TLP:CLEAR

CYBERSECURITY ADVISORY

Authored by:

Product ID: AA23-059A

February 28, 2023

CISA Red Team Shares Key Findings to Improve Monitoring and Hardening of Networks

SUMMARY

In 2022, the Cybersecurity and Infrastructure Security Agency (CISA) conducted a red team assessment (RTA) at the request of a large critical infrastructure organization with multiple geographically separated sites. The team gained persistent access to the organization's network, moved laterally across the organization's multiple geographically separated sites, and eventually gained access to systems adjacent to the organization's sensitive business systems (SBSs). Multifactor authentication (MFA) prompts prevented the team from achieving access to one SBS, and the team was unable to complete its viable plan to compromise a second SBSs within the assessment period.

Actions to take today to harden your local environment:

- **Establish a security baseline** of normal network activity; tune network and host-based appliances to detect anomalous behavior.
- **Conduct regular assessments** to ensure appropriate procedures are created and can be followed by security staff and end users.
- **Enforce phishing-resistant MFA** to the greatest extent possible.

Despite having a mature cyber posture, the organization did not detect the red team's activity throughout the assessment, including when the team attempted to trigger a security response.

CISA is releasing this Cybersecurity Advisory (CSA) detailing the red team's tactics, techniques, and procedures (TTPs) and key findings to provide network defenders of critical infrastructure organizations proactive steps to reduce the threat of similar activity from malicious cyber actors. This CSA highlights the importance of collecting and monitoring logs for unusual activity as well as

---

All organizations should report incidents and anomalous activity to CISA's 24/7 Operations Center at report@cisa.gov. When available, please include the following information regarding the incident: date, time, and location of the incident; type of activity; number of people affected; type of equipment used for the activity; the name of the submitting company or organization; and a designated point of contact

continuous testing and exercises to ensure your organization's environment is not vulnerable to compromise, regardless of the maturity of its cyber posture.

CISA encourages critical infrastructure organizations to apply the recommendations in the Mitigations section of this CSA—including conduct regular testing within their security operations center—to ensure security processes and procedures are up to date, effective, and enable timely detection and mitigation of malicious activity.

TECHNICAL DETAILS

Note: This advisory uses the MITRE ATT&CK® for Enterprise framework, version 12. See the appendix for a table of the red team's activity mapped to MITRE ATT&CK tactics and techniques.

Introduction

CISA has authority to, upon request, provide analyses, expertise, and other technical assistance to critical infrastructure owners and operators and provide operational and timely technical assistance to Federal and non-Federal entities with respect to cybersecurity risks. (See generally 6 U.S.C. §§ 652[c][5], 659[c][6].) After receiving a request for a red team assessment (RTA) from an organization and coordinating some high-level details of the engagement with certain personnel at the organization, CISA conducted the RTA over a three-month period in 2022.

During RTAs, a CISA red team emulates cyber threat actors to assess an organization's cyber detection and response capabilities. During Phase I, the red team attempts to gain and maintain persistent access to an organization's enterprise network while avoiding detection and evading defenses. During Phase II, the red team attempts to trigger a security response from the organization's people, processes, or technology.

The "victim" for this assessment was a large organization with multiple geographically separated sites throughout the United States. For this assessment, the red team's goal during Phase I was to gain access to certain sensitive business systems (SBSs).

Phase I: Red Team Cyber Threat Activity

Overview

The organization's network was segmented with both logical and geographical boundaries. CISA's red team gained initial access to two organization workstations at separate sites via spearphishing emails. After gaining access and leveraging Active Directory (AD) data, the team gained persistent access to a third host via spearphishing emails. From that host, the team moved laterally to a misconfigured server, from which they compromised the domain controller (DC). They then used forged credentials to move to multiple hosts across different sites in the environment and eventually gained root access to all workstations connected to the organization's mobile device management (MDM) server. The team used this root access to move laterally to SBS-connected workstations. However, a multifactor authentication (MFA) prompt prevented the team from achieving access to one SBS, and Phase I ended before the team could implement a seemingly viable plan to achieve access to a second SBS.

Initial Access and Active Directory Discovery

The CISA red team gained initial access [TA0001] to two workstations at geographically separated sites (Site 1 and Site 2) via spearphishing emails. The team first conducted open-source research [TA0043] to identify potential targets for spearphishing. Specifically, the team looked for email addresses [T1589.002] as well as names [T1589.003] that could be used to derive email addresses based on the team's identification of the email naming scheme. The red team sent tailored spearphishing emails to seven targets using commercially available email platforms [T1585.002]. The team used the logging and tracking features of one of the platforms to analyze the organization's email filtering defenses and confirm the emails had reached the target's inbox.

The team built a rapport with some targeted individuals through emails, eventually leading these individuals to accept a virtual meeting invite. The meeting invite took them to a red team-controlled domain [T1566.002] with a button, which, when clicked, downloaded a "malicious" ISO file [T1204]. After the download, another button appeared, which, when clicked, executed the file.

Two of the seven targets responded to the phishing attempt, giving the red team access to a workstation at Site 1 (Workstation 1) and a workstation at Site 2. On Workstation 1, the team leveraged a modified SharpHound collector, `ldapsearch`, and command-line tool, `dsquery`, to query and scrape AD information, including AD users [T1087.002], computers [T1018], groups [T1069.002], access control lists (ACLs), organizational units (OU), and group policy objects (GPOs) [T1615].
Note: SharpHound is a BloodHound collector, an open-source AD reconnaissance tool. Bloodhound has multiple collectors that assist with information querying.

There were 52 hosts in the AD that had `Unconstrained Delegation` enabled and a `lastlogon` timestamp within 30 days of the query. Hosts with `Unconstrained Delegation` enabled store Kerberos ticket-granting tickets (TGTs) of all users that have authenticated to that host. Many of these hosts, including a Site 1 SharePoint server, were Windows Server 2012R2. The default configuration of Windows Server 2012R2 allows unprivileged users to query group membership of local administrator groups.

The red team queried parsed Bloodhound data for members of the SharePoint admin group and identified several standard user accounts with administrative access. The team initiated a second spearphishing campaign, similar to the first, to target these users. One user triggered the red team's payload, which led to installation of a persistent beacon on the user's workstation (Workstation 2), giving the team persistent access to Workstation 2.

Lateral Movement, Credential Access, and Persistence

The red team moved laterally [TA0008] from Workstation 2 to the Site 1 SharePoint server and had `SYSTEM` level access to the Site 1 SharePoint server, which had `Unconstrained Delegation` enabled. They used this access to obtain the cached credentials of all logged-in users—including the New Technology Local Area Network Manager (NTLM) hash for the SharePoint server account. To obtain the credentials, the team took a snapshot of `lsass.exe` [T1003.001] with a tool called nanodump, exported the output, and processed the output offline with Mimikatz.

The team then exploited `Unconstrained Delegation` to perform an NTLM-relay attack and steal the DC's TGT. Specifically, the team used the Sharepoint server's machine NTLM hash and DFSCoerce's python script (`DFSCoerce.py`) to prompt DC authentication to the server, and they

captured the incoming DC TGT using Rubeus [T1550.002], [T1557.001]. (DFSCoerce is used for NTLM relay attacks; it abuses Microsoft's Distributed File System [MS-DFSNM] protocol to relay authentication against an arbitrary server.[1])

The team then used the TGT to harvest advanced encryption standard (AES)-256 hashes via `DCSync` [T1003.006] for the `krbtgt` account and several privileged accounts—including domain admins, workstation admins, and a system center configuration management (SCCM) service account (SCCM Account 1). The team used the `krbtgt` account hash throughout the rest of their assessment to perform golden ticket attacks [T1558.001] in which they forged legitimate TGTs. The team also used the `asktgt` command to impersonate accounts they had credentials for by requesting account TGTs [T1550.003].

The team first impersonated the SCCM Account 1 and moved laterally to a Site 1 SCCM distribution point (DP) server (SCCM Server 1) that had direct network access to Workstation 2. The team then moved from SCCM Server 1 to a central SCCM server (SCCM Server 2) at a third site (Site 3). Specifically, the team:

1. Queried the AD using Lightweight Directory Access Protocol (LDAP) for information about the network's sites and subnets [T1016]. This query revealed all organization sites and subnets broken down by classless inter-domain routing (CIDR) subnet and description.
2. Used LDAP queries and domain name system (DNS) requests to identify recently active hosts.
3. Listed existing network connections [T1049] on SCCM Server 1, which revealed an active Server Message Block (SMB) connection from SCCM Server 2.
4. Attempted to move laterally to the SCCM Server 2 via `AppDomain` hijacking, but the HTTPS beacon failed to call back.
5. Attempted to move laterally with an SMB beacon [T1021.002], which was successful.

The team also moved from SCCM Server 1 to a Site 1 workstation (Workstation 3) that housed an active server administrator. The team impersonated an administrative service account via a golden ticket attack (from SCCM Server 1); the account had administrative privileges on Workstation 3. The user employed a KeePass password manager that the team was able to use to obtain passwords for other internal websites, a kernel-based virtual machine (KVM) server, virtual private network (VPN) endpoints, firewalls, and another KeePass database with credentials. The server administrator relied on a password manager, which stored credentials in a database file. The red team pulled the decryption key from memory using KeeThief and used it to unlock the database [T1555.005].

At the organization’s request, the red team confirmed that SCCM Server 2 provided access to the organization’s sites because firewall rules allowed SMB traffic to SCCM servers at all other sites.

The team moved laterally from SCCM Server 2 to an SCCM DP server at Site 5 and from the SCCM Server 1 to hosts at two other sites (Sites 4 and 6). The team installed persistent beacons at each of these sites. Site 5 was broken into a private and a public subnet and only DCs were able to cross that boundary. To move between the subnets, the team moved through DCs. Specifically, the team moved from the Site 5 SCCM DP server to a public DC; and then they moved from the public DC to the

private DC. The team was then able to move from the private DC to workstations in the private subnet.

The team leveraged access available from SCCM 2 to move around the organization's network for post-exploitation activities (See Post-Exploitation Activity section).

See Figure 1 for a timeline of the red team's initial access and lateral movement showing key access points.

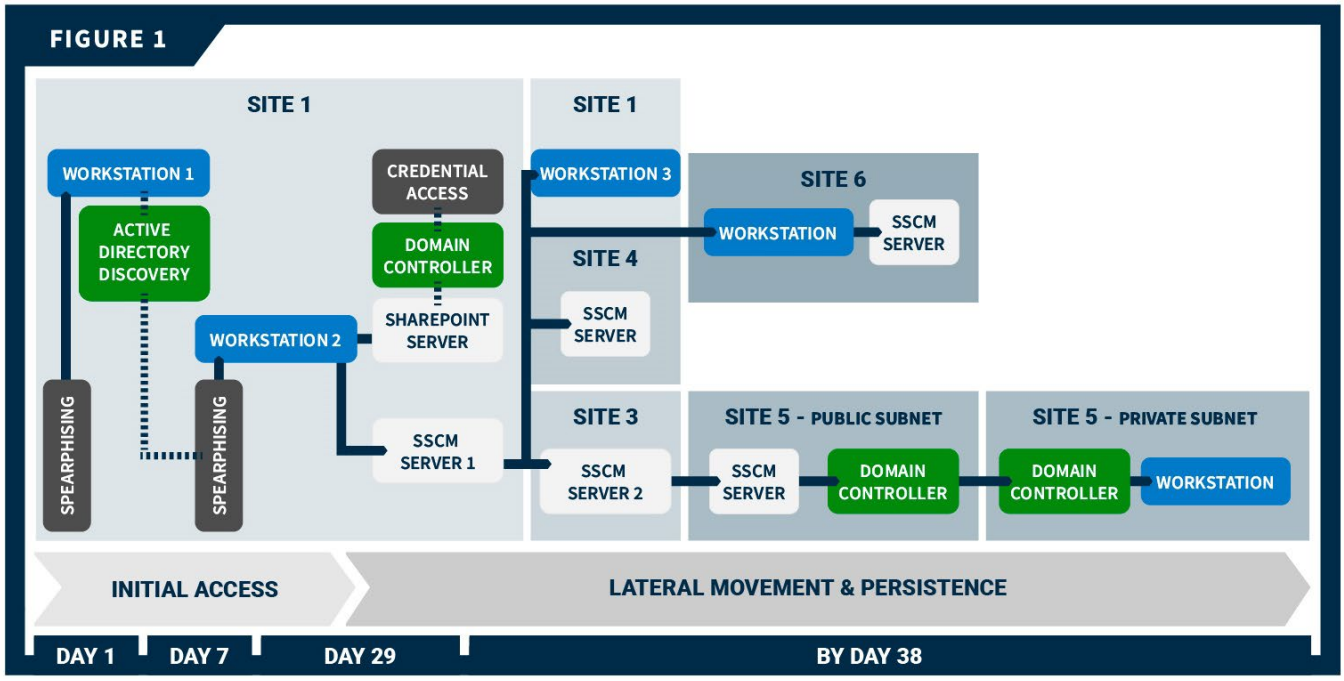

Figure 1: Red Team Cyber Threat Activity: Initial Access and Lateral Movement

While traversing the network, the team varied their lateral movement techniques to evade detection and because the organization had non-uniform firewalls between the sites and within the sites (within the sites, firewalls were configured by subnet). The team's primary methods to move between sites were `AppDomainManager` hijacking and dynamic-link library (DLL) hijacking [T1574.001]. In some instances, they used Windows Management Instrumentation (WMI) Event Subscriptions [T1546.003].

The team impersonated several accounts to evade detection while moving. When possible, the team remotely enumerated the local administrators group on target hosts to find a valid user account. This technique relies on anonymous SMB pipe binds [T1071], which are disabled by default starting with Windows Server 2016. In other cases, the team attempted to determine valid accounts based on group name and purpose. If the team had previously acquired the credentials, they used `asktgt` to impersonate the account. If the team did not have the credentials, they used the golden ticket attack to forge the account.

Post-Exploitation Activity: Gaining Access to SBSs

With persistent, deep access established across the organization's networks and subnetworks, the red team began post-exploitation activities and attempted to access SBSs. Trusted agents of the

organization tasked the team with gaining access to two specialized servers (SBS 1 and SBS 2). The team achieved root access to three SBS-adjacent workstations but was unable to move laterally to the SBS servers:

- Phase I ended before the team could implement a plan to move to SBS 1.
- An MFA prompt blocked the team from moving to SBS 2, and Phase I ended before they could implement potential workarounds.

However, the team assesses that by using Secure Shell (SSH) session socket files (see below), they could have accessed any hosts available to the users whose workstations were compromised.

Plan for Potential Access to SBS 1

Conducting open-source research [1591.001], the team identified that SBS 1 and 2 assets and associated management/upkeep staff were located at Sites 5 and 6, respectively. Adding previously collected AD data to this discovery, the team was able to identify a specific SBS 1 admin account. The team planned to use the organization's mobile device management (MDM) software to move laterally to the SBS 1 administrator's workstation and, from there, pivot to SBS 1 assets.

The team identified the organization's MDM vendor using open-source and AD information [T1590.006] and moved laterally to an MDM distribution point server at Site 5 (MDM DP 1). This server contained backups of the MDM MySQL database on its `D:` drive in the `Backup` directory. The backups included the encryption key needed to decrypt any encrypted values, such as SSH passwords [T1552]. The database backup identified both the user of the SBS 1 administrator account (USER 2) and the user's workstation (Workstation 4), which the MDM software remotely administered.

The team moved laterally to an MDM server (MDM 1) at Site 3, searched files on the server, and found plaintext credentials [T1552.001] to an application programming interface (API) user account stored in PowerShell scripts. The team attempted to leverage these credentials to browse to the web login page of the MDM vendor but were unable to do so because the website directed to an organization-controlled single-sign on (SSO) authentication page.

The team gained root access to workstations connected to MDM 1—specifically, the team accessed Workstation 4—by:

1. Selecting an MDM user from the plaintext credentials in PowerShell scripts on MDM 1.
2. While in the MDM MySQL database,
    a. Elevating the selected MDM user's account privileges to administrator privileges, and
    b. Modifying the user's account by adding `Create Policy` and `Delete Policy` permissions [T1098], [T1548].
3. Creating a policy via the MDM API [T1106], which instructed Workstation 4 to download and execute a payload to give the team interactive access as `root` to the workstation.
4. Verifying their interactive access.
5. Resetting permissions back to their original state by removing the policy via the MDM API and removing `Create Policy` and `Delete Policy` and administrator permissions and from the MDM user's account.

While interacting with Workstation 4, the team found an open SSH socket file and a corresponding `netstat` connection to a host that the team identified as a bastion host from architecture documentation found on Workstation 4. The team planned to move from Workstation 4 to the bastion host to SBS 1. **Note:** A SSH socket file allows a user to open multiple SSH sessions through a single, already authenticated SSH connection without additional authentication.

The team could not take advantage of the open SSH socket. Instead, they searched through SBS 1 architecture diagrams and documentation on Workstation 4. They found a security operations (SecOps) network diagram detailing the network boundaries between Site 5 SecOps on-premises systems, Site 5 non-SecOps on-premises systems, and Site 5 SecOps cloud infrastructure. The documentation listed the SecOps cloud infrastructure IP ranges [T1580]. These "trusted" IP addresses were a public `/16` subnet; the team was able to request a public IP in that range from the same cloud provider, and Workstation 4 made successful outbound SSH connections to this cloud infrastructure. The team intended to use that connection to reverse tunnel traffic back to the workstation and then access the bastion host via the open SSH socket file. However, Phase 1 ended before they were able to implement this plan.

Attempts to Access SBS 2

Conducting open-source research, the team identified an organizational branch [T1591] that likely had access to SBS 2. The team queried the AD to identify the branch's users and administrators. The team gathered a list of potential accounts, from which they identified administrators, such as `SYSTEMS ADMIN` or `DATA SYSTEMS ADMINISTRATOR`, with technical roles. Using their access to the MDM MySQL database, the team queried potential targets to (1) determine the target's last contact time with the MDM and (2) ensure any policy targeting the target's workstation would run relatively quickly [T1596.005]. Using the same methodology as described by the steps in the Plan for Potential Access to SBS 1 section above, the team gained interactive root access to two Site 6 SBS 2-connected workstations: a software engineering workstation (Workstation 5) and a user administrator workstation (Workstation 6).

The Workstation 5 user had bash history files with what appeared to be SSH passwords mistyped into the bash prompt and saved in bash history [T1552.003]. The team then attempted to authenticate to SBS 2 using a similar tunnel setup as described in the Access to SBS 1 section above and the potential credentials from the user's bash history file. However, this attempt was unsuccessful for unknown reasons.

On Workstation 6, the team found a `.txt` file containing plaintext credentials for the user. Using the pattern discovered in these credentials, the team was able to crack the user's workstation account password [T1110.002]. The team also discovered potential passwords and SSH connection commands in the user's bash history. Using a similar tunnel setup described above, the team attempted to log into SBS 2. However, a prompt for an MFA passcode blocked this attempt.

See figure 2 for a timeline of the team's post exploitation activity that includes key points of access.

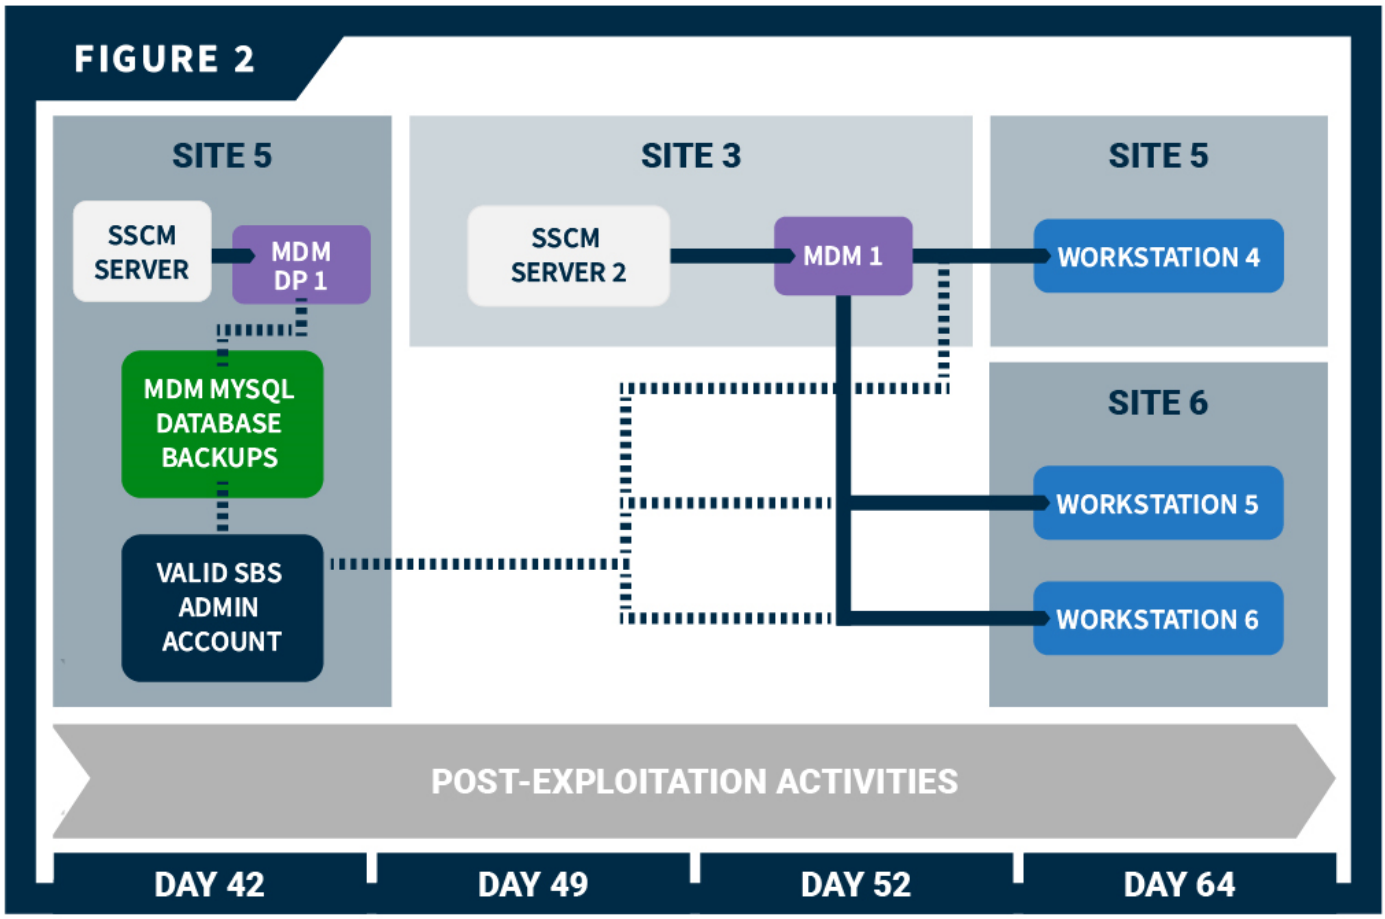

Figure 2: Red Team Cyber Threat Activity: Post Exploitation

Command and Control

The team used third-party owned and operated infrastructure and services [T1583] throughout their assessment, including in certain cases for command and control (C2) [TA0011]. These included:

- Cobalt Strike and Merlin payloads for C2 throughout the assessment. **Note:** Merlin is a post-exploit tool that leverages HTTP protocols for C2 traffic.
    - The team maintained multiple Cobalt Strike servers hosted by a cloud vendor. They configured each server with a different domain and used the servers for communication with compromised hosts. These servers retained all assessment data.
- Two commercially available cloud-computing platforms.
    - The team used these platforms to create flexible and dynamic redirect servers to send traffic to the team's Cobalt Strike servers [T1090.002]. Redirecting servers make it difficult for defenders to attribute assessment activities to the backend team servers. The redirectors used HTTPS reverse proxies to redirect C2 traffic between the target organization's network and the Cobalt Strike team servers [T1071.002]. The team encrypted all data in transit [T1573] using encryption keys stored on team's Cobalt Strike servers.

- A cloud service to rapidly change the IP address of the team's redirecting servers in the event of detection and eradication.
- Content delivery network (CDN) services to further obfuscate some of the team's C2 traffic.
    - This technique leverages CDNs associated with high-reputation domains so that the malicious traffic appears to be directed towards a reputation domain but is actually redirected to the red team-controlled Cobalt Strike servers.
    - The team used domain fronting [T1090.004] to disguise outbound traffic in order to diversify the domains with which the persistent beacons were communicating. This technique, which also leverages CDNs, allows the beacon to appear to connect to third-party domains, such as nytimes.com, when it is actually connecting to the team's redirect server.

Phase II: Red Team Measurable Events Activity

The red team executed 13 measurable events designed to provoke a response from the people, processes, and technology defending the organization's network. See Table 1 for a description of the events, the expected network defender activity, and the organization's actual response.

Table 1: Measurable Events

Measurable Event	Description	MITRE ATT&CK Technique(s)	Expected Detection Points	Expected Network Defender Reactions	Reported Reactions
Internal Port Scan	Launch scan from inside the network from a previously gained workstation to enumerate ports on target workstation, server, and domain controller system(s).	• Network Service Discovery [T1046]	• Network Monitoring and Analysis Tools • Intrusion Detection or Prevention Systems • Endpoint Protection Platform	• Detect target hosts and ports • Identify associated scanning process • Analyze scanning host once detected • Develop response plan	None
Comprehensive Active Directory and Host Enumeration	Perform AD enumeration by querying all domain objects from the DC; and enumerating trust relationships within the AD Forest, user accounts, and current session information from every domain computer (Workstation and Server).	• Domain Trust Discovery [T1482] • Account Discovery: Dom ain Account [T1087.002] • System Owner/User Discovery [T1033] • Remote System Discovery [T1018]	• Network Monitoring and Analysis Tools • Intrusion Detection or Prevention Systems • Endpoint Protection Platform	• Detect target hosts and ports • Identify associated scanning process • Analyze scanning host once detected • Develop response plan	Collection process stopped before completion. Host isolated and sent for forensics.
Data Exfiltration—1 GB of Data	Send a large amount (1 GB) of mock sensitive information to an external system over various protocols, including ICMP,	• Exfiltration Over Alternative Protocol [T1048]	• Network Monitoring and Analysis Tools • Intrusion Detection or Prevention Systems • Endpoint Protection Platform	• Detect target hosts and ports • Identify associated scanning process • Analyze scanning host once detected	None

Measurable Event	Description	MITRE ATT&CK Technique(s)	Expected Detection Points	Expected Network Defender Reactions	Reported Reactions
	DNS, FTP, and/or HTTP/S.			• Develop response plan	
Malicious Traffic Generation—Workstation to External Host	Establish a session that originates from a target Workstation system directly to an external host over a clear text protocol, such as HTTP.	• Application Layer Protocol [T1071]	• Intrusion Detection or Prevention Systems • Endpoint Protection Platform • Windows Event Logs	• Detect and Identify source IP and source process of enumeration • Analyze scanning host once detected • Develop response plan	None
Active Directory Account Lockout	Lock out several administrative AD accounts	• Account Access Removal [T1531]	• Windows Event Logs • End User Reporting	• Detect and Identify source IP and source process of exfiltration • Analyze host used for exfiltration once detected Develop response plan	None
Local Admin User Account Creation (workstation)	Create a local administrator account on a target workstation system.	• Create Account: Local Account [T1136.001] • Account Manipulation [T1098]	• Intrusion Detection or Prevention Systems • Endpoint Protection Platform • Web Proxy Logs	• Detect and identify source IP and source process of malicious traffic • Investigate destination IP address • Triage compromised host • Develop response plan	None
Local Admin User Account Creation (server)	Create a local administrator account on a target server system.	• Create Account: Local Account [T1136.001] • Account Manipulation [T1098]	• Windows Event Logs	• Detect account creation • Identify source of change • Verify change with system owner • Develop response plan	None

Measurable Event	Description	MITRE ATT&CK Technique(s)	Expected Detection Points	Expected Network Defender Reactions	Reported Reactions
Active Directory Account Creation	Create AD accounts and add it to domain admins group	• Create Account: Domain Account [T1136.002] • Account Manipulation [T1098]	• Windows Event Logs	• Detect account creation • Identify source of change • Verify change with system owner • Develop response plan	None
Workstation Admin Lateral Movement—Workstation to Workstation	Use a previously compromised workstation admin account to upload and execute a payload via SMB and Windows Service Creation, respectively, on several target Workstations.	• Valid Accounts: Domain Accounts [T1078.002] • Remote Services: SMB/Windows Admin Shares, Sub-technique [T1021.002] • Create or Modify System Process: Windows Service [T1543.003]	• Windows Event Logs	• Detect account compromise • Analyze compromised host • Develop response plan	None
Domain Admin Lateral Movement—Workstation to Domain Controller	Use a previously compromised domain admin account to upload and execute a payload via SMB and Windows Service Creation, respectively, on a target DC.	• Valid Accounts: Domain Accounts [T1078.002] • Remote Services: SMB/Windows Admin Shares,	• Windows Event Logs	• Detect account compromise • Triage compromised host • Develop response plan	None

Measurable Event	Description	MITRE ATT&CK Technique(s)	Expected Detection Points	Expected Network Defender Reactions	Reported Reactions
		Sub-technique [T1021.002] • Create or Modify System Process: Windows Service [T1543.003]			
Malicious Traffic Generation—Domain Controller to External Host	Establish a session that originates from a target Domain Controller system directly to an external host over a clear text protocol, such as HTTP.	• Application Layer Protocol [T1071]	• Intrusion Detection or Prevention Systems • Endpoint Protection Platform • Web Proxy Logs	• Detect and identify source IP and source process of malicious traffic • Investigate destination IP address • Triage compromised host Develop response plan	None
Trigger Host-Based Protection—Domain Controller	Upload and execute a well-known (e.g., with a signature) malicious file to a target DC system to generate host-based alerts.	• Ingress Tool Transfer [T1105]	• Endpoint Protection Platform • Endpoint Detection and Response	• Detect and identify source IP and source process of malicious traffic • Investigate destination IP address • Triage compromised host • Develop response plan	Malicious file was removed by antivirus
Ransomware Simulation	Execute simulated ransomware on multiple Workstation systems to simulate a ransomware attack. Note: This technique does NOT encrypt files on the target system.	N/A	• End User Reporting	• Investigate end user reported event • Triage compromised host • Develop response Plan	Four users reported event to defensive staff

Findings

Key Issues

The red team noted the following key issues relevant to the security of the organization's network. These findings contributed to the team's ability to gain persistent, undetected access across the organization's sites. See the Mitigations section for recommendations on how to mitigate these issues.

- **Insufficient host and network monitoring.** Most of the red team's Phase II actions failed to provoke a response from the people, processes, and technology defending the organization's network. The organization failed to detect lateral movement, persistence, and C2 activity via their intrusion detection or prevention systems, endpoint protection platform, web proxy logs, and Windows event logs. Additionally, throughout Phase I, the team received no deconflictions or confirmation that the organization caught their activity. Below is a list of some of the higher risk activities conducted by the team that were opportunities for detection:
  - Phishing
  - Lateral movement reuse
  - Generation and use of the golden ticket
  - Anomalous LDAP traffic
  - Anomalous internal share enumeration
  - `Unconstrained Delegation` server compromise
  - `DCSync`
  - Anomalous account usage during lateral movement
  - Anomalous outbound network traffic
  - Anomalous outbound SSH connections to the team's cloud servers from workstations
- **Lack of monitoring on endpoint management systems.** The team used the organization's MDM system to gain root access to machines across the organization's network without being detected. Endpoint management systems provide elevated access to thousands of hosts and should be treated as high value assets (HVAs) with additional restrictions and monitoring.
- **KRBTGT never changed.** The Site 1 `krbtgt` account password had not been updated for over a decade. The `krbtgt` account is a domain default account that acts as a service account for the key distribution center (KDC) service used to encrypt and sign all Kerberos tickets for the domain. Compromise of the `krbtgt` account could provide adversaries with the ability to sign their own TGTs, facilitating domain access years after the date of compromise. The red team was able to use the `krbtgt` account to forge TGTs for multiple accounts throughout Phase I.
- **Excessive permissions to standard users.** The team discovered several standard user accounts that have local administrator access to critical servers. This misconfiguration allowed the team to use the low-level access of a phished user to move laterally to an `Unconstrained Delegation` host and compromise the entire domain.
- **Hosts with `Unconstrained Delegation` enabled unnecessarily.** Hosts with `Unconstrained Delegation` enabled store the Kerberos TGTs of all users that authenticate to that host, enabling actors to steal service tickets or compromise `krbtgt` accounts and

perform golden ticket or "silver ticket" attacks. The team performed an NTLM-relay attack to obtain the DC's TGT, followed by a golden ticket attack on a SharePoint server with `Unconstrained Delegation` to gain the ability to impersonate any Site 1 AD account.

- **Use of non-secure default configurations.** The organization used default configurations for hosts with Windows Server 2012 R2. The default configuration allows unprivileged users to query group membership of local administrator groups. The red team used and identified several standard user accounts with administrative access from a Windows Server 2012 R2 SharePoint server.

Additional Issues

The team noted the following additional issues.

- **Ineffective separation of privileged accounts**. Some workstations allowed unprivileged accounts to have local administrator access; for example, the red team discovered an ordinary user account in the local admin group for the SharePoint server. If a user with administrative access is compromised, an actor can access servers without needing to elevate privileges. Administrative and user accounts should be separated, and designated admin accounts should be exclusively used for admin purposes.
- **Lack of server egress control.** Most servers, including domain controllers, allowed unrestricted egress traffic to the internet.
- **Inconsistent host configuration**. The team observed inconsistencies on servers and workstations within the domain, including inconsistent membership in the local administrator group among different servers or workstations. For example, some workstations had "Server Admins" or "Domain Admins" as local administrators, and other workstations had neither.
- **Potentially unwanted programs**. The team noticed potentially unusual software, including music software, installed on both workstations and servers. These extraneous software installations indicate inconsistent host configuration (see above) and increase the attack surfaces for malicious actors to gain initial access or escalate privileges once in the network.
- **Mandatory password changes enabled**. During the assessment, the team keylogged a user during a mandatory password change and noticed that only the final character of their password was modified. This is potentially due to domain passwords being required to be changed every 60 days.
- **Smart card use was inconsistent across the domain**. While the technology was deployed, it was not applied uniformly, and there was a significant portion of users without smartcard protections enabled. The team used these unprotected accounts throughout their assessment to move laterally through the domain and gain persistence.

Noted Strengths

The red team noted the following technical controls or defensive measures that prevented or hampered offensive actions:

- **The organization conducts regular, proactive penetration tests and adversarial assessments** and invests in hardening their network based on findings.

  - The team was unable to discover any easily exploitable services, ports, or web interfaces from more than three million external in-scope IPs. This forced the team to resort to phishing to gain initial access to the environment.
  - Service account passwords were strong. The team was unable to crack any of the hashes obtained from the 610 service accounts pulled. This is a critical strength because it slowed the team from moving around the network in the initial parts of the Phase I.
  - The team did not discover any useful credentials on open file shares or file servers. This slowed the progress of the team from moving around the network.
- **MFA was used for some SBSs**. The team was blocked from moving to SBS 2 by an MFA prompt.
- **There were strong security controls and segmentation for SBS systems.** Direct access to SBS were located in separate networks, and admins of SBS used workstations protected by local firewalls.

MITIGATIONS

CISA recommends organizations implement the recommendations in Table 2 to mitigate the issues listed in the Findings section of this advisory. These mitigations align with the Cross-Sector Cybersecurity Performance Goals (CPGs) developed by CISA and the National Institute of Standards and Technology (NIST). The CPGs provide a minimum set of practices and protections that CISA and NIST recommend all organizations implement. CISA and NIST based the CPGs on existing cybersecurity frameworks and guidance to protect against the most common and impactful threats, tactics, techniques, and procedures. See CISA's Cross-Sector Cybersecurity Performance Goals for more information on the CPGs, including additional recommended baseline protections.

Table 2: Recommendations to Mitigate Identified Issues

Issue	Recommendation

Insufficient host and network monitoring	• **Establish a security baseline of normal network traffic and tune network appliances** to detect anomalous behavior [CPG 3.1]. Tune host-based products to detect anomalous binaries, lateral movement, and persistence techniques.
	  ○ Create alerts for Windows event log authentication codes, especially for the domain controllers. This could help detect some of the pass-the-ticket, DCSync, and other techniques described in this report.
	  ○ From a detection standpoint, focus on identity and access management (IAM) rather than just network traffic or static host alerts.

Issue	Recommendation
	▪ Consider who is accessing what (what resource), from where (what internal host or external location), and when (what day and time the access occurs). ▪ Look for access behavior that deviates from expected or is indicative of AD abuse. • **Reduce the attack surface by limiting the use of legitimate administrative pathways and tools** such as PowerShell, PSExec, and WMI, which are often used by malicious actors. CISA recommends selecting one tool to administer the network, ensuring logging is turned on [CPG 3.1], and disabling the others. • **Consider using "honeypot" service principal names (SPNs)** to detect attempts to crack account hashes [CPG 1.1]. • **Conduct regular assessments** to ensure processes and procedures are up to date and can be followed by security staff and end users. o Consider using red team tools, such as SharpHound, for AD enumeration to identify users with excessive privileges and misconfigured hosts (e.g., with `Unconstrained Delegation` enabled). • **Ensure all commercial tools deployed in your environment are regularly tuned** to pick up on relevant activity in your environment.
Lack of monitoring on endpoint management systems	• **Treat endpoint management systems as HVAs** with additional restrictions and monitoring because they provide elevated access to thousands of hosts.
KRBTGT never changed	• **Change the `krbtgt` account password on a regular schedule** such as every 6 to 12 months or if it becomes compromised. Note that this password change must be carefully performed to effectively change the credential without breaking AD functionality. The password must be changed twice to effectively invalidate the old credentials. However, the required waiting period between resets must be greater than the maximum lifetime period of Kerberos tickets, which is 10 hours by default. See Microsoft's KRBTGT account maintenance considerations guidance for more information.

Issue	Recommendation

Excessive permissions to standard users and ineffective separation of privileged accounts	• Implement the principle of least privilege:
	  ○ **Grant standard user rights for standard user tasks** such as email, web browsing, and using line-of-business (LOB) applications.
	  ○ **Periodically audit standard accounts** and minimize where they have privileged access.
	  ○ **Periodically Audit AD permissions** to ensure users do not have excessive permissions and have not been added to admin groups.
	  ○ **Evaluate which administrative groups should administer which servers/workstations**. Ensure group members administrative accounts instead of standard accounts.
	  ○ **Separate administrator accounts from user accounts** [CPG 1.5]. Only allow designated admin accounts to be used for admin purposes. If an individual user needs administrative rights over their workstation, use a separate account that does not have administrative access to other hosts, such as servers.
	• **Consider using a privileged access management (PAM) solution** to manage access to privileged accounts and resources [CPG 3.4]. PAM solutions can also log and alert usage to detect any unusual activity and may have helped stop the red team from accessing resources with admin accounts. **Note:** password vaults associated with PAM solutions should be treated as HVAs with additional restrictions and monitoring (see below).
	• **Configure time-based access for accounts set at the admin level and higher**. For example, the just-in-time (JIT) access method provisions privileged access when needed and can support enforcement of the principle of least privilege, as well as the Zero Trust model. This is a process in which a network-wide policy is set in place to automatically disable administrator accounts at the AD level when the account is not in direct need. When individual users need the account, they submit their requests through an automated process that enables access to a system but only for a set timeframe to support task completion.

Issue	Recommendation
Hosts with `Unconstrained Delegation` enabled	• **Remove `Unconstrained Delegation` from all servers**. If `Unconstrained Delegation` functionality is required, upgrade operating systems and applications to leverage other approaches (e.g., `constrained delegation`) or explore whether systems can be retired or further isolated from the enterprise. CISA recommends Windows Server 2019 or greater. • **Consider disabling or limiting NTLM and WDigest Authentication** if possible, including using their use as criteria for prioritizing updates to legacy systems or for segmenting the network. Instead use more modern federation protocols (SAML, OIDC) or Kerberos for authentication with AES-256 bit encryption [CPG 3.4]. • If NTLM must be enabled, **enable Extended Protection for Authentication (EPA)** to prevent some NTLM-relay **attacks, and implement SMB signing** to prevent certain adversary-in-the-middle and pass-the-hash attacks CPG 3.4]. See Microsoft Mitigating NTLM Relay Attacks on Active Directory Certificate Services (AD CS) and Microsoft Overview of Server Message Block signing for more information.
Use of non-secure default configurations	• **Keep systems and software up to date** [CPG 5.1]. If updates cannot be uniformly installed, update insecure configurations to meet updated standards.
Lack of server egress control	• **Configure internal firewalls and proxies to restrict internet traffic from hosts that do not require it**. If a host requires specific outbound traffic, consider creating an allowlist policy of domains.

Large number of credentials in a shared vault	• **Treat password vaults as HVAs** with additional restrictions and monitoring [CPG 3.4]:
	  ○ If on-premise, require MFA for admin and apply network segmentation [CPG 1.3]. Use solutions with end-to-end encryption where applicable [CPG 3.3].
	  ○ If cloud-based, evaluate the provider to ensure use of strong security controls such as MFA and end-to-end encryption [CPG 1.3, 3.3].

Issue	Recommendation
Inconsistent host configuration	• **Establish a baseline/gold-image for workstations and servers** and deploy from that image [CPG 2.5]. Use standardized groups to administer hosts in the network.
Potentially unwanted programs	• **Implement software allowlisting** to ensure users can only install software from an approved list [CPG 2.1]. • Remove unnecessary, extraneous software from servers and workstations.
Mandatory password changes enabled	• **Consider only requiring changes for memorized passwords in the event of compromise**. Regular changing of memorized passwords can lead to predictable patterns, and both CISA and the National Institute of Standards and Technology (NIST) recommend against changing passwords on regular intervals.

Additionally, CISA recommends organizations implement the mitigations below to improve their cybersecurity posture:

- **Provide users with regular training and exercises**, specifically related to phishing emails [CPG 4.3]. Phishing accounts for majority of initial access intrusion events.
- **Enforce phishing-resistant MFA** to the greatest extent possible [CPG 1.3].
- Reduce the risk of credential compromise via the following:
    - **Place domain admin accounts in the protected users group** to prevent caching of password hashes locally; this also forces Kerberos AES authentication as opposed to weaker RC4 or NTLM.
    - **Implement Credential Guard for Windows 10 and Server 2016** (Refer to Microsoft: Manage Windows Defender Credential Guard for more information). For Windows Server 2012R2, enable Protected Process Light for Local Security Authority (LSA).
    - **Refrain from storing plaintext credentials in scripts** [CPG 3.4]. The red team discovered a PowerShell script containing plaintext credentials that allowed them to escalate to admin.
- **Upgrade to Windows Server 2019 or greater and Windows 10 or greater**. These versions have security features not included in older operating systems.

As a long-term effort, **CISA recommends organizations prioritize implementing a more modern, Zero Trust network architecture** that:

- Leverages secure cloud services for key enterprise security capabilities (e.g., identity and access management, endpoint detection and response, policy enforcement).
- Upgrades applications and infrastructure to leverage modern identity management and network access practices.

- Centralizes and streamlines access to cybersecurity data to drive analytics for identifying and managing cybersecurity risks.
- Invests in technology and personnel to achieve these goals.

CISA encourages organizational IT leadership to ask their executive leadership the question: Can the organization accept the business risk of NOT implementing critical security controls such as MFA? Risks of that nature should typically be acknowledged and prioritized at the most senior levels of an organization.

VALIDATE SECURITY CONTROLS

In addition to applying mitigations, CISA recommends exercising, testing, and validating your organization's security program against the threat behaviors mapped to the MITRE ATT&CK for Enterprise framework in this advisory. CISA recommends testing your existing security controls inventory to assess how they perform against the ATT&CK techniques described in this advisory.

To get started:

1. Select an ATT&CK technique described in this advisory (see Table 3).
2. Align your security technologies against the technique.
3. Test your technologies against the technique.
4. Analyze your detection and prevention technologies' performance.
5. Repeat the process for all security technologies to obtain a set of comprehensive performance data.
6. Tune your security program, including people, processes, and technologies, based on the data generated by this process.

CISA recommends continually testing your security program, at scale, in a production environment to ensure optimal performance against the MITRE ATT&CK techniques identified in this advisory.

RESOURCES

See CISA's RedEye tool on CISA's GitHub page. RedEye is an interactive open-source analytic tool used to visualize and report red team command and control activities. See CISA's RedEye tool overview video for more information.

REFERENCES

[1] Bleeping Computer: New DFSCoerce NTLM Relay attack allows Windows domain takeover

APPENDIX: MITRE ATT&CK TACTICS AND TECHNIQUES

See Table 3 for all referenced red team tactics and techniques in this advisory. **Note:** activity was from Phase I unless noted.

Table 3: Red Team ATT&CK Techniques for Enterprise

Reconnaissance		
Technique Title	**ID**	**Use**
Gather Victim Identity Information: Email Addresses	T1589.002	The team found employee email addresses via open-source research.
Gather Victim Identify Information: Employee Names	T1589.003	The team identified employee names via open-source research that could be used to derive email addresses.
Gather Victim Network Information: Network Security Appliances	T1590.006	The team identified the organization's MDM vendor and leveraged that information to move laterally to SBS-connected assets.
Gather Victim Org Information	T1591	The team conducted open-source research and identified an organizational branch that likely had access to an SBS asset.
Gather Victim Org Information: Determine Physical Locations	T1591.001	The team conducted open-source research to identify the physical locations of upkeep/management staff of selected assets.
Search Open Technical Databases: Scan Databases	T1596.005	The team queried an MDM SQL database to identify target administrators who recently connected with the MDM.

Resource Development		
Technique Title	**ID**	**Use**
Acquire Infrastructure	T1583	The team used third-party owned and operated infrastructure throughout their assessment for C2.
Establish Accounts: Email Accounts	T1585.002	The team used commercially available email platforms for their spearphishing activity.
Obtain Capabilities: Tool	T1588.002	The team used the following tools: • Cobalt Strike and Merlin payloads for C2. • KeeThief to obtain a decryption key from a KeePass database • Rubeus and DFSCoerce in an NTLM relay attack
Initial Access		
Technique Title	**ID**	**Use**
Phishing: Spearphishing Link	T1566.002	The team sent spearphishing emails with links to a red-team-controlled domain to gain access to the organization's systems.
Execution		
Technique Title	**ID**	**Use**
Native API	T1106	The team created a policy via the MDM API, which downloaded and executed a payload on a workstation.
User Execution	T1204	Users downloaded and executed the team's initial access payloads after clicking buttons to trigger download and execution.

Persistence		
Technique Title	**ID**	**Use**
Account Manipulation	T1098	The team elevated account privileges to administrator and modified the user's account by adding `Create Policy` and `Delete Policy` permissions. During Phase II, the team created local admin accounts and an AD account; they added the created AD account to a domain admins group.
Create Account: Local Account	T1136.001	During Phase II, the team created a local administrator account on a workstation and a server.
Create Account: Domain Account	T1136.002	During Phase II, the team created an AD account.
Create or Modify System Process: Windows Service	T1543.003	During Phase II, the team leveraged compromised workstation and domain admin accounts to execute a payload via Windows Service Creation on target workstations and the DC.
Event Triggered Execution: Windows Management Instrumentation Event Subscription	T1546.003	The team used WMI Event Subscriptions to move laterally between sites.
Hijack Execution Flow: DLL Search Order Hijacking	T1574.001	The team used DLL hijacking to move laterally between sites.
Privilege Escalation		
Technique Title	**ID**	**Use**
Abuse Elevation Control Mechanism	T1548	The team elevated user account privileges to administrator by modifying the user's account via adding `Create Policy` and `Delete Policy` permissions.

Defense Evasion		
Technique Title	**ID**	**Use**
Valid Accounts: Domain Accounts	T1078.002	During Phase II, the team compromised a domain admin account and used it to laterally to multiple workstations and the DC.
Credential Access		
Technique Title	**ID**	**Use**
OS Credential Dumping: LSASS Memory	T1003.001	The team obtained the cached credentials from a SharePoint server account by taking a snapshot of `lsass.exe` with a tool called nanodump, exporting the output and processing the output offline with Mimikatz.
OS Credential Dumping: DCSync	T1003.006	The team harvested AES-256 hashes via `DCSync`.
Brute Force: Password Cracking	T1110.002	The team cracked a user's workstation account password after learning the user's patterns from plaintext credentials.
Unsecured Credentials	T1552	The team found backups of a MySQL database that contained the encryption key needed to decrypt SSH passwords.
Unsecured Credentials: Credentials in Files	T1552.001	The team found plaintext credentials to an API user account stored in PowerShell scripts on an MDM server.
Unsecured Credentials: Bash History	T1552.003	The team found bash history files on a Workstation 5, and the files appeared to be SSH passwords saved in bash history.
Credentials from Password Stores: Password Managers	T1555.005	The team pulled credentials from a KeePass database.
Adversary-in-the-middle: LLMNR/NBT-NS Poisoning and SMB Relay	T1557.001	The team leveraged Rubeus and DFSCoerce in a NTLM relay attack to obtain the DC's TGT from a host with `Unconstrained Delegation` enabled.

Steal or Forge Kerberos Tickets: Golden Ticket	T1558.001	The team used the acquired `krbtgt` account hash throughout their assessment to forge legitimate TGTs.
Steal or Forge Kerberos Tickets: Kerberoasting	T1558.003	The team leveraged Rubeus and DFSCoerce in a NTLM relay attack to obtain the DC's TGT from a host with `Unconstrained Delegation` enabled.
Discovery		
Technique Title	**ID**	**Use**
System Network Configuration Discovery	T1016	The team queried the AD for information about the network's sites and subnets.
Remote System Discovery	T1018	The team queried the AD, during phase I and II, for information about computers on the network.
System Network Connections Discovery	T1049	The team listed existing network connections on SCCM Server 1 to reveal an active SMB connection with server 2.
Permission Groups Discovery: Domain Groups	T1069.002	The team leveraged `ldapsearch` and `dsquery` to query and scrape active directory information.
Account Discovery: Domain Account	T1087.002	The team queried AD for AD users (during Phase I and II), including for members of a SharePoint admin group and several standard user accounts with administrative access.
Cloud Infrastructure Discovery	T1580	The team found SecOps network diagrams on a host detailing cloud infrastructure boundaries.
Domain Trust Discovery	T1482	During Phase II, the team enumerated trust relationships within the AD Forest.
Group Policy Discovery	T1615	The team scraped AD information, including GPOs.
Network Service Discovery	T1046	During Phase II, the team enumerated ports on target systems from a previously compromised workstation.

System Owner/User Discovery	T1033	During Phase II, the team enumerated the AD for current session information from every domain computer (Workstation and Server).
Lateral Movement		
Technique Title	**ID**	**Use**
Remote Services: SMB/Windows Admin Shares	T1021.002	The team moved laterally with an SMB beacon. During Phase II, they used compromised workstation and domain admin accounts to upload a payload via SMB on several target Workstations and the DC.
Use Alternate Authentication Material: Pass the Hash	T1550.002	As part of a NTLM relay attack, the team used a server's NTLM hash and `DFSCoerce.py` to prompt DC authentication to the server, and they captured the incoming DC TGT using Rubeus.
Pass the Ticket	T1550.003	The team used the `asktgt` command to impersonate accounts for which they had credentials by requesting account TGTs.
Command and Control		
Technique Title	**ID**	**Use**
Application Layer Protocol	T1071	The team remotely enumerated the local administrators group on target hosts to find valid user accounts. This technique relies on anonymous SMB pipe binds, which are disabled by default starting with Server 2016. During Phase II, the team established sessions that originated from a target Workstation and from the DC directly to an external host over a clear text protocol.
Application Layer Protocol: Web Protocols	T1071.001	The team's C2 redirectors used HTTPS reverse proxies to redirect C2 traffic.
Application Layer Protocol: File Transfer Protocols	T1071.002	The team used HTTPS reverse proxies to redirect C2 traffic between target network and the team's Cobalt Strike servers.

Encrypted Channel	T1573	The team's C2 traffic was encrypted in transit using encryption keys stored on their C2 servers.
Ingress Tool Transfer	T1105	During Phase II, the team uploaded and executed well-known malicious files to the DC to generate host-based alerts.
Proxy: External Proxy	T1090.002	The team used redirectors to redirect C2 traffic between the target organization's network and the team's C2 servers.
Proxy: Domain Fronting	T1090.004	The team used domain fronting to disguise outbound traffic in order to diversify the domains with which the persistent beacons were communicating.
Impact		
Technique Title	**ID**	**Use**
Account Access Removal	T1531	During Phase II, the team locked out several administrative AD accounts.