Automatic disruption of human-operated attacks through containment of compromised user accounts | Microsoft Security Blog

By Microsoft Threat Intelligence

Published: 2023-10-11 · Archived: 2026-04-05 16:53:33 UTC

Our experience and insights from real-world incidents tell us that the swift containment of compromised user accounts is key to disrupting hands-on-keyboard attacks, especially those that involve human-operated ransomware. In these attacks, lateral movement follows initial access as the next critical stage for attackers to advance their objective of targeting valuable assets and sensitive data. Successful lateral movement depends on attackers' ability to compromise user accounts and elevate permissions: our observations of attacks show that all human-operated ransomware attacks where ransomware deployment was successful involve attackers gaining access to a domain admin-level account or local administrator passwords.

Attackers compromise user accounts through numerous and diverse means, including techniques like credential dumping, keylogging, and brute-forcing. Poor credential hygiene could very quickly lead to the compromise of domain admin-level accounts, which could allow attackers to access domain resources and devices, and completely take over the network. Based on incidents analyzed by Microsoft, it can take only a single hop from the attacker's initial access vector to compromise domain admin-level accounts. For instance, an attacker can target an over-privileged service account configured in an outdated and vulnerable internet-facing server.

Highly privileged user accounts are arguably the most important assets for attackers. Compromised domain admin-level accounts in environments that use traditional solutions provide attackers with access to Active Directory and could subvert traditional security mechanisms. In addition to compromising existing accounts, attackers have adopted the creation of additional dormant, highly privileged user accounts as persistence mechanisms.

Identifying and containing these compromised user accounts, therefore, prevents attacks from progressing, even if attackers gain initial access. This is why, as announced today, we added user containment to the automatic attack disruption capability in Microsoft Defender for Endpoint, a unique and innovative defense mechanism that stops human-operated attacks in their tracks. User containment prevents a compromised user account from accessing endpoints and other resources in the network, limiting attackers' ability to move laterally regardless of the account's Active Directory state or privilege level. It is automatically triggered by high-fidelity signals indicating that a compromised user account is being used in an ongoing attack. With user containment, even compromised domain admin accounts cannot help attackers access other devices in the network.

In this blog we will share our analysis of real-world incidents and demonstrate how automatic attack disruption protected our customers by containing compromised user accounts. We then explain how this capability fits in our automatic attack disruption strategy and how it works under the hood.

User containment stops Storm-1567 attack, prevents Akira ransomware encryption

In early June 2023, an industrial engineering organization was the target of a human-operated attack by an Akira ransomware operator tracked by Microsoft as Storm-1567. Akira is a ransomware strain first observed by Microsoft in March 2023 and has features common to other ransomware payloads like the use of ChaCha encryption algorithm, PowerShell, and Windows Management Instrumentation (WMI). Microsoft assesses that Akira is most likely a closed ransomware offering and not openly marketed as ransomware as a service.

In this attack, the threat actor leveraged devices that were not onboarded to Microsoft Defender for Endpoint for most of the attack stages, a defense evasion tactic we've seen in other attacks. While visibility by our endpoint solution could have blocked the attack earlier in the attack chain and helped to protect the organization's devices much sooner, Defender for Endpoint nonetheless successfully prevented the ransomware stage, protecting all onboarded devices in the organization from getting encrypted.

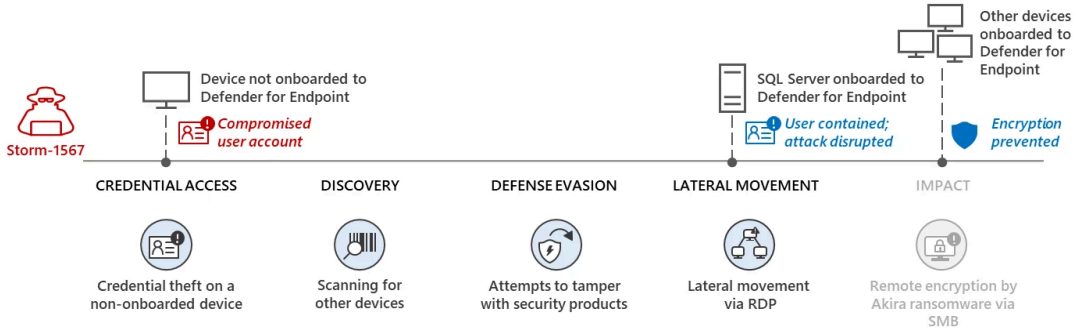

Figure 1. Storm-1567 attempt to encrypt devices

Based on our analysis, after gaining access to the network, the threat actor started preparing to encrypt devices by scanning, attempting to tamper with security products, conducting lateral movement using Remote Desktop Protocol (RDP), and other anomalous activities. It should be noted that the activities were conducted on a Sunday evening, a time when SOC teams might be at a limited capacity. Most of these activities were done on Windows Server devices, including SQL Servers onboarded to Microsoft Defender for Endpoint. These activities were highly anomalous compared to routine activity in the customer's network and therefore triggered multiple alerts.

Microsoft Defender for Endpoint's next-generation protection capabilities detected and prevented several attacker activities, prompting the attackers to try tampering with the security product. However, tamper protection was enabled in the environment, so these attempts were not successful. Meanwhile, Microsoft 365 Defender correlated signals from multiple Defender products, identified the malicious activity, and incriminated – that is, determined as malicious with high confidence – the associated compromised assets, including a user account the attackers used.

Approximately half an hour after activity began, attackers leveraged the compromised user account and attempted to encrypt devices remotely via Server Message Block (SMB) protocol from a device not onboarded to Microsoft Defender for Endpoint. Because of the earlier incrimination, the compromised user account was contained, and the devices onboarded to Defender for Endpoint were protected from encryption attempts.

Later the same day, the attackers repeated the same malicious sequences by pivoting to other compromised user accounts, attempting to bypass attack disruption protection. Defender for Endpoint was again able to protect onboarded devices from encryption over the network. In this incident, automatic attack disruption's ability to contain additional compromised user accounts demonstrated unique and innovative impact for endpoint and identity security, helping to protect all devices onboarded to Defender for Endpoint from the attack.

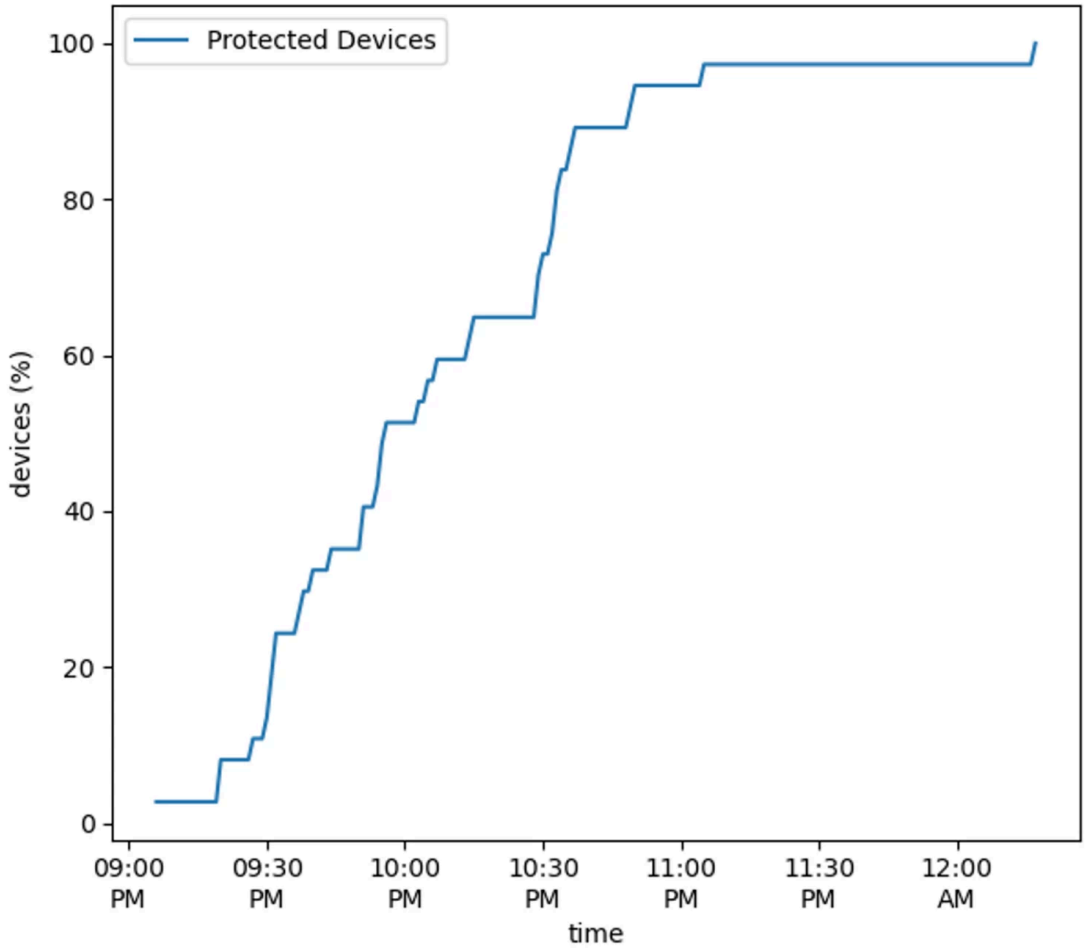

Figure 2. Chart showing remote encryption attempts being blocked on devices onboarded to Microsoft Defender for Endpoint as the attack progresses

User containment stops lateral movement in human-operated campaign

In early August 2023, Microsoft Defender for Endpoint automatically disrupted a human-operated attack early in the attack chain by containing the compromised user account prior to any impact, saving a medical research lab from what could have been a large-scale attack. The first indication of the attack was observed at roughly 4:00 AM local time on a Friday, when attackers, operating from a device not onboarded to Defender for Endpoint, initiated a remote password reset for the default domain administrator account. This account wasn't active on any device onboarded to Microsoft Defender for Endpoint in the months prior to the intrusion. We infer that the account credentials were likely expired, and that the attackers found the stale password hashes belonging to the account by using commodity credential theft tools like Mimikatz on a device not-onboarded to Microsoft Defender for Endpoint. Expired credentials, while often not seen as a security risk, could still be abused and could allow attackers to update an account's password.

Minutes after the administrator account password was reset, the attackers started scanning the network for accessible shares and enumerated other account and domain configurations using SMB-accessible services. This scan and all subsequent malicious activities originated from the same non-onboarded device and compromised administrator account.

Parallel to the network scan, the threat actor initiated an RDP session to a SQL Server, attempting to tamper with security products on the server and running a variety of credential theft and domain discovery tools.

At this point, the compromised administrator account was incriminated based on cumulative signals from the Defender for Endpoint-onboarded SQL server and the account's anomalous activity. Automatic attack disruption was triggered and the compromised account was contained. All devices in the organization that supported the user containment feature immediately blocked SMB access from the compromised user account, stopping the discovery operations and preventing the possibility of subsequent lateral movement.

Following the initial containment of the attack through automatic attack disruption, the SOC was then able to take additional critical remediation actions to expand the scope of the disruption and evict the attackers from the network. This included terminating the attackers' sessions on two compromised servers and disabling the compromised domain administrator account at the Active Directory-level.

While user containment is automatic for devices onboarded to Defender for Endpoint, this incident demonstrates the importance of active engagement of the SOC team after the automatic attack disruption action to fully evict the attackers from the environment. It also shows that onboarding devices to Microsoft Defender for Endpoint improves the overall capability to detect and disrupt attacks within the network sooner, before high-privileged user accounts are compromised.

In addition, as of September 2023, user containment also supports terminating active RDP sessions, in addition of blocking new attempted connections, a critical first step in evicting attackers from the network. Disabling compromised user accounts at the Active Directory-level is already supported by automatic attack disruption through integration with Defender for Identity. In this particular incident, the customer was not using Defender for Identity, but this case highlights the stronger defenses as a result of cross-domain visibility.

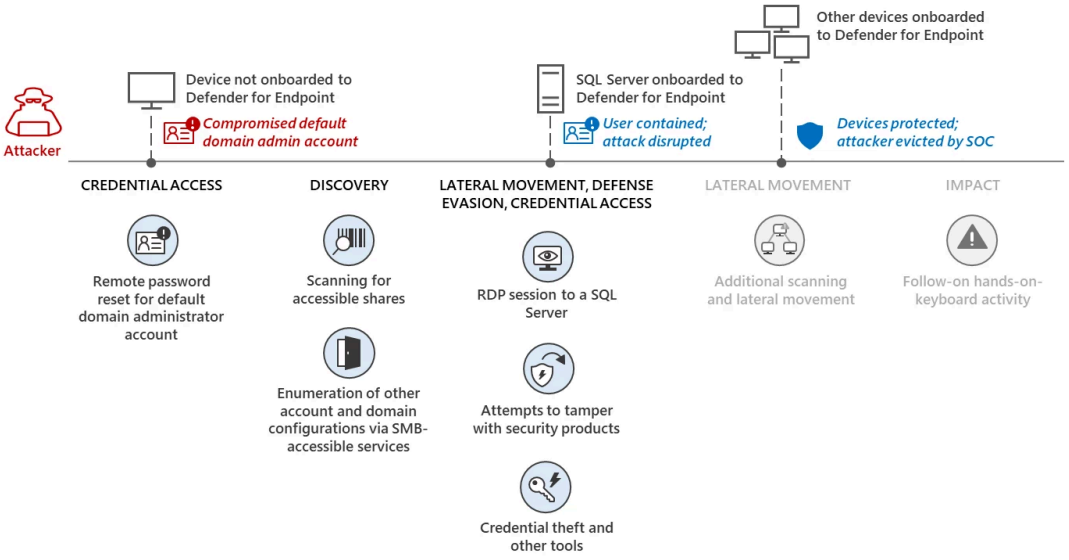

Figure 3. Attack chain of human-operated campaign that targeted a medical research lab

Protecting against compromised user accounts through automatic containment

As demonstrated by the incidents we described above, unlike commodity malware infection, human-operated attacks are driven by humans with hands-on-keyboard access to the network who make decisions at every stage of their attack. Attack patterns vary depending on what attackers find in the target network. Protecting against such highly skilled, profit-driven, and determined adversaries is not trivial. These attackers leverage key principles of on-premises Active Directory environments, which provide an active domain administrator account unlimited access to domain resources. Once attackers obtain accounts with sufficient privileges, they can conduct malicious activities like lateral movement or data access using legitimate administrative tools and protocols.

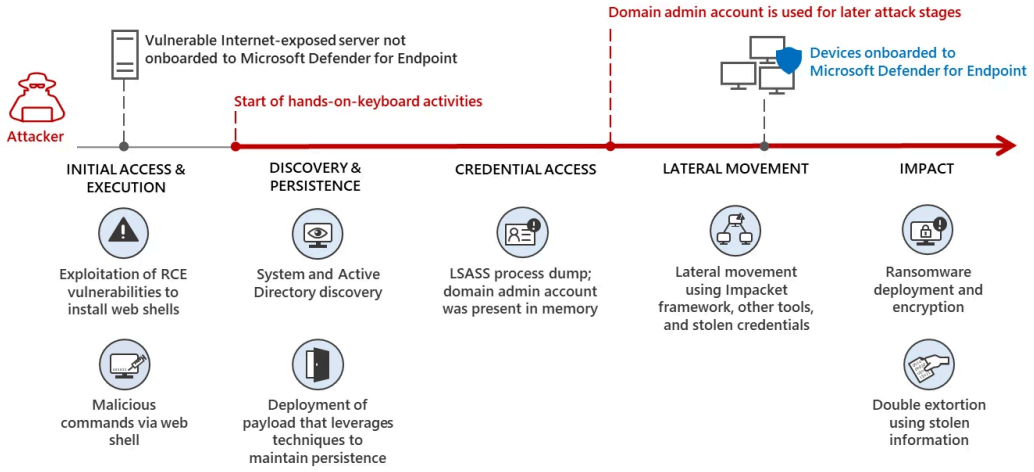

Figure 4. An example of a malicious activity of compromised user accounts in a human-operated ransomware attack

At Microsoft, we understand that to better defend our customers against such highly motivated attackers, a multi-layer defense approach must be used for an optimal security protection solution across endpoints and identities. More importantly, this solution should prioritize organization-wide protection, rather than protecting only a single endpoint. Motivated attackers search for security weaknesses and prioritize compromising unprotected devices. As a result, assuming that initial attack stages have occurred, with potentially at least a few compromised user accounts, is critical for developing security defenses for later attack stages. Using key assumptions and principles of on-premises Active Directory environments, a security-first mindset means limiting the access of even the most privileged user accounts to mitigate security risks.

The automatic attack disruption capability contains user accounts by creating a boundary between healthy onboarded devices and compromised user accounts and devices. It works in a decentralized nature: a containment policy distributed to all onboarded devices across the organization enables each Microsoft Defender for Endpoint client to protect the device against any compromised account, even an account belonging to the Domain Admins group.

This decentralized approach avoids some of the pitfalls of centralized manual or automatic controls, such as disabling an account in Active Directory, which possesses a single point of failure as it can be overridden by the attacker who may already have compromised domain controllers. The virtual security boundary set to contain the user is implemented by controls that were tailored to disrupt attacker activity during various attack stages, including lateral movement, credential theft, and impact such as remote encryption or deployment of ransomware payload. The actual set of controls triggered to contain a user might vary depending on the attack scenario and stage, and includes:

1. Sign-in restriction: This is the most aggressive control in containing a user account. When this control is triggered, devices will deny all or some types of sign-ins by a compromised account. This control takes effect immediately and is effective regardless of the account's state (i.e., active or disabled) in the authority it belongs to. This control can block most attacker capabilities, but in cases where an attacker had already authenticated to device before a compromise was identified, the other controls might still be required to block the attack.
2. Intercepting SMB activity: Attack disruption can contain a user by denying inbound file system access from a remote origin, limiting the attacker's ability to remotely steal or destroy valuable data. Notably, this control can prevent or limit ransomware encryption over SMB. It can also block lateral movement methods that include a payload being created on a remote device, including PsExec and similar tools.
3. Filtering RPC activity: Attack disruption can selectively restrict compromised users' access to remote procedure call (RPC) interfaces that attackers often leverage during attacks. Attackers abuse RPC-based protocols for a variety of goals such credential theft (DCsync and DPAPI), privilege escalation ("PetitPotam", Print Spooler), discovery (server & workstation services), and lateral movement (remote WMI, scheduled tasks, and services). Blocking such activities can contain an attack before the attacker gains a strong foothold in the network or can deny the ability to capitalize on such a foothold during the impact stage.
4. Disconnecting or terminating active sessions: In case a compromised account had already gained a foothold on the device, when attack disruption is triggered, it can disconnect or terminate sessions previously initiated by the account. This control differs from the others in this list as it's effective against already compromised devices, protecting against any additional malicious activity by the attacker. Once a session is terminated, attackers are locked out of the device by the sign-in restriction control. This is specifically critical in stopping attacks earlier in the attack chain, disrupting and containing attacks before reaching impact stage.

The user containment capability is part of the existing protections provided by solutions within Microsoft 365 Defender. As we described in this blog, this capability correlates high-fidelity signals from multiple Defender products to incriminate malicious entities with high confidence and then immediately contain them to automatically disrupt ongoing attacks, including the pre-ransomware and encryption stages in human-operated attacks.

To benefit from this capability, organizations need only to onboard devices to Microsoft Defender for Endpoint. As more devices are onboarded, the scope of disruption is larger and the level of protection is higher. And as more Defender products are used in the organization, the visibility is wider and the effectiveness of the solution is greater. This also lowers the risk of attackers taking advantage of unprotected devices as launch pads for attacks.

Automatic attack disruption represents an innovative solution designed to increase defenses against the increasingly more sophisticated threat of hands-on-keyboard attacks, especially human-operated ransomware. This capability is informed by threat intelligence and insights from investigations and analysis of threats and actors in the cybercrime economy, and reflects our commitment to provide industry-best protections for our customers.

Edan Zwick, Amir Kutcher, Charles-Edouard Bettan, Yair Tsarfaty, Noam Hadash

Further reading

Learn how Microsoft Defender for Endpoint stops human-operated attacks.

For more information, read our documentation on the automatic attack disruption capability.

For the latest security research from the Microsoft Threat Intelligence community, check out the Microsoft Threat Intelligence Blog: https://aka.ms/threatintelblog.

To get notified about new publications and to join discussions on social media, follow us at https://twitter.com/MsftSecIntel.
