Scattered Spider x RansomHub: A New Partnership

By ReliaQuest Threat Research Team 24 October 2024

Published: 2024-10-24 · Archived: 2026-04-05 22:07:31 UTC

Editor's note: James Xiang *and* Hayden Evans *contributed to this blog.*

Key Points

- In October 2024, ReliaQuest responded to an intrusion affecting a manufacturing sector customer. We identified "Scattered Spider" to be behind the incident. This English-speaking collective previously served as an affiliate for ransomware group "ALPHV" and now partners with "RansomHub."
- The attacker gained initial access to two employee accounts by carrying out social engineering attacks on the organization's help desk twice. Within six hours, the attacker began encrypting the organization's systems.
- To maintain persistence, Scattered Spider leveraged the organization's ESXi environment to create a virtual machine (VM). This concealed their attack until the environment was encrypted and backups were sabotaged.
- Implementing comprehensive measures to mitigate social engineering techniques, such as restricting SharePoint permissions and hardening ESXi environments, can reduce the attack surface and decrease the likelihood of threat actors achieving their objectives.

What Happened?

In October 2024, ReliaQuest investigated an intrusion for a customer in the manufacturing sector. We attributed the incident with high confidence to "Scattered Spider," an English-speaking collective acting as an affiliate for the ransomware group "RansomHub."

Scattered Spider previously targeted telecommunications firms, likely to support its SIM-swapping activities that facilitate account takeovers. Lately, it's shifted focus to extorting large organizations by collaborating with ransomware groups, aiming for higher financial returns.

Our investigation uncovered Scattered Spider's tactics, techniques, and procedures (TTPs), including a unique method of gaining initial access to organizations. Leveraging its English proficiency, the collective uses social engineering for initial access.

In this incident, the attacker convinced the organization's help desk to reset the Chief Financial Officer's (CFO) account credentials. After discovering that the CFO's account lacked the permissions required for further pivoting, the attacker repeated the social engineering tactic to compromise a domain administrator account. With this

privileged access, they created a virtual machine (VM) within the ESXi environment, evading security tools like endpoint detection and response (EDR). They then deployed a RansomHub encryptor to impact a critical ESXi environment in just **six hours**.

In this report, we explore Scattered Spider's evolution from low-level cybercrimes to partnering with ransomware groups to target major organizations. We'll break down the TTPs observed in the incident and offer practical advice to help organizations understand, investigate, and mitigate similar threats.

Scattered Spider Teams Up with RansomHub

Active since at least May 2022, Scattered Spider (aka "UNC3944," "Octo Tempest") is a collective of at least one thousand English-speaking threat actors linked to the cybercriminal network known as The Community or The Com. Operating across forums and Telegram groups, The Com engages in attacks that require social engineering such as SIM swapping, swatting, carding, and identity fraud. Members of this community buy and sell social engineering services to one another to facilitate these illicit activities (see Figure 1).

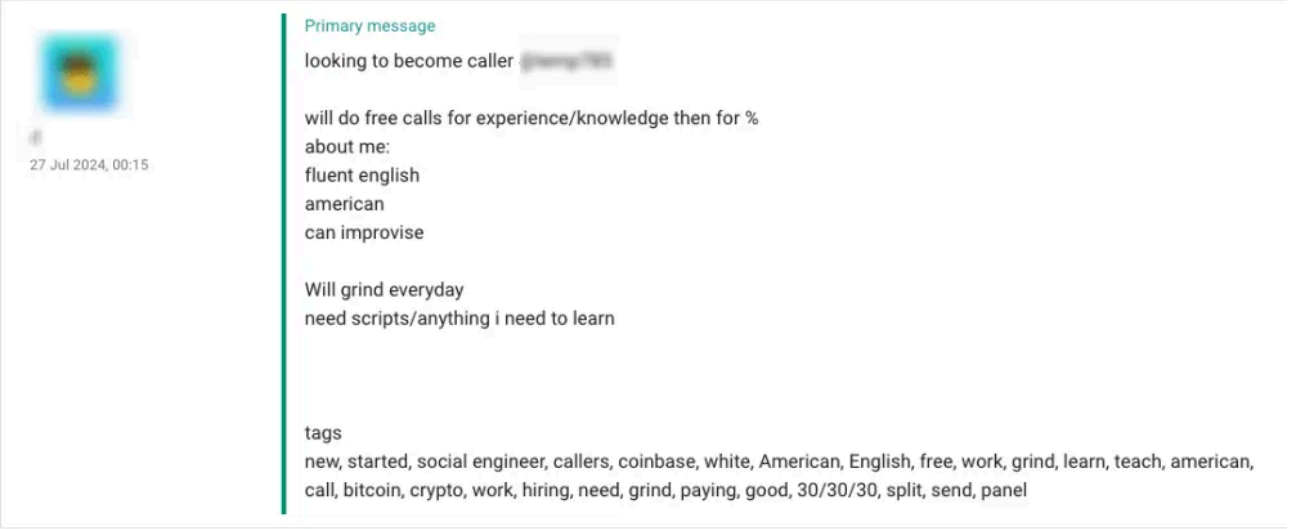

Figure 1: Telegram user offers social engineering services

Scattered Spider members likely refined their social engineering skills through these activities, now using them alongside Russia-linked ransomware groups to target organizations for financial gain.

Since at least August 2023, Scattered Spider has been collaborating with ransomware-as-a-service (RaaS) groups. Initially an affiliate for "ALPHV" (aka "BlackCat"), Scattered Spider gained notoriety by attacking multiple US-based casinos. In February 2024, ALPHV conducted an exit scam against its affiliates and disbanded, leaving them searching for new partners.

That same month, a new ransomware group, RansomHub, began recruiting affiliates (see Figure 2). RansomHub offered an enticing deal, keeping just 10% of attack profits for malware developers and leaving affiliates with 90%. Since June 2024, security researchers have detected intrusions leading to the deployment of the RansomHub malware, which featured tactics typical of Scattered Spider, suggesting the group is now a RansomHub affiliate.

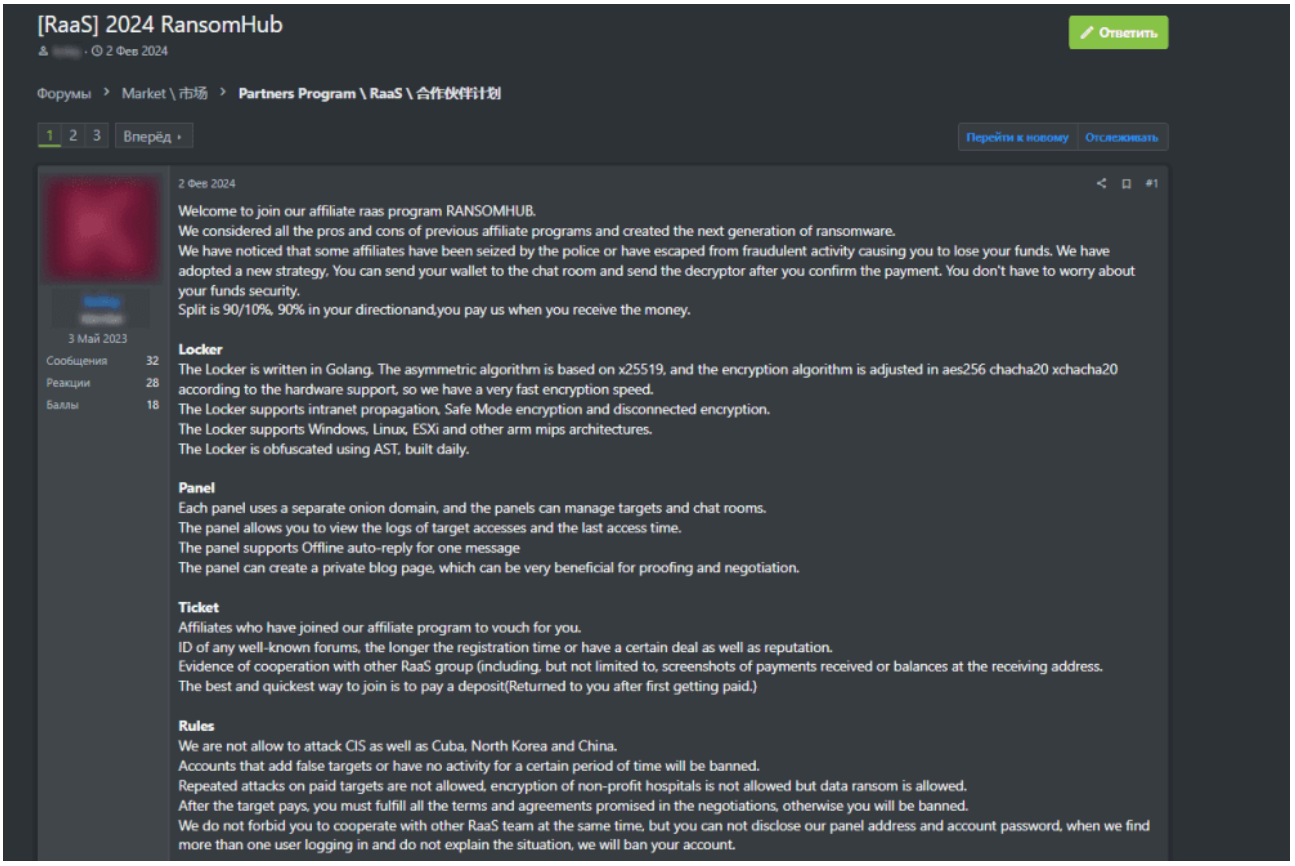

Figure 2: RansomHub advertises affiliate program

The potent combination of RansomHub's lucrative incentives and Scattered Spider's sophisticated social engineering poses a significant threat. Companies across all sectors must rigorously evaluate their security measures to ensure resilient defenses against such attacks, particularly as we anticipate that other adversaries will likely adopt Scattered Spider's effective techniques.

Attack Analysis

During our investigation, we observed the following noteworthy behaviors in this incident:

- **Persistent Social Engineering:** Consistent with Scattered Spider's typical initial access method, the threat actor in this incident gained initial access by social engineering the organization's help desk to compromise the CFO's account. When this account lacked the sufficient permissions, the threat actor used the same social engineering tactic on the help desk again to gain access and compromise a domain admin account.
- **Telecom Infrastructure Abuse:** The threat actor used Verizon IPv6 addresses to access the network, leveraging telecommunications infrastructure with a clean reputation to bypass security controls.
- **ESXi Defense Evasion:** The threat actor spun up their own VM in the victim's ESXi environment to carry out a wide range of adversarial actions such as lateral movement, credential dumping, and data exfiltration.
- **Rapid Time to Impact:** The adversary compromised two accounts within an hour of calling the help desk, accessed the virtual environment in under two hours, and encrypted systems in just over six hours. They

maintained access for roughly ten hours by moving from the organization’s identity and cloud solutions to their on-premises environment.

Given the speed and simultaneous actions in this event, we assess with high confidence that multiple individuals facilitated the attack. The following timeline provides a breakdown of each step during the incident.

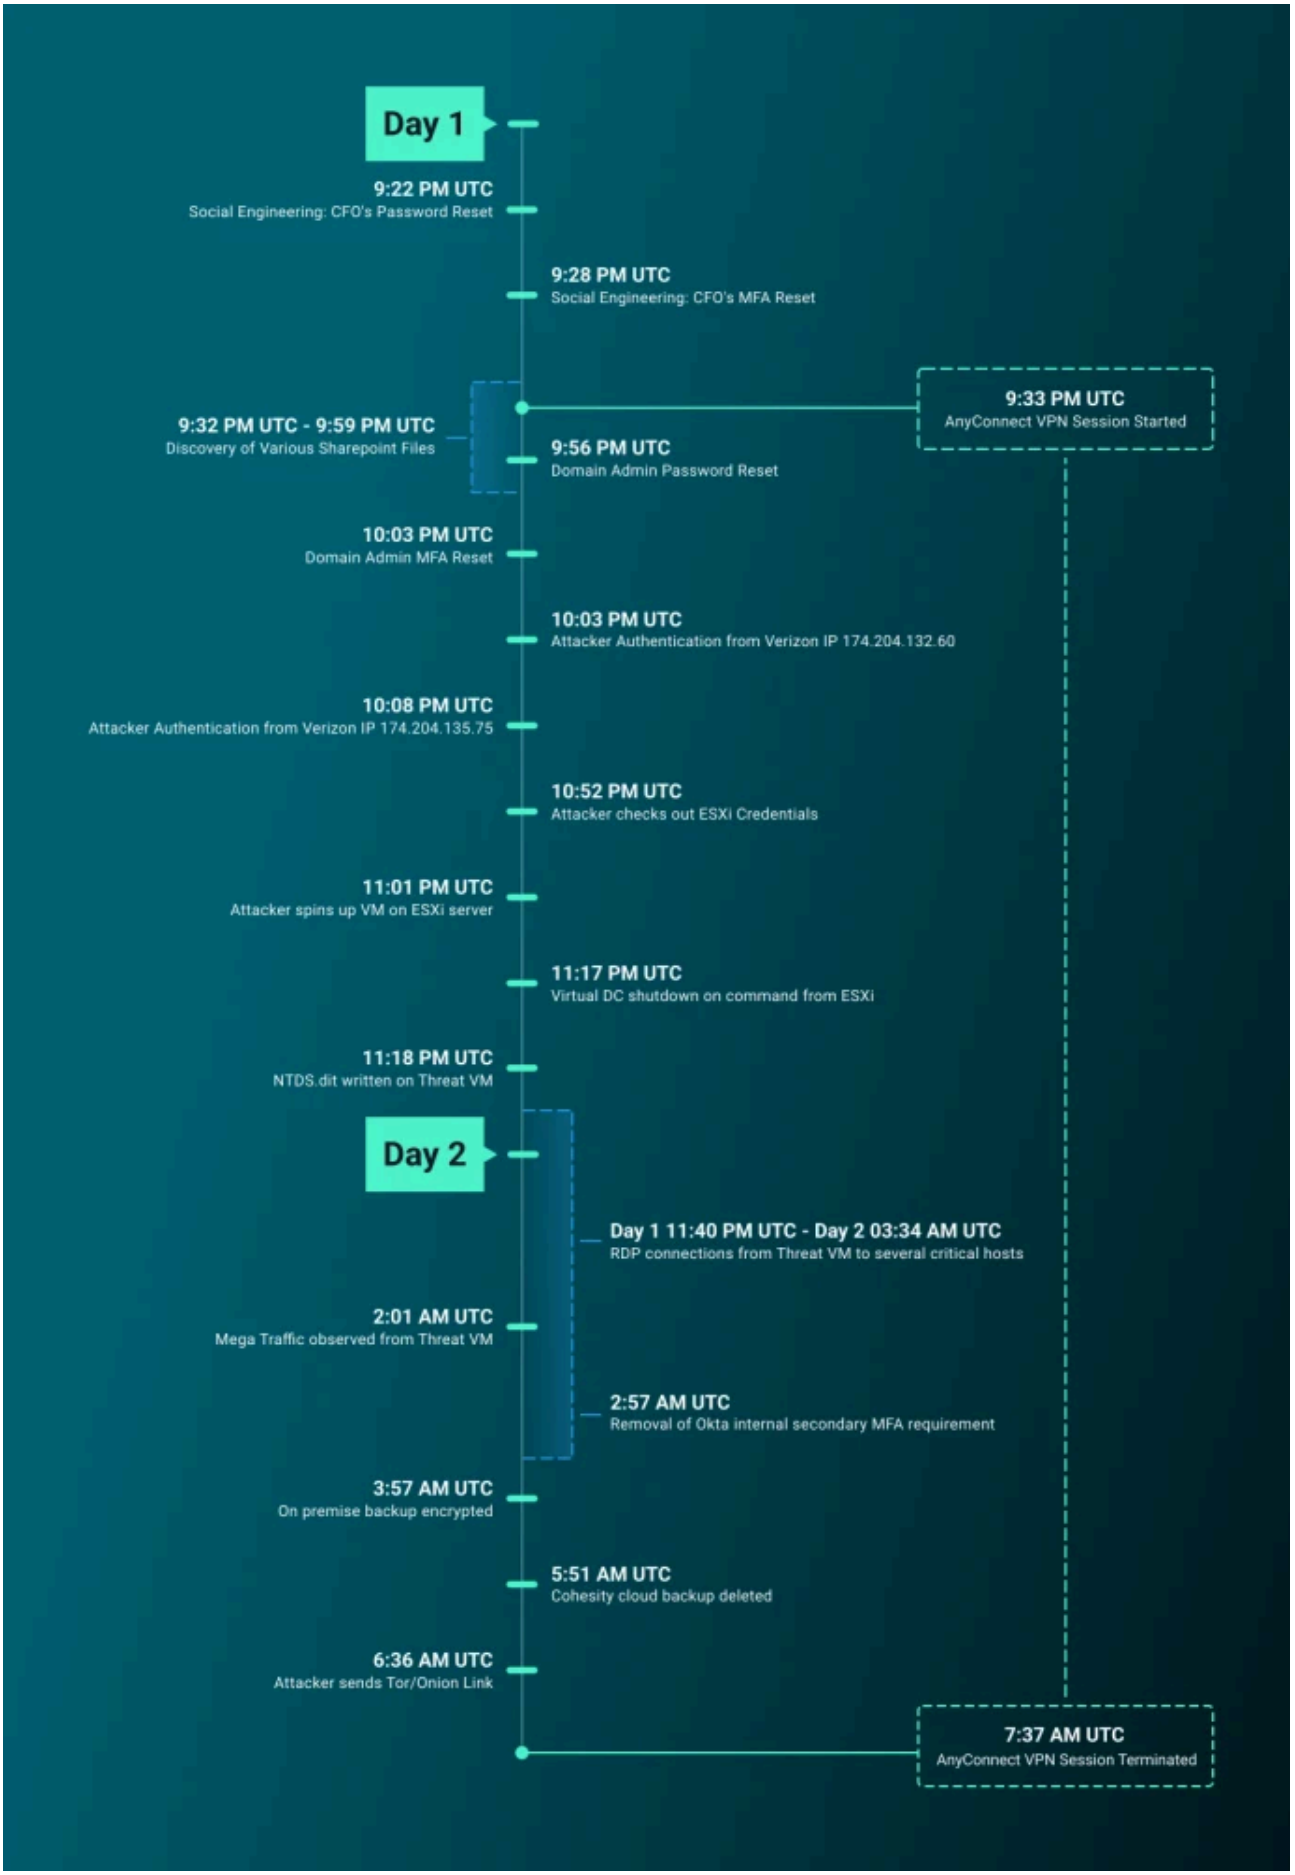
Day 1
9:22 PM UTC
Social Engineering: CFO's Password Reset
9:28 PM UTC
Social Engineering: CFO's MFA Reset
9:33 PM UTC
AnyConnect VPN Session Started
9:32 PM UTC - 9:59 PM UTC
Discovery of Various Sharepoint Files
9:56 PM UTC
Domain Admin Password Reset
10:03 PM UTC
Domain Admin MFA Reset
10:03 PM UTC
Attacker Authentication from Verizon IP 174.204.132.60
10:08 PM UTC
Attacker Authentication from Verizon IP 174.204.135.75
10:52 PM UTC
Attacker checks out ESXi Credentials
11:01 PM UTC
Attacker spins up VM on ESXi server
11:17 PM UTC
Virtual DC shutdown on command from ESXi
11:18 PM UTC
NTDS.dit written on Threat VM
Day 2
Day 1 11:40 PM UTC - Day 2 03:34 AM UTC
RDP connections from Threat VM to several critical hosts
2:01 AM UTC
Mega Traffic observed from Threat VM
2:57 AM UTC
Removal of Okta internal secondary MFA requirement
3:57 AM UTC
On premise backup encrypted
5:51 AM UTC
Cohesity cloud backup deleted
6:36 AM UTC
Attacker sends Tor/Onion Link
7:37 AM UTC
AnyConnect VPN Session Terminated

Figure 3: Scattered Spider attack timeline

Social Engineering: Fool Me Once, Fool Me Twice

To gain initial access to the target network, the threat actor called the organization's IT help desk and persuaded staff to reset the CFO's account password. They then made a second call to another help desk employee, convincing them to reset the multifactor authentication (MFA) controls on the CFO's account. This allowed the attacker to enroll their own SMS device, which was later identified as a voice over IP (VOIP) Google Voice phone number: (971) 444-5872.

The attacker now had access to the user's Okta account. Given that Okta is a single sign-on (SSO) solution, the threat actor was able to access all Okta applications provisioned to the CFO. The following Okta payload shows how the threat actor sent MFA requests to their own SMS device.


```
{"actor":{"id":"Redacted","type":"User","alternateId":"CFO@organization.com","displa
yName":"CFO","detailEntry":null},"client":{"userAgent":{"rawUserAgent":"Mozilla/5.0(
WindowsNT10.0;Win64;x64)AppleWebKit/537.36(KHTML,likeGecko)Chrome/129.0.0.0Safari/53
7.36","os":"Windows10","browser":"CHROME"},"zone":"null","device":"Computer","id":nu
ll,"ipAddress":"174.204.132.60","geographicalContext":{"city":"NewYork","state":"New
York","country":"UnitedStates","postalCode":"10128","geolocation":{"lat":40.7809,"lo
n":-
73.9502}}},"device":null,"authenticationContext":{"authenticationProvider":null,"cre
dentialProvider":null,"credentialType":null,"issuer":null,"interface":null,"authenti
cationStep":0,"rootSessionId":"Redacted","externalSessionId":"Redacted"},"displayMes
sage":"SendsecondfactorauthSMS","eventType":"system.sms.send_factor_verify_message",
"outcome":{"result":"SUCCESS","reason":null},"published":"Day1T21:29:14.518Z","secur
ityContext":{"asNumber":6167,"asOrg":"verizon","isp":"verizon","domain":null,"isProx
y":false},"severity":"INFO","debugContext":{"debugData":{"behaviors":"{NewGeo-
Location=NEGATIVE,NewDevice=POSITIVE,NewIP=NEGATIVE,NewState=NEGATIVE,NewCountry=NEG
ATIVE,Velocity=NEGATIVE,NewCity=NEGATIVE}","requestUri":"/api/v1/authn/factors/Redac
ted","transactionId":"Redacted","url":"/api/v1/authn/factors/smsRedacted/verify/rese
nd?","isSmsHookFailover":"false","countryCodeIso2":"US","phoneNumber":"+1_971-444-
XXXX","authnRequestId":"Redacted","countryCallingCode":"1","requestId":"Redacted","d
tHash":"Redacted","smsProvider":"TELESIGN","risk":"{reasons=AnomalousDevice,level=HI
GH}","threatSuspected":"false"}},"legacyEventType":"core.user.sms.message_sent.facto
r","transaction":{"type":"WEB","id":"Redacted","detail":{}},"uuid":"Redacted","versi
on":"0","request":{"ipChain":[{"ip":"174.204.132.60","geographicalContext":{"city":"
NewYork","state":"NewYork","country":"UnitedStates","postalCode":"10128","geolocatio
n":{"lat":40.7809,"lon":-
73.9502}},"version":"V4","source":null}]},"target":[{"id":"Redaced","type":"User","a
lternateId":"CFO@organization.com","displayName":"CFO","detailEntry":null},{"id":"Re
dacted","type":"MobilePhone","alternateId":"+19714445872","displayName":"+1971444587
2","detailEntry":null}]}
```


Next, the threat actor set their sights on Thycotic—a password vault housing organizational secrets (passwords), including those for privileged accounts. They attempted to access it through its Okta tile but couldn't progress as the CFO's account lacked the sufficient permissions. Undeterred, they searched across the organization's SharePoint, which we will explore further below, and pinpointed a domain administrator account to gain the elevated privileges they needed. This isn't the first time we've seen Scattered Spider target password managers. As previously reported, the collective has previously used the same tactic in other intrusions.

Having identified a new target account, the threat actor made another call to the help desk and requested a password reset for the domain administrator account, which also carried Okta Super Administrator privileges. This account had access to Thycotic through Okta and could self-assign any Okta apps due to its privileged role. In

both instances, the help desk failed to follow the firm's standard operating procedures (SOPs), resulting in the password and MFA information for the domain administrator account falling into the hands of the threat actor.

Human errors are bound to happen, which is why it's vital to implement technical controls alongside SOPs to reinforce organizational policies and prevent inadvertent mistakes. For instance, change request controls should require secondary authorization before resetting credentials for privileged or executive accounts.

Attacker Infrastructure

During this early access phase of the intrusion, we uncovered some intriguing aspects of the attacker's infrastructure. These insights not only helped us identify the threat actor but also provided valuable intelligence about their operations.

- **Verizon IPv4 for Okta Access:** These IP addresses appeared in the organization's logs as Verizon IPv4 addresses with a clean reputation, meaning they would not trigger any suspicious IP detections. Later in the attack, the attacker's activity was associated with Verizon IPv6 addresses. While it's unclear why the IPv4 showed up initially, we do know that Okta doesn't currently support IPv6. Therefore, this could be IPv6 reverting to IPv4 when it isn't supported. Notably, we've seen Scattered Spider using Verizon IP addresses in several other intrusions.
- **Verizon IPv6 for EntraID and O365 Access:** When pivoting to EntraID and O365, the attacker's activity was associated with IPv6 addresses on Verizon's network. Most threat intelligence sources that feed into SIEM correlations don't support IPv6 addresses, making this a clever method to bypass detections. Additionally, IPv6 addresses under certain conditions can circumvent risky sign-in and location-based conditional access policies.
- **Two Attackers, One Compromised Account (High Confidence):** We observed authentications from two different Verizon IP addresses just minutes after the initial account compromise, each requesting MFA separately for the same account. Different user agents were used and separate actions were simultaneously performed within the compromised account. Parsing these user agents shows two different browsers were used: a Chrome browser was tied to one IP address and a Firefox browser was tied to the other IP address, indicating that two threat actors accessed the same compromised account from two different hosts.
- **Scattered Spider's Use of Cellular Hotspots (Medium Confidence):** In nearly every intrusion we've investigated and attributed to Scattered Spider, we've identified mobile providers as their primary infrastructure. The constant shift between IPv4 and IPv6 is a known fallback capability of hotspot devices, suggesting that such devices were used in this incident.

Accessing SharePoint Secrets and Breaking into SentinelOne

In this incident, the threat actor accessed several SharePoint files via the CFO's account to gather information for lateral movement. Notably, in a previous attack carried out by Scattered Spider, we observed the collective abuse SharePoint access and knowledge article repositories. In the current incident, the following files were accessed:

Guide to Working from Home.jpg	Remote Access to Your Computer.jpg
Citrix Login.docx	What Requires an Access Request.pdf
VPN and Multifactor Authentication Guide.pdf	New VPN Setup Instructions.docx
Logmein Prerequisites.pdf	IT Administrative Access.docx
Install Cisco AnyConnect Client.pdf	ESXi Server Refresh Project.xslx
Change Password via Okta.pdf	Engineering Password Vault Utility.pdf

Among all the files, those detailing the IT organization structure almost certainly facilitated in targeting the domain administrator account.

Once the threat actor had access to the domain administrator account, they retrieved additional files relating to backups and key network infrastructure. Since this domain administrator was also an Okta Super Admin, the attacker was also able to access several additional IT applications through Okta SSO, including:

Okta Admin Console	SentinelOne
Thycotic Prod	Cohesity Helios
Microsoft Office 365	Solarwinds
LogMeIn	

The following Okta payload shows the threat actor impersonated a user with SentinelOne access (more details below):

```
{"actor":{"id":"Redacted","type":"User","alternateId":"DomainAdmin@organization.com"
,"displayName":"Domain,Admin","detailEntry":null},"client":{"userAgent":{"rawUserAge
nt":"Mozilla/5.0(WindowsNT10.0;Win64;x64)AppleWebKit/537.36(KHTML,likeGecko)Chrome/8
6.0.4240.75Safari/537.36Edg/86.0.622.38","os":"Windows10","browser":"CHROMIUM_EDGE"}
,"zone":"null","device":"Computer","id":null,"ipAddress":"OrgIPRedacted","geographic
alContext":{"city":"OrgCity","state":"OrgState","country":"UnitedStates","postalCode
":"OrgCode","geolocation":{"lat":Redacted,"Redacted"}}},"device":null,"authentication
Context":{"authenticationProvider":null,"credentialProvider":null,"credentialType":n
ull,"issuer":null,"interface":null,"authenticationStep":0,"rootSessionId":"Redacted"
,"externalSessionId":"Redacted"},"displayMessage":"Changeuser'sapplicationusername",
"eventType":"application.user_membership.change_username","outcome":{"result":"SUCCE
SS","reason":null},"published":"Day2T02:13:23.697Z","securityContext":{"asNumber":Re
dact,"asOrg":"Redact","isp":"Redacted.","domain":null,"isProxy":false},"severity":"I
NFO","debugContext":{"debugData":{"requestId":"Redacted","dtHash":"Redacted","reques
tUri":"/api/v1/apps/Redacted/users/Redacted","oldDisplayName":"DomainAdmin@organizat
ion.com","url":"/api/v1/apps/Redacted/users/Redacted?"}},"legacyEventType":"app.gene
ric.config.app_username_update","transaction":{"type":"WEB","id":"Redacted","detail"
:{}},"uuid":"Redacted","version":"0","request":{"ipChain":[{"ip":"OrgIP","geographic
alContext":{"city":"OrgCity","state":"OrgState","country":"UnitedStates","postalCode
":"OrgCode","geolocation":{"lat":Redacted,"lon":Redacted}},"version":"V4","source":n
ull}]},"target":[{"id":"unknown","type":"AppUser","alternateId":"S1Analyst@organizat
ion.com","displayName":"Domain,Admin","detailEntry":null},{"id":"Redacted","type":"A
ppInstance","alternateId":"SentinelOne","displayName":"SentinelOne","detailEntry":nu
ll},{"id":"Redacted","type":"User","alternateId":"DomainAdmin@organization.com","dis
playName":"Domain,Admin","detailEntry":null}]}
```


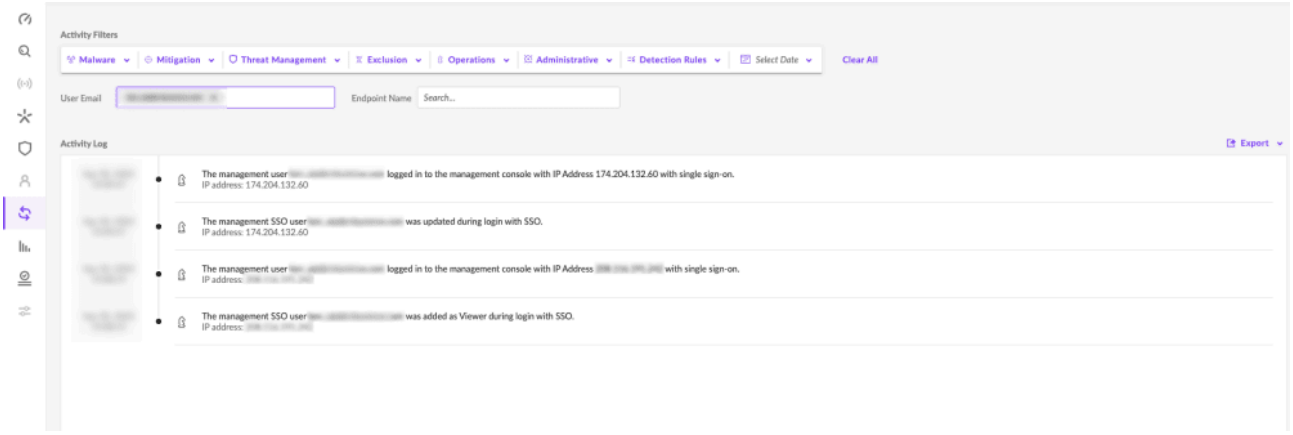

Figure 4: SentinelOne access gained by threat actor

Authentication Manipulation

With access to the Okta Super Admin account, the threat actor manipulated the authentication process in several ways:

- **Removing Secondary MFA:**The attacker disabled secondary MFA for several critical applications through Okta. Despite having already registered their own MFA device, this step was likely intended to simplify and ensure continuous access to the applications.
- **Targeting SentinelOne:**The attacker specifically targeted SentinelOne, an EDR application, but initially lacked access. By impersonating a user with SentinelOne permissions, they granted access to the domain administrator account, allowing them to access the SentinelOne console with view permissions.

- **Resetting MFA:**The attacker reset MFA on three additional accounts in the environment to perform specific functions associated with each one. All these accounts had administrative access to different applications due to their roles (e.g., the Database Administrator had access to Snowflake). We also noticed the same phone number being reused for MFA. The following Thycotic log shows the VMWare ESXi host secret being accessed by the compromised domain admin.

```
Day 1 18:52:13 Redacted CEF:0|Thycotic Software|Secret
Server|11.7.000043|10004|SECRET - VIEW|2|msg=[[SecretServer]] Event: [Secret]
Action: [View] By User: Organization\DomainAdmin Item name: VMWare ESXi host (Item
Id: 1234) Container name: OrgInfrastructure (Container Id: 00) Details: Account
Name: root suid=00 suser=Organization\DomainAdmin cs4=Organization\DomainAdmin
cs4Label=suser Display Name src=174.204.132.60 rt=Day 1 2024 22:51:53 fname=VMWare
ESXi host fileType=Secret fileId=1234 cs3Label=Folder cs3=OrgInfrastructure
```

To gain credentialed access, the actor enumerated several “secrets” (passwords) in the Thycotic password vault. A secret containing ESXi admin credentials allowed them to gain access to the on-premises environment.

Exploiting VPN and ESXi for Undercover Operations

At this stage of the attack, visibility was lost as unmanaged devices were used. However, by working closely with a partner forensics team, we recovered several key events:

- **AnyConnect VPN Session:**The attacker initiated an AnyConnect virtual private network (VPN) session using the CFO’s account from the Verizon IP address 174.204.132[.]60, gaining access to the on-premises environment.
- **ESXi Host Access:**From the VPN device, the attacker checked out the VMware ESXi host credentials from Thycotic and logged into the ESXi server. They then created a new VM on the ESXi host, likely to evade endpoint-based detection, as the new VM wouldn’t have any logging available.
- **Network Connections:**From the attacker’s VM, several network connections were made to on-premises domain controllers. Further investigation revealed that the attacker used Remote Desktop Protocol (RDP) to access multiple servers, including domain controllers, SQL servers, and backup servers.

Our investigation uncovered an NTDS.dit file on the attacker’s VM. This file is the “crown jewel of the domain,” containing the core elements of Windows Active Directory and enabling the extraction of password hashes for any domain user. Despite having full EDR coverage on all domain controllers and SQL servers, no EDR alerts fired during the intrusion. We theorize that this is because the events occurred at the hypervisor level. Logging and EDR telemetry are collected at the operating system (OS) layer, where we observed very little activity. Based on the available facts, we formed the following concrete narrative around the events surrounding the NTDS.dit file write:

- The domain controllers were registered in SentinelOne as being virtualized through VMware ESXi.
- With access to the ESXi host, the attacker could mount or copy the virtual hard drive of the virtual domain controller.
- After mounting the virtual hard drive onto their VM, the attacker copied over the NTDS.dit file.

Forensics pinpointed the exact time the NTDS.dit file was written to the threat actor's VM. Endpoint logs show that just seconds earlier, the virtualized domain controller had been shut down via a default command from the compromised ESXi server. This shutdown command was identified through the parent process vmtoolsd.exe, which is used to delegate commands from the vCenter/ESXi server to individual VMs.

The shutdown was a critical step for the attacker, allowing them to mount the virtual hard drive. Once mounted, copying the NTDS.dit file and dumping hashes is straightforward and can be executed on the attacker's VM host without raising any suspicion. The following log shows the shutdown process being initiated on the domain controller.

```
<13>Day 1 16:19:40 DomainCtrller AgentDevice=WindowsLog AgentLogFile=Security
PluginVersion=WC.MSEVEN6.10.1.1.30 Source=Microsoft-Windows-Security-Auditing
Computer=DomainCtrller OriginatingComputer=Redacted User= Domain= EventID=4688
EventIDCode=4688 EventType=8 EventCategory=13312 RecordNumber=Redacted
TimeGenerated=Redacted TimeWritten=Redacted Level=LogAlways
Keywords=AuditSuccess Task=SE_ADT_DETAILEDTRACKING_PROCESSCREATION Opcode=Info
Message=A new process has been created. Creator Subject: Security ID: NT
AUTHORITY\SYSTEM Account Name: DomainCtrller$ Account Domain: ENT Logon
ID: 0x3E7 Target Subject: Security ID: \NULL SID Account Name: - Account
Domain: - Logon ID: 0x0 Process Information: New Process ID: 0x1a5c New
Process Name: C:\Windows\System32\cmd.exe Token Elevation Type: %%1936 Mandatory
Label: Mandatory Label\System Mandatory Level Creator Process ID: 0x8cc Creator
Process Name: C:\Program Files\VMware\VMware Tools\vmtoolsd.exe Process Command
Line: C:\Windows\system32\cmd.exe /c ""C:\Program Files\VMware\VMware
Tools\poweroff-vm-default.bat""
```

The hypothesized sequence of events is important, especially if domain controllers in an enterprise environment are virtualized. Infosec teams may have a false sense of security when critical servers are equipped with EDR technology and redundant logging. However, if an attacker gains access to the underlying hypervisor or cloud service hosting the virtual server, OS-level visibility cannot be relied upon, and teams must have other defense measures in place to prevent further damaging consequences.

New Tactic: Demanding Ransom Through Teams

The threat actor carried out a double extortion attack: they encrypted the ESXi environment and exfiltrated data. They further targeted the organization's backup solutions, encrypting on-premises backups and deleting cloud backups.

For the local data backup server, the attacker used the open-source disk encryption tool VeraCrypt. For the cloud-based backup solution, they used Okta to access Cohesity, an enterprise data backup and security solution, and deleted associated storage accounts.

With their extensive access, the threat actor exfiltrated several gigabytes of data, transferring it from their VM to an IP address owned by Mega Cloud, a frequently abused cloud storage service.

Notably, we also observed a novel ransom note technique. Traditionally, threat actors leave a message named README on every host after a successful encryption event. However, in this attack, after encrypting hosts and exfiltrating data, the attacker sent a Microsoft Teams message from the compromised domain admin account, containing an Onion link for the ransom demand. They also sent an email titled "Urgent Update on Cyber Attack"

from the same account. Below is the payload of the Microsoft Teams message sent from the compromised domain admin.


```
{"AppAccessContext":{"IssuedAtTime":"Day21T06:23:06","UniqueTokenId":"Redacted"},"Cr
eationTime":"Day2T06:33:23","Id":"Redacted","Operation":"MessageCreatedHasLink","Org
anizationId":"Redacted","RecordType":25,"UserKey":"Redacted","UserType":0,"Version":
1,"Workload":"MicrosoftTeams","ClientIP":"2600:1017:b80d:9a48:adf0:5809:b514:1b5b","
UserId":"DomainAdmin@organization.com","ChatThreadId":"Redacted@unq.gbl.spaces","Com
municationType":"OneOnOne","ExtraProperties":[{"Key":"TimeZone","Value":"America/
"},{"Key":"OsName","Value":"windows"},{"Key":"OsVersion","Value":"NT10.0"},{"Key":"C
ountry","Value":"us"},{"Key":"ClientName","Value":"skypeteams"},{"Key":"ClientVersio
n","Value":"1415/24081700421"},{"Key":"ClientUtcOffsetSeconds","Value":"-
25200"}],"MessageId":"Redacted","MessageVersion":"Redacted","ParticipantInfo":{"HasF
oreignTenantUsers":false,"HasGuestUsers":false,"HasOtherGuestUsers":false,"HasUnauth
enticatedUsers":false,"ParticipatingDomains":[],"ParticipatingSIPDomains":[],"Partic
ipatingTenantIds":["Redacted"]},"ResourceTenantId":"Redacted","ItemName":"Redacted@u
nq.gbl.spaces","MessageURLs":["http://lwrw6gh7ckfe62batalzvozqp3cgkukp5v2jpidvrisk7t
p7uluef4id.onion/","http://torproject.org/)"]}
```


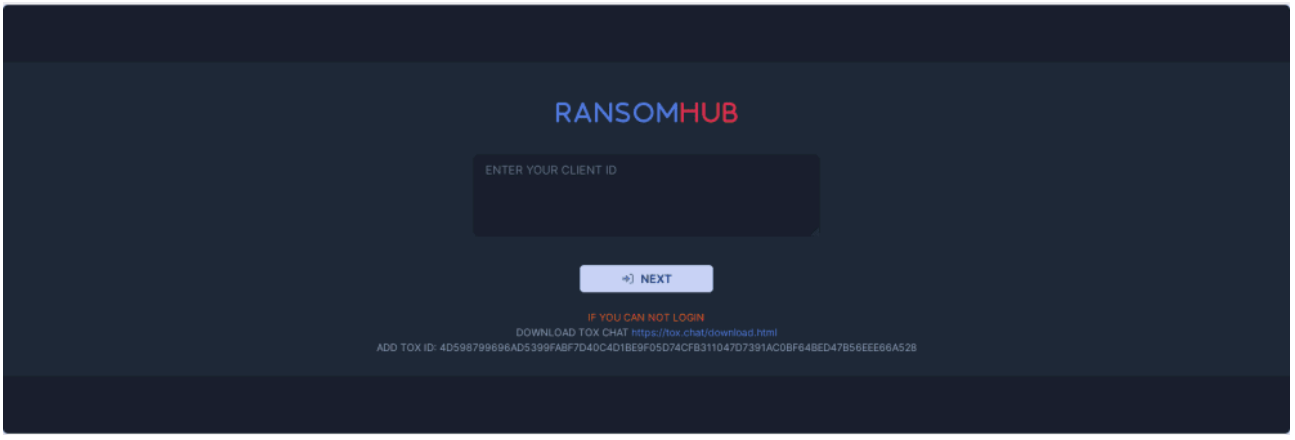

Figure 5: Ransom negotiation portal sent to customer

What Lies Ahead

This incident sheds light on the ongoing partnership between Scattered Spider and RansomHub and offers several other insights into the future threat landscape.

First, the involvement of threat actors skilled in social engineering demonstrates a demand for fluent English speakers to collaborate with ransomware affiliates. Second, RansomHub is attracting talented adversaries, suggesting it will continue to be the most dominant ransomware group, after recently surpassing former leading group "LockBit." Third, financially motivated attackers are getting much better at pressuring organizations into paying ransoms by targeting critical virtual infrastructure, such as ESXi and backups, while successfully evading detection.

These projections, supported by our observations and thorough investigations, emphasize the evolving tactics and growing sophistication of ransomware affiliates, highlighting the increasingly complex cyber risks that organizations face.

Cross-Language Collaboration

Previously, Russia-aligned threat actors have been reluctant to collaborate with English-speaking counterparts due to language and cultural barriers, law enforcement concerns, and trust issues. They often enjoy greater operational freedom within Russia, provided they don't target organizations in the Commonwealth of Independent States (CIS), China, and North Korea. They also perceive English-speaking counterparts as having poorer operational security practices and being less capable of executing sophisticated attacks, which could increase the risk of exposure and possible arrest for all parties involved.

Despite the arrests of three alleged Scattered Spider members in 2024, Russia-aligned adversaries continue to see the value of partnering with English speakers. For example, in June 2024, a user on the Russian-language forum XSS commented on the arrests: "One example of why it's still worth working with English speakers, but also important to keep in mind how quickly they can be caught." This sentiment is echoed by the many forum posts and replies requesting or offering calling services in English. For example, in July 2024, an XSS post was created advertising English calling services (see Figure 6) to Russia-linked threat actors and recruiting more English speakers due to growing demand: "As a result of the expansion of my business, I am actively searching for a competent English-language caller. I'm prepared to take them on permanently or for project work for %." The post continues, "Apart from English-language calling, we also offer calling services in Spanish, French, and German."

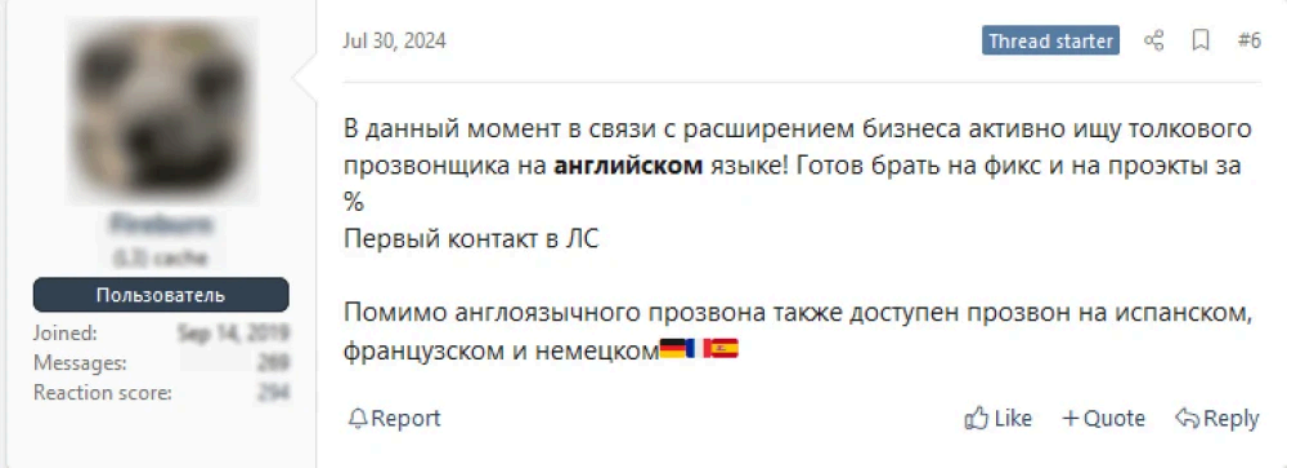

Figure 6: XSS user advertises calling services

In September 2024, the same user emphasized the importance of targeting specific employees within an organization, stating, "A reminder that for more effective corporate calls, you need the contact details for C-suite management level, legal, or finance departments." Forum users responded to the advertisement, some with specific requests, such as, "Hello, I need a native speaker, strictly American Floridian accent. Can you help?" These listings highlight the high demand for calling services and social engineering in multiple languages, particularly English.

Likely due to its effectiveness, affiliates of ransomware groups, including those beyond RansomHub, have also been using English-speaking callers to target organizations. For example, in May 2024, we identified an ongoing social engineering campaign in which affiliates of the "Black Basta" ransomware group made calls as IT personnel to target employees.

We forecast with high confidence that ransomware affiliates will almost certainly continue using English speakers for social engineering attacks in the long term (beyond one year). To mitigate this risk, organizations should implement stringent help desk procedures and standard IT interaction processes. Additionally, investing in regular employee training is critically important to raise awareness about the significant risk of social engineering and maintain high levels of vigilance.

The Rise of RansomHub

Although RansomHub only became active in Q1 2024, it quickly gained dominance—surpassing previously prominent groups like LockBit and "Play" by mid-year (see Figure 7), just as we forecasted in Q1 2024. This rapid rise is attributed to law enforcement action taken against LockBit and the disbandment of ALPHV, which led affiliates to gravitate towards RansomHub, enticed particularly by their lucrative 90/10 profit split.

This profit-sharing structure has also attracted more advanced adversaries, including members of Scattered Spider, who are likely working together with Russia-linked threat actors. Scattered Spider's social engineering skills complement the network-compromising expertise of their Russia-linked counterparts, making their collaboration particularly effective. The 90/10 profit split results in higher income for both Scattered Spider and Russia-linked groups, attracting them to RansomHub for sustained collaboration.

Given RansomHub's favorable positioning and its expert affiliates, we forecast with high confidence that RansomHub will remain the dominant ransomware group in the mid-term (between three months and one year).

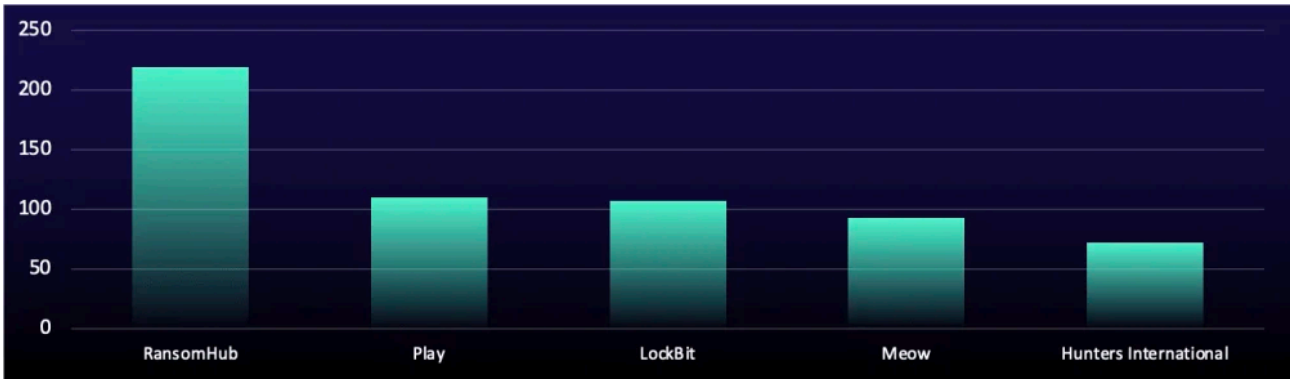

Figure 7: Victims named by most active ransomware groups, July 1, 2024 to October 18, 2024

Why ESXi Will Remain a Prime Target

Adversaries have targeted ESXi servers since at least 2021, as demonstrated by the "Defray777" and "Darkside" ransomware variants deployed by "Sprite Spider" and "Carbon Spider," respectively. This trend is likely due to two main reasons. First, VMware is a leading vendor in virtualization, which means that developing a specific ransomware variant provides many targeting opportunities. Second, ESXi servers host multiple VMs on a single server that commonly run critical applications and services. This creates a single point of failure that disrupts essential business operations and heightens organizations' urgency to resolve the issue. This enables attackers to quickly achieve maximum impact and apply pressure on victims to, as a result, increase the likelihood of a ransom payment.

Additionally, as observed in this incident, attackers are becoming increasingly aware of security controls such as EDR. This indicates that adversaries are likely to evade detection by limiting interaction with physical systems, thereby avoiding detection or log generation. The growing mentions of “ESXi” on the Russian-language ransomware-focused forum RAMP (see Figure 8) also reflects threat actors’ heightened focus on ESXi servers due to their potential profitable returns.

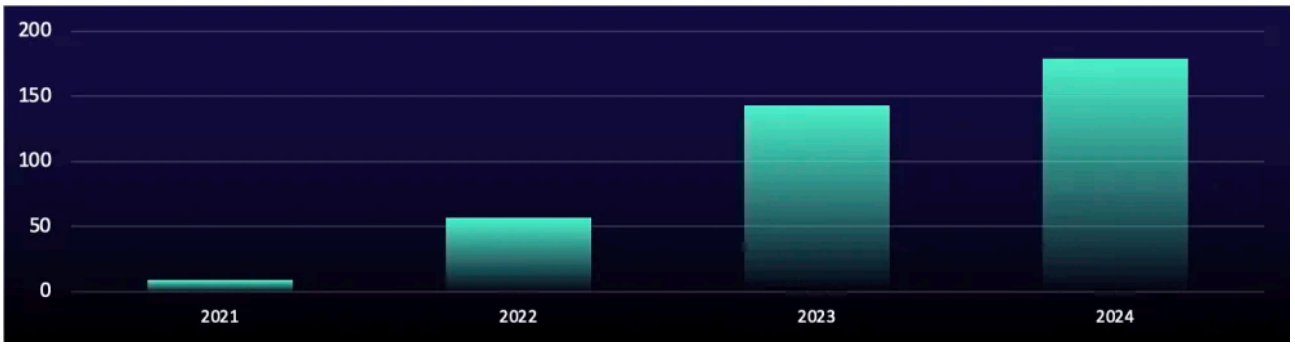

Figure 8: ESXi mentions on RAMP from January 2021 to October 2024

We anticipate with high confidence that, in the long term, ransomware developers will continue to create variants targeting ESXi servers in response to increased demand from affiliates. Additionally, affiliates will persist in targeting ESXi servers by deploying ransomware to halt business operations and maximize potential profits from ransom payments. Furthermore, these virtual systems will continue to be abused for evasion as they offer an alternative for directly targeting physical systems that are likely to generate detections and log malicious activity.

What ReliaQuest Is Doing

For the fastest remediation, organizations should implement automated incident response, such as enabling GreyMatter Automated Response Playbooks, to automatically contain threats, reducing mean time to contain (MTTC) and halting the adversary’s progress. Alternatively, organizations can set GreyMatter Response Playbooks to “RQ Approved” to allow our analyst team to handle remediation actions. This speeds up containment while requiring a ReliaQuest analyst’s discretion to execute the Response Playbook. Note that certain Playbooks, such as “Isolate Host,” can be set to require phone approval to avoid business disruption.

Terminate Active Sessions and Reset Passwords: Scattered Spider gains initial access through social engineering attacks, deceiving help desks into resetting a targeted user’s password. Enabling these Playbooks can revoke any established malicious sessions and force a password reset, effectively cutting off the attacker’s access.

Disable User: If an account is suspected to be compromised, this GreyMatter Response Playbook will disable the affected account, revoking the adversary’s access and preventing further advancement toward gaining sensitive information and deploying encryption.

Block IP: This Playbook blocks IP addresses using associated technologies like EDR or a firewall. While not a long-term solution, this Playbook should be executed alongside the account remediation plays to revoke the attacker’s access, as IP addresses can easily be changed.

Next Steps: Enhancing Your Defense

This incident offers important lessons that organizations should review in accordance with their own existing processes and technical controls to harden their defensive measures against similar attacks. The threat actor bypassed MFA and help desk policies by social engineering the help desk twice to access high-priority accounts. They were aware of the organization's security controls, leading them to minimize endpoint interactions and generate minimal logs, evading EDR detection as activities occurred at the hypervisor layer.

Organizations across all sectors should consider the following recommendations to strengthen their security posture against these techniques, which are likely to remain prevalent and be adopted by other adversaries.

Shield Your Network from Initial Access Threats

- **Avoid Using SMS Messaging for MFA:** This method is vulnerable to SIM-swapping attacks, which allow adversaries to intercept one-time password codes and gain unauthorized access to accounts. Instead, consider using more secure MFA methods such as authenticator apps or hardware tokens.
- **Mitigating Social Engineering Attacks:** Implement video calls with ID verification and callbacks to verified phone numbers from the employee directory. However, be aware that callbacks are not resistant to SIM-swapping attacks. Attackers often use publicly available information or information obtained via a data breach, such as addresses, social security numbers, and answers to common security questions, to deceive verification processes. These extra steps create a robust multistep verification process and help deter attackers.
- **Conduct Social Engineering Assessments:** These assessments should focus on testing help desk policies, educating employees on recognizing social engineering attacks, and evaluating established procedures. Regular testing ensures that controls are adequate and prepares staff to effectively identify and respond to social engineering attempts.
- **Implement Client-Based Conditional Access Policies:** These policies should require a certificate on the host machine performing the VPN authentication. This control restricts an attacker's ability to authenticate to the network, even if credentials are compromised.
- **Restrict SharePoint Permissions:** In this event, the adversary used the CFO's account to access sensitive resources on SharePoint, including network diagrams, ESXi documentation, and IT organization charts. This information likely provided insights for further attacks, such as targeting the domain administrator account and the virtual environment. To counter this, reduce permissions of sensitive files in SharePoint so only employees who require access can view them.

Reinforce Your VMware ESXi Defenses

- **Ensure Virtualization Systems are Up to Date:** Vulnerable systems can be exploited to escalate privileges, allowing attackers greater access to deploy malicious software like ransomware.
- **Implement vCenter Network Access Control:** Create a network allowlist using the vCenter Server Appliance Firewall. The allowlist permits only trusted traffic to access the VSphere environment, preventing an attacker from accessing the virtual environment if their traffic originates from an untrusted host.

- **Implement ESXi Smart-Card Authentication:** This authentication control restricts access to the ESXi environment even if an administrator's credentials and MFA are compromised, and it replaces the VSphere authentication process with a smart card and PIN. This prevents adversaries from using a compromised privileged account to create, modify, or shut down ESXi hosts.

Resilient Protection from ReliaQuest

This incident highlights the rapid pace at which advanced attackers like Scattered Spider can move through environments, taking just six hours from initial access to impact in this event. They are also increasingly targeting systems that are essential for business operations, as this allows for quicker and more lucrative financial gains.

To counter the increased speed and precision achieved by these adversaries, defenders must minimize their MTTC threats during incident response. Reducing MTTC is crucial for preventing a full-blown attack, as it decreases dwell time and halts the attack before further damage can occur. For more resilient protection, ReliaQuest's security operations platform, GreyMatter, leverages advanced AI to reduce threat response times and lower MTTC to under 5 minutes, helping you to improve your overall security posture.
