

MITRE ATT&CK v17 ESXi Matrix – ZeroLock® Mapping – In-Depth Guide

Tactic	Technique ID	Technique Name	Description / ESXi Relevance	Potential Impact	ZeroLock Feature(s)
Initial Access	T1190	Exploit Public-Facing Application	Adversaries may exploit vulnerabilities in public-facing ESXi components like OpenSLP, vSphere, or vCenter servers that are exposed to the internet.	Grants attackers an initial foothold/attack surface on a highly sensitive system.	Lockdown Rules, Virtual Patching
Initial Access	T1078	Valid Accounts	Adversaries may use stolen or default credentials to log into ESXi hosts or vCenter systems via SSH or web interfaces, enabling stealthy access without triggering alarms or using malware.	Grants attackers an initial foothold/attack surface on a highly sensitive system.	SSH MFA
Execution	T1059	Command and Scripting Interpreter	Attackers may abuse command-line tools like Bash or SH on ESXi hosts to execute malicious scripts or commands—often via SSH, cron jobs, or reverse shells—bypassing conventional controls and using native hypervisor functionality. ESXi subtechnique T1059.012 - Hypervisor CLI: Adversaries may exploit hypervisor CLIs to run malicious commands, leveraging their broad control over both the hypervisor and its guest VMs.	Allows attackers to execute actions on compromised systems.	Lockdown Rules, Application Filtering
Execution	T1675	ESXi Administration Command	Adversaries may abuse Guest Operations APIs (e.g., StartProgramInGuest) to run commands or transfer files on ESXi-hosted VMs through VMware Tools, enabling direct control from the hypervisor.	Allows attackers to execute actions on compromised systems.	SSH MFA, Lockdown Rules
Execution	T1053	Scheduled Task/Job	On ESXi systems, adversaries may use cron or systemd timers to schedule malicious scripts for persistent or privileged execution, often mimicking legitimate admin behavior to evade detection.	Allows attackers to execute actions on compromised systems.	AI-Behavioral Detection, Lockdown Rules, Application Filtering
Persistence	T1098	Account Manipulation	Adversaries may manipulate ESXi user accounts—modifying permissions, SSH keys, or access policies—to establish persistence or escalate privileges across the virtual infrastructure.	Allows attackers to maintain a foothold in the compromised environment as they broaden their attack surface.	SSH MFA, AI-Behavioral Detection, Lockdown Rules
Persistence	T1037	Boot or Logon Initialization Scripts	Attackers may create or modify initialization scripts (e.g., rc.local, init.d) to execute malicious payloads during boot or login, maintaining persistent access with elevated privileges.	Allows attackers to maintain a foothold in the compromised environment as they broaden their attack surface.	Lockdown Rules, Application Filtering
Persistence	T1554	Compromise Host Software Binary	Adversaries may replace or modify trusted binaries on ESXi—such as OpenSSH or management utilities—to embed persistent backdoors.	Allows attackers to maintain a foothold in the compromised environment as they broaden their attack surface.	Application Filtering, Lockdown Rules
Persistence	T1136	Create Account	Adversaries may create new user accounts directly on the ESXi host or vCenter system to maintain persistent access without needing to rely on malware or exploit re-entry.	Allows attackers to maintain a foothold in the compromised environment as they broaden their attack surface.	SSH MFA, Lockdown Rules
Persistence	T1053	Scheduled Task/Job	On ESXi systems, adversaries may use cron or systemd timers to schedule malicious scripts for persistent or privileged execution, often mimicking legitimate admin behavior to evade detection.	Allows attackers to maintain a foothold in the compromised environment as they broaden their attack surface.	AI-Behavioral Detection, Lockdown Rules, Application Filtering
Persistence	T1505	Server Software Component	Includes ESXi specific subtechnique T1505.006 - vSphere Installation Bundles: attackers may install malicious vSphere Installation Bundles (VIBs) or other plug-in components to extend server behavior on ESXi.	Allows attackers to maintain a foothold in the compromised environment as they broaden their attack surface.	Application Filtering, Lockdown Rules

Persistence	T1078	Valid Accounts	Adversaries may use stolen or default credentials to log into ESXi hosts or vCenter systems via SSH or web interfaces, enabling stealthy access without triggering alarms or using malware.	Allows attackers to maintain a foothold in the compromised environment as they broaden their attack surface.	SSH MFA
Privilege Escalation	T1098	Account Manipulation	On ESXi, attackers may manipulate existing accounts—e.g., changing SSH keys, escalating roles, or resetting passwords for default admin users—to maintain control or escalate privileges without triggering new account creation alerts.	Allows attackers to maintain a foothold in the compromised environment as they broaden their attack surface.	SSH MFA, AI-Behavioral Detection, Lockdown Rules
Privilege Escalation	T1037	Boot or Logon Initialization Scripts	Attackers may use or modify startup scripts (e.g., /etc/rc.local, init.d, or systemd units) on ESXi or Linux-based hypervisor hosts to ensure malware executes automatically at boot or login—often with root or elevated privileges.	Allows attackers to maintain a foothold in the compromised environment as they broaden their attack surface.	Lockdown Rules, Application Filtering
Privilege Escalation	T1611	Escape to Host	Adversaries may exploit vulnerabilities in ESXi or misconfigured containers to escape virtual machines and gain control over the hypervisor or host system. This could allow access to other VMs or administrative control of the host.	Allows attackers to maintain a foothold in the compromised environment as they broaden their attack surface.	Lockdown Rules, Virtual Patching
Privilege Escalation	T1053	Scheduled Task/Job	On ESXi systems, adversaries may use cron or systemd timers to schedule malicious scripts for persistent or privileged execution, often mimicking legitimate admin behavior to evade detection.	Allows attackers to maintain a foothold in the compromised environment as they broaden their attack surface.	AI-Behavioral Detection, Lockdown Rules, Application Filtering
Privilege Escalation	T1078	Valid Accounts	Adversaries may use stolen or default credentials to log into ESXi hosts or vCenter systems via SSH or web interfaces, enabling stealthy access without triggering alarms or using malware.	Allows attackers to maintain a foothold in the compromised environment as they broaden their attack surface.	SSH MFA
Defense Evasion	T1140	Deobfuscate/Decode Files or Information	Adversaries may attempt to hide or reassemble malicious payloads on ESXi systems by decoding or deobfuscating files. This can enable hidden persistence, script execution, or data staging without immediate detection.	Enables attacker to avoid detection from defensive actors or systems.	Application Filtering, Lockdown Rules
Defense Evasion	T1480	Execution Guardrails	Constrain malware activity to only execute on systems with specific environment variables—like ESXi folder paths, system language, or domain names. Adversaries may rely on specific paths or identifiers to ensure payloads only run on hypervisor targets.	Enables attacker to avoid detection from defensive actors or systems.	Application Filtering, Lockdown Rules
Defense Evasion	T1222	File and Directory Permissions Modification	Adversaries may alter file permissions on ESXi hosts to access protected configurations, logs, or core hypervisor binaries. This includes manipulating ACLs or symbolic links that point to remote paths or backup volumes—especially impactful in ransomware scenarios where attackers target mounted datastores or shared folders.	Enables attacker to avoid detection from defensive actors or systems.	AI-Behavioral Detection, Lockdown Rules
Defense Evasion	T1564	Hide Artifacts	Adversaries may use built-in OS tools and ESXi admin functions to conceal malicious processes, files, or configurations—evading detection from traditional endpoint and file-based defenses. This can involve hiding backdoors in startup scripts, scheduled tasks, or abusing virtual machine metadata.	Enables attacker to avoid detection from defensive actors or systems.	Application Filtering, Lockdown Rules
Defense Evasion	T1562	Impair Defenses	Adversaries may disable or weaken defenses such as logging, monitoring, and AV tools—especially those that operate within virtual environments. This includes altering host firewall rules, suppressing log files, or tampering with scheduled tasks and registry entries that monitor hypervisor integrity.	Enables attacker to avoid detection from defensive actors or systems.	Lockdown Rules, AI-Behavioral Detection

Defense Evasion	T1070	Indicator Removal	Adversaries may delete or alter logs, files, or registry keys to erase their presence—compromising forensic evidence and limiting detection. This can include clearing logs like <code>/var/log/hostd.log</code> , modifying audit settings, or wiping entries that reflect unauthorized activity.	Enables attacker to avoid detection from defensive actors or systems.	Lockdown Rules, Virtual Patching
Defense Evasion	T1036	Masquerading	Masquerading techniques on ESXi may involve renaming binaries (like <code>esxcli</code> , <code>dcui</code> , or <code>hostd</code>) to hide malicious intent, disguising malicious scripts or processes with legitimate-sounding names, or placing payloads in trusted directories such as <code>/etc/vmware/</code> or <code>/bin</code> .	Enables attacker to avoid detection from defensive actors or systems.	Lockdown Rules, Application Filtering
Defense Evasion	T1027	Obfuscated Files or Information	On ESXi, obfuscated payloads may arrive as encoded shell scripts, Base64-encoded binaries, or split archive files. These could be hidden in authorized directories (e.g., <code>/tmp/</code> , <code>/scratch/</code> , or <code>/var/core/</code>) and are commonly decoded on-device using native tools like <code>base64</code> , <code>dd</code> , or <code>openssl</code> . Obfuscation also extends to CLI commands run via SSH, where attackers use shell tricks or encoded flags to bypass logging and monitoring tools.	Enables attacker to avoid detection from defensive actors or systems.	SSH MFA, Application Filtering, Lockdown Rules
Defense Evasion	T1078	Valid Accounts	Adversaries may use stolen or default credentials to log into ESXi hosts or vCenter systems via SSH or web interfaces, enabling stealthy access without triggering alarms or using malware.	Enables attacker to avoid detection from defensive actors or systems.	SSH MFA
Credential Access	T1110	Brute Force	ESXi systems are often exposed through SSH or accessed via management interfaces like vSphere or API endpoints. Gaining shell access to ESXi through brute-forced credentials gives adversaries the control they need to shut down VMs, encrypt data stores, and laterally move across hosts.	Allows attacker to steal authenticated access or account privileges.	SSH MFA
Discovery	T1087	Account Discovery	Adversaries who gain a foothold in a virtualized environment often try to enumerate user accounts and roles (especially administrative or service accounts). This data enables subsequent credential attacks, privilege escalation, and lateral movement—especially if default accounts or weakly protected credentials are used for hypervisor management (e.g., SSH, vCenter access, or service scripts).	Allows attackers to uncover important information about the host or network.	Lockdown Rules, Virtual Patching
Discovery	T1083	File and Directory Discovery	Adversaries search VM directories (e.g., <code>/vmfs/volumes/</code>) to locate and stage high-value files like <code>.vmdk</code> and config backups for encryption, deletion, or exfiltration.	Allows attackers to uncover important information about the host or network.	Lockdown Rules, Virtual Patching
Discovery	T1654	Log Enumeration	Adversaries may access <code>/var/log/</code> and other ESXi log directories to review authentication activity, system changes, or security event traces—especially to track incident response or find gaps in defenses.	Allows attackers to uncover important information about the host or network.	Lockdown Rules, Virtual Patching
Discovery	T1057	Process Discovery	Attackers may use <code>esxcli system process list</code> or <code>ps</code> to view all active processes—seeking insight into running services, security tools, or backup agents to avoid or terminate.	Allows attackers to uncover important information about the host or network.	Lockdown Rules, Virtual Patching
Discovery	T1018	Remote System Discovery	Adversaries may use <code>esxcli network diag ping</code> , <code>parse /etc/hosts</code> , or <code>scan /vmfs/volumes</code> paths to discover adjacent systems or virtual infrastructure for lateral movement.	Allows attackers to uncover important information about the host or network.	Lockdown Rules, Virtual Patching
Discovery	T1518	Software Discovery	Adversaries may run <code>esxcli software vib list</code> , <code>parse /etc/vmware/</code> , or query VM agent logs to learn what hypervisor versions or software packages are installed—often used to identify vulnerable or unpatched systems.	Allows attackers to uncover important information about the host or network.	Lockdown Rules, Virtual Patching

Discovery	T1082	System Information Discovery	Attackers may run <code>esxcli system version get</code> , <code>esxcli hardware</code> , or <code>access /etc/vmware-release</code> to fingerprint the ESXi host—revealing patch levels, version details, and architecture used for privilege escalation or lateral movement planning.	Allows attackers to uncover important information about the host or network.	Lockdown Rules, Virtual Patching
Discovery	T1016	System Network Configuration Discovery	Adversaries often use commands like <code>esxcli network nic list</code> , <code>esxcli network ip interface ipv4 get</code> , or <code>examine /etc/networking/</code> configurations to collect MAC addresses, IP ranges, and interface details for lateral movement or infrastructure targeting.	Allows attackers to uncover important information about the host or network.	Lockdown Rules, Virtual Patching
Discovery	T1049	System Network Connections Discovery	Adversaries use <code>esxcli network ip connection list</code> to enumerate active connections to/from an ESXi host—revealing services, sessions, or lateral movement targets. This can help them understand traffic flows, isolate targets, or identify management interfaces and backdoors.	Allows attackers to uncover important information about the host or network.	Lockdown Rules, Virtual Patching
Discovery	T1124	System Time Discovery	On ESXi, accurate time is often synchronized via NTP, so understanding this configuration can also help attackers disrupt logging or evade correlation-based detection systems.	Allows attackers to uncover important information about the host or network.	Lockdown Rules, Virtual Patching
Discovery	T1673	Virtual Machine Discovery	Adversaries who compromise an ESXi host can run commands to enumerate active VMs, identify high-value targets, or plan lateral movement.	Allows attackers to uncover important information about the host or network.	Lockdown Rules, Virtual Patching
Lateral Movement	T1210	Exploitation of Remote Services	Once an adversary gains initial access, they may exploit unpatched vulnerabilities in remote services to move laterally to other systems or to the hypervisor itself. VMware vCenter and ESXi have historically been targeted for this (e.g., CVE-2021-21985, CVE-2022-22960).	Gives the attacker the ability to move from the compromised system to another component of the network.	Lockdown Rules, Virtual Patching
Lateral Movement	T1570	Lateral Tool Transfer	Adversaries may move tools or payloads between internal systems using file-sharing protocols or CLI utilities (e.g., <code>scp</code> , <code>curl</code> , <code>ftp</code> , <code>rsync</code> , etc.). On ESXi, this can include using <code>esxcli</code> , shared folders, or remote management tools to move binaries between hypervisors or from host to VM.	Gives the attacker the ability to move from the compromised system to another component of the network.	Lockdown Rules, Application Filtering
Lateral Movement	T1021	Remote Services	Adversaries may use stolen credentials to access ESXi hosts remotely via SSH or management interfaces like vCenter. Once connected, they can execute commands or pivot deeper into the virtual infrastructure.	Gives the attacker the ability to move from the compromised system to another component of the network.	SSH MFA, Lockdown Rules
Collection	T1005	Data from Local System	Adversaries may search for and collect sensitive data stored locally on an ESXi host, including configuration files, logs, or virtual machine data (like <code>.vmdk</code> and <code>.vmsn</code> files in <code>/vmfs/volumes</code>). This information can be used for further exploitation or exfiltration.	Allows attacker to correlate and extract relevant data simply as a consolidated package.	Lockdown Rules, Virtual Patching
Collection	T1074	Data Staged	Adversaries may consolidate collected data on an ESXi host before exfiltrating it, often using built-in shell tools to organize VM-related files or logs into a staging directory. On ESXi, this could include aggregating <code>.vmdk</code> , <code>.vmx</code> , or log files within <code>/vmfs/volumes/</code> for easier extraction.	Allows attacker to correlate and extract relevant data simply as a consolidated package.	Lockdown Rules, Virtual Patching
Command and Control	T1071	Application Layer Protocol	Adversaries may use common application layer protocols like SSH, RDP, or SMB to blend command-and-control traffic into legitimate network activity. On ESXi, SSH is frequently targeted, as it's often enabled for remote management and can be abused for stealthy communication between compromised nodes.	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	SSH MFA, Lockdown Rules
Command and Control	T1132	Data Encoding	Adversaries may encode C2 traffic using formats like Base64 or hex to evade detection. On ESXi, encoded payloads may be delivered through SSH sessions or embedded in legitimate-looking commands, making them harder to distinguish without deep inspection.	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	Lockdown Rules, Virtual Patching

Command and Control	T1001	Data Obfuscation	Adversaries may mask C2 traffic on ESXi using techniques like junk padding, misleading headers, or encoding within legitimate protocols (e.g., hiding commands in HTTP cookies or SSH sessions). Since ESXi environments often rely on CLI and SSH for management, malicious traffic can blend in unless closely inspected.	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	Lockdown Rules, Virtual Patching
Command and Control	T1568	Dynamic Resolution	Adversaries may leverage dynamic DNS or algorithmically generated domains (DGA) to sustain C2 connections in ESXi environments, especially when traditional static indicators are blocked. This technique is particularly relevant given ESXi's typical reliance on static IPs and DNS records for management—making dynamic shifts harder to detect in a hypervisor-focused security model.	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	Lockdown Rules, Virtual Patching
Command and Control	T1573	Encrypted Channel	Adversaries may encrypt their C2 traffic using custom or standard algorithms (e.g., TLS) to bypass detection mechanisms. In ESXi environments, encrypted traffic to/from hypervisors or between malicious management interfaces can help attackers evade scrutiny—especially when using trusted ports or protocols (e.g., HTTPS, SSH).	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	Lockdown Rules, Virtual Patching
Command and Control	T1008	Fallback Channels	Adversaries may maintain backup C2 mechanisms to ensure persistent communication if the primary channel is blocked or fails. In ESXi environments, fallback channels may leverage open outbound ports or alternate protocols from compromised VMs or management interfaces to silently reestablish contact—bypassing traditional detections.	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	Lockdown Rules, Virtual Patching
Command and Control	T1665	Hide Infrastructure	Adversaries may disguise their C2 infrastructure using proxy chains, VPNs, spoofed IP ranges, or deceptive domains to avoid detection and prolong persistence. Attackers might exploit trusted hosting or filter traffic from security tools to evade sandboxing and detection, especially when targeting hypervisor management interfaces or vCenter.	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	Lockdown Rules, Virtual Patching
Command and Control	T1105	Ingress Tool Transfer	Adversaries may transfer tools into an ESXi environment using native utilities like wget, curl, or scp, which are available on many ESXi builds. These tools are often used to stage ransomware payloads or lateral movement scripts directly on the hypervisor, bypassing endpoint defenses.	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	Application Filtering, Lockdown Rules
Command and Control	T1104	Multi-Stage Channels	Adversaries targeting ESXi may use multi-stage command and control to evade detection. A lightweight first-stage implant might collect system data or deploy additional scripts, with a second-stage payload providing full control over the hypervisor. Each stage may operate over distinct channels, further complicating detection and response efforts.	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	Application Filtering, Lockdown Rules
Command and Control	T1095	Non-Application Layer Protocol	Adversaries may use non-application layer protocols like ICMP, TCP, or VMCI to establish stealthy C2 communication. In ESXi environments, VMCI (Virtual Machine Communication Interface) provides a unique opportunity for attackers to bypass traditional network monitoring by enabling guest-to-host communication that never leaves the physical machine—making it effectively invisible to most detection tools like tcpdump or Wireshark.	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	Lockdown Rules, Virtual Patching
Command and Control	T1571	Non-Standard Port	Adversaries may send command and control traffic over uncommon port-protocol combinations—like HTTP over TCP 9999 or HTTPS over port 8080—to evade security controls that rely on default port expectations. This helps malware blend in with legitimate network traffic or bypass basic firewall rules. In virtual environments like ESXi, where east-west traffic may not be deeply inspected, this behavior can fly under the radar if not specifically monitored.	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	Lockdown Rules, Virtual Patching

Command and Control	T1572	Protocol Tunneling	Adversaries encapsulate malicious C2 traffic within other protocols (e.g., HTTP, HTTPS, DNS, SSH) to bypass traditional network defenses and reach restricted systems. This can also hide communications by mimicking trusted traffic. Attackers could potentially tunnel VM-to-host or host-to-external traffic—especially through overlooked or unmanaged channels like VMCI or management APIs.	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	Lockdown Rules, Virtual Patching
Command and Control	T1090	Proxy	Adversaries use proxies to relay traffic through intermediary systems—either externally or between internal systems—to disguise their origin, evade detection, or maintain stealthy access. This could involve compromised guest VMs being turned into relay nodes or proxy tools deployed directly on the host.	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	Lockdown Rules, Virtual Patching
Command and Control	T1102	Web Service	Adversaries may abuse legitimate web services—like Dropbox, Google Drive, GitHub, or OneDrive—to host payloads or manage command and control traffic. In an ESXi context, this tactic poses a risk if malicious components (e.g., scripts or second-stage payloads) are fetched from such services using utilities like curl, wget, or VM guest integrations.	Adversarial Persistence within infrastructure environment, enabling further command and control remaining undetected by defenders.	Application Filtering, Lockdown Rules
Exfiltration	T1030	Data Transfer Size Limits	Adversaries may exfiltrate sensitive data in small chunks to fly under the radar of traffic thresholds or anomaly-based detection systems. This technique is relevant in ESXi environments if attackers leverage CLI utilities or scripts to split and exfiltrate logs, configuration files, or virtual disk snapshots.	Sensitive data loss (Including but not limited to customer data, credentials, financial data, etc).	Application Filtering, Lockdown Rules
Exfiltration	T1048	Exfiltration Over Alternative Protocol	This technique involves data exfiltration using a different protocol than the primary C2 channel. On ESXi, attackers might use tools like curl, scp, or even custom scripts to export configuration files, VM snapshots, or logs via FTP, HTTP, or DNS. These channels may bypass existing detection that’s focused on a different protocol.	Sensitive data loss (Including but not limited to customer data, credentials, financial data, etc).	Application Filtering, Lockdown Rules
Exfiltration	T1041	Exfiltration Over C2 Channel	This technique involves attackers piggybacking on an existing command and control (C2) session to exfiltrate data—embedding the stolen data into outbound C2 traffic. This could mean sensitive logs, configuration files, or snapshot data are quietly funneled out using the same channel established for remote access or monitoring, making detection much harder.	Sensitive data loss (Including but not limited to customer data, credentials, financial data, etc).	Lockdown Rules, Virtual Patching
Exfiltration	T1567	Exfiltration Over Web Service	Adversaries bypass traditional detection by uploading sensitive data to legitimate external web services (e.g., Google Drive, Office 365, Mega, or Telegram). This might include logs, configurations, VM images, or snapshot data discreetly transferred via common platforms that are often allowed through corporate firewalls—making the activity blend in with normal traffic.	Sensitive data loss (Including but not limited to customer data, credentials, financial data, etc).	Lockdown Rules, Virtual Patching
Impact	T1531	Account Access Removal	Adversaries may deliberately disable, delete, or modify user accounts on ESXi systems to block access, delay incident response, or prepare for destructive actions like ransomware encryption. This could include removing local shell accounts, disabling SSH logins, or modifying credentials via esxcli.	Devastating effects on targeted endpoint, usually resulting in significant damage to IT resources. Includes the denial/degradation of enterprise services.	SSH MFA, AI-Behavioral Detection, Lockdown Rules
Impact	T1485	Data Destruction	Adversaries may target ESXi environments to irreversibly destroy critical files, virtual machine data, or configuration files. Techniques may include overwriting VMDK files with garbage data or deleting VM directories entirely via shell access. These actions may be performed directly using rm or esxcli, or indirectly through tools dropped on the host.	Devastating effects on targeted endpoint, usually resulting in significant damage to IT resources. Includes the denial/degradation of enterprise services.	SSH MFA, Automated Remediation, AI-Behavioral Detection

Impact	T1486	Data Encrypted for Impact	Ransomware actors increasingly target VMware ESXi environments, encrypting entire virtual machines (VMDK, VMSD, and VMX files) to disrupt business continuity at scale. Because ESXi lacks built-in file-level protections and many security agents don't run natively on the hypervisor, adversaries exploit it as a high-impact chokepoint.	Devastating effects on targeted endpoint, usually resulting in significant damage to IT resources. Includes the denial/degradation of enterprise services.	AI-Behavioral Detection, Automated Remediation
Impact	T1491	Defacement	Defacement on ESXi or related infrastructure often appears in the form of login screen modifications, ransom messages injected into host banners, or tampered admin UIs. This type of attack aims to intimidate, humiliate, or pressure victims and is frequently paired with ransomware campaigns.	Targeted effects on endpoint, usually resulting in potential reputational damage.	Lockdown Rules, Automated Remediation
Impact	T1490	Inhibit System Recovery	Adversaries often delete VM snapshots to prevent recovery, using commands like vim-cmd vmsvc/snapshot.removeall. In more destructive cases, adversaries may also delete the VMDKs themselves or encrypt backup volumes stored on accessible storage.	Devastating effects on targeted endpoint, usually resulting in significant damage to IT resources. Includes the denial/degradation of enterprise services.	AI-Behavioral Detection, Automated Remediation, Lockdown Rules
Impact	T1489	Service Stop	This technique is especially impactful in VMware ESXi environments, where adversaries may shut down virtual machines or kill running VM processes to enable encryption or destruction of files that would otherwise be locked.	Devastating effects on targeted endpoint, usually resulting in significant damage to IT resources. Includes the denial/degradation of enterprise services.	Lockdown Rules, AI-Behavioral Detection
Impact	T1529	System Shutdown/Reboot	Adversaries may reboot or shut down ESXi hosts or guest VMs as a final step to lock users out, interrupt operations, or complete the effects of prior destructive techniques like encryption or snapshot deletion. Commands like esxcli vm process kill, vim-cmd vmsvc/power.off, or esxcli system shutdown are often abused to forcibly terminate systems.	Devastating effects on targeted endpoint, usually resulting in significant damage to IT resources. Includes the denial/degradation of enterprise services.	Lockdown Rules, AI-Behavioral Detection, Automated Remediation