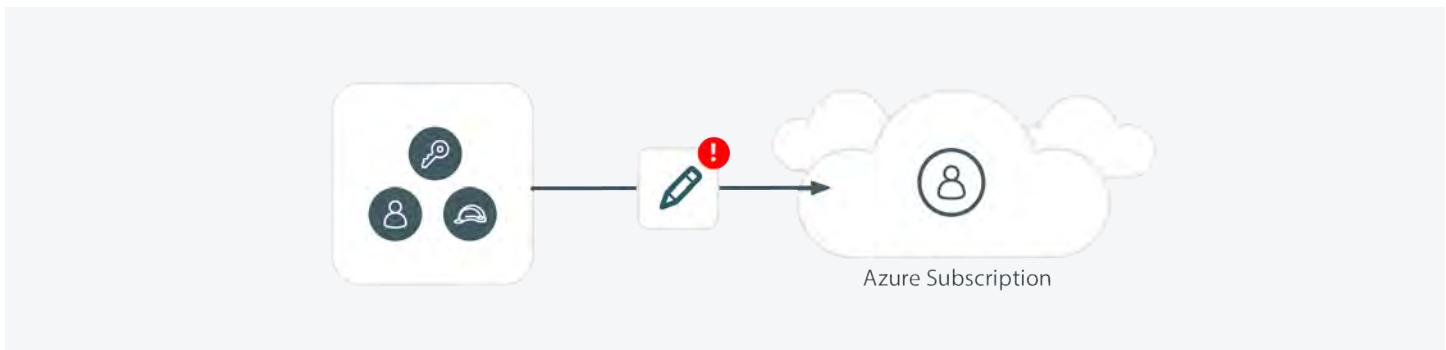
- An attacker can hinder the defenses of their victim compromises the victim's ability to respond and also has the ability to evade detection.

Steps to Verify

- Investigate the Principal which performed the actions for other signs of malicious activity.
- Review security policy to determine if the removal of the privileged role is sanctioned.
- If review indicates possible malicious actions or high-risk modifications:
 - Disable credentials associated with this alert.
 - Regrant privileges within the Resource Group as necessary to restore visibility.
 - Perform a comprehensive investigation to determine initial compromise and the scope of impacted resources.

Azure Subscription Admin Privilege Granting

Lateral Movement



MITRE | ATT&CK®

T1548 Abuse Elevation Control

T1098 Account Manipulation

Triggers

- Assignment of a highly permissive role to an entity at the Subscription scope level

Possible Root Causes

- An attacker is changing the permissions of an entity to enable them to leverage those permissions to gain additional or persistent access to the environment.
- An administrator has been granted a highly permissive role to enable them complete access to the environment.

Business Impact

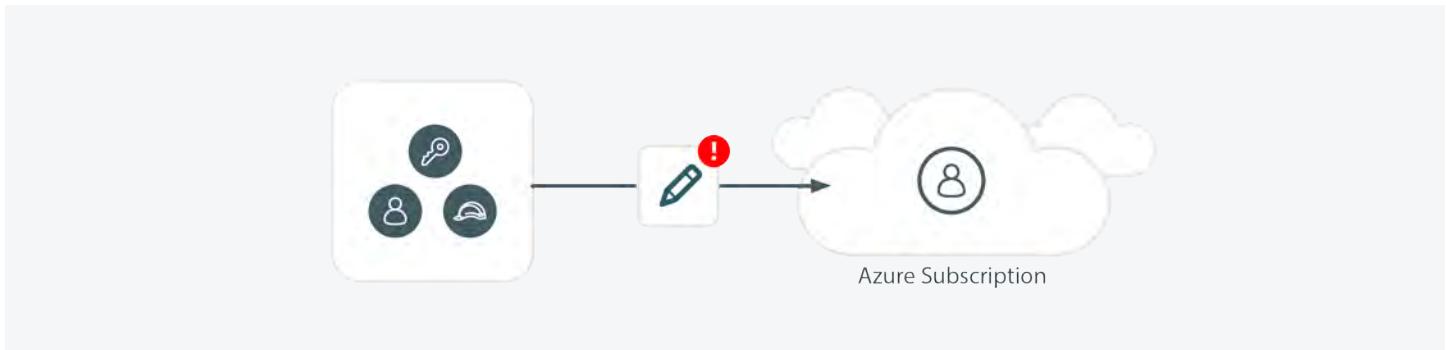
- Lateral movement may indicate that an adversary has established a foothold in the environment and is progressing towards their objective, increasing the risk of material impact.

Steps to Verify

- Investigate the Principal which performed the actions for other signs of malicious activity.
- Review whether this account should have the assigned level of privilege for their normal duties.
- If review indicates possible malicious actions or high-risk configurations:
 - Revert any configuration changes.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and scope of impacted resources.

Azure Subscription Admin Role Unassigned

Lateral Movement



MITRE | ATT&CK®

T1531 Account Access Removal

Triggers

- Removal of a highly permissive role assigned to an entity at the Subscription scope level

Possible Root Causes

- An attacker is attempting to isolate access to a Subscription by removing a legitimate administrator. This is done in order to impair defenses and disrupt logging visibility.
- An administrator is performing authorized changes to permissions.

Business Impact

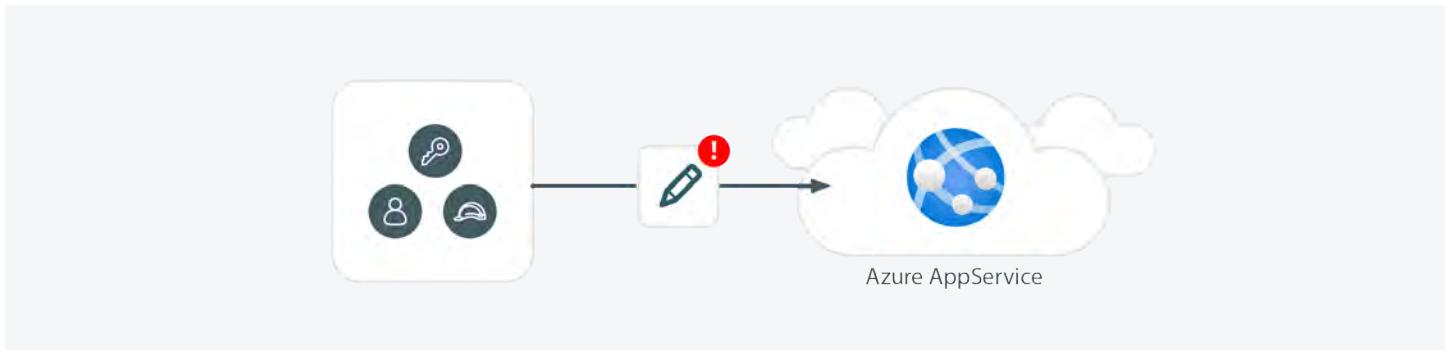
- An attacker who is able to hinder the defenses of their victim compromises the victim's ability to respond and also has the ability to evade detection.

Steps to Verify

- Investigate the Principal which performed the actions for other signs of malicious activity.
- Review security policy to determine if the removal of the privileged role is sanctioned.
- If review indicates possible malicious actions or high-risk modifications:
 - Disable credentials associated with this alert.
 - Regrant privileges within the Subscription as necessary to restore visibility.
 - Perform a comprehensive investigation to determine initial compromise and the scope of impacted resources.

Azure Suspect App Service Deployment Activity

Lateral Movement



MITRE | ATT&CK®

[T1525 Implant Internal Image](#)

[T1546 Event Triggered Execution](#)

[T1574 Hijack Execution Flow](#)

Triggers

- Modification of existing Azure Function App Service Deployment Slot with unusual parameters or logic.
- Unusual changes to App Service Deployment Slot triggers.
- Unusual creation or modification of Azure Functions Deployment Slot by an unexpected or unauthorized user/service principal, potentially indicating malicious activity.

Possible Root Causes

- A compromised principal account is attempting unauthorized modifications
- A developer creating or modifying an Azure Function App Service runtime version.
- Automated deployment scripts not previously used are updating the function code or configuration.
- A developer creating a new function or modifying an existing one as part of a legitimate development process.

Business Impact

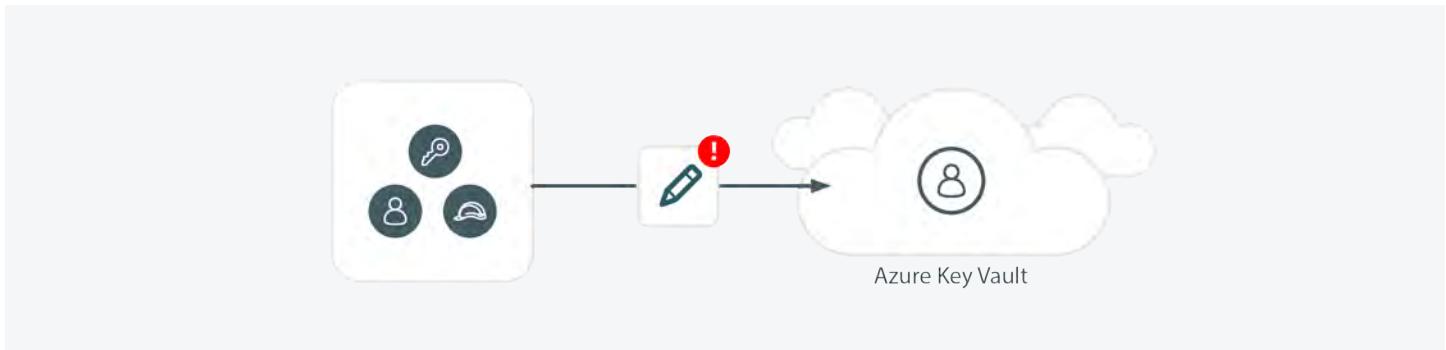
- Exposure of sensitive data through unauthorized access or data leaks.
- Security vulnerabilities exploited due to misconfigured functions.
- Unplanned changes to business logic or workflows.
- Potential data breaches, unauthorized access to sensitive resources, disruption of critical business services, and reputational damage.

Steps to Verify

- Review the Azure Activity Logs for the suspicious event, focusing on the user/service principal and the created or modified App Service Deployment slots.
- Investigate the user's or service principal's permissions and access levels within Azure.
- Verify if other security alerts or notifications were triggered around the time of the suspicious event.
- Check the Azure Function Deployment Slot code for signs of malicious activity.
- Consult with Azure administrators, security teams, and relevant stakeholders to determine the cause and scope of the incident.

Azure Suspect Key Vault Privilege Granting

Lateral Movement



MITRE | ATT&CK®

T1548 Abuse Elevation Control

T1098 Account Manipulation

Triggers

- Assignment of a privileged Key Vault role to an entity for a given Key Vault.

Possible Root Causes

- An attacker is assigning permissions to an entity to enable them to leverage those permissions to gain additional or persistent access to the environment.
- An administrator has been granted a highly permissive role to enable them complete access to the environment.

Business Impact

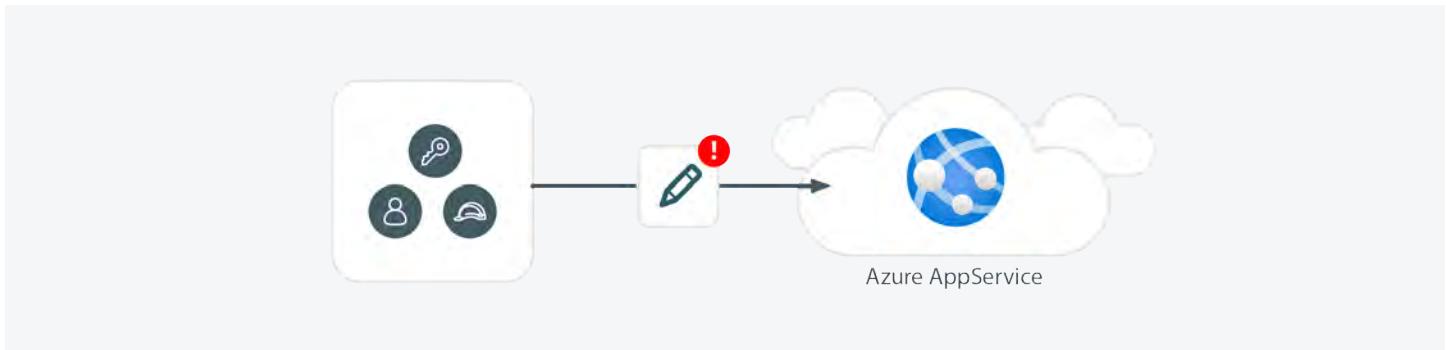
- Lateral movement may indicate that an adversary has established a foothold in the environment and is progressing towards their objective, increasing the risk of material impact.

Steps to Verify

- Investigate the Principal which performed the actions for other signs of malicious activity.
- Review whether this account should have the assigned level of privilege for their normal duties.
- If review indicates possible malicious actions or high-risk configurations:
 - Revert any configuration changes.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and scope of impacted resources.

Azure Suspicious App Service Creation or Modification

Lateral Movement



MITRE | ATT&CK®

[T1525 Implant Internal Image](#)

[T1546 Event Triggered Execution](#)

[T1574 Hijack Execution Flow](#)

Triggers

- Modification of existing Azure Function App Service with unusual parameters or logic.
- Unusual changes to App Service triggers, bindings, or storage accounts.
- Unusual creation or modification of Azure Functions by an unexpected or unauthorized user/service principal, potentially indicating malicious activity.

Possible Root Causes

- A compromised principal account is attempting unauthorized modifications.
- A developer creating or modifying an Azure Function App Service runtime version.
- Automated deployment scripts not previously used are updating the function code or configuration.
- A developer creating a new function or modifying an existing one as part of a legitimate development process.

Business Impact

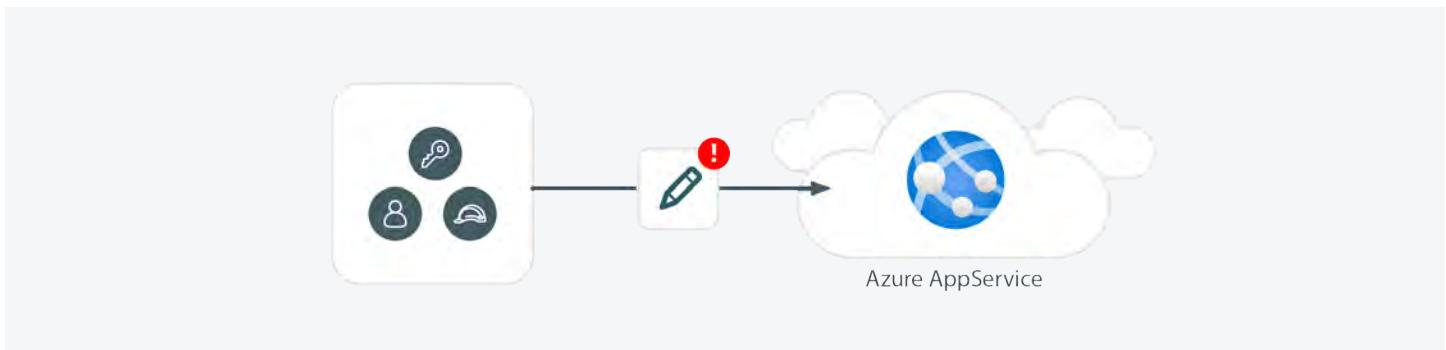
- Exposure of sensitive data through unauthorized access or data leaks.
- Security vulnerabilities exploited due to misconfigured functions or storage accounts.
- Unplanned changes to business logic or workflows.
- Potential data breaches, unauthorized access to sensitive resources, disruption of critical business services, and reputational damage.

Steps to Verify

- Review the Azure Activity Logs for the suspicious event, focusing on the user/service principal and the created or modified App Service.
- Investigate the user's or service principal's permissions and access levels within Azure.
- Verify if other security alerts or notifications were triggered around the time of the suspicious event.
- Check the Azure Function code for signs of malicious activity.
- Consult with Azure administrators, security teams, and relevant stakeholders to determine the cause and scope of the incident.

Azure Suspicious App Service Credential Download

Lateral Movement



MITRE | ATT&CK®

[T1555 Credentials from Password Stores](#)

[T1552 Unsecured Credentials](#)

Triggers

- Unusual access or unexpected changes made to App Service configurations.
- Creation or modification of deployment credentials for a Web App
- Use of 'az cli' or azure PowerShell commands to retrieve or modify deployment credentials.

Possible Root Causes

- Unusual access or unexpected changes made to App Service configurations.
- Creation or modification of deployment credentials for a Web App
- Use of 'az cli' or azure PowerShell commands to retrieve or modify deployment credentials.

Business Impact

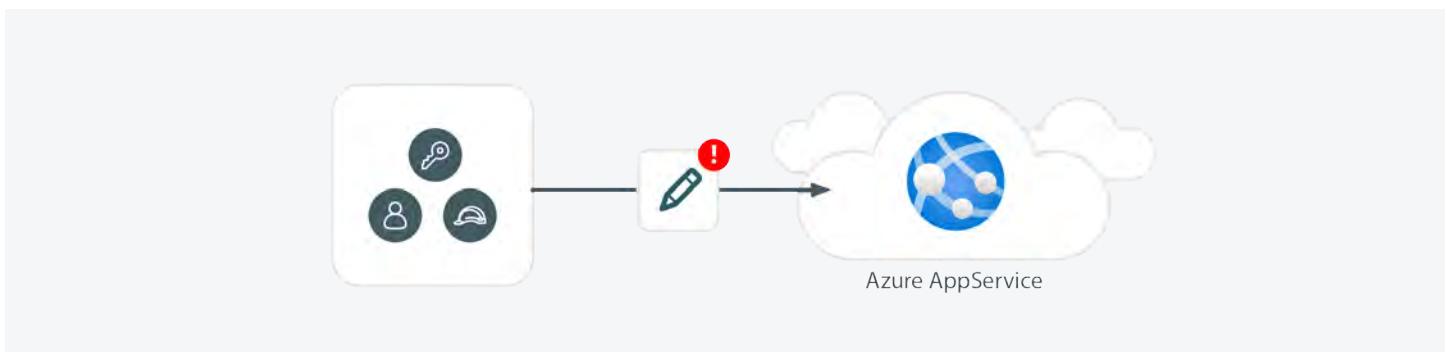
- Unusual access or unexpected changes made to App Service configurations.
- Creation or modification of deployment credentials for a Web App
- Use of 'az cli' or azure PowerShell commands to retrieve or modify deployment credentials.

Steps to Verify

- Unauthorized access to intellectual property (e.g., source code, configuration files)
- Potential data breaches, unauthorized access to sensitive resources, disruption of critical business services, and reputational damage.
- Exposure of sensitive data through unauthorized access or data leaks.
- Service disruptions or downtime due to unauthorized changes made to the Web App
- Compliance and regulatory issues due to inadequate security controls.

Azure Suspicious App Service Deployment Configuration Download

Lateral Movement



MITRE | ATT&CK®

[T1555 Credentials from Password Stores](#)

[T1552 Unsecured Credentials](#)

Triggers

- Unusual access or unexpected changes made to App Service deployment configurations.
- API requests to retrieve deployment configuration “credentials” from a specific slot.

Possible Root Causes

- Unauthorized user or service principal accessing App Service credentials.
- Malicious code execution within an Azure Function instance, leading to credential exposure.
- Misconfigured security settings or permissions on the deployment slot.
- Insider threat or compromised account with elevated privileges accessing App Service deployment configuration credentials
- Developers legitimately downloading App Service Slot credentials for debugging or testing purposes.
- IT administrators accessing App Service configurations for maintenance or troubleshooting.
- Automated scripts or tools used by authorized personnel to manage App Service deployments.

Business Impact

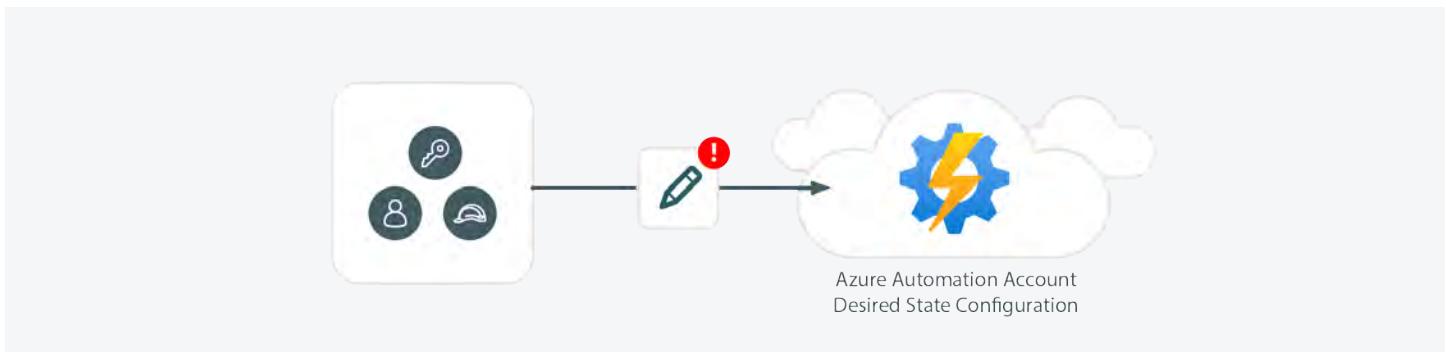
- Unauthorized access to intellectual property (e.g. configuration files)
- Potential data breaches and unauthorized access to sensitive information (e.g., API keys, connection strings).
- Compromise of App Service environment and potential DoS attacks or resource exhaustion.
- Compliance and regulatory issues due to inadequate security controls.

Steps to Verify

- Investigate the specific deployment slot and instance involved in the suspicious activity.
- Review the Azure Activity Logs for API requests and changes made to configuration or environment variables within the affected slot.
- Check for any recent updates or changes in App Service settings, permissions, or access control.
- Run a vulnerability assessment and penetration testing on the affected App Service environment.
- Validate user and service principal permissions to identify potential security issues.
- Review monitoring logs for any signs of suspicious activity or anomalies within the affected slot.

Azure Suspicious Automation DSC Execution

Lateral Movement



MITRE | ATT&CK®

[T1059 Command and Scripting Interpreter](#)

[T1098 Account Manipulation](#)

[T1071 Application Layer Protocol](#)

Triggers

- Unusual changes to VM configuration or Hybrid Environment using Azure Automation Account with DSC (Desired State Configuration).
- Modification of existing DSC configurations with unusual parameters or logic.
- Unusual creation or modification of DSC resources by an unexpected or unauthorized user/service principal.
- Possible misuse of DSC for malicious purposes, such as lateral movement or privilege escalation.

Possible Root Causes

- The VM or hybrid machine has been compromised by an attacker who is using DSC to maintain persistence and evade detection.
- DSC has been misconfigured by a user, leading to unintended behavior and potential security issues.
- An unauthorized user or service principal has gained access to the Azure environment and is using DSC for malicious purposes.
- Automated deployment scripts not previously used are updating the DSC configurations.

Business Impact

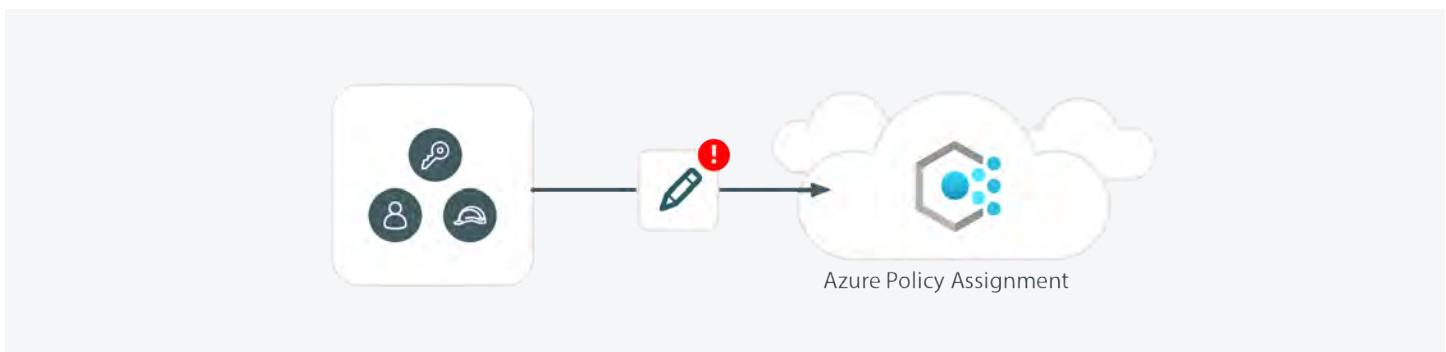
- Exposure of sensitive data through unauthorized access or data leaks
- Security vulnerabilities exploited due to misconfigured resources or DSC configurations
- Unplanned changes to business logic or workflows
- Potential data breaches, unauthorized access to resources, disruption of critical business services, and reputational damage

Steps to Verify

- Review Azure Activity Logs: Investigate the user/service principal and the affected VM for suspicious activity.
- Audit DSC Configurations for Azure Automation Accounts: Review the DSC configurations for unusual changes or modifications. To view the DSC service:
 - Navigate to the 'Automation Accounts' service in Azure.
 - Identify the Automation Account associated to the DSC service.
 - DSC can be found under the 'Configuration Management' tab for the selected Automation Account.
- Check Network Traffic: Examine network traffic to and from the VM to detect any potential lateral movement or communication with malicious actors.
- Conduct a Security Audit: Perform a thorough security audit of the affected VM and its resources to identify any vulnerabilities or misconfigurations.
- Consult with Azure Administrators: Work with Azure administrators, security teams, and relevant stakeholders to determine the cause and scope of the incident.

Azure Suspicious Policy Assignment

Lateral Movement



MITRE | ATT&CK®

T1546 Event Triggered Execution

Triggers

- Policy assignment by a principal (user or service account) who usually does not perform such actions.

Possible Root Causes

- A compromised principal attempting unauthorized policy assignments.
- A legitimate principal with elevated privileges performing actions outside their normal scope of work.
- Human error or misconfiguration leading to unintended policy assignments.

Business Impact

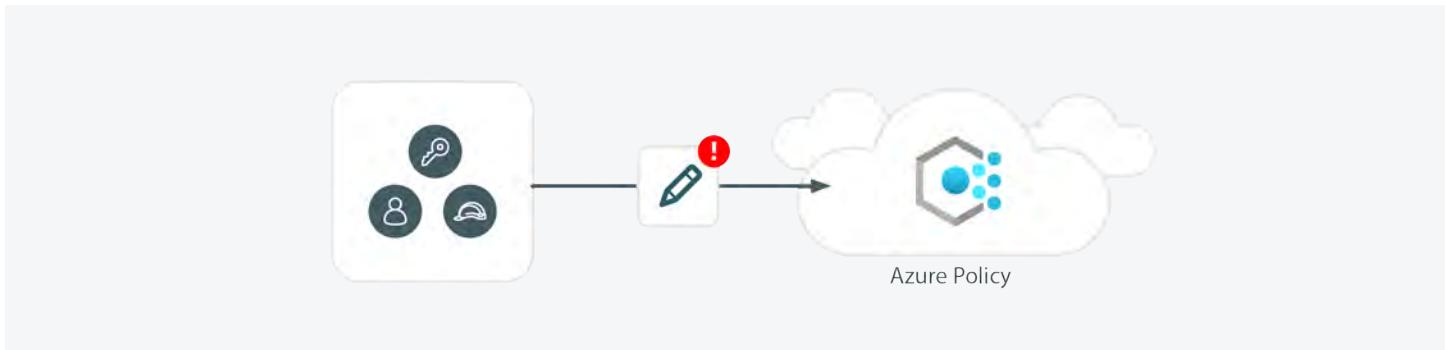
- Enforcement of Non-Compliant Policies: Assignment of policies that enforce configurations not aligned with the organization's security or compliance standards.
- Disruption of Operations: Assignment of restrictive policies that could block or disrupt critical operations.
- Introduction of Security Vulnerabilities: Assignment of policies that inadvertently weaken the security posture of the environment.

Steps to Verify

- Review Azure Activity Logs for the principal's previous activities related to policy assignments.
- Investigate the context around the policy assignment—such as time of day, location, and associated resource changes—to determine if the activity is legitimate.
- If review indicates possible malicious actions or high-risk modifications:
 - Revert modifications made by the principal.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and the scope of impacted resources.

Azure Suspicious Policy Creation or Modification

Lateral Movement



MITRE | ATT&CK®

T1546 Event Triggered Execution

T1059 Command and Scripting Interpreter

Triggers

- Policy creation or modification by a principle who usually does not perform such actions.

Possible Root Causes

- A compromised principal account is attempting unauthorized modifications
- Misconfiguration or human error leading to unintended policy changes.
- A legitimate user is performing actions outside their normal scope of work.

Business Impact

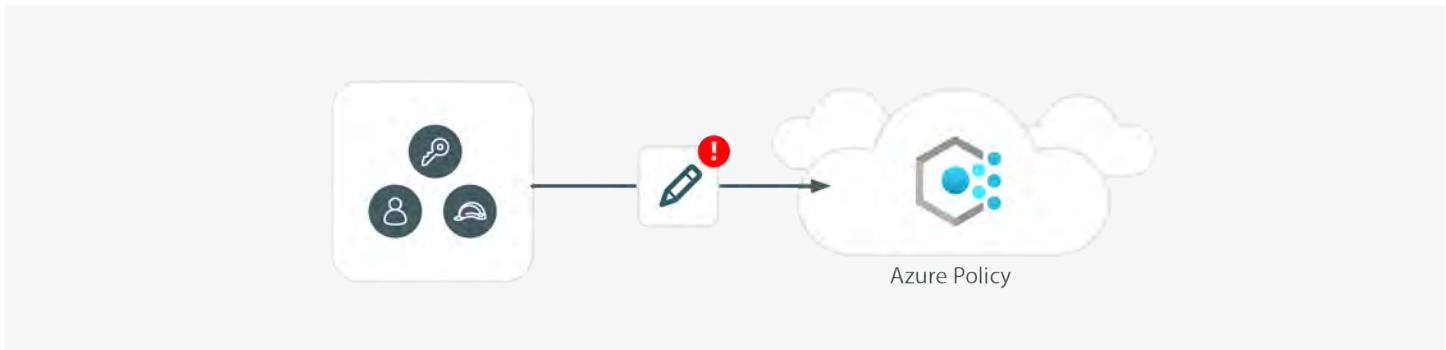
- Modification to Azure policies can lead to a number of consequences depending on the policy effect:
 - DeployIfNotExists: Deployment of resources that introduce vulnerabilities or malicious configurations.
 - Append: Application of tags or settings, weakening security postures.
 - Modify: Direct alteration of resource configurations, leading to misconfigurations.

Steps to Verify

- Check Azure activity logs for the user's previous activities related to policy management.
- Investigate the context around the action, such as associated resource changes to determine if the activity is legitimate.
- If review indicates possible malicious actions or high-risk modifications:
 - Revert modifications made by the principal.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and the scope of impacted resources.

Azure Suspicious Policy Remediation Task

Lateral Movement



MITRE | ATT&CK®

[T1546 Event Triggered Execution](#)

[T1059 Command and Scripting Interpreter](#)

Triggers

- Creation of an Azure Policy remediation task by a user who usually does not perform such actions.

Possible Root Causes

- A compromised principal is attempting to create a remediation task to exploit Azure resources.
- A legitimate principal with elevated privileges is performing actions outside their typical scope of work.

Business Impact

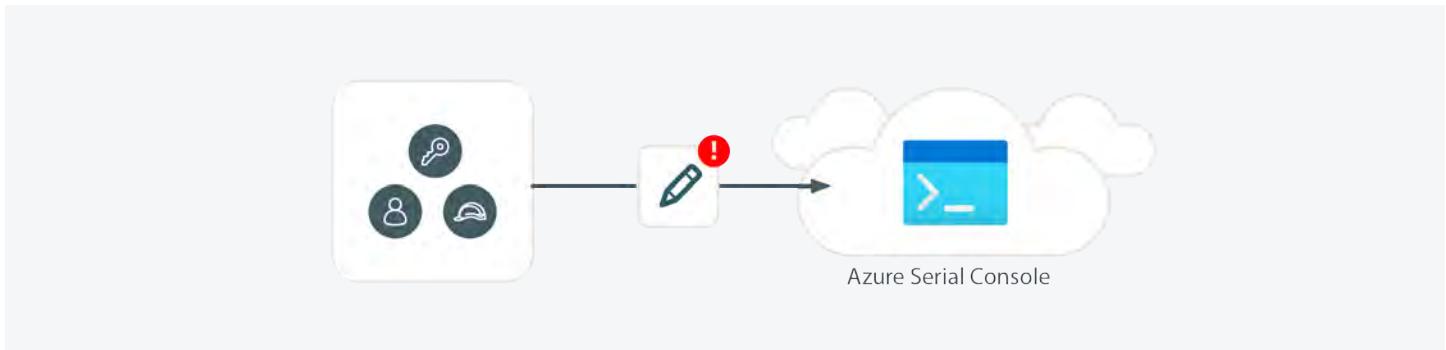
- A remediation task could automatically deploy unwanted resources, such as virtual machines, potentially creating backdoors or exposing data.
- Automatic application of configurations could weaken security, such as enabling additional network ports or changing compliance-critical settings

Steps to Verify

- Review Azure Activity Logs for the user's previous activities to verify their typical behavior.
- Investigate the context surrounding the remediation task, such as associated policy changes or triggered events, to determine whether the action is legitimate.
- Identify any unexpected resource deployments or configuration changes triggered by the task.

Azure Suspicious Serial Console Usage

Lateral Movement



MITRE | ATT&CK®

T1651 Cloud
Administration Command

Triggers

- The Serial Console is accessed by an unexpected or unauthorized user/service principal.
- Unusual frequency of Serial Console connections in a short window.
- Serial Console connections are made from an unusual location or IP address.

Possible Root Causes

- An attacker has gained access to a user account or service principal with Contributor or higher permissions and is using it to access the Serial Console on a virtual machine (VM) or virtual machine scale set (VMSS) instance.
- A user account or service principal has excessive permissions, allowing them to access the Serial Console on a VM or virtual machine scale set instance they shouldn't have access to.
- An authorized user is intentionally accessing the Serial Console for malicious purposes, such as data theft or system manipulation.
- A legitimate user has accidentally configured their account or service principal with excessive permissions, allowing them to access the Serial Console on a VM or virtual machine scale set instance.

Business Impact

- Exposed sensitive data due to unauthorized access.
- Security vulnerabilities exploited due to misconfigured permissions.
- Unplanned changes to business logic or workflows.
- Potential data breaches, unauthorized access to resources, disruption of critical business services, and reputational damage.

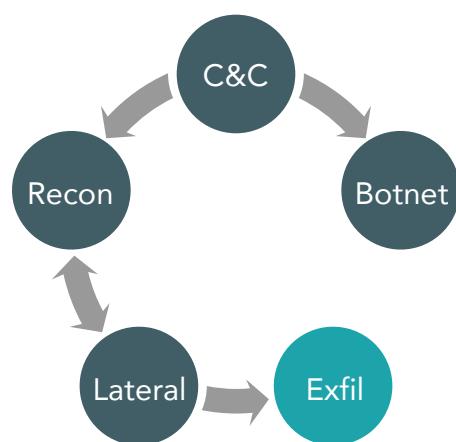
Steps to Verify

- Investigate User Activity: Review the Azure Activity Logs for the suspicious Serial Console usage event, focusing on the user/service principal who accessed the console.
- Verify Permissions: Investigate the user's or service principal's permissions and access levels within Azure to determine if they have excessive privileges.
- Analyze Connection Parameters: Examine the parameters used in the Serial Console connection to identify potential security risks such as data theft or system manipulation.
- Consult with Azure Administrators: Consult with Azure administrators, security teams, and relevant stakeholders to determine the cause and scope of the incident.

Category

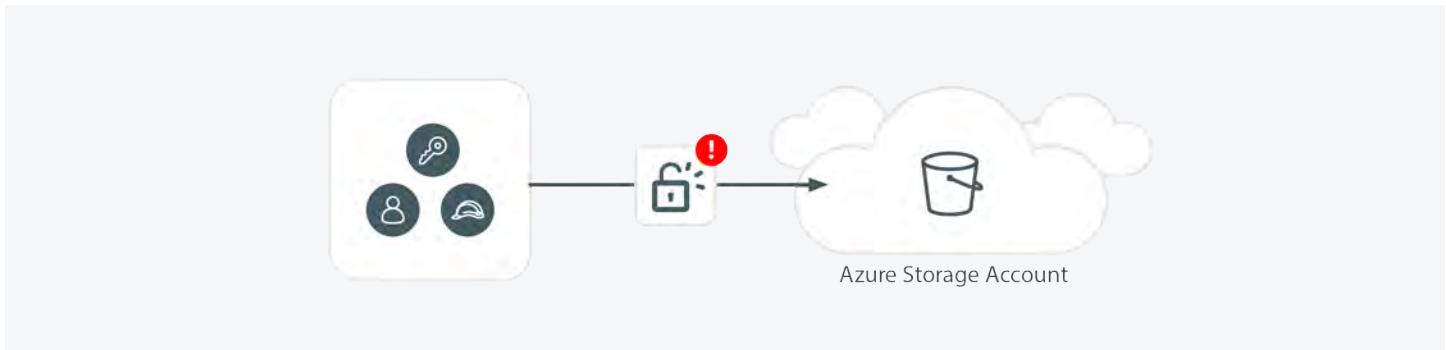
Exfiltration

- An attacker with access to the Azure environment is collecting and removing data from the environment
- Attackers after gaining access and gaining sufficient permissions will steal high value data
- Attackers will take actions to modify services and identities to ensure continued access
- Attackers will leverage credentials within their defined permissions but for non-intended purposes to further the attacker's objective
- Actions are associated with [Execution](#), [Persistence](#), [Privilege Escalation](#), [Defense Evasion](#), [Lateral Movement](#), [Credential Access](#) MITRE Tactics



Azure Suspect Public Storage Account Change

Exfiltration



MITRE | ATT&CK®

T1578 Modify Cloud Compute Infrastructure

T1530 Data from Cloud Storage

Triggers

- An entity was observed enabling public access to a given Storage Account

Possible Root Causes

- An attacker may be scanning and modifying configurations around a storage account to enable data exfiltration.
- An IT misconfiguration may have been made by an authorized administrator which could weaken the posture around a storage account and promote the risk of data loss.
- An administrator or automated task is making authorized modifications to controls associated with a storage account.

Business Impact

- Malicious or unintentional weakening of security posture controls around storage accounts are commonly associated with data loss.

Steps to Verify

- Investigate the account context that made the change for other signs of malicious activity.
- Investigate for data loss.
- Verify if the storage account in question is authorized for public access.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

Azure Suspicious Disk Download

Exfiltration



MITRE | ATT&CK®

[T1567 Exfiltration Over Web Service](#)

[T1039 Data from Network Shared Drive](#)

[T1074 Data Staged](#)

[T1020 Automated Exfiltration](#)

Triggers

- An identity created a Shared Access Signature (SAS) URL link to gain access to an Azure resource.

Possible Root Causes

- An attacker may be using SAS URLs to exfiltrate data.
- An authorized user is using SAS URLs to access a resource as they start new projects, back up data or access files to support their job function.

Business Impact

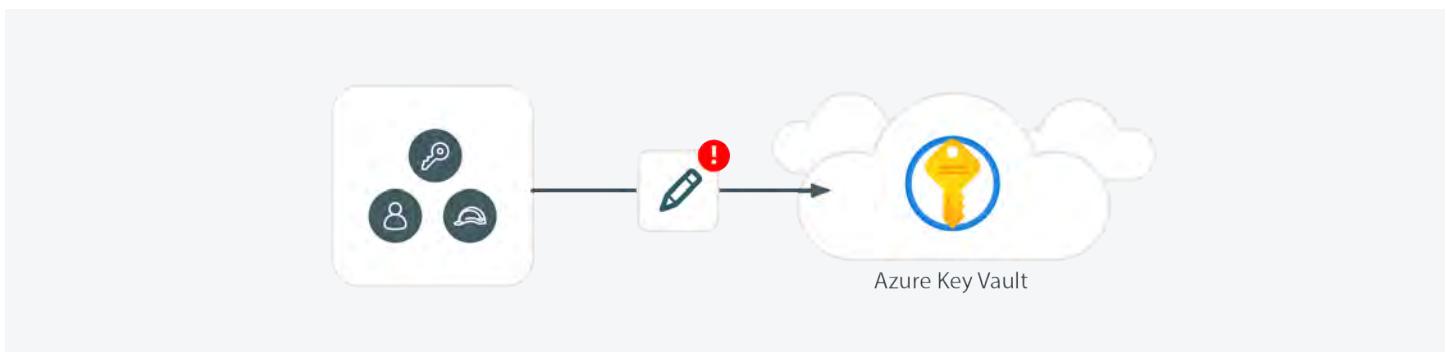
- Ability to access a resource and exfiltrate from it is often the last stage of a security compromise.
- Exfiltration of sensitive business data may lead to loss of control of company secrets and intellectual property.

Steps to Verify

- Investigate the account context that leveraged the SAS URL for other signs of malicious activity which may indicate their account has been compromised.
- Investigate for data loss. Review the details and contents of the resources to assess risk.
- If review indicates possible malicious actions, disable credentials associated with this alert then perform a comprehensive investigation.

Azure Suspicious Key Vault Extraction

Exfiltration



MITRE | ATT&CK®

[T1555 Credentials from Password Stores](#)

[T1552 Unsecured Credentials](#)

[T1649 Steal or Forge Authentication Certificates](#)

Triggers

- Unusual volumes of secret, certificate, or key exports from Azure Key Vault within a short time period.
- A new type of credential was accessed from within the Key Vault which is unusual for the identity.
- Key Vault objects are accessed by accounts or service principals that haven't previously accessed these resources.

Possible Root Causes

- An attacker may have gained access to a privileged account and is extracting secrets or keys from Key Vault.
- Scripts or tools may be extracting secrets or certificates as part of backup processes or regulatory compliance.
- Security or compliance teams performing tests or audits that involve accessing multiple Key Vault items.

Business Impact

- Exposed secrets or certificates may allow attackers to impersonate trusted services.
- Stolen credentials could enable lateral movement or privilege escalation within the cloud environment.
- Potential of significant financial impact if extracted keys or secrets are used to deploy or manipulate resources.

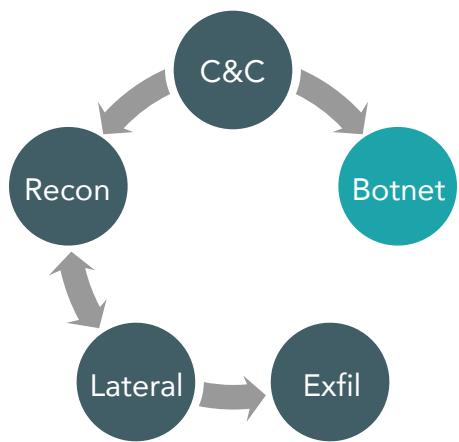
Steps to Verify

- Check Key Vault resource logs for patterns of access to secrets, certificates, or keys, focusing on the frequency, timing, and source IPs of access.
- Verify if the accessing account or service principal is showing unusual behavior and that permissions on the Key Vault haven't been recently modified.
- Actions if confirmed suspicious:
 - Change or disable affected secrets, certificates, or keys to prevent further unauthorized use.
 - Implement temporary network restrictions for unusual IPs or regions involved in the suspicious activity.
 - Review the Key Vault for any additional unauthorized activities and ensure that security policies are properly enforced.
 - Notify security teams and document the event in the incident management system for tracking and follow-up actions.

Category

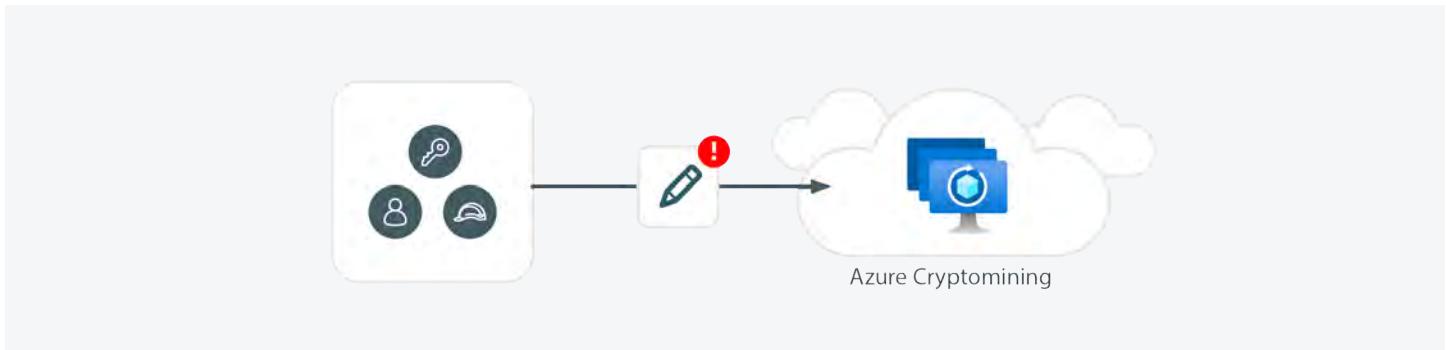
Botnet

- An attacker with access to the Azure environment is leveraging Azure infrastructure for financial gain
- Actions are associated with the [Impact](#) MITRE Tactic



Azure Cryptomining

Botnet



MITRE | ATT&CK®

[T1496 Resource Hijacking](#)

[T1583 Acquire Infrastructure](#)

Triggers

- An attacker is leveraging compromised credentials to create powerful compute instances (virtual machines) for use in a cryptomining attack.
- High powered VM instances are being created to enable compute intensive operations to occur in support of authorized business activity.

Possible Root Causes

- An attacker is creating VMs powerful enough to mine cryptocurrency in the Azure environment.
- Legitimate creation a large and/or expensive VMs as part of normal operations. A common example of this could be the creation of VMs to host an AI deployment.

Business Impact

- Creation of large and/or expensive VMs will have a significant impact on the organization's cloud budget, especially if the VMs remain undetected & are allowed to run for a prolonged period of time.

Steps to Verify

- Investigate the identity for signs of other malicious activity.
- Ensure that any large and/or expensive VMs were intentionally created.

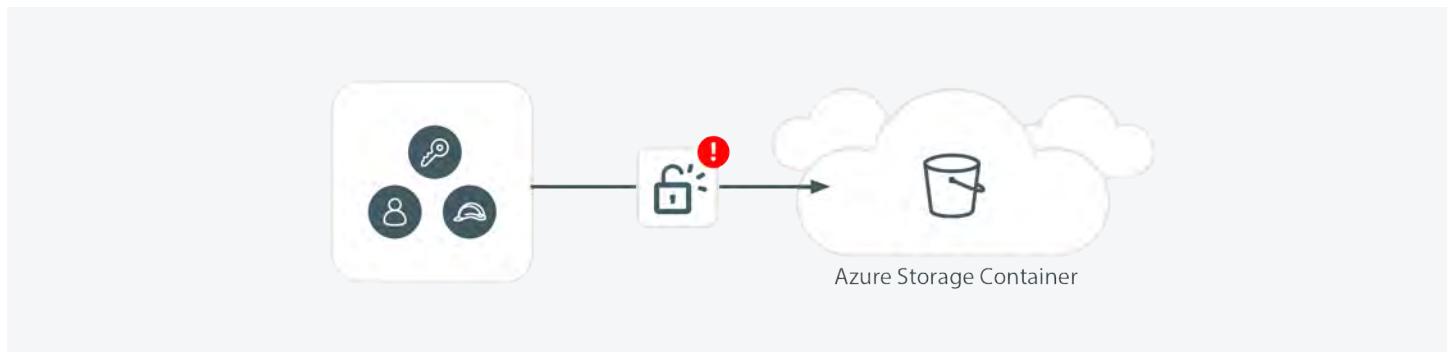
Category

Info

- Reports on new and novel events without directly impacting scoring.
- New and novel events occur normally in most network and cloud environments and in most cases are not directly linked to threats.
- Awareness of new and novel events support better situational awareness and provide additional context when observed with kill chain alerts.

Azure Public Storage Container Change

Info



MITRE | ATT&CK®

N/A

Triggers

- The settings for a storage container were changed. This may indicate an attempt to enable public access.

AWS Detections



Kingpin Technology



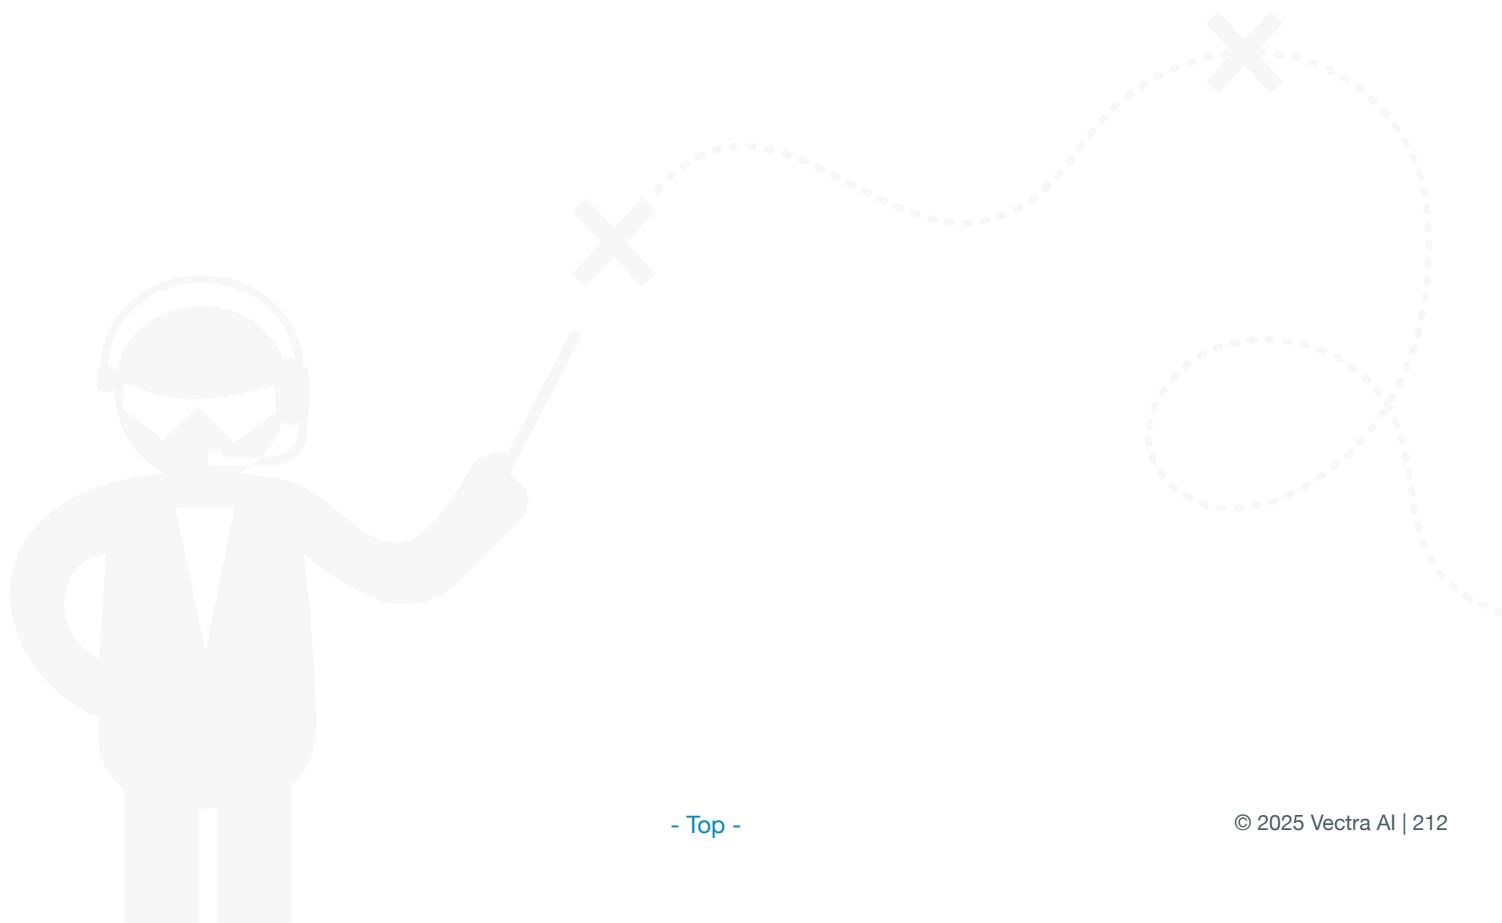
Vectra AI attributes all of our detections to actionable User identities such as IAM Users, SAML Users, an External Account, or AWS Services.

This is a complex problem because users in AWS are encouraged to assume other roles to perform actions, and actively discouraged to perform actions as the account they logged in with. In some cases, users will even assume roles after assuming a role in order to be able to perform certain actions. Our dedicated team of Data Scientists use advanced machine learning techniques to attribute any activity up to the original actor based on logged activity across your AWS account. When you see any AWS detections in our product, you will be able to see a chain of roles assumed by the actor before performing their action, which will explain how this user assumed this role.

Category

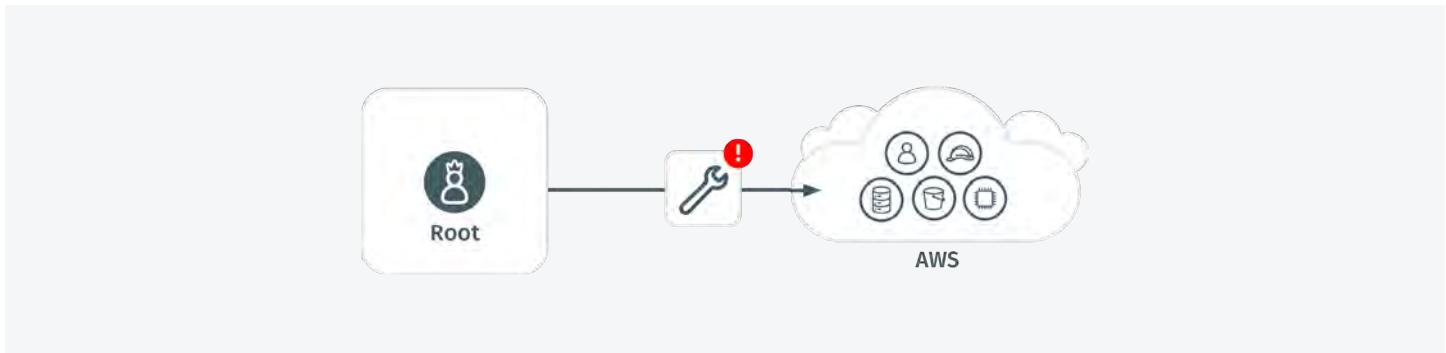
Command & Control

- An attacker is controlling an AWS account to orchestrate their attack against AWS infrastructure and services
- Attackers will access credential through various methods
- Attackers will control accounts in order to then perform reconnaissance and lateral movement in the name of achieving an object like data exfiltration or resource impact
- Actions are associated with [Command and Control](#), [Initial Access](#), [Reconnaissance](#) MITRE Tactics



AWS Root Credential Usage

Command & Control



MITRE | ATT&CK®

T1078 Valid Accounts

Triggers

- An action was taken by the root account.

Possible Root Causes

- An attacker has compromised the root account and is using the unfettered access it grants to further their attack.
- Administrators are using the root account for normal activities, which is against best practices and should not be done.

Business Impact

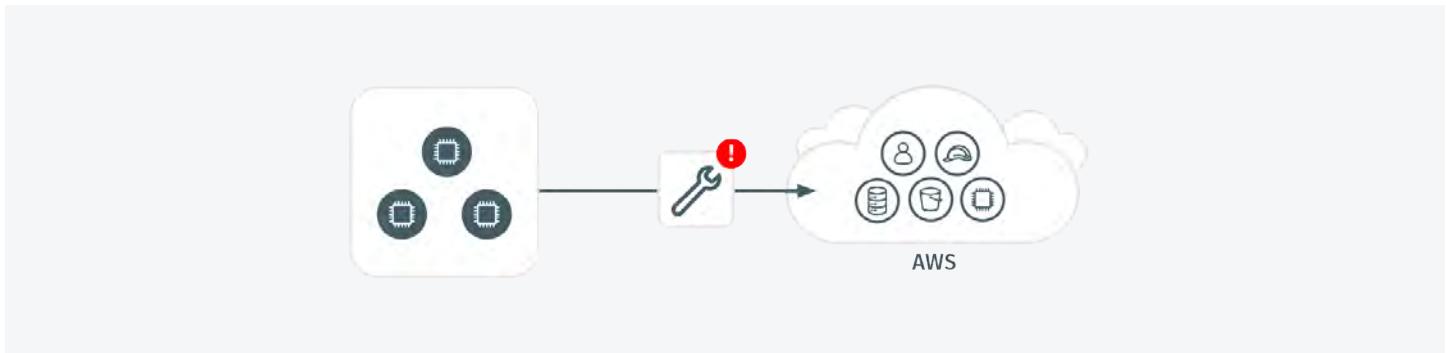
- Malicious use of the root account indicates significant opportunity for negative impact to organizational assets, services, and data to include disruptive impact and sensitive data loss.
- Misuse of the root account by admins for routine activities greatly elevates the risk of accidental damage or disruption.

Steps to Verify

- Review the activity completed by the root account for indications of malicious activity.
- Validate with the team responsible for administering AWS that they used the root account for an authorized activity.

AWS Suspicious Credential Usage

Command & Control



MITRE | ATT&CK®

T1078 Valid Accounts

Triggers

- EC2 generated temporary credential used outside of EC2.

Possible Root Causes

- An attacker has extracted a temporary credential from an EC2 instance and is using it to further their attack.
- An application is using temporary credential generation via EC2s in an unusual way.

Business Impact

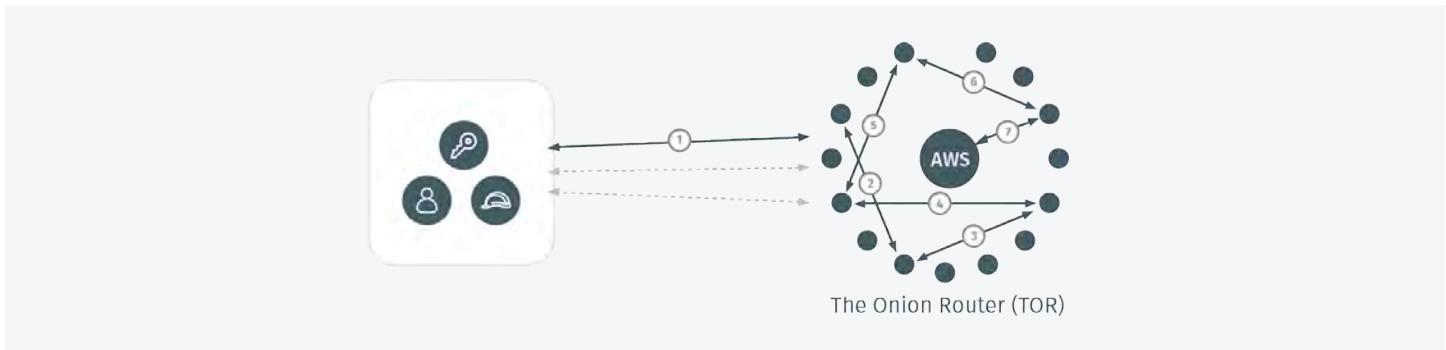
- Attackers may use temporary credentials as a means of maintaining persistent command and control in an environment, which increases the risk of data loss or impacted assets and services.

Steps to Verify

- Review the actions being undertaken by the credential after the identified activity and potential risk posed by that access.
- Discuss with the EC2 instance owners to determine if the use of instance generated temporary keys outside of EC2 is known and legitimate.
- If the review determines there is a high risk to data or the environment, disable the credentials and perform a comprehensive investigation.

AWS TOR Activity

Command & Control



MITRE | ATT&CK®

T1090 Proxy

Triggers

- A credential was observed accessing the environment from a known anonymized (TOR) exit node.

Possible Root Causes

- An attacker is using an anonymizing proxy like TOR to obfuscate details of their source connection or make an investigation more difficult by using multiple source IP addresses.
- A user may be intentionally using TOR to circumvent restrictions preventing access to the resources in question, such as those applied by the country they are accessing from.

Business Impact

- Attackers identified under this detection are actively operating within the environment while maintaining some level of operational security by obfuscating their source details.
- Attackers operating using TOR will reduce the ability of teams to connect identified attacker behavior with other behaviors not yet identified since it enables the attacker to regularly change the source detail of their connections while undertaking operations within the environment.
- Authorized users that have adopted TOR may be in violation of IT Policies and be placing organizational assets at risk.

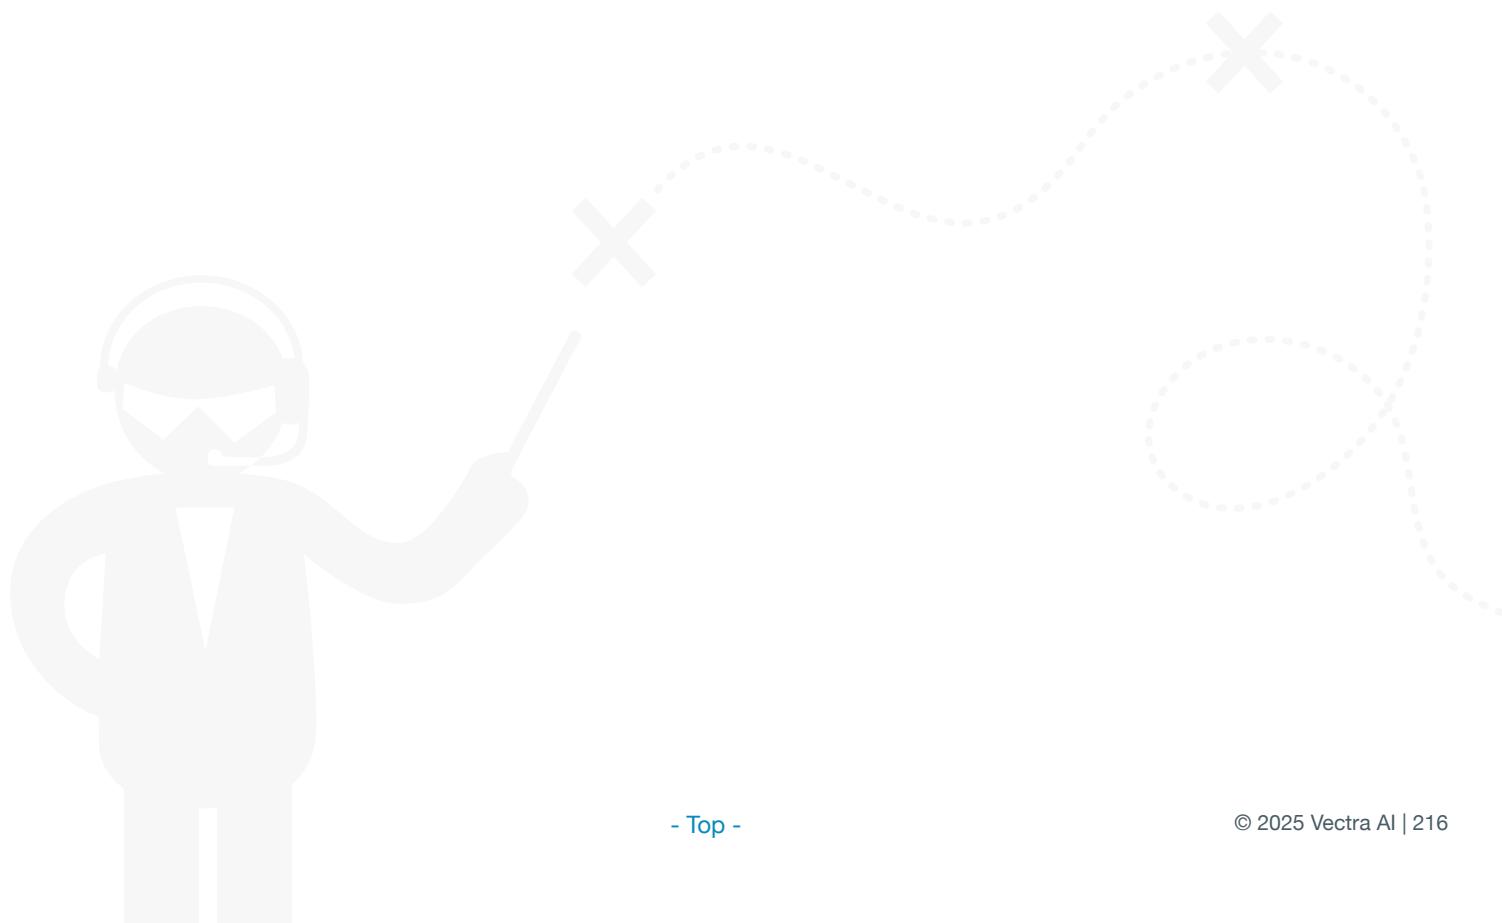
Steps to Verify

- Review the actions being undertaken by the user after the identified activity and potential risk posed by that access
- Review security policy to determine if the use of TOR is allowed.
- Discuss with the user to determine if the use of TOR is known and legitimate.
- If the review determines there is a high risk to data or the environment, disable the credentials and perform a comprehensive investigation.

Category

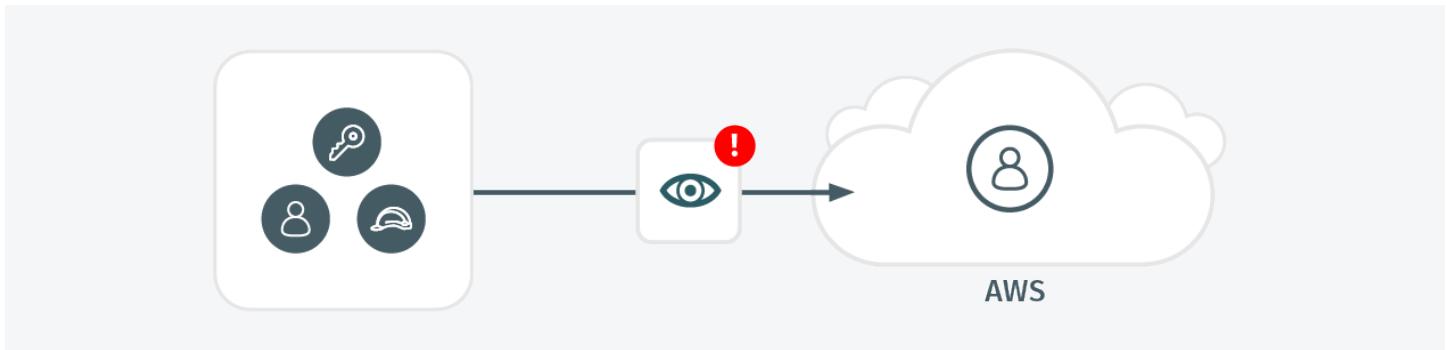
Reconnaissance

- An attacker is surveying and learning about AWS infrastructure
- Attackers with control of an account will look to identify paths to their final objectives
- Attackers will probe AWS services in-order to find credentials to gain additional access
- Attackers will identify services in order to collect, exfiltrate, or impact data
- Actions are associated with the [Discovery](#) MITRE Tactic



AWS External Network Discovery

Reconnaissance



MITRE | ATT&CK®

T1049 System Network
Connections Discovery

Triggers

- An AWS control-plane API was observed programmatically enumerating the configuration details associated with Cloud-Native network integrations such as VPC Peering Connections, VPN Connections and/or DirectConnect Gateways.

Possible Root Causes

- An attacker may be actively enumerating how external networks are connected into the environment, in order to further their attack.
- An administrator may intentionally be enumerating network configurations as part of their normal duties.

Business Impact

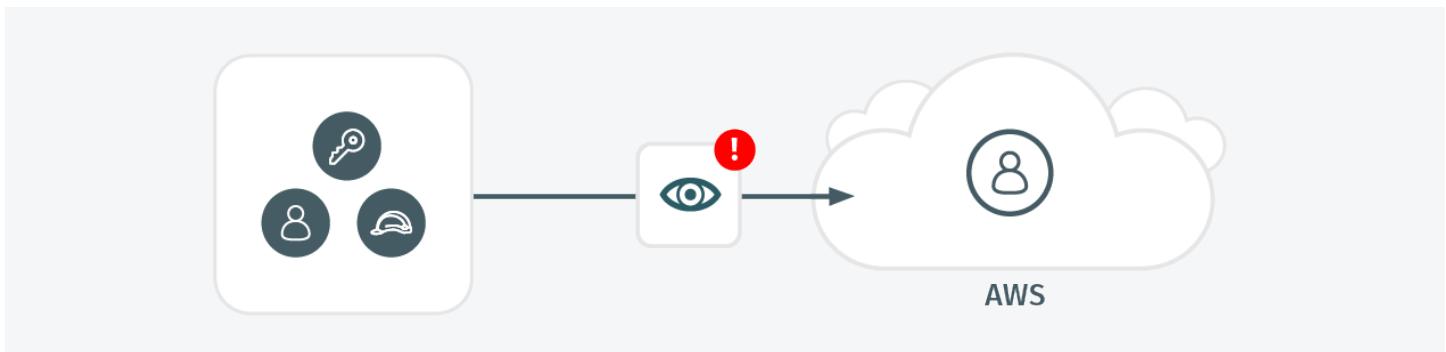
- Reconnaissance may indicate the presence of an adversary gaining details necessary to support additional malicious activities within the environment. A successful attack may yield information that can be used by an adversary to mount a campaign against any external, connected network

Steps to Verify

- Investigate the Principal that performed the action for other signs of malicious activity.
- Investigate if any modifications were made to the enumerated DirectConnect Gateways, VPN or VPC Peering Connections configurations.
- Validate that any changes were authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration change:
 - Revert configuration change.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and the scope of impacted resources.

AWS Network Configuration Discovery

Reconnaissance



MITRE | ATT&CK®

T1526 Cloud Service Discovery

Triggers

- An AWS control-plane API was observed programmatically enumerating the configuration details associated with the Virtual Private Network (VPC) such as Network Interfaces, Gateways, Network ACLs and Route Tables.

Possible Root Causes

- An attacker may be actively enumerating how networks in the environment are configured in order to further their attack.
- An administrator may intentionally be enumerating network configurations as part of their normal duties.

Business Impact

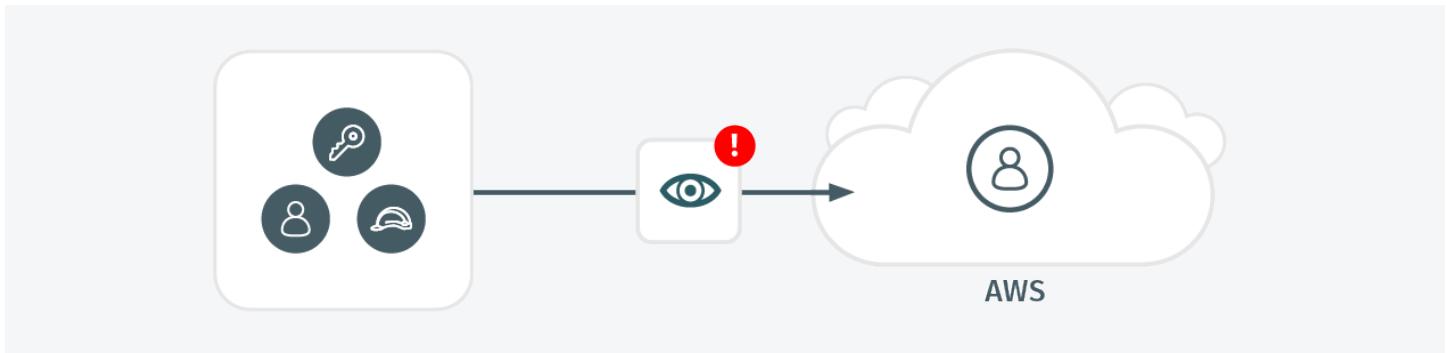
- Recon may indicate the presence of an adversary gaining details necessary to support additional malicious activities within the environment. A successful attack may yield information that can be used by an adversary to mount a campaign within the AWS Environment.

Steps to Verify

- Investigate the Principal that performed the action for other signs of malicious activity.
- Investigate if any modifications were made to the enumerated VPCs such as changes to Network Interfaces, Gateways, Network ACLs or Routing Tables.
- Validate that any changes were authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration change:
 - Revert configuration change.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and the scope of impacted resources.

AWS Organization Discovery

Reconnaissance



MITRE | ATT&CK®

T1580 Cloud Infrastructure Discovery

T1614 System Location Discovery

Triggers

- A user lists AWS account aliases via ListAliases or retrieves details for the AWS organization via DescribeOrganization

Possible Root Causes

- An attacker is enumerating details on the AWS organization to further their attack planning and next steps.
- An administrator or user is retrieving organization details as part of their normal duties.
- Automation in the environment is collecting these details to support additional activities.

Business Impact

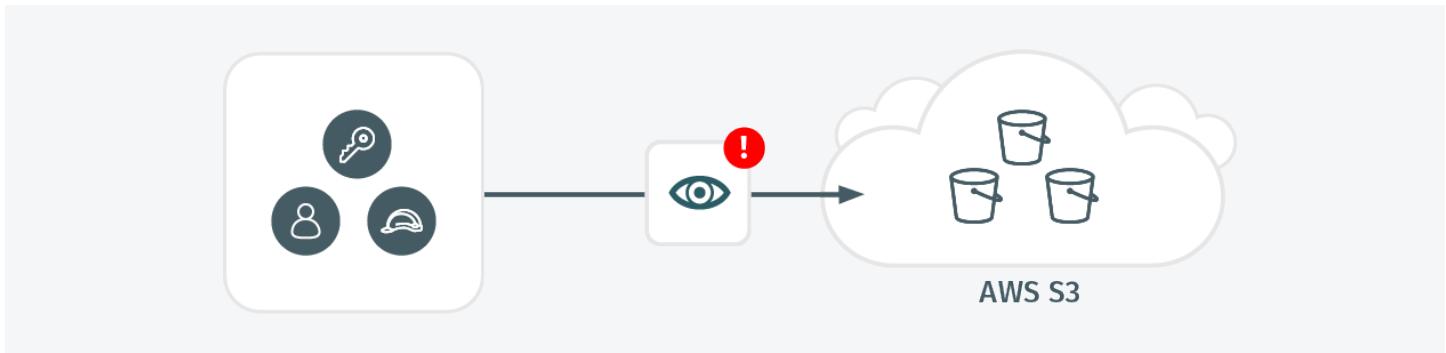
- Recon may indicate the presence of an adversary gaining details necessary to support additional malicious activities within the environment.

Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS S3 Enumeration

Reconnaissance



MITRE | ATT&CK®

T1526 Cloud Service Discovery

Triggers

- Credential was observed performing a set of anomalous API requests that can be associated with the discovery or subsequent phases of an attack.

Possible Root Causes

- An attacker may be actively looking for privilege escalation opportunities.
- A security or IT service may intentionally be enumerating these APIs for monitoring reasons.

Business Impact

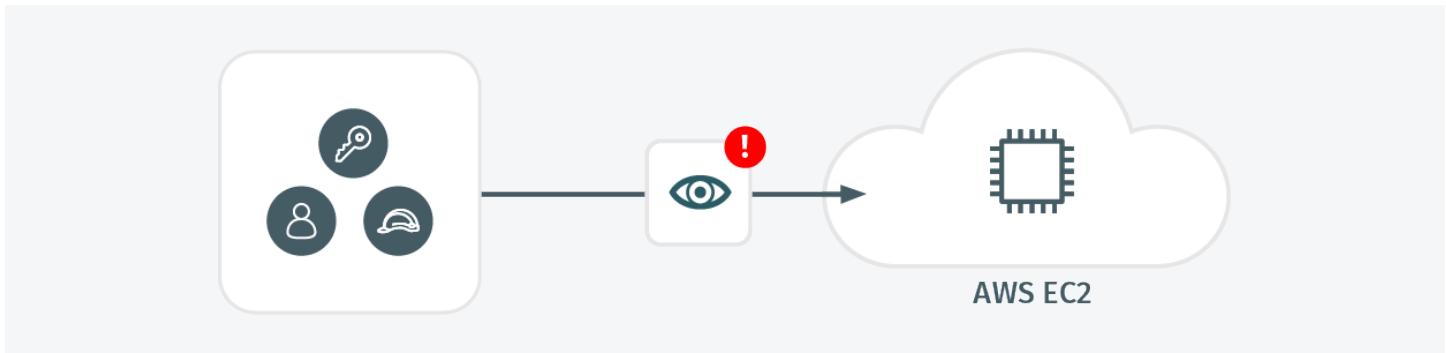
- Privilege escalation may indicate the presence of an adversary that is modifying permissions to progress towards an objective.

Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Validate that any modifications are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Suspect Credential Access from EC2

Reconnaissance



MITRE | ATT&CK®

T1552 Unsecured Credentials

Triggers

- A set of AWS control plane APIs commonly used to search EC2 user data on EC2 resources for credentials was invoked in an unusual way that may be associated with a potential attack.

Possible Root Causes

- An attacker is searching for credentials inside of the EC2 user data to pivot in the environment.
- An authorized administrator is performing an unusual activity commonly associated with attack progression.

Business Impact

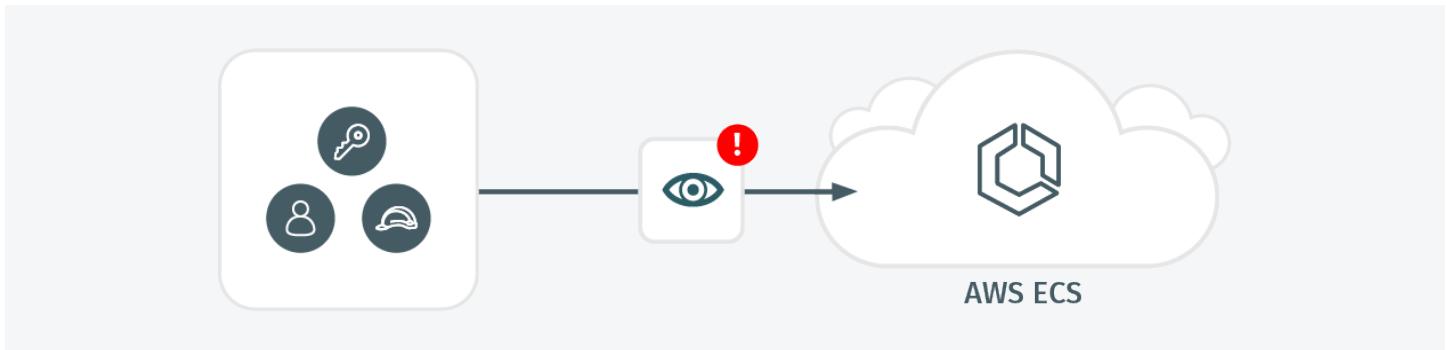
- Lateral movement may indicate that an adversary has established a foothold in the environment and is progressing towards their objective, increasing the risk of material impact.

Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Validate that any modifications are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Suspect Credential Access from ECS

Reconnaissance



MITRE | ATT&CK®

T1552 Unsecured Credentials

Triggers

- Credential was observed performing a set of API requests to retrieve a broad range of container configuration details which may further their attack through the leak of credentials or other data about the environment.

Possible Root Causes

- An attacker may be actively looking for privilege escalation opportunities.
- A security or IT service may intentionally be enumerating these APIs for monitoring or configuration management reasons.

Business Impact

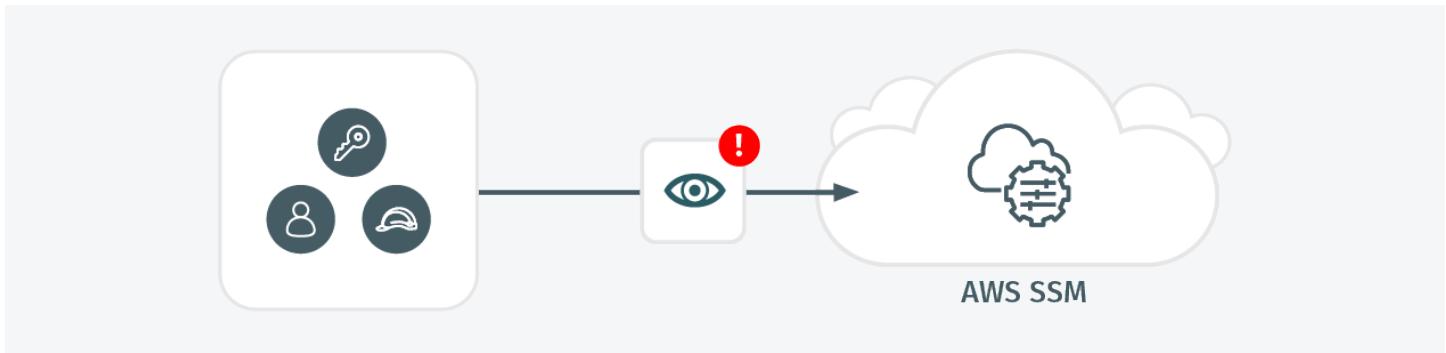
- Stolen credentials allow an adversary to leverage authorized services and APIs to extend their attack which can be difficult for traditional security solutions to detect.
- Abused credentials are typically associated with impactful attacks, and if unmitigated may increase the likelihood that an adversary may inflict a loss of data or service availability.

Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Validate that any modifications are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Suspect Credential Access from SSM

Reconnaissance



MITRE | ATT&CK®

T1552 Unsecured Credentials

Triggers

- Credential was observed performing a set of API requests to list and then retrieve parameters within the AWS parameter store.

Possible Root Causes

- An attacker may be actively looking for privilege escalation opportunities.
- A security or IT service may intentionally be enumerating these APIs for monitoring or configuration management reasons.

Business Impact

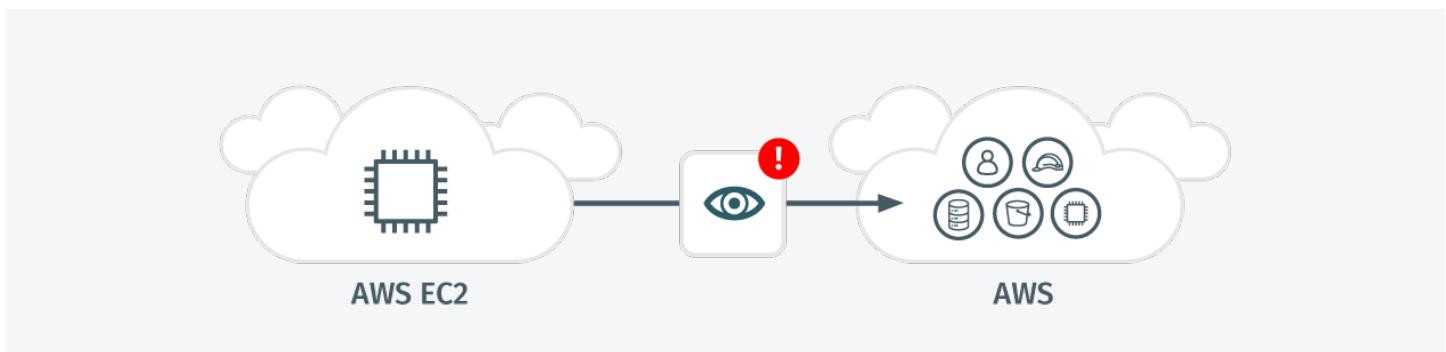
- Stolen credentials allow an adversary to leverage authorized services and APIs to extend their attack which can be difficult for traditional security solutions to detect.
- Abused credentials are typically associated with impactful attacks, and if unmitigated may increase the likelihood that an adversary may inflict a loss of data or service availability.

Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Validate that parameters requested do not contain sensitive details, such as credentials. If they do, investigate those credentials for potential malicious use.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Suspect Discovery from EC2 Instance

Reconnaissance



MITRE | ATT&CK®

T1069.003 Permission Groups Discovery: Cloud Groups
T1087 Account Discovery
T1580 Cloud Infrastructure Discovery

Triggers

- A set of AWS control-plane APIs were invoked from an EC2 instance which enumerated the configuration details about cloud environment.

Possible Root Causes

- An attacker may have compromised an EC2 instance and be actively looking for opportunities to escalate their permissions by enumerating details of the environment such as the IAM Users, Roles, S3 Buckets and Logging configuration.
- An application hosted on the EC2 instance maybe be intentionally programmed to invoke AWS control-plane APIs associated with environment reconnaissance leveraging the Instance Profile credentials available on the instance.

Business Impact

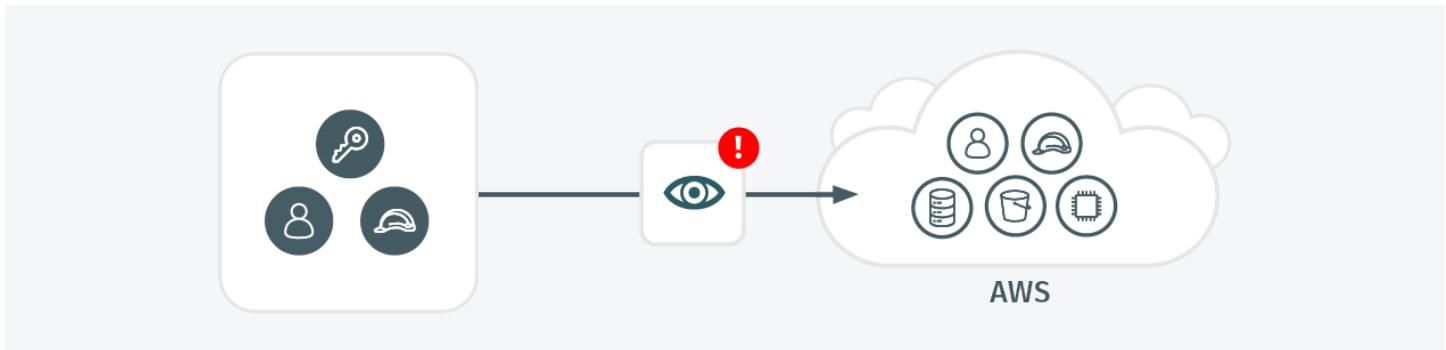
- Reconnaissance of the environment from an EC2 instance may indicate an adversary has compromised the EC2 instance and is gaining details necessary to support additional malicious activities within the environment.

Steps to Verify

- Investigate the Instance Profile that performed the action for other signs of malicious activity.
- Identify the EC2 instances associated with the suspicious Instance Profile and review them for other signs of malicious activity.
- If any modifications were made to the environment by the suspicious Instance Profile, validate they are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration:
 - Revert any configuration changes.
 - Disable credentials associated with the Instance Profile.
 - Perform a comprehensive investigation to determine initial compromise and scope of impacted resources.

AWS Suspect Escalation Reconnaissance

Reconnaissance



MITRE | ATT&CK®

T1069 Permission Groups
Discovery

Triggers

- Credential was observed performing a set of unusual API requests that can be associated with the discovery or subsequent phase of an attack.

Possible Root Causes

- An attacker may be actively looking for privilege escalation opportunities,
- A security or IT service may intentionally be enumerating these APIs for monitoring reasons.

Business Impact

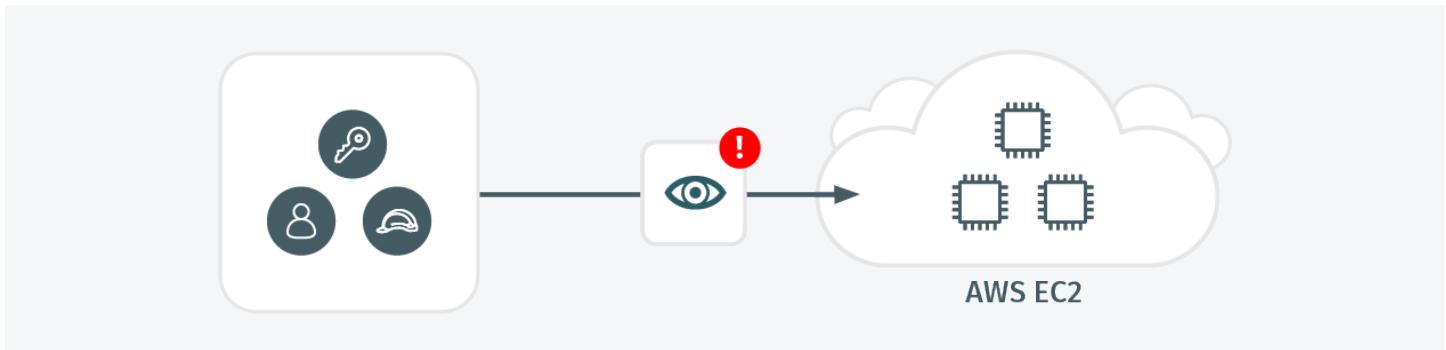
- Privilege escalation may indicate the presence of an adversary that is modifying permissions to progress towards an objective.

Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Validate that any modifications are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Suspicious EC2 Enumeration

Reconnaissance



MITRE | ATT&CK®

T1526 Cloud Service Discovery

Triggers

- Credential was observed performing a set of anomalous API requests that can be associated with the discovery or subsequent phases of an attack.

Possible Root Causes

- An attacker may be actively looking for privilege escalation opportunities.
- A security or IT service may intentionally be enumerating these APIs for monitoring reasons.

Business Impact

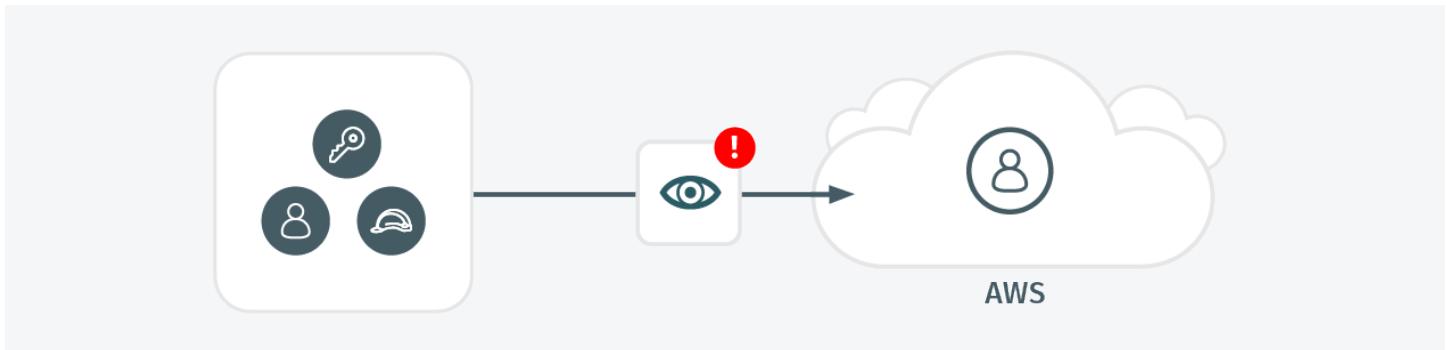
- Privilege escalation may indicate the presence of an adversary that is modifying permissions to progress towards an objective.

Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Validate that any modifications are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS User Permissions Enumeration

Reconnaissance



MITRE | ATT&CK®

T1069 Permission Groups
Discovery

Triggers

- Credential was observed performing a set of unusual API requests that can be associated with the discovery or subsequent phase of an attack.

Possible Root Causes

- An attacker may be actively looking for privilege escalation opportunities,
- A security or IT service may intentionally be enumerating these APIs for monitoring reasons.

Business Impact

- Privilege escalation may indicate the presence of an adversary that is modifying permissions to progress towards an objective.

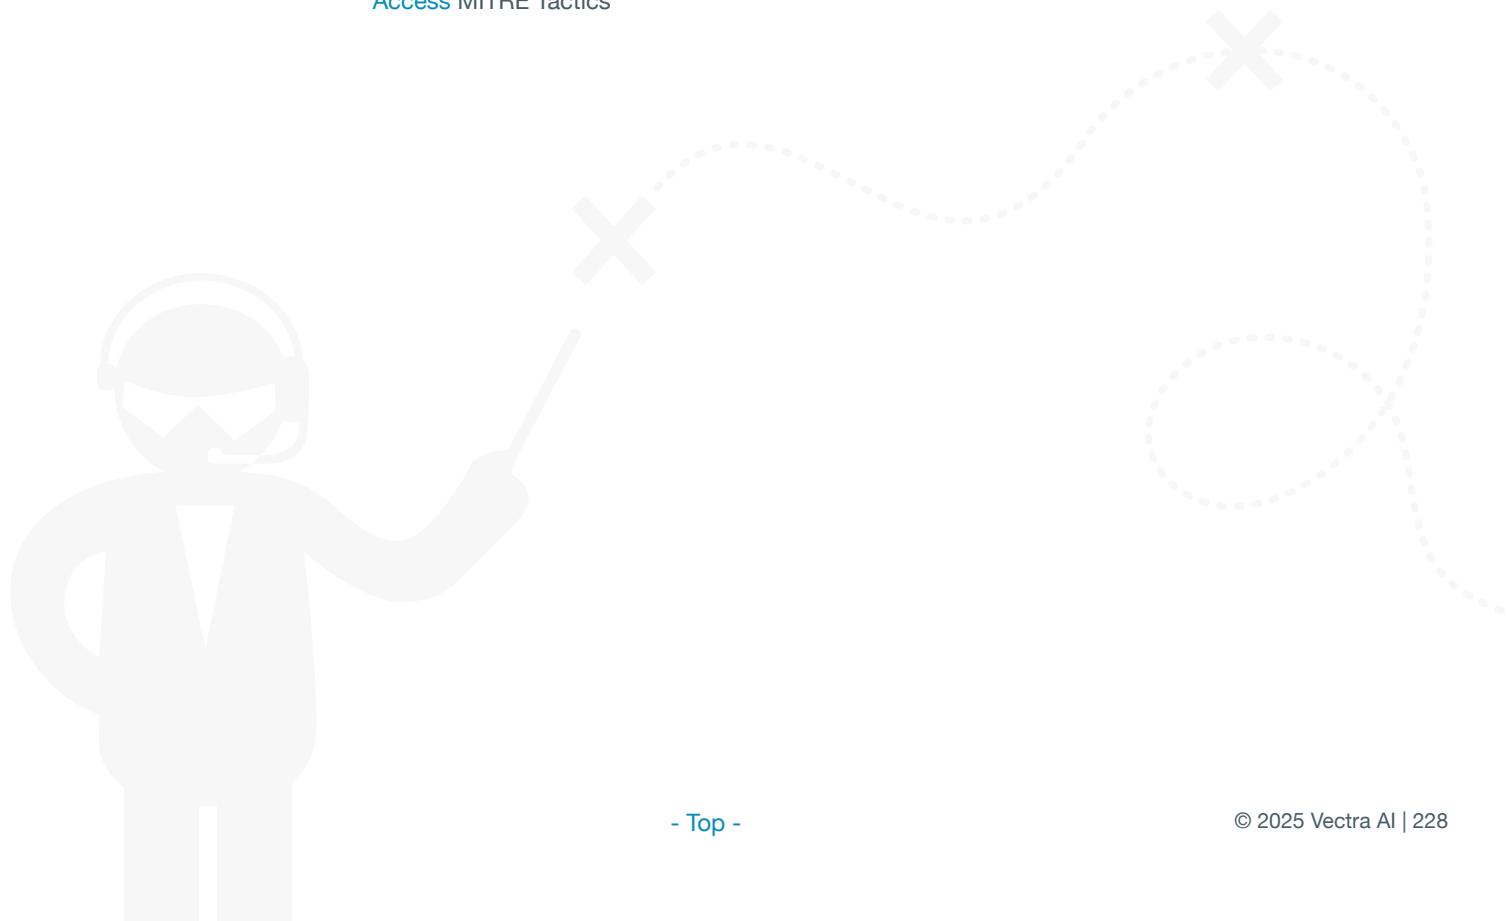
Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Validate that any modifications are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

Category

Lateral Movement

- An attacker in the AWS environment is spreading and taking actions that ensure continuous undetected access
- Attackers after gaining access to the credentials and discovering the environment will propagate and solidify their access
- Attackers will take actions to modify services and identifies to ensure continued access
- Attackers will leverage credentials within their defined permissions but for non-intended purposes to further the attacker's objective
- Actions are associated with [Execution](#), [Persistence](#), [Privilege Escalation](#), [Defense Evasion](#), [Lateral Movement](#), [Credential Access](#) MITRE Tactics



AWS Attack Tools

Lateral Movement



MITRE | ATT&CK®

T1562

Triggers

- An AWS control plane API was invoked with a well known, off-the-shelf attack tool that has a User-Agent signature.

Possible Root Causes

- An attacker is utilizing an AWS attack tool to as a method for interacting with your cloud environment and in support of a malicious campaign.
- An authorized party is utilizing an AWS Attack Tool to aide in auditing the environment or performing a penetration test.

Business Impact

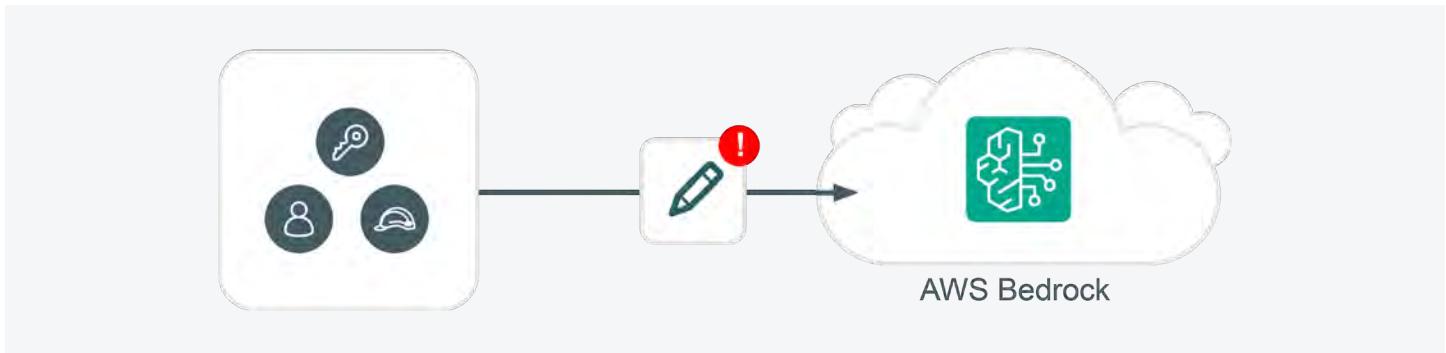
- Attack Tools being used in your environment could indicate the presence of a malicious actor. Activity could represent any phase of an attack on your cloud environment.

Steps to Verify

- Investigate the Principal that performed the actions for other signs of malicious activity.
- Determine if the Principal is authorized to use the tool.
- If review indicates possible malicious actions or high-risk configuration:
 - Revert any configuration change
 - Disable credentials associated with this alert
 - Perform a comprehensive investigation to determine initial compromise and scope of impacted resources

AWS Bedrock Logging Configuration Disabled

Lateral Movement



MITRE | ATT&CK®

T1562

Triggers

- Disable prompt logging for all AWS Bedrock models in an AWS region.

Possible Root Causes

- An attacker attempting to suppress visibility into model usage after gaining unauthorized access.
- Legitimate operational change made without coordination (e.g., for cost optimization).

Business Impact

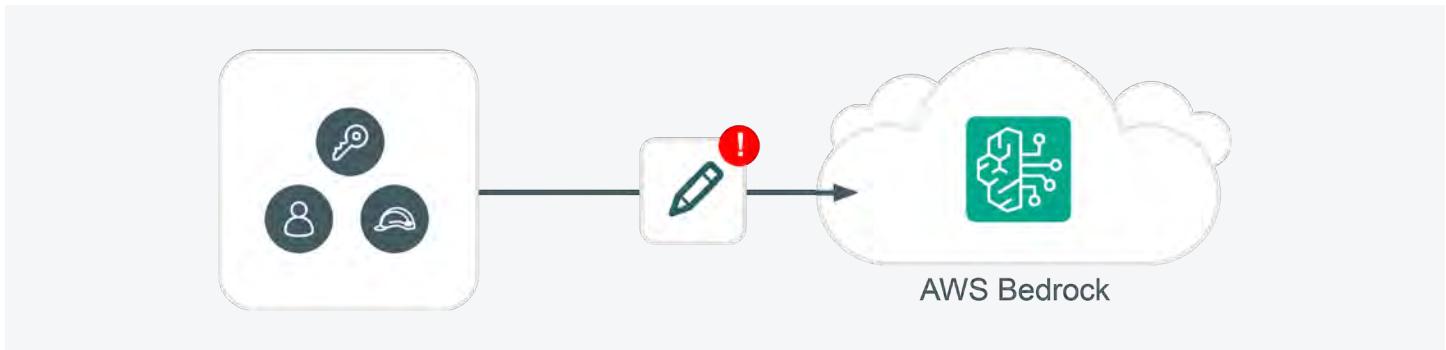
- Eliminates logging for all model prompts and completions in the region, severely limiting ability to audit actions and perform post-incident analysis.
- May enable malicious prompt injection, data extraction, or abuse of the model without detection.

Steps to Verify

- Review the CloudTrail entry to confirm which principal issued the request.
- Investigate whether this principal has a history of administrative actions or if this is anomalous behavior.
- Correlate with recent AWS Bedrock activity (e.g., model invocation, enablement) to assess potential misuse.

AWS Bedrock Novel Model Enabled

Lateral Movement



MITRE | ATT&CK®

AML.T0008

Triggers

- Detection of an identity enabling an AWS Bedrock model who has no prior history of performing similar actions.

Possible Root Causes

- A compromised principal is enabling Bedrock AI model capabilities for illicit gain.
- A legitimate user or developer enabling Bedrock for the first time.

Business Impact

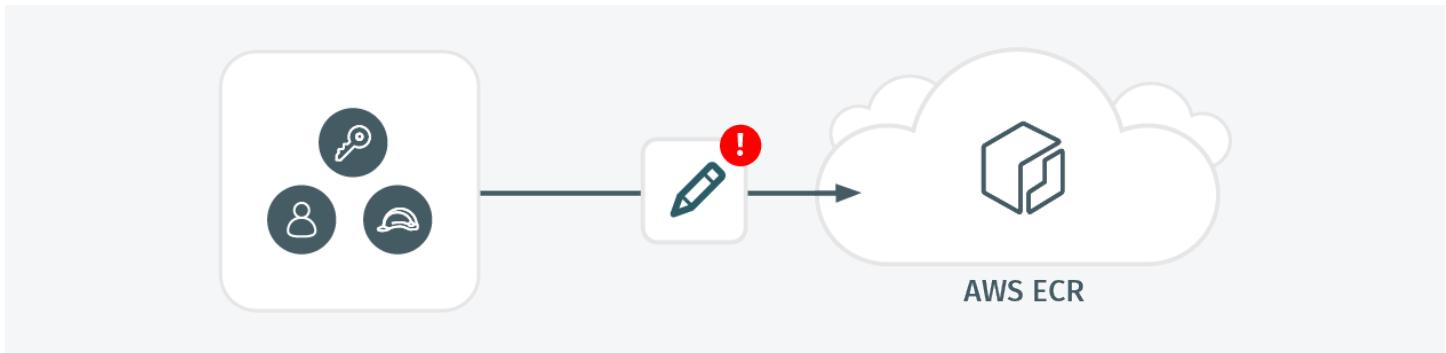
- Enabling a Bedrock model allows for inference usage, which can lead to high compute charges. If a model is enabled unintentionally or by an unauthorized user, it can result in unexpected spend on the organization's AWS bill.
- Enabling a model opens up powerful generative capabilities. If this is done by an unapproved user or workload, it may lead to misuse of the model to generate inappropriate content or leak sensitive information.

Steps to Verify

- Check whether this user is authorized to use Bedrock and whether the model type aligns with their role.
- Investigate whether logging is enabled and whether proper guardrails are configured.

AWS ECR Hijacking

Lateral Movement



MITRE | ATT&CK®

T1525 Implant Internal Image

Triggers

- After enumerating ECR repositories and enumerating the images within those repositories, the attacker requests an authorization token for an image.

Possible Root Causes

- An attacker is inserting a backdoor into an existing image.
- An ECR administrator is making an authorized change to the image.

Business Impact

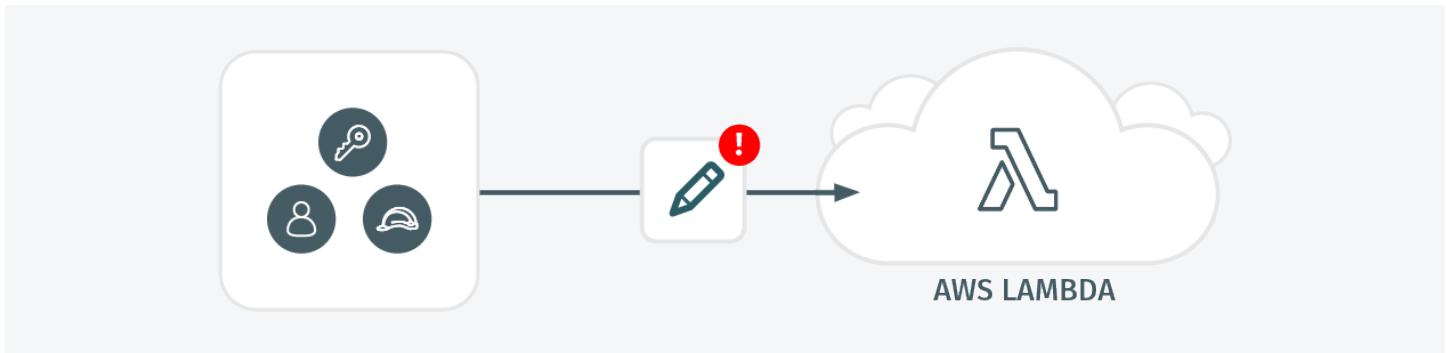
- Lateral movement may indicate that an adversary has established a foothold in the environment and is progressing towards their objective, increasing the risk of material impact.
- An inserted backdoor may provide hidden access persistence within the environment, allowing attackers to return to the environment after eviction.

Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Validate that any modifications are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Lambda Hijacking

Lateral Movement



MITRE | ATT&CK®

T1525 Implant Internal Image

Triggers

- After enumerating Lambda functions and IAM roles, create a Lambda function, and add a new rule to that Lambda function.

Possible Root Causes

- An attacker is creating a Lambda function that serves as a backdoor into the environment.
- An administrator is creating a Lambda function with a trigger for legitimate reasons.

Business Impact

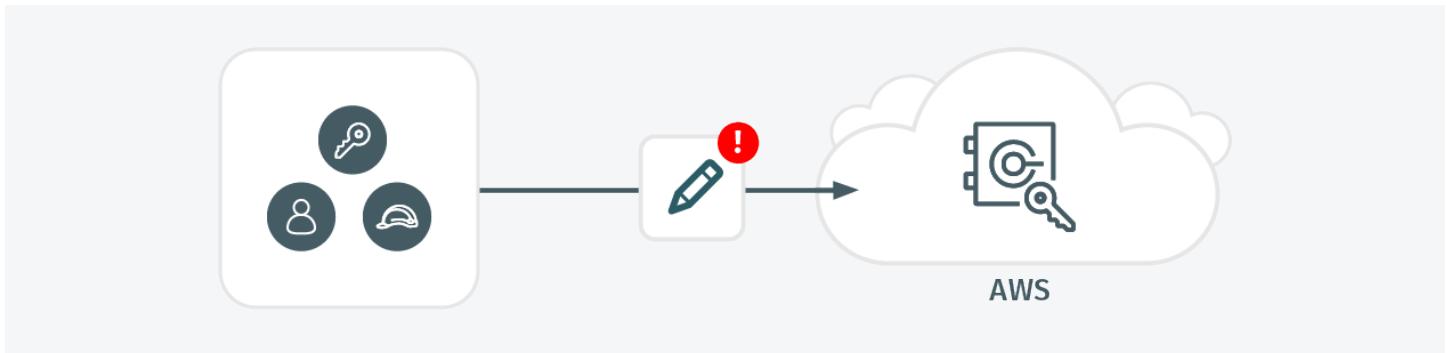
- Lateral movement may indicate that an adversary has established a foothold in the environment and is progressing towards their objective, increasing the risk of material impact.
- An inserted backdoor may provide hidden access persistence within the environment, allowing attackers to return to the environment after eviction.

Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Validate that any modifications are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Logging Disabled

Lateral Movement



MITRE | ATT&CK®

T1562 Impair Defenses

Triggers

- Disable or delete CloudTrail logging within a region where the logging is already enabled.

Possible Root Causes

- An attacker has deleted CloudTrail logs to hide their tracks and/or has deleted the logs to prevent investigation of their historical activities.
- An administrator has disabled CloudTrail logging as part of normal changes to the environment.

Business Impact

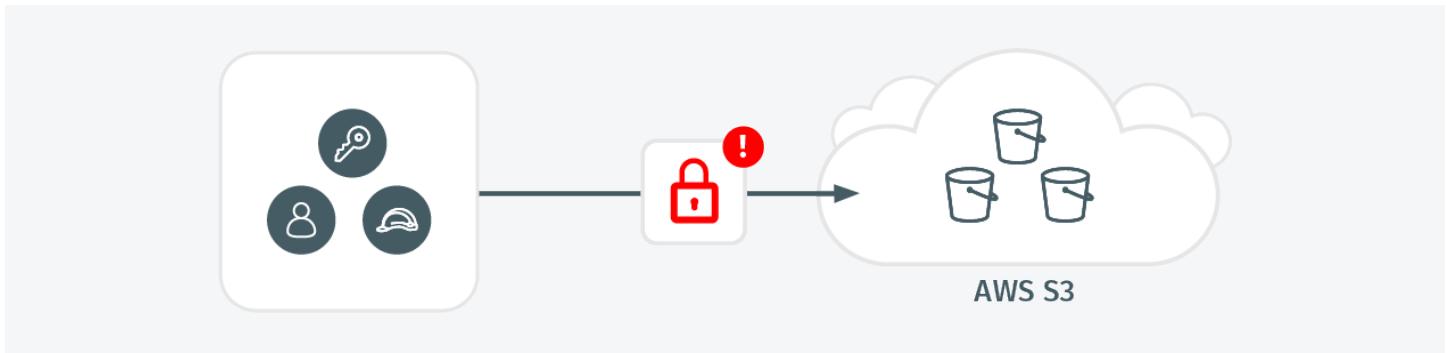
- Inability to detect future attacks, investigate future or historical attacks, or audit activity within the environment.
- Increased risk of activity that may negatively impact the business going unnoticed.

Steps to Verify

- Review the actions being undertaken by the user after the identified activity and potential risk posed by that access in regions where logging remains (if any).
- Review security policy to determine if the removal of logging capabilities is allowed.
- Discuss with the user to determine if the activity is known and legitimate.
- If the review determines there is a high risk to data or the environment, disable the credentials and perform a comprehensive investigation.

AWS Ransomware S3 Activity

Lateral Movement



MITRE | ATT&CK®

T1486 Data Encrypted for Impact

Triggers

- A large number of S3 objects were copied in a way that may indicate the encryption phase of ransomware activity in the environment.

Possible Root Causes

- An attacker leveraging AWS APIs to encrypt S3 objects with the goal of demanding a ransom for the key to decrypt.
- Security or IT operations are manipulating and encrypting S3 objects in bulk as part of normal operations.

Business Impact

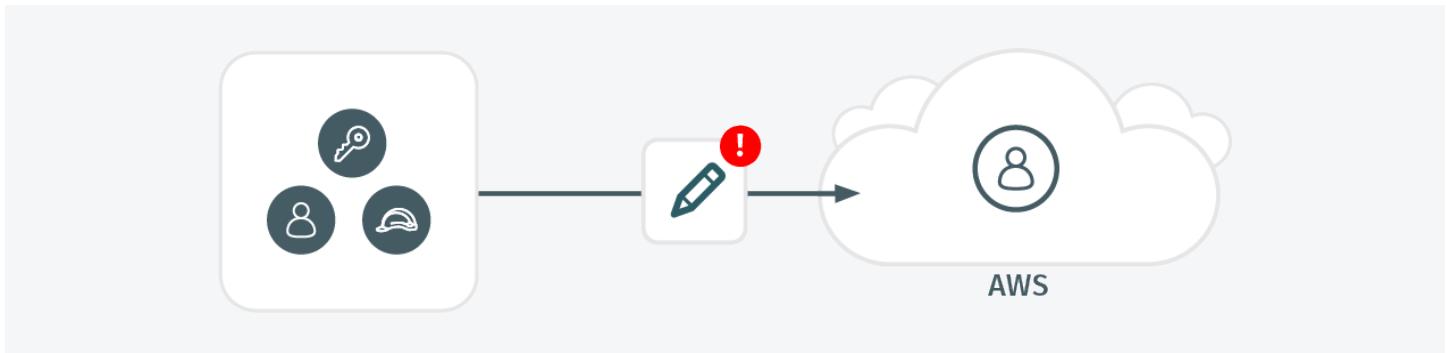
- Ransomware attacks directly impact access to the organization's data and are popular among attackers due to the possibility of a quick transition from attack to monetization.
- After files have been encrypted, the attacker will ask the organization to pay a ransom in return for a promise to provide the encryption key which would allow the files to be decrypted.
- Even if an organization is willing to pay the ransom, there is no guarantee that the encryption key will be provided by the attacker or that the decryption process will work.

Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Validate that any modifications are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration, disable credential associated with this alert then perform a comprehensive investigation.

AWS Root MFA Disabled

Lateral Movement



MITRE | ATT&CK®

T1078 Valid Accounts

Triggers

- The root user was observed disabling or deactivating its MFA device.

Possible Root Causes

- An attacker has compromised the root user and has disabled MFA in order to further solidify control over the root account
- An AWS root user steward has disabled MFA during routine account management such as the migration of root account ownership.

Business Impact

- Malicious use of the root user indicates significant opportunity for negative impact to organizational assets, services, and data to include disruptive impact and sensitive data loss.
- Disabling MFA on the root user greatly elevates the risk of compromise of the root user and the AWS account as a whole.

Steps to Verify

- Investigate the AWS root user steward and if the disabling of MFA was a part of routine, authorized activity.
- Review the activity completed by the root user for indications of malicious activity.
- Ensure MFA is reactivated on the root account.

AWS S3 Objection Deletion

Lateral Movement



MITRE | ATT&CK®

T1485

Triggers

- AWS control-plane APIs were observed suspiciously downloading and then deleting S3 objects or their historical versions impacting data availability.

Possible Root Causes

- An attacker is utilizing compromised credentials to exfiltrate and destroy data for extortion.
- Legitimate scheduled or routine bucket maintenance processes are performing large-scale object deletion.

Business Impact

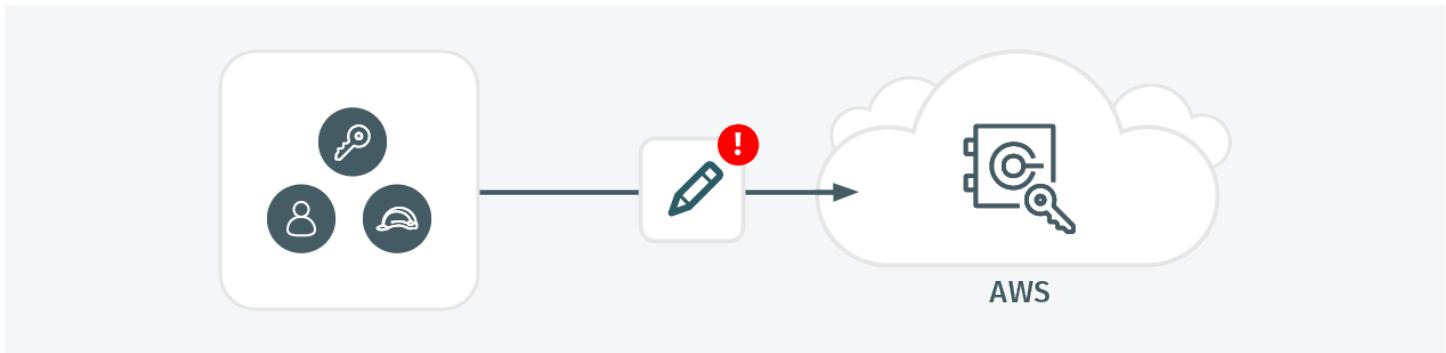
- Ransomware attacks directly impact access to the organization's data and are popular among attackers due to the possibility of a quick transition from attack to monetization. This may lead to operational disruption or permanent data loss.
- After files have been deleted, the attacker will ask the organization to pay a ransom in return for a promise to restore the files.

Steps to Verify

- Investigate the Principal that performed the actions for other signs of malicious activity.
- Examine the CloudTrail logs for suspicious API operation patterns involving 'DeleteObject' events.
- Review the logs for other suspicious enumeration activity such as 'ListObjectVersions' and 'ListObjects'.
- Validate bucket policies, automated scripts, or lifecycle configurations to rule out legitimate activities.
- If review indicates possible malicious actions or high-risk configuration:
 - Revert any configuration changes.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and scope of impacted resources.

AWS Security Tools Disabled

Lateral Movement



MITRE | ATT&CK®

T1562 Impair Defenses

Triggers

- Credential was observed performing a set of API requests capable of disabling native AWS security measures.

Possible Root Causes

- Attackers are attempting to disable or downgrade AWS security mechanisms to blind defenders or to enable further malicious activities without the risk of detection.
- A security or IT service may intentionally be disabling security tools while troubleshooting problems.

Business Impact

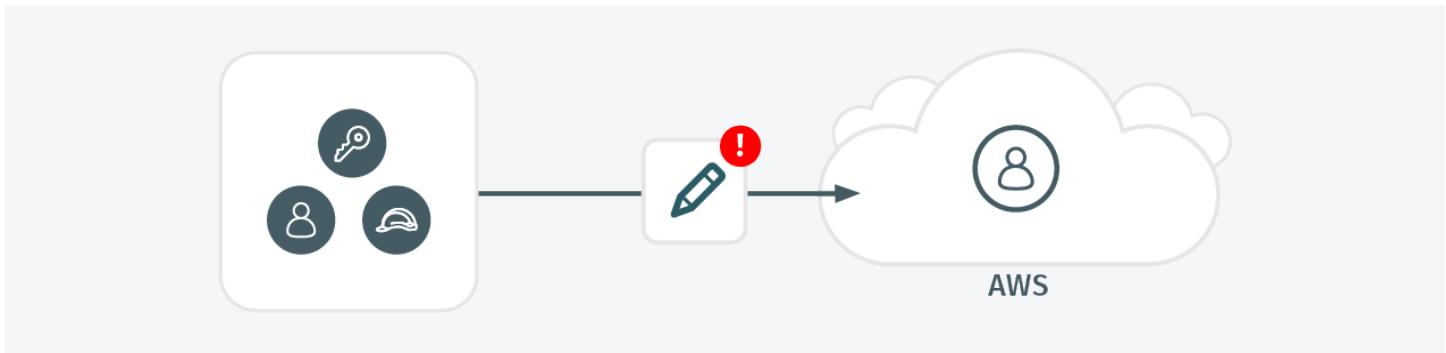
- Attackers who have successfully degraded, disabled, or bypassed security controls can more easily progress towards their objectives.
- Unintentional disabling of security controls increases the potential impact of both present and future attacks against the organization.

Steps to Verify

- Review if this configuration is expected and appropriate in light of any available compensating controls.
- If this is a temporary configuration for troubleshooting purposes, confirm it has been reenabled once that troubleshooting is complete.

AWS Suspect Admin Privilege Granting

Lateral Movement



MITRE | ATT&CK®

T1078 Valid Accounts

T1098 Account Manipulation

Triggers

- Apply a highly permissive inline policy (i.e. ":" or "*:*") to a user, role, or group.

Possible Root Causes

- An attacker is changing the permissions of a user, role, or group to enable them to leverage those permissions to gain additional or persistent access to the environment.
- An administrator has been granted highly permissive policies to enable them complete access to the environment.

Business Impact

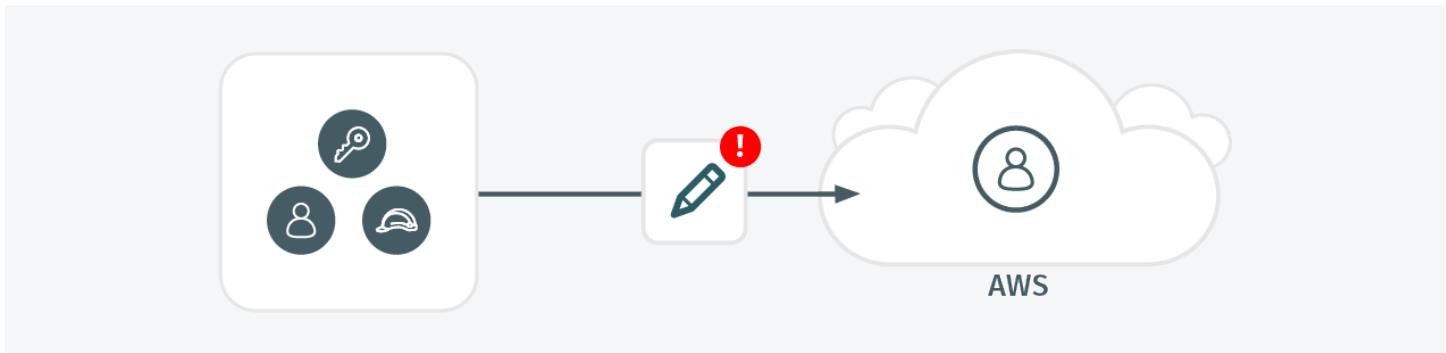
- Lateral movement may indicate that an adversary has established a foothold in the environment and is progressing towards their objective, increasing the risk of material impact.

Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Review whether this account should have access to the console for their normal duties.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Suspect Console Pivot

Lateral Movement



MITRE | ATT&CK®

[T1538 Cloud Service Dashboard](#)

[T1098 Account Manipulation](#)

Triggers

- An account enumerates users or obtains details on their own account, after which they request a token for console login and use that token to login to the console.

Possible Root Causes

- An attacker is pivoting from the AWS API to the AWS management console to continue their attack progression.
- An administrator has started to use the AWS management console in an unusual way.

Business Impact

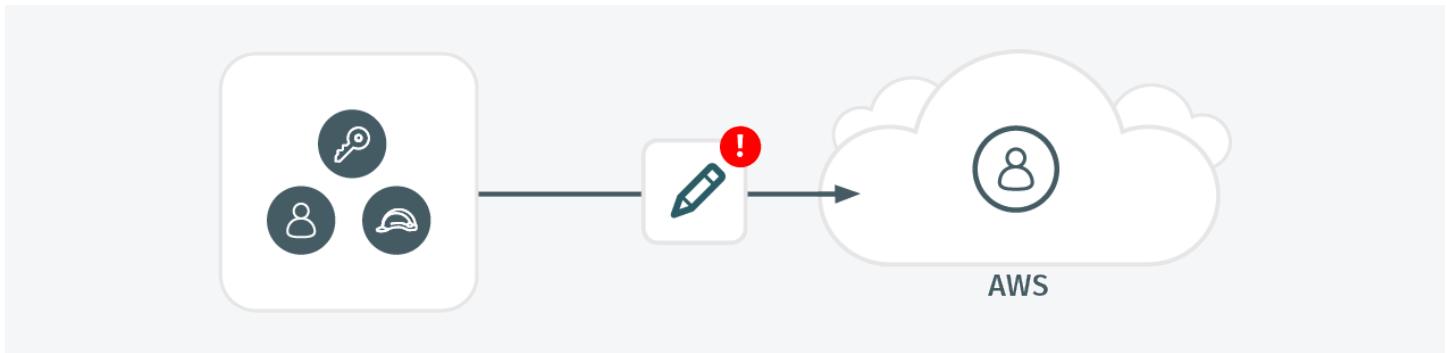
- Lateral movement may indicate that an adversary has established a foothold in the environment and is progressing towards their objective, increasing the risk of material impact.

Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Review whether this account should have access to the console for their normal duties.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Suspect Login Profile Manipulation

Lateral Movement



MITRE | ATT&CK®

T1098 Account Manipulation

Triggers

- A source AWS account modifies the login profile of a target account, following which the target account accesses the AWS console.

Possible Root Causes

- An attacker is enabling access to the console for credentials they have access to, to further their attack.
- An administrator has enabled console access for another user within the environment.

Business Impact

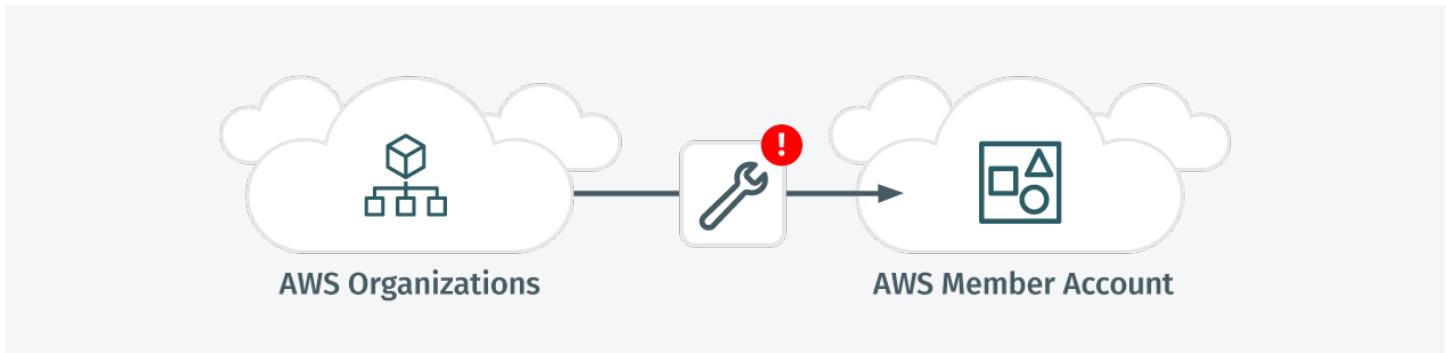
- Lateral movement may indicate that an adversary has established a foothold in the environment and is progressing towards their objective, increasing the risk of material impact.

Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Validate that any modifications are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Suspect Organization Exit

Lateral Movement



MITRE | ATT&CK®

T1562.008 Impair Defenses:
Disable or Modify Cloud
Logs

Triggers

- An AWS control-plane API was invoked in an attempt to leave the AWS Organization in which the target account is a member.

Possible Root Causes

- An attacker is attempting to leave the AWS organization in which the target account is a member. This is done in order to evade restrictions and disrupt logging visibility.
- An administrator or automated task is performing authorized account migration activities.

Business Impact

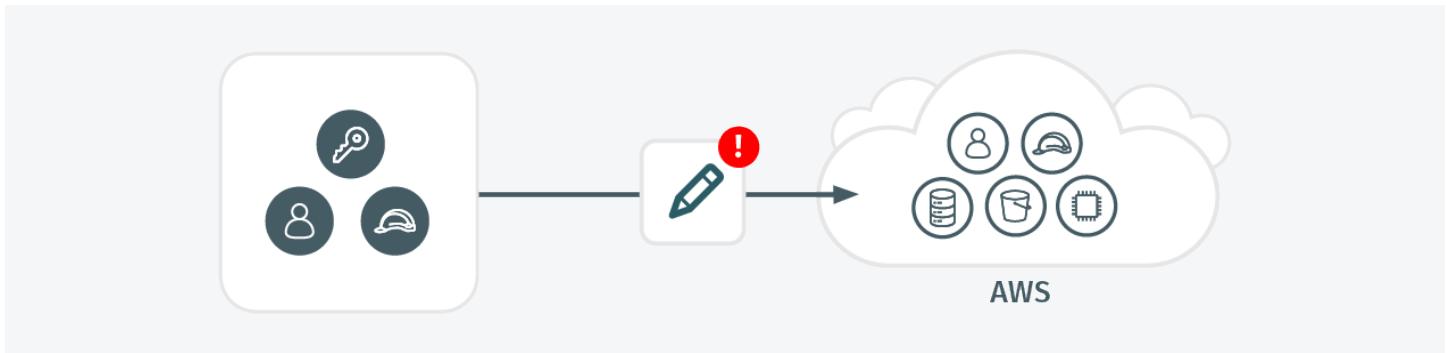
- An attacker who is able to hinder the defenses of their victim also has the ability to evade detection.
- If an attacker is able to successfully remove a targeted AWS account from its AWS Organization:
 - Guardrails such as Service Control Policies (SCP) will be lifted leading to an increased risk of malicious activity in the account.
 - Logging may be interrupted and as a result there would be an increased risk of malicious activity in the account going unnoticed.

Steps to Verify

- Investigate the Principal which performed the actions for other signs of malicious activity.
- Review security policy to determine if the removing the Member Account from the Organization is allowed.
- If review indicates possible malicious actions or high-risk modifications:
 - Disable credentials associated with this alert.
 - Invite the Member Account to re-join the Organization.
 - Establish control over the email inbox of the Member Account Root User in order to approve the invitation to re-join the Organization.
 - Perform a comprehensive investigation to determine initial compromise and the scope of impacted resources.
 - Create a Service Control Policies (SCP) preventing Member Accounts from leaving the Organization.

AWS Suspect Privilege Escalation

Lateral Movement



MITRE | ATT&CK®

T1098 Account Manipulation

Triggers

- Credential was observed performing a set of unusual API requests that enumerate privileges, following which a modification of privileges was observed which may be indicative of a privilege escalation occurring within the environment.

Possible Root Causes

- An attacker has attempted to escalate privileges within the environment.
- An account misconfiguration has weakened IAM protections associated with resource authorizations.
- A security service, administrator, or other automation completed these actions as part of normal environment operation.

Business Impact

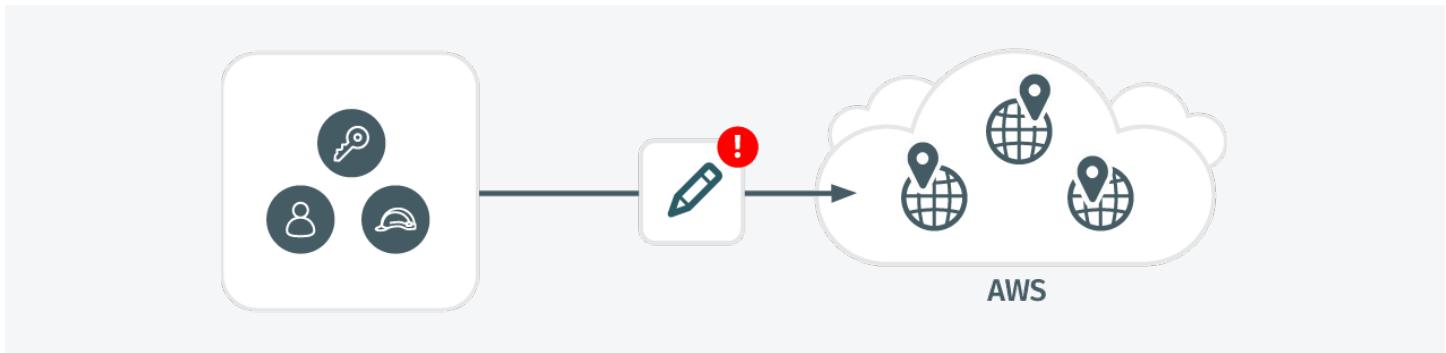
- Privilege escalation may indicate the presence of an adversary that is modifying permissions to progress towards an objective.
- IT misconfigurations may act to increase the risk of impact to assets, data, or services.

Steps to Verify

- Investigate the account context that made the change for other signs of malicious activity.
- Validate that the modifications are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Suspect Region Activity

Lateral Movement



MITRE | ATT&CK®

T1098 Defense Evasion

Triggers

- An AWS control-plane API was observed being invoked in a geographic service region which is not normally used in your cloud environment.

Possible Root Causes

- An attacker is operating out of an unused geographic service region in order to evade detection.
- An administrator or automated task is creating resources in a previously unused geographic service region, as a part of normal activities.

Business Impact

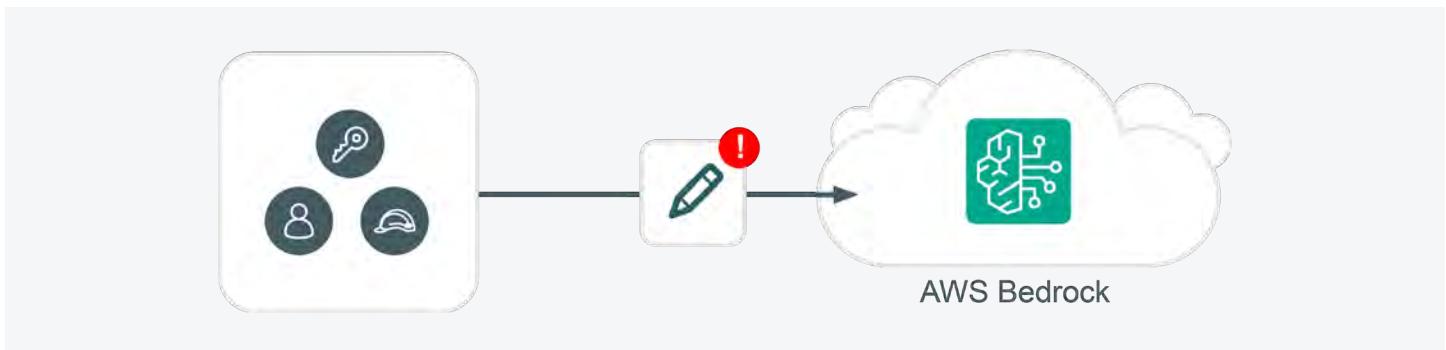
- If logging and monitoring is not inclusive of all regions, a successful attack within an unused region would evade detection capabilities.
- An organization could be hindered from both investigating historical activity and auditing any current activity within the region.

Steps to Verify

- Investigate the Principal which performed the actions for other signs of malicious activity.
- Discuss with the user to determine if the activity is known and legitimate.
- Review Logging and Monitoring coverage for the region to ensure it is adequate.
- Review available internal documentations such as a Cloud Security Policy for details on approved AWS service regions.
- If review indicates possible malicious actions or high-risk configuration:
 - Revert any configuration changes
 - Disable credentials associated with this alert
 - Perform a comprehensive investigation to determine initial compromise and scope of impacted resources

AWS Suspicious Bedrock Activity

Lateral Movement



MITRE | ATT&CK®

AML.T0040

Triggers

- Detection of an identity enabling and then immediately invoking an AWS Bedrock model who has no prior history of performing similar actions.

Possible Root Causes

- Compromised credentials being used to enable and exploit AWS Bedrock for inference, shifting cost to the victim.
- Legitimate usage by a new developer or service.

Business Impact

- Indicates potential misuse of high-cost AI inference services at the victim's expense.
- Could lead to generation of harmful or policy-violating content via LLMs.

Steps to Verify

- Review the CloudTrail events for both model enablement and subsequent invocation.
- Investigate the identity involved in the sequence of activities.
- Check if prompt logging was enabled to inspect prompt content and context surrounding this behavior.

AWS Suspicious S3 Batch Deletion

Lateral Movement



MITRE | ATT&CK®

T1485

Triggers

- AWS control-plane APIs were observed suspiciously downloading and then mass-deleting S3 objects or their historical versions impacting data availability.

Possible Root Causes

- An attacker is utilizing compromised credentials to exfiltrate and destroy data for extortion.
- Legitimate scheduled or routine bucket maintenance processes are performing large-scale object deletion.

Business Impact

- Ransomware attacks directly impact access to the organization's data and are popular among attackers due to the possibility of a quick transition from attack to monetization. This may lead to operational disruption or permanent data loss.
- After files have been deleted, the attacker will ask the organization to pay a ransom in return for a promise to restore the files.

Steps to Verify

- Investigate the Principal that performed the batch delete action for other signs of malicious activity.
- Examine the CloudTrail logs for suspicious API operation patterns involving 'DeleteObjects' events.
- Review the logs for other suspicious enumeration activity such as ListObjectVersions and ListObjects.
- Validate bucket policies, automated scripts, or lifecycle configurations to rule out legitimate activities.
- If review indicates possible malicious actions or high-risk configuration:
 - Revert any configuration changes.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and scope of impacted resources.

AWS Suspicious S3 Encryption

Lateral Movement



MITRE | ATT&CK®

T1486

Triggers

- An AWS control-plane API was observed encrypting a large number of s3 objects and specified either an external KMS key (SSE-KMS) or a client-controlled key (SSE-C) to perform the server-side encryption operations.

Possible Root Causes

- An attacker is leveraging AWS APIs to encrypt and overwrite s3 objects in a way that may indicate the exfiltration phase of ransomware activity in the environment, with the goal of demanding a ransom.
- Legitimate automated scripts or processes are performing scheduled or bulk data encryption as part of security compliance or internal data protection practices.

Business Impact

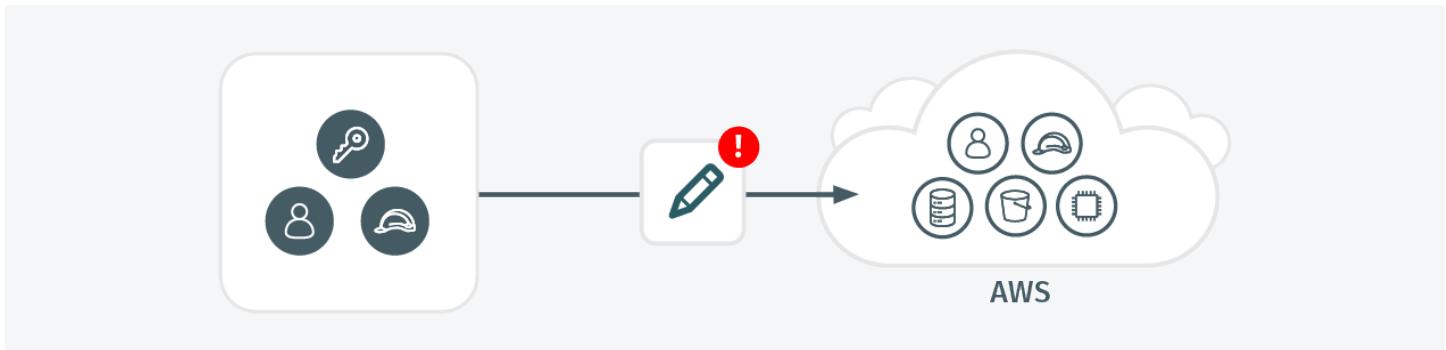
- Ransomware attacks directly impact access to the organization's data and are popular among attackers due to the possibility of a quick transition from attack to monetization. This may lead to operational disruption due to data unavailability.
- After files have been encrypted, the attacker will ask the organization to pay a ransom in return for a promise to provide access to the decryption key.
- Even if an organization is willing to pay the ransom, there is no guarantee that access to the decryption key will be provided by the attacker.

Steps to Verify

- Investigate the Principal that performed the actions for other signs of malicious activity.
- Inspect the CloudTrail logs for suspicious SSE-C or SSE-KMS encrypted CopyObject or PutObject operations.
- If the encryption type is SSE-KMS, investigate the AWS Account in which the KMS key resides. Determine if this Account is authorized to perform cryptographic operations on s3 objects.
- Review the logs for other suspicious enumeration activity such as ListObjectVersions and ListObjects.
- Validate that any movement and encryption of data in s3 is authorized, given the purpose and policies governing the s3 objects.
- If review indicates possible malicious actions or high-risk configuration:
 - Revert any configuration changes.
 - Disable credentials associated with this alert.
 - Review bucket versioning status and assess data recovery capabilities.
 - Perform a comprehensive investigation to determine initial compromise and scope of impacted resources.

AWS User Hijacking

Lateral Movement



MITRE | ATT&CK®

T1098 Account Manipulation

Triggers

- After enumerating users in the environment, add an access key to another user in the environment.

Possible Root Causes

- An attacker is expanding access to additional users within the environment.
- Authorized IT Automation is using access keys to interact on behalf of other users within the environment.

Business Impact

- Lateral movement may indicate that an adversary has established a foothold in the environment and is progressing towards their objective, increasing the risk of material impact.

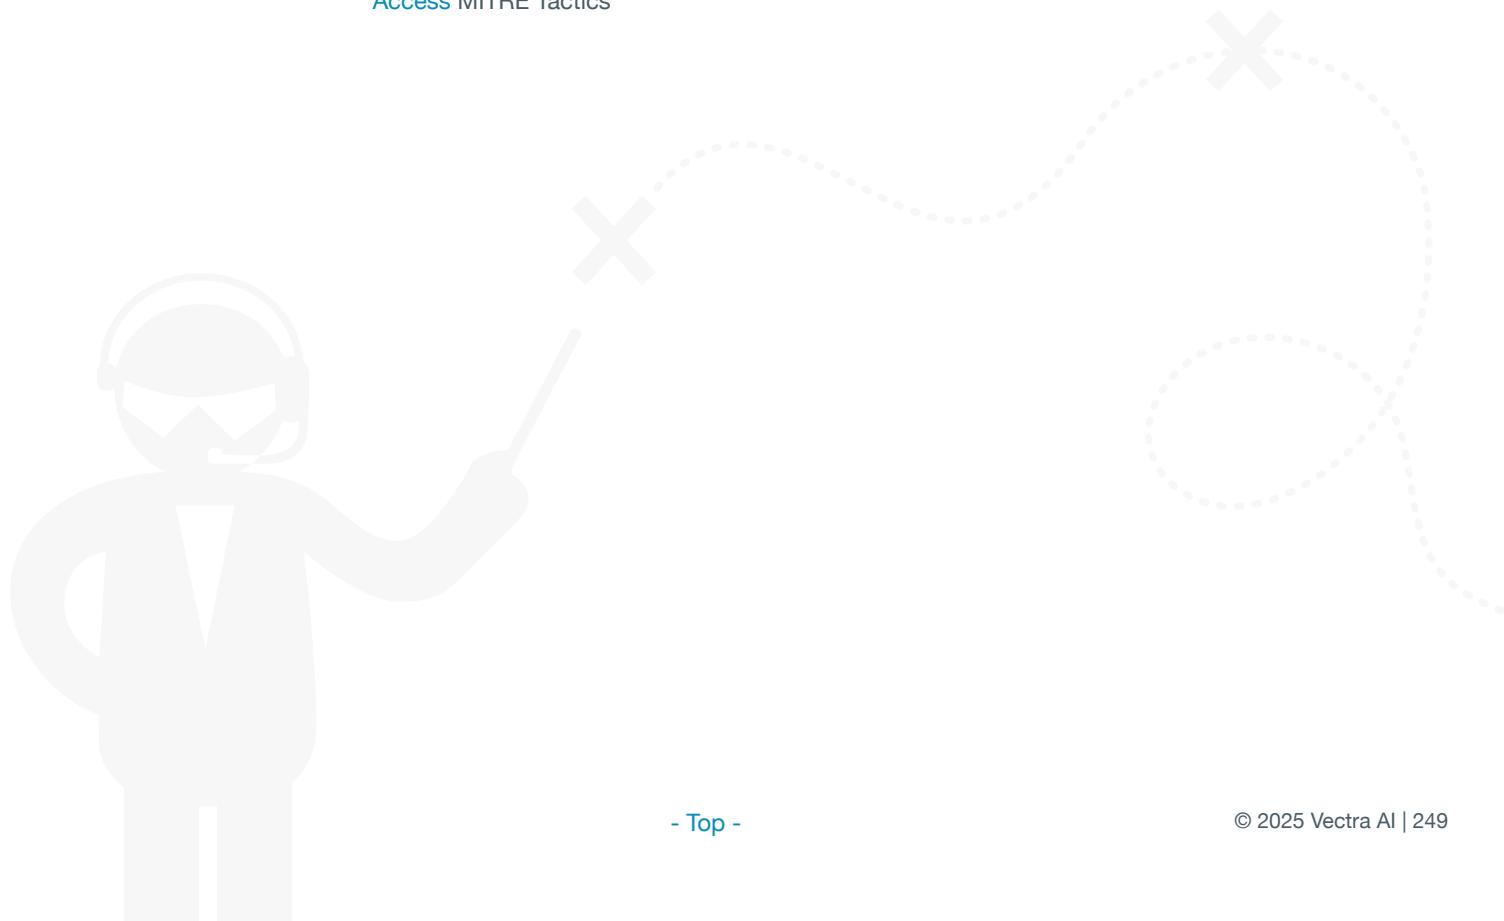
Steps to Verify

- Investigate the account context that performed the action for other signs of malicious activity.
- Validate that any modifications are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

Category

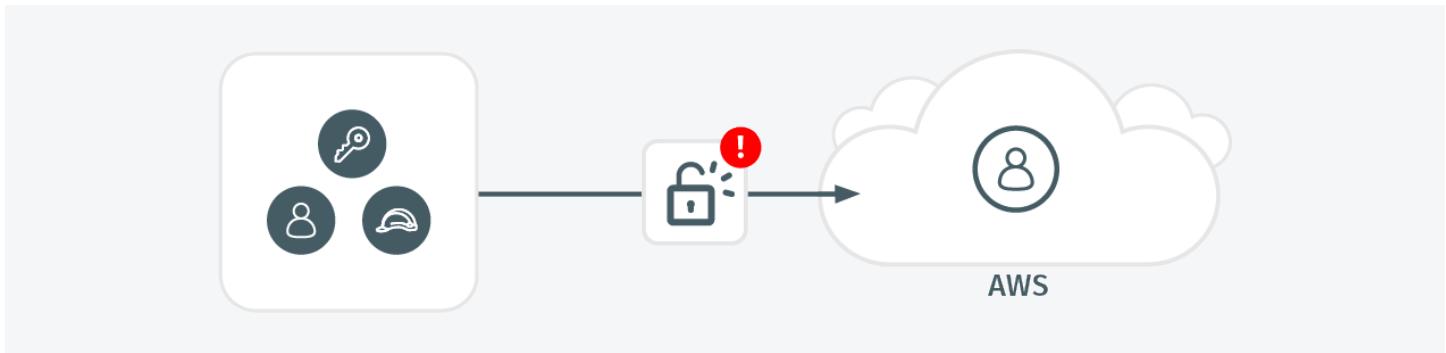
Exfiltration

- An attacker with access to the AWS environment is collecting and removing data from the environment
- Attackers after gaining access and gaining sufficient permissions will steal high value data
- Attackers will take actions to modify services and identifies to ensure continued access
- Attackers will leverage credentials within their defined permissions but for non-intended purposes to further the attacker's objective
- Actions are associated with [Execution](#), [Persistence](#), [Privilege Escalation](#), [Defense Evasion](#), [Lateral Movement](#), [Credential Access](#) MITRE Tactics



AWS Suspect External Access Granting

Exfiltration



MITRE | ATT&CK®

T1078 Valid Accounts

T1098 Account Manipulation

Triggers

- A credential was observed enabling external access to AWS resources through an IAM role.

Possible Root Causes

- An attacker may be creating a means of accessing data from a separate AWS account.
- A sanctioned third-party security or IT service may be granted access to AWS resources in order to perform normal activities.

Business Impact

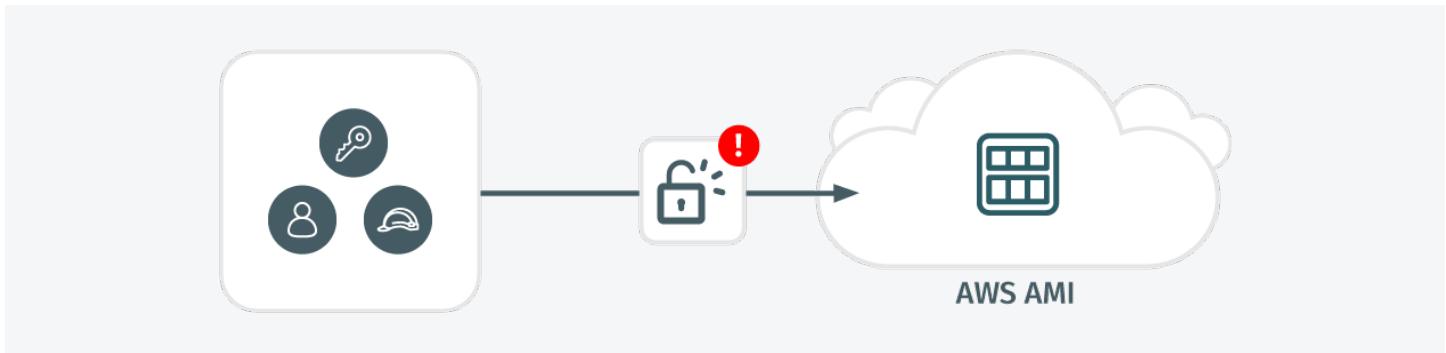
- Once an adversary achieves persistent access, they've established the opportunity to stage subsequent phases of an attack.

Steps to Verify

- Validate that the access is authorized, given the purpose and policies governing these resources.
- If review indicates possible malicious actions or high-risk configuration, delete the created IAM role and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Suspect Public AMI Change

Exfiltration



MITRE | ATT&CK®

T1537 Transfer Data to Cloud Account

Triggers

- An AWS control-plane API was invoked which modified the launch permissions of an Amazon Machine Image (AMI) granting either an unknown, external account or the public the ability to launch an EC2 instance from the image.

Possible Root Causes

- An attacker may be exfiltrating data contained in the Amazon Machine Image (AMI) by sharing it externally and launching an instance from the stolen template.
- An authorized administrator may be performing a backup, disaster recovery activities or sharing the image in order to coordinate troubleshooting efforts.

Business Impact

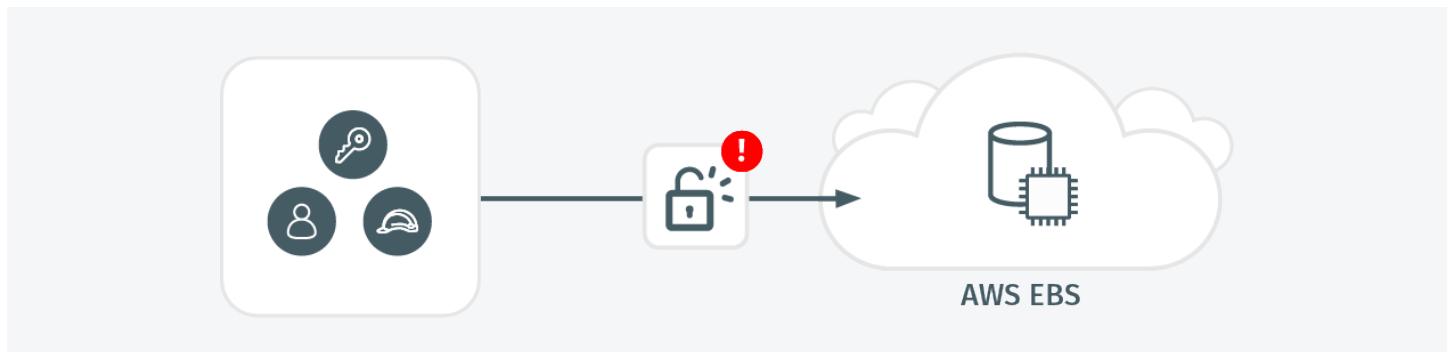
- Exfiltration of Amazon Machine Images (AMI) by an attacker may expose details that support further attack progression. An impacted organization may incur data loss, impacting the confidentiality of sensitive information contained in the impacted Amazon Machine Images (AMIs).

Steps to Verify

- Investigate the Principal that performed the actions for other signs of malicious activity.
- Investigate the affected AMI for potential data loss.
- Validate that any modifications to AMI launch permissions are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration:
 - Revert any configuration changes.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and scope of impacted resources.

AWS Suspect Public EBS Change

Exfiltration



MITRE | ATT&CK®

[T1213 Data from Information-
Repositories](#)

[T1530 Data From Cloud
Storage Object](#)

[T1537 Transfer Data to Cloud
Account](#)

Triggers

- A credential was observed performing a set of AWS control plane API actions related to exfiltration EC2 snapshots.

Possible Root Causes

- An attacker may be actively looking for privilege escalation opportunities
- A security or IT service may intentionally be enumerating these APIs for monitoring reasons.

Business Impact

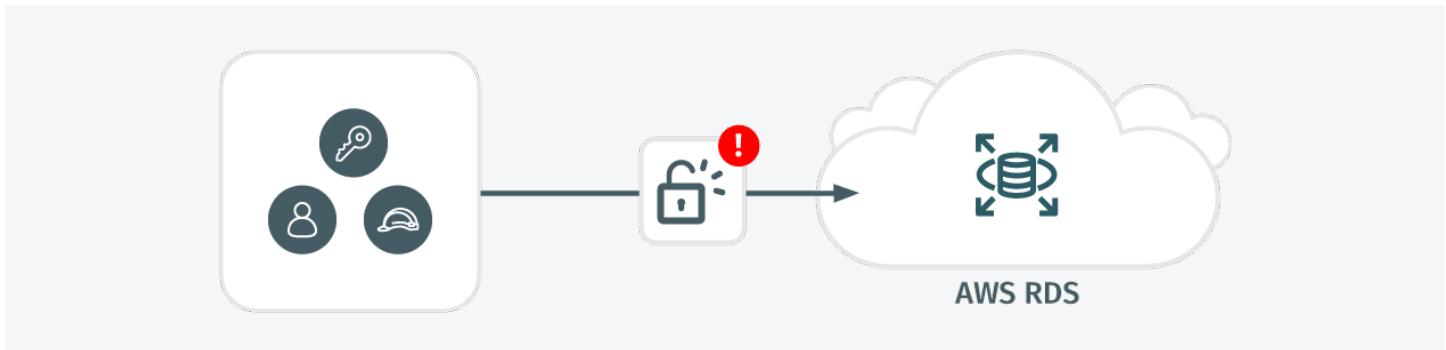
- Exfiltration by an attacker of EC2 snapshots may expose details that support further attack progression, or lead to data loss.

Steps to Verify

- Investigate the account context that performed this action for other signs of malicious activity.
- Investigate for data loss.
- If review indicates possible malicious actions or high-risk configuration, revert applicable configurations and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Suspect Public RDS Change

Exfiltration



MITRE | ATT&CK®

[T1537 Transfer Data to Cloud Account](#)

Triggers

- An AWS control-plane API was invoked which modified the attributes of a RDS snapshot granting either an unknown, external account or the public the ability to restore a RDS database from the snapshot.

Possible Root Causes

- An attacker may be exfiltrating data contained in the RDS database by sharing a snapshot externally.
- An authorized administrator may be performing backup or disaster recovery activities or sharing snapshots in order to coordinate troubleshooting efforts.

Business Impact

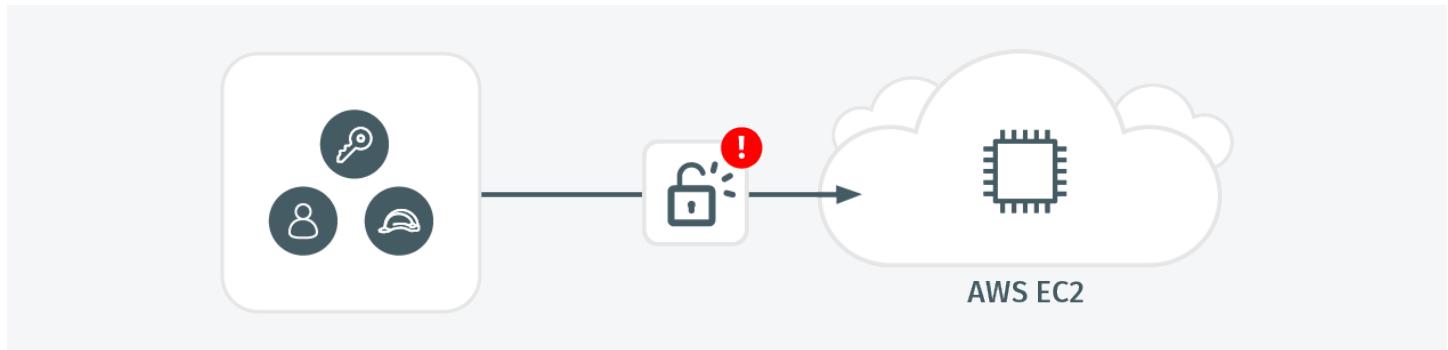
- Exfiltration of RDS database snapshots by an attacker may expose details that support further attack progression. An impacted organization may incur data loss, impacting the confidentiality of sensitive information contained in the impacted RDS database.

Steps to Verify

- Investigate the Principal that performed the actions for other signs of malicious activity.
- Investigate for potential data loss.
- Validate that any modifications to snapshot attributes are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration:
 - Revert any configuration changes.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and scope of impacted resources.

AWS Suspect Public EC2 Change

Exfiltration



MITRE | ATT&CK®

[T1213 Data from Information-
Repositories](#)

[T1530 Data From Cloud
Storage Object](#)

[T1537 Transfer Data to Cloud
Account](#)

[T1578 Modify Cloud Compute
Infrastructure](#)

Triggers

- After enumerating the existing security group policies, the ingress policy for an EC2 instance is modified.

Possible Root Causes

- An attacker is enabling external access to an EC2 instance to maintain persistence.
- An EC2 instance is exposed to external access as a part of its normal operation.

Business Impact

- Once an adversary achieves persistent access, they've established the opportunity to stage subsequent phases of an attack.

Steps to Verify

- Validate that any modifications are authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

AWS Suspect Public S3 Change

Exfiltration



MITRE | ATT&CK®

[T1213 Data from Information-
Repositories](#)

[T1530 Data From Cloud
Storage Object](#)

[T1537 Transfer Data to Cloud
Account](#)

Triggers

- A credential was observed suspiciously invoking a set of S3 APIs that permits public access to a given bucket.

Possible Root Causes

- An attacker may be scanning and maliciously modifying configurations around an S3 bucket to enable data exfiltration.
- An IT misconfiguration may have been made by an authorized user which could weaken the posture around an S3 bucket and promote the risk of data loss.
- An internal tool is scanning the buckets for security reasons.

Business Impact

- Malicious or unintentional weakening of security posture controls around S3 buckets are commonly associated with data loss.

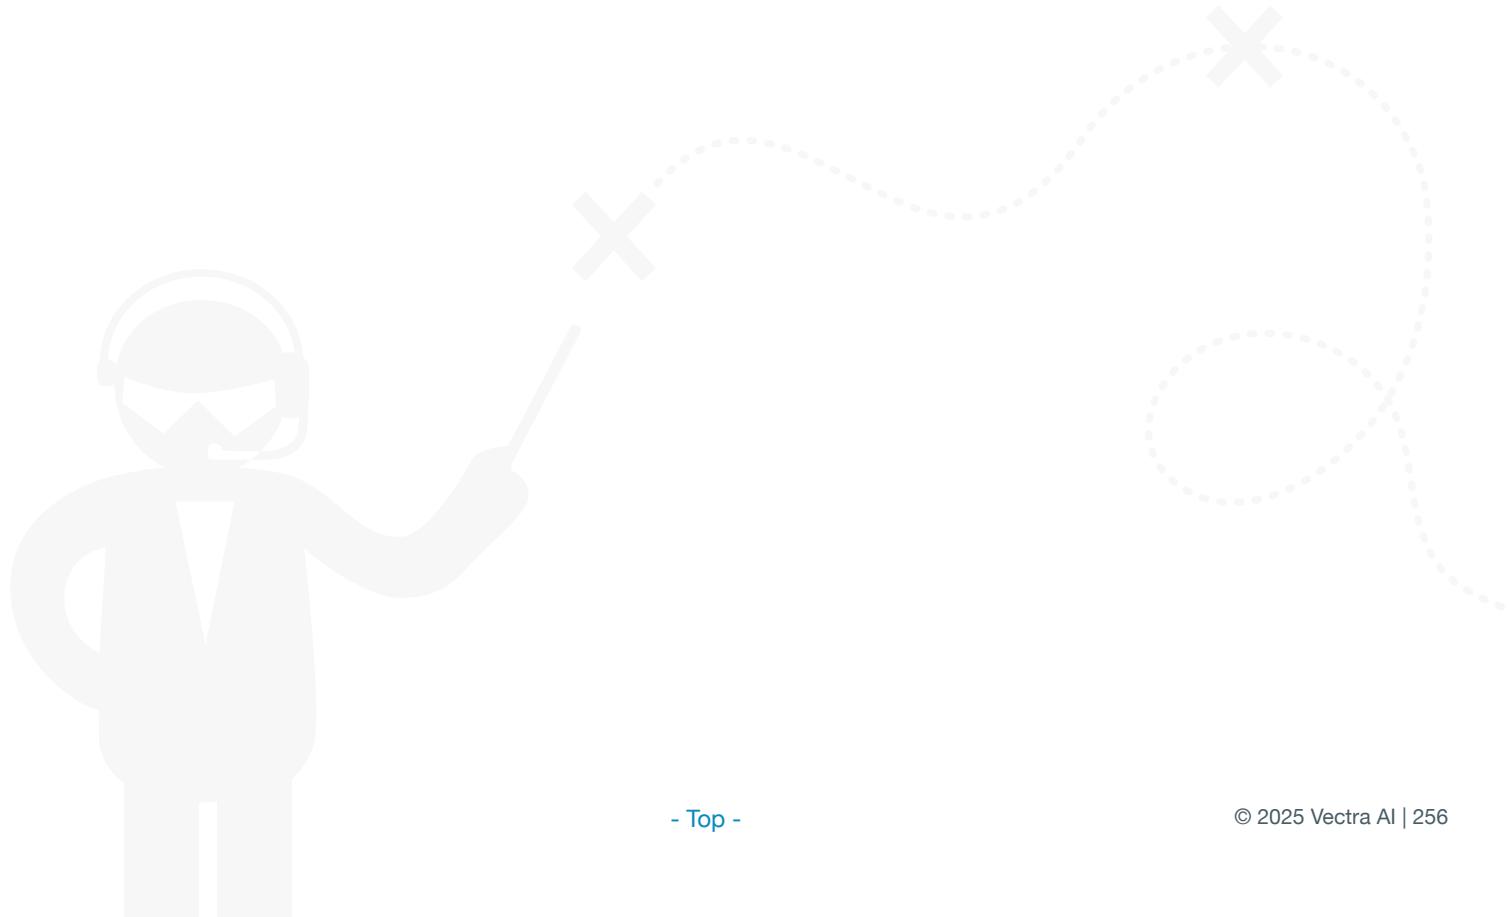
Steps to Verify

- Investigate the account context that made the change for other signs of malicious activity.
- Investigate for data loss.
- Verify if the S3 bucket in question is authorized for public access.
- If review indicates possible malicious actions or high-risk configuration, revert configuration and disable credentials associated with this alert then perform a comprehensive investigation.

Category

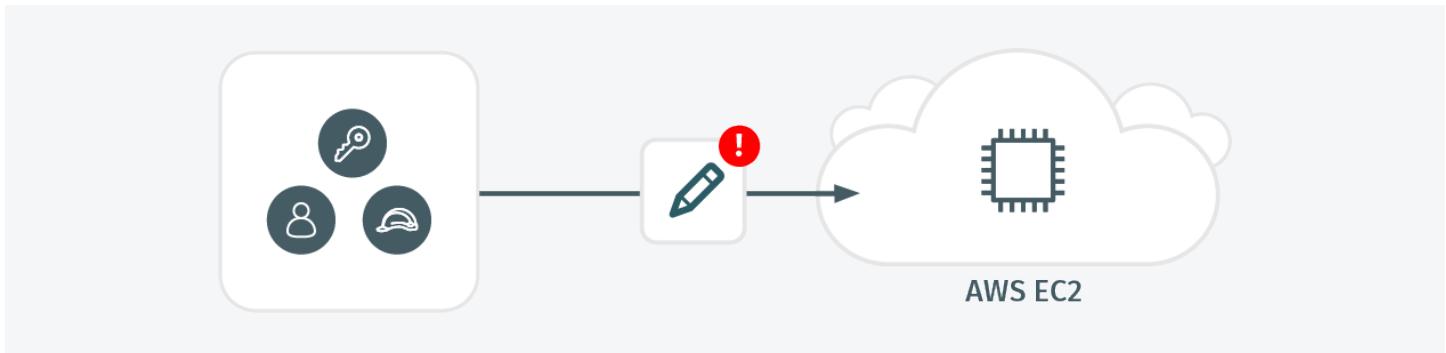
Botnet

- An attacker with access to the AWS environment is leveraging AWS infrastructure for financial gain
- Actions are associated with the [Impact](#) MITRE Tactic



AWS Cryptomining

Botnet



MITRE | ATT&CK®

T1496 Resource Hijacking

Triggers

- An AWS control-plane API was observed creating new EC2 instances that are ideally suited for cryptomining operations.

Possible Root Causes

- An attacker is leveraging a compromised principal, EC2 instance and/or token to create powerful EC2 instances for use in cryptomining.
- Internal users, infrastructure and applications create highly powered EC2 instances to enable compute intensive operations to occur in support of authorized business activity.

Business Impact

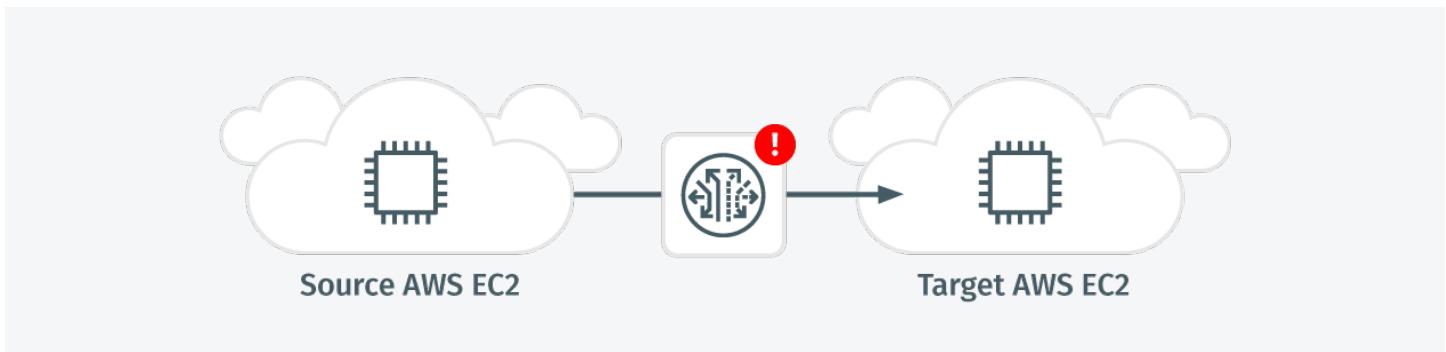
- High powered EC2 instances utilized for cryptomining result in significant costs billed to the organization that owns the AWS account.

Steps to Verify

- Investigate the principal and determine if this resource is authorized to create new, high-powered, EC2 instances.
- Investigate the newly created EC2 instances to determine their purpose and ensure they are not malicious.
- If review indicates possible malicious actions:
 - Perform a comprehensive investigation to determine where initial compromise occurred on the source principal.
 - Revoke the compromised principal's credentials.
 - Remediate any resources and accounts that were created by the compromised principal.

AWS Suspect Traffic Mirror Creation

Botnet



MITRE | ATT&CK®

T1040 Network Sniffing

Triggers

- An AWS control-plane API was invoked, which leveraged an EC2 instance as a traffic mirroring target. This suggests a malicious network traffic session will be created, mirroring traffic to the target EC2 instance.

Possible Root Causes

- A malicious actor is mirroring network traffic to an attacker controlled EC2 in order to steal credentials like passwords and further pivot into the environment.
- An administrator may have intentionally configured an EC2 as a traffic mirroring target as part of normal operations.

Business Impact

- Malicious traffic mirroring can be extremely impactful as the traffic moving within VPCs is frequently unencrypted. This is common due to the cloud network design practice of terminating SSL/TLS encryption at load balancers.
- Stolen credentials sniffed from a network can further an attack campaign, impacting the confidentiality of data stored on impacted systems.
- When confidentiality of data is affected, there may be regulatory or compliance implications for the business.

Steps to Verify

- Investigate the Principal that performed the actions for other signs of malicious activity.
- Validate that the creation of the traffic mirroring target is authorized, given the purpose and policies governing this resource.
- Review CloudTrail logs to determine if a traffic mirroring session was established and is authorized, given the purpose and policies governing this resource.
- If review indicates possible malicious actions or high-risk configurations were made:
 - Revert any configuration changes.
 - Terminate any traffic mirroring session created by the Principal.
 - Disable credentials associated with this alert.
 - Perform a comprehensive investigation to determine initial compromise and if network traffic from the source EC2 instance was encrypted in transit.

Category

Info

- AWS Bedrock Model Enabled: Detects every instance when an AWS Bedrock foundational model is enabled, as this action is uncommon and may have cost or security implications.