

Ransomware: Defense in Depth with VMware

Table of contents

Introduction	3
What is ransomware?	3
Ransomware distribution scenarios	4
Ransomware distributed through websites or email	4
Ransomware distributed through an infected USB drive	4
Ransomware distributed through a targeted attack	5
NIST controls for ransomware	5
Protect	6
Detect	7
Respond	7
VMware solutions for ransomware	8
Defense in depth	8
Cloud visibility	9
End-user computing solutions	10
Protect	11
Detect	12
Respond	14
Private cloud solutions	15
Protect	15
Detect	18
Respond	20
Multi-cloud solutions	21
CloudHealth Secure State	21
VMware Carbon Black Cloud	22
Modern applications	22
Protect the build process	23
Reduce the attack surface after deployment	23
The role of service mesh	24
Summary	26
Authors	26
Citations	26

Introduction

Cyberattacks increased significantly in 2020, and ransomware is one of the most recurring and devastating attacks. Criminal enterprises using ransomware have collected billions of dollars from infected companies and individuals. The average attack results in companies spending hundreds of thousands of dollars to recover. This cost increases when enterprises pay the ransom. Ransomware attacks typically do not focus on a single industry and are common across all enterprises across the globe. In the United States, the government is now considering enforcing fines for enterprises paying ransoms to discourage the paying of ransom during a ransomware attack. This threat has done little to stop ransomware proliferation. With offerings such as ransomware as a service, it is easier to sabotage an organization by encrypting their systems and data, and deleting or encrypting the organization's backups to negate an established recovery strategy.

All organizations must bolster their defenses against ransomware and other cyberattacks immediately. However, the security technology landscape is complex and presents many challenges in implementing a strong security practice. The VMware intrinsic security approach simplifies the toolsets and processes required to secure end users, private clouds, and public clouds. Intrinsic security is a fundamentally different approach to securing businesses that unifies security and IT teams, and empowers organizations with deep context and insights that accelerate how they identify risk and prevent, detect and respond to threats.

In this white paper, VMware documents solutions by incorporating the National Institute of Standards and Technology (NIST) controls specific to protect, detect and respond functions contained in the NIST Cybersecurity Framework. These controls are further detailed in the NIST SP 1800 series of documents that address data integrity. This paper uses a defense in depth approach and outlines the VMware components and solutions necessary to protect users working from anywhere as well as an organization's resources in a private or public cloud.

What is ransomware?

Ransomware is a type of malware that attempts to deny access to a user's or organization's data, usually by encrypting the data with a cryptographic key known only to the hacker who deployed the malware. The organization's data is held hostage until the ransom is paid. Once ransomware enters a system, it begins encrypting files or complete file systems. It blocks user access until requests for payments, which are often displayed in warning messages, are fulfilled. Unfortunately, there is no guarantee that the cryptographic keys needed to break the encryption will be provided upon payment.

This malware typically enters through malicious downloads, email links, social network messages, and websites. More recently, ransomware has been distributed through aggressive worms and targeted brute force attacks against public-facing remote access services, such as the Remote Desktop Protocol (RDP). Once the end user has executed the malicious content, which often masquerades as legitimate files, the encryption takes place, and a message is displayed demanding a ransom. The ransom note usually threatens permanently losing access to their data, and publicly releasing intellectual property or embarrassing content.

Numerous criminal enterprises using ransomware target companies, nonprofits, governments and schools of all sizes. Some of the most prominent ransomware organizations are DoppelPaymer, Egregor/Maze, NetWalker, REvil and Ryuk. While these criminal enterprises use various strains of ransomware, they have common attack vectors for compromise, such as brute force attempts at public-facing services including RDP, the exploitation of outdated public-facing web software, and known vulnerabilities that may have not been remediated.

Ransomware distribution scenarios

This paper focuses on three specific ransomware scenarios and how VMware solutions protect, detect and respond to attacks against end-user computing, private cloud, and multi-cloud environments.

Ransomware distributed through websites or email

While working from anywhere, a user mistakenly downloads ransomware from an external web server or a link in an email. When the user executes this malicious software, it generates a cryptographic key that is sent back to the external web server. The malware then utilizes a privilege escalation exploit to propagate across the network. The malicious software deletes local backups, shadow copies (Windows), and Time Machine backups (macOS); encrypts files on the machines to which it propagated; and demands payment in exchange for decryption of these files.




FIGURE 1: Ransomware distributed through websites and email.

Ransomware distributed through an infected USB drive

A user finds an unmarked USB device and inserts it into their system. The USB device contains malicious software that may run automatically or with user interaction. The malicious software encrypts the master boot record (MBR) of the machine it is inserted into, then scans the local network looking for vulnerable systems to propagate to. After malware has spread, it reboots the infected systems, and the user is presented with a ransom request.

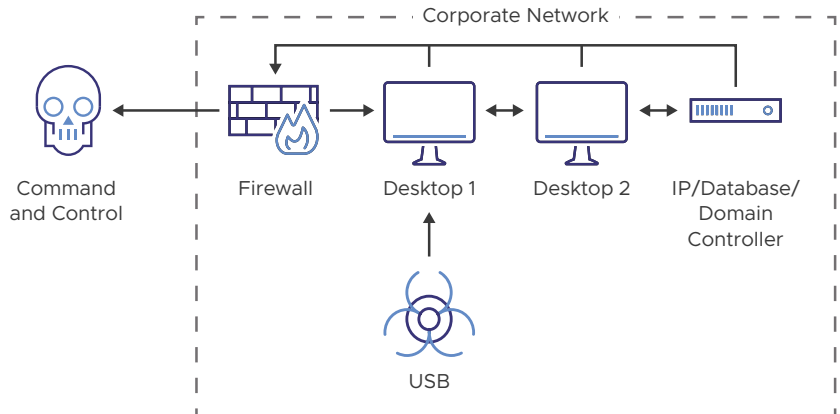


FIGURE 2: Ransomware distributed through an infected USB drive.

Ransomware distributed through a targeted attack

The actors obtain valid user credentials through a brute force attack on publicly accessible RDP. This user's credentials are for a VMware administrator who has reused these credentials on multiple VMware vSphere® hosts. The attacker is able to enable a secure shell (SSH) on the vSphere hosts and upload ransomware to encrypt the VMDK files for all the virtual servers on these hosts. A ransom request is left on the vSphere hosts.

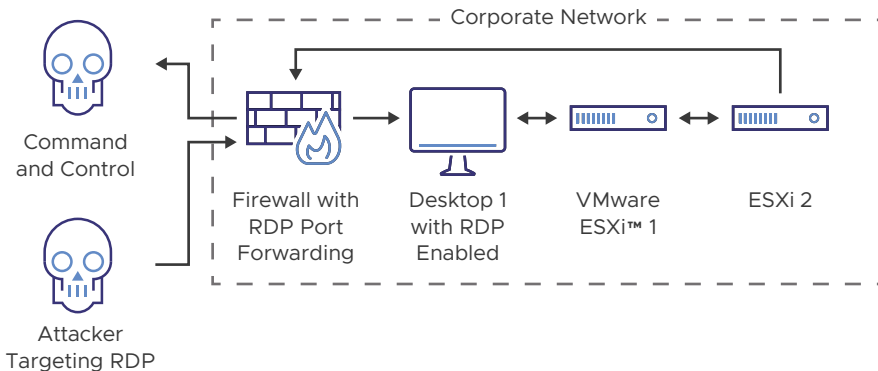


FIGURE 3: Ransomware distributed through a targeted attack.

The solutions sections in this paper are dedicated to a defense in depth strategy for end-user computing, private cloud, and multi-cloud environments. Ransomware propagating through websites or email, an infected USB drive, or a targeted attack are addressed in the [End-user computing solutions](#) section. If the attacker has gained access beyond end users into a private or public cloud network, the capabilities documented will be necessary to protect, detect and respond to ransomware in the [Private cloud solutions](#) and [Multi-cloud solutions](#) sections.

NIST controls for ransomware

The VMware solutions in this paper follow the functions, categories and subcategories documented in the NIST publications addressing data integrity:

- SP 1800-25 – Identifying and Protecting Assets Against Ransomware and Other Destructive Events
- SP 1800-26 – Detecting and Responding to Ransomware and Other Destructive Events
- SP 1800-11 – Recovering from Ransomware and Other Destructive Events

These NIST publications do not focus solely on ransomware. For the purposes of this paper, only the scenarios and controls specific to ransomware within these three publications are used in documenting ransomware-specific capabilities in VMware solutions.

The NIST approach to cybersecurity is documented in the NIST Cybersecurity Framework. The framework core contains five functions:¹

- Identify – develop an organizational understanding to manage cybersecurity risk to systems, people, assets, data, and capabilities.
- Protect – develop and implement appropriate safeguards to ensure delivery of critical services.
- Detect – develop and implement appropriate activities to identify the occurrence of a cybersecurity incident.
- Respond – develop and implement appropriate activities to take action regarding a detected cybersecurity incident.
- Recover – develop and implement appropriate activities to maintain plans for resilience and to restore any capabilities or services that were impaired due to a cybersecurity event.

To learn more about the framework, read the [Framework for Improving Critical Infrastructure Cybersecurity](#).

The following subsections specify the categories and subcategories in SP 1800-25, SP 1800-26, and SP 1800-11 supported by VMware solutions.

Protect^{2,3}

Identity management, authentication, and access control:

- PR.AC-1 – Identities and credentials are issued, managed, verified, revoked, and audited for authorized devices, users, and processes

Data security:

- PR.DS-1 – Data-at-rest is protected
- PR.DS-6 – Integrity checking mechanisms are used to verify software, firmware, and information integrity

Information protection, processes, and procedures:

- PR.IP-1 – A baseline configuration of information technology/industrial control systems is created and maintained incorporating security principles (e.g., concept of least functionality)
- PR.IP-3 – Configuration change control processes are in place
- PR.IP-4 – Backups of information are conducted, maintained, and tested periodically
- PR.IP-9 – Response plans (Incident Response and Business Continuity) and recovery plans (Incident Recovery and Disaster Recovery) are in place and managed

Maintenance:

- PR.MA-1 – Maintenance and repair of organizational assets are performed and logged, with approved and controlled tools
- Protective technology:
- PR.PT-1 – Audit/log records are determined, documented, implemented, and reviewed in accordance with policy

Detect^{2,3}

Anomalies and events:

- DE.AE-1 – A baseline of network operations and expected data flows for users and systems is established and managed
- DE.AE-2 – Detected events are analyzed to understand attack targets and methods
- DE.AE-3 – Event data are collected and correlated from multiple sources and sensors
- DE.AE-4 – Impact of events is determined
- DE.AE-5 – Incident alert thresholds are established

Security continuous monitoring:

- DE.CM-1 – The network is monitored to detect potential cybersecurity events
- DE.CM-3 – Personnel activity is monitored to detect potential cybersecurity events
- DE.CM-4 – Malicious code is detected
- DE.CM-5 – Unauthorized mobile code is detected
- DE.CM-7 – Monitoring for unauthorized personnel, connections, devices, and software is performed

Detection processes:

- DE.DP-2 – Detection activities comply with all applicable requirements

Respond^{2,3}

Response planning:

- RS.RP-1 – Response plan is executed during or after an incident

Communications:

- RS.CO-2 – Incidents are reported consistent with established criteria

Analysis:

- RS.AN-1 – Notifications from detection systems are investigated
- RS.AN-2 – The impact of the incident is understood
- RS.AN-3 – Forensics are performed
- RS.AN-4 – Incidents are categorized consistent with response plans

Mitigation:

- RS.MI-1 – Incidents are contained
- RS.MI-2 – Incidents are mitigated
- RS.MI-3 – Newly identified vulnerabilities are mitigated or documented as accepted risks

While the NIST publications do not directly address the recover function for ransomware, VMware does provide a backup and recovery solution well positioned to securely recover from a ransomware event. This solution is documented in the [Multi-cloud solutions](#) section. The ransomware capabilities in this paper should be implemented in addition to a complete cybersecurity defense in depth approach. Defense in depth best practices are documented in the next section.

VMware solutions for ransomware

Defense in depth

Defense in depth is defined as deploying multiple layers of defense across endpoints and public and private clouds to protect an organization from cybersecurity events. In this section, recommendations are documented on how to implement defense in depth to better secure organizations from ransomware attacks. While not specifically addressing the issue of ransomware, these recommendations for a defense in depth posture will provide a strong security foundation. It is recommended to follow these defense in depth best practices:

- Awareness and training program – End users are top targets. Everyone in the organization must be aware of the threat of ransomware and how it is delivered.
- Email scanning – Content scanning and email filtering should be used to detect threats before they reach end users.
- Spam filters – These prevent phishing emails from reaching end users. Also, spam filters authenticate inbound emails using technologies such as Sender Policy Framework and DomainKeys Identified Mail.
- Block ads – Ransomware is often distributed through malicious ads served when visiting certain sites. Blocking ads can reduce that risk.
- Firewalls – These software or hardware appliances control network traffic through access or deny policies or rules. These rules include denylisting or allowlisting IP addresses, MAC addresses, and ports. There are also application-specific firewalls, such as web application firewalls (WAFs) and secure email gateways, that focus on detecting malicious activity directed at a particular application.⁴
- Intrusion detection or prevention systems (IDS/IPS) – An IDS sends an alert when malicious network traffic is detected, whereas an IPS attempts to prevent and alert on identified malicious activity on the network or a user's workstation. These solutions base recognition of attacks on signatures of known malicious network activity.⁴
- Endpoint detection and response (EDR) – This is software or agents that reside on the client system (e.g., a user's laptop or mobile phone) and provide antivirus protection, alert, detection, analysis, threat triage, and threat intelligence capabilities. These solutions run on rulesets (i.e., signatures or firewall rules) or heuristics (i.e., detection of anomalous or malicious behaviors).⁴
- Network segmentation – This is the practice of splitting a network into multiple sub-networks designed around business needs and technology requirements. This might include different sub-networks for executives, finance, operations, and human resources. Depending on the level of security required, these networks may not be able to communicate directly. Segmentation is often accomplished through the use of network switches or firewall rules.⁴
- Inspect east-west internal traffic – This provides anomaly detection of certificates when traffic is encrypted.
- Inspect north-south traffic – This detects command and control (C&C) traffic by using threat intelligence to identify malicious IPs, domains and more.
- Scan network artifacts – This dynamically analyzes file behaviors for threats by using AI to detect malicious code.
- Log aggregation – This collects logs from all critical devices, security controls, and endpoints in a central location for correlation and analysis.
- The principle of least privilege – This requires policy and technical controls to only assign users, systems, and processes access to resources (networks, systems and files) that are absolutely necessary to perform their assigned function.⁴

- Strong passwords – These are a critical authentication mechanism in information security. Modern password guidance involves using multifactor authentication for any account of value, using a phrase with multiple words, and not reusing passwords.⁴
- Patch management – This is the process of applying updates to an operating system, software, hardware or plug-in. Often, these patches address identified vulnerabilities that could allow cyberthreat actors unauthorized access to information systems or networks.⁴
- Back up data regularly – This verifies the integrity of backups and tests the restoration process to ensure it's working.
- Secure your offline backups – This ensures backups are not permanently connected to the computers and networks they are backing up.
- Conduct an annual penetration test and vulnerability assessment.

Cloud visibility

Log aggregation and traffic inspection are key concepts for defense in depth. VMware provides a number of solutions that perform these functions for end-user computing, private clouds, and multi-cloud environments: VMware vRealize® Network Insight™, vRealize Log Insight™, and vRealize Operations™.

VMware vRealize Network Insight

VMware vRealize Network Insight delivers intelligent operations for software-defined networking and security. Through its innate abilities to discover applications, observe network traffic, or plan micro-segmentation, vRealize Network Insight can be crucial when architecting the security strategy of an organization. Security personnel can audit any changes to the network infrastructure, understand the topology of applications to identify any risks, or analyze traffic after an incident to isolate an attack. Recommendations for VMware NSX® firewall rules and security group memberships can easily be extracted from the solution. Moreover, vRealize Network Insight makes these processes repeatable, so the same security patterns can be applied across the entire infrastructure.

vRealize Network Insight contributes to preparing for a more adequate reaction to ransomware attacks. By analyzing east-west and north-south traffic, it provides valuable information about areas of security improvements, such as missing firewall rules that can allow ransomware to spread or leak information to the outside. Two major capabilities are identifying objects that should not belong in the environment and searching for actors accessing resources they should not be able to access.

VMware vRealize Log Insight

VMware vRealize Log Insight offers heterogenous and highly scalable log management with intuitive dashboards and sophisticated analytics. It integrates with vSphere, NSX, VMware vSAN™, vRealize Operations, vRealize Automation™, and *VMware Horizon®*. vRealize Log Insight also integrates with many third-party solutions, such as F5, Extreme Networks, HPE OneView, Infoblox, and more. By collecting logs from across the infrastructure and utilizing analytics features, administrators are able to correlate various events that are difficult to associate otherwise. This capability gives tremendous opportunities during attack investigations and makes the process of log retrieval from months ago a simple and quick action. If logs are needed for auditing purposes, tracking a certain incident among various platforms and solutions is achievable quickly with customizable dashboards. Alerts will be triggered when security events happen.

Trusted indicators of ransomware attacks are log messages notifying about lost or prevented access to some resources. For example, when an application running on a server is prevented access to its database and this event is logged, vRealize Log Insight filters can quickly alert administrators and show the event on a dashboard. By leveraging filtering and field extraction, investigators are able to pinpoint the origin of the attack. Investigators can trace the spread of an attack to see how it behaves in a test environment.

VMware vRealize Operations

VMware vRealize Operations delivers insight about past events and conditions as well as the future state of the infrastructure. These capabilities lower risk, add predictability, and increase operational efficiency. By collecting information from various infrastructure and application platforms, it provides a holistic picture of how components interact with each other, where vulnerabilities might occur, and how to improve defenses. Customizable alerts and automatic remediation in vRealize Operations enables security administrators to ensure proper and adequate reaction during emergencies. vRealize Operations further reduces risk by enforcing IT regulatory standards for VMware Cloud Foundation™ and VMware Cloud™ on AWS with integrated compliance and automated drift remediation. It ensures the environment's adherence to standards such as the Payment Card Industry (PCI), the Health Insurance Portability and Accountability Act (HIPAA), or the Sarbanes-Oxley Act (SOX).

Similar to vRealize Log Insight, vRealize Operations can create an alert when a resource's status or behavior has been altered. vRealize Operations integrated compliance enables administrators to look for any changes in a system dealing with privacy information, both in terms of OS settings as well as the underlying infrastructure. Inventorying critical systems and showing the results in dashboards and reports allows administrators to easily control what applications are running. Backup of data is crucial, especially when ransomware is a risk. With the vast integrations available, vRealize Operations can alert on and self-remediate any problems with backups.

Going offline when a ransomware attack emerges is not an option for most organizations, and vRealize Operations helps by running what-if scenarios and budgeting public cloud costs in case of emergencies.

End-user computing solutions

As organizations move to a future-ready workforce with their employees working from anywhere, the need to protect their laptop and desktop endpoint devices from ransomware has never been greater. The protection offered by a corporate firewall has gone away, and these endpoints are often connected directly to the internet. Also, new risks are present, such as a family member casually using the employee's laptop to browse the web and accidentally infecting it.

VMware has the ability to defend both physical endpoints and virtual desktops against ransomware attacks with a modern, cloud-first, platform approach utilizing [VMware Workspace ONE® UEM](#), [VMware Carbon Black Cloud™](#), [VMware Workspace ONE Intelligence™](#), and [Horizon](#). This intrinsic security solution is called [VMware Workspace Security™](#).

Horizon virtual desktops defend against ransomware attacks by being isolated from the user's endpoint. Also, non-persistent desktops are wiped after each use to prevent propagation. While an unprotected user might be affected by ransomware on their physical endpoint, this wouldn't transfer into the secure Horizon solution. Horizon deployments are protected by the defense methods detailed in the [Private cloud solutions](#) and [Multi-cloud solutions](#) sections. [Secure cloning](#) from VMware Carbon Black Cloud ensures policies are inherited from the golden image, regardless of where the virtual desktop infrastructure is deployed. Horizon also uses policies to block USB drives or just specific types, such as mass storage in virtual desktops. Monitoring and access control of USB storage are also available for physical devices using VMware Carbon Black Cloud.

The following subsections provide a detailed look at defending against ransomware attacks on physical Windows 7, Windows 10, and macOS endpoints.

Protect

The first step is to protect endpoints by installing the VMware Carbon Black Cloud Sensor. When enrolling Windows 10 and macOS endpoints, the sensor is automatically installed, always running, and always up to date. There is nothing for an end user to do here, and this is a key part of VMware Workspace Security.

For Windows 7 endpoints, a management tool can install the sensor, or the end user can install it manually.

The sensor runs continuously on the endpoint, which can be enforced by Workspace ONE UEM and can't be removed by the end user, even with administrator privileges, when enforced by VMware Carbon Black Cloud. VMware Carbon Black Cloud offers a number of ways to protect against ransomware in real time. These methods provide visibility into, detection of, and blocking of ransomware attacks along the kill chain. They include monitoring for malicious signatures, advanced detections, machine learning, and behavioral analysis.

To enable this protection, VMware Carbon Black Cloud collects a large amount of data on every process execution. This includes:

- USB mass storage devices attached
- USB mass storage devices' metadata
- Parent and child relationships between processes
- Cross-process events
- Inbound and outbound network connections
- File modifications
- Registry modifications
- Module loads
- Script loads

The telemetry that VMware Carbon Black Cloud collects is available in the console UI and can be forwarded to a security information and event management (SIEM) system or data lake for further analysis.

VMware Carbon Black Cloud provides all these capabilities with a single, lightweight endpoint sensor and a single console UI without scheduled scans by monitoring the activity of all binary executions in real time. VMware Carbon Black Cloud detects and blocks ransomware and other more advanced attacks without impacting the performance of the endpoint.

With VMware Carbon Black Cloud, organizations maintain visibility, protection and control over their employees' endpoints, even if they are disconnected from the corporate network.

Detect

Distribution scenario 1: Ransomware distributed through websites or email

Known malicious IP addresses can be detected and alerted on by VMware Carbon Black Cloud, giving an initial indication that an attack is underway. This detection happens regardless of whether the endpoint is on the corporate network or not. The known malicious URLs or IP addresses can be identified using out-of-the-box threat intelligence from VMware Carbon Black Cloud that includes feeds from Facebook ThreatExchange and AlienVault. In addition to the VMware Carbon Black Cloud threat feed, organizations can ingest their own threat feed into VMware Carbon Black Cloud.

Communication from the compromised machine back to C&C would be detected as an anomalous activity by the network traffic analysis (NTA) capability of NSX Network Detection and Response™. NSX Distributed IDS/IPS™ would use signature-based detections to uncover ransomware in the network traffic before it reached the endpoint. NSX Advanced Threat Analyzer™ can also sandbox potential malware to detect its malicious payload.

If the IP address is unknown and the malicious payload is downloaded, then VMware Carbon Black Cloud can stop destructive actions early on in the kill chain, as described in the [Detection techniques](#) subsection.

Any lateral movement from the compromised machine to another or communication with C&C servers could be detected by NSX Distributed IDS/IPS and NTA.

Distribution scenario 2: Ransomware distributed through an infected USB drive

Workspace ONE UEM restriction profiles can block the use of attaching USB mass storage to endpoints.

If USB devices are not completely blocked, all unapproved USB mass storage devices can be blocked and an alert generated according to the policy by VMware Carbon Black Cloud. USB mass storage devices can be approved by product ID, vendor ID, or serial number.

If the device was in a policy that allowed all USB mass storage devices and the malicious binary was executed, then VMware Carbon Black Cloud can stop destructive actions early on in the kill chain, as described in the [Detection techniques](#) subsection.

Any lateral movement from the compromised machine to another or communication with C&C servers could be detected by NSX Distributed IDS/IPS and NTA.

Distribution scenario 3: Ransomware distributed through a targeted attack

An attacker targeting RDP from outside the corporate network can be detected by identifying known malicious IP addresses from the threat feed. This attack is alerted on by VMware Carbon Black Cloud, giving an initial indication that an attack is underway.

If the attacker copies their tools and ransomware to the endpoint they are connected to, then VMware Carbon Black Cloud can stop destructive actions early on in the kill chain, as described in the [Detection techniques](#) subsection.

Any lateral movement from the compromised machine to another or communication with C&C servers could be detected by NSX Distributed IDS/IPS and NTA.

Detection techniques

All new binaries deployed to the end-user device are checked against malicious signatures and a reputation database in VMware Carbon Black Cloud, which is made up of more than 30 threat feeds. If the binary is known to be malicious, it can be blocked through policy. The ransomware binaries can also be detected by NSX Advanced Threat Analyzer, which leverages the most advanced network sandbox with a unique isolation and inspection environment that emulates the entire host—including CPU, system memory, and all input and output devices—while interacting with malware to safely analyze all behaviors. NSX Advanced Threat Analyzer provides the most complete malware analysis of artifacts traversing your data center and enables accurate detection and prevention of advanced threats, including ransomware and zero-day attacks.

VMware Carbon Black Cloud also integrates with the Microsoft Antimalware Scan Interface (AMSI) to leverage high-fidelity, behavior-based detections created by the VMware Threat Analysis Unit™ to block ransomware in its tracks by detecting the exploit kits commonly used by strains of ransomware, such as Cobalt Strike, PowerShell Empire, and Metasploit. By detecting these exploit kits, VMware Carbon Black Cloud can shut down the ransomware attack before it is able to do any damage.

VMware Carbon Black Cloud uses machine learning to conduct static analysis of the binary to determine if it is suspected to be malware. This static analysis looks at the metadata embedded in the binary as well as human-readable strings for indicators that it is malicious. If the binary is judged to be suspected malware, it can be blocked through policy.

If the binary has not been blocked at this point, then it will be allowed to start executing. However, the VMware Carbon Black Cloud detections and blocking do not stop there. Ransomware will typically attempt to delete local backups, such as shadow copies on Windows or Time Machine backups on macOS, to prevent users from easily recovering from the attack. VMware Carbon Black Cloud monitors access to utilities to delete these backups and can block these attempts by policy and terminate the calling process.

VMware Carbon Black Cloud Sensor writes canary files in strategic locations on the file system of an endpoint. If a binary attempts to modify or encrypt these files, the binary will be terminated according to policy.

Also, part of the behavioral analysis is a heuristic detection for ransomware that looks at the number and type of files being accessed in a short time window. This detection will terminate the offending process according to policy. This capability is currently only available on macOS endpoints.

Organizations can also create custom detections in the form of watchlists. These are searches that are run against all new data seen by VMware Carbon Black Cloud. A watchlist is made up of any number of reports, and the reports are made up of one or more indicators of compromise (IoCs). IoCs can be rated by severity from 1 to 10, with 10 being the most severe. Watchlists can be set to create alerts when they detect a hit.

The opt-in cloud analysis feature sends the binary to VMware Carbon Black Cloud and then to our partner, Avira, for sandboxing. This feature is only currently available for Windows endpoints.

The opt-in unified binary store collects a single copy of all newly seen binaries in an organization's infrastructure. Alerts can be created in the form of watchlists for YARA hits. YARA is a tool used by researchers and analysts to identify and classify malware through complex signatures, but blocking files identified by YARA is not currently supported.

NSX can apply Layer 7 distributed firewall policies to all desktops and workloads to allow only specific ports/protocols and drop the remaining traffic. Lateral movement can be greatly reduced by blocking desktops from communicating to other desktops and segmenting communication to specific workloads. The NSX Identity Firewall can also be leveraged for remote desktop users to only permit certain users from accessing appropriate workloads.

Respond

Once a ransomware attack is discovered and/or blocked, the security team is automatically notified so they can ensure the attack does not spread from the impacted endpoint and that no data is exfiltrated.

If the attack is determined to be serious, then the quarantine feature in VMware Carbon Black Cloud can isolate the endpoint, which prevents it from communicating anywhere except for VMware Carbon Black Cloud. The quarantined endpoint will still collect and report telemetry, so there are no blind spots in the investigation.

When using Workspace ONE Intelligence automation as part of VMware Workspace Security, this quarantining process is automatic for Windows 10 and macOS endpoints. This automation can also perform other useful actions, such as tagging the endpoint as at risk in Workspace ONE UEM; notifying the end user by email; notifying the security team in *Slack*, *Microsoft Teams*, and the like; and creating an incident ticket in *ServiceNow*, *Remedy*, or other IT service management tools. The notification contains a direct URL to the attacked endpoint in the VMware Carbon Black Cloud console so it can be immediately investigated.

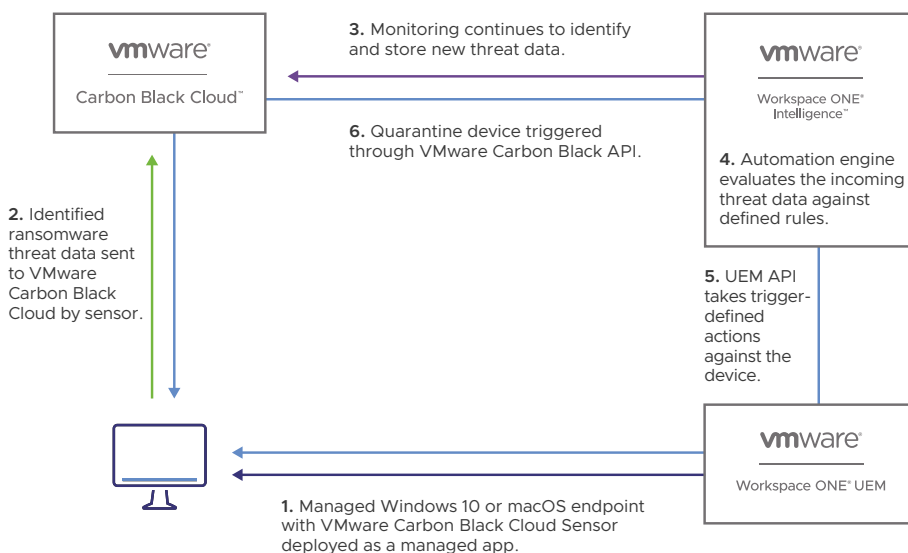


FIGURE 4: VMware Workspace Security automatically quarantines an endpoint when ransomware is detected.

VMware Carbon Black Cloud enables security teams to obtain a remote, secure shell on an endpoint to gather additional forensic artifacts and take remedial actions, such as:

- Pushing and executing scripts or third-party tools to the endpoint
- Running any OS commands
- Downloading files
- Killing processes
- Querying and modifying the registry
- Deleting files
- Dumping memory

This shell can be obtained even if the endpoint is in quarantine.

VMware Carbon Black Cloud can also automatically delete known malware in a configurable window of one day to four months. This known malware will not be allowed to execute before it is deleted according to policy. This window is to help security teams ensure that false positives are not being deleted.

VMware Carbon Black Cloud classifies malicious activity by MITRE ATT&CK technique IDs (TIDs). It also generates an alert visualization to help quickly understand the scope of the attack, establish the root cause, and prioritize response.

Finally, actions that can be taken in the VMware Carbon Black Cloud UI can also be taken by leveraging the API. The VMware Carbon Black Cloud API is a robust, open, two-way API that allows security teams to script remedial actions, create automation, and leverage other VMware products and third-party integrations.

If the endpoint is attached to the corporate network, NSX can assist in the response. NSX classifies malicious activity by MITRE ATT&CK stage. It also generates a dynamic intrusion blueprint and a detailed timeline of a threat to quickly understand the scope of the attack and prioritize response.

Private cloud solutions

A private cloud is a cloud infrastructure provisioned for exclusive use by a single organization, often comprised of multiple business units. This section documents VMware solutions and capabilities specific to private clouds that allow organizations to protect, detect and respond to ransomware. This section focuses specifically on vSphere, VMware Carbon Black Cloud, and NSX. vRealize Network Insight, vRealize Log Insight, and vRealize Automation—as documented in the [Defense in depth](#) section—are also critical to detecting ransomware events.

Protect

An organization must first focus on protecting their private cloud based on vSphere. This includes implementing security best practices in addition to deploying and configuring VMware solutions and features. Ransomware attackers target vSphere hosts or individual virtual machines (VMs) after gaining access to the environment, as outlined in the [Ransomware distribution scenarios](#) section. It is important to protect vSphere hosts and VMs as part of a defense in depth strategy. The following information should be implemented in all private clouds to improve an organization's security posture.

Protect vSphere from ransomware

The proper configuration of vSphere is critical to protecting the private cloud from ransomware attacks. VMware vCenter Server® must be deployed to a separate management network. Access to the VMware vCenter® management network must only be allowed from a jump server. Multifactor authentication (MFA) must be implemented for the jump server. Connections to the jump server must only be allowed from a specific IP range by a solution such as the NSX micro-segmentation rules. Disable all non-necessary ports and protocols in the vSphere environment by default. One of the most important services that must be disabled is SSH on vSphere hosts. After deploying the vSphere hosts, implement the configuration items from the vSphere Security Configuration Guide, and enable Secure Boot and vSphere Trust Authority™ on vSphere hosts. MFA must also be implemented for accounts with access to vCenter.

Once the vSphere hosts are deployed and configured as documented in the previous paragraph, a complete backup and recovery solution must also be implemented for the vSphere environment. Many backup and restore solutions support vSphere and VMs. It is recommended to implement a solution that supports MFA. The backup storage solution should support version control. Organizations must implement different Active Directory accounts for vSphere administrations and backup administrators. Write-once, read-many (WORM) devices must be used for immutable storage. Utilizing an air-gapped backup service is critical as ransomware attackers first delete backups to prevent organizations from being able to recover. Figure 5 shows the seven major concepts for deploying and running ransomware-resistant vSphere.

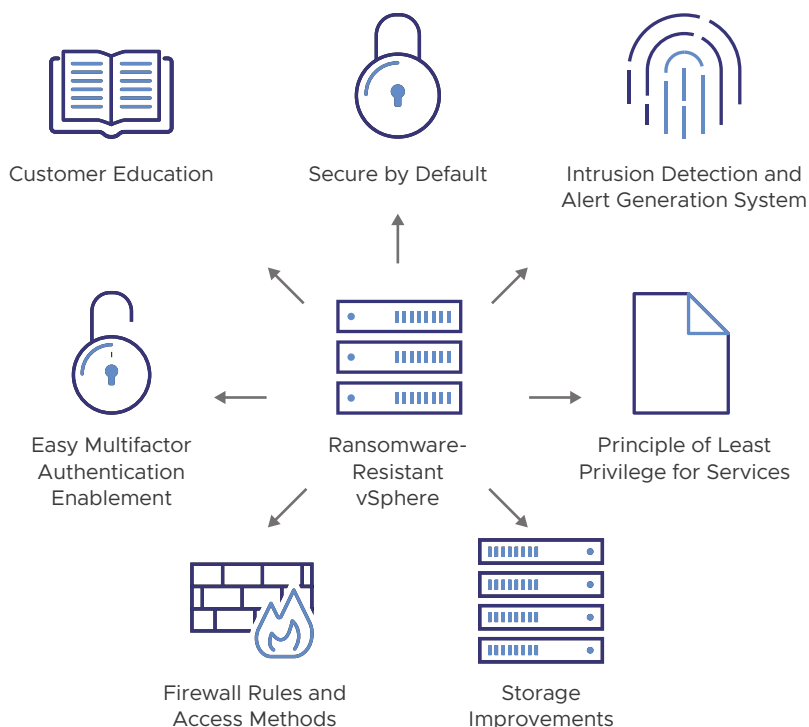


FIGURE 5: The seven concepts for deploying and running a ransomware-resistant vSphere.

Once the private cloud based on vSphere is deployed, additional best practices must be followed to protect the environment. These best practices include leveraging industry standards to protect the vSphere environment, applying security updates, limiting access, and protecting VMs.

Industry standards and benchmarks

There are several industry-accepted standards and benchmarks that an organization can use to develop a hardened vSphere image. The Center for Internet Security (CIS) provides benchmarks that can be applied against a vSphere host to develop a hardened image. Additionally, NIST has provided general guidance in SP 800-53 – Security and Privacy Controls for Information Systems and Organizations. The Security Technical Implementation Guide (STIG) can also be used as a guideline to harden vSphere. The guidance in the industry-standard documents can be used to protect vSphere and align to an organization’s specific security and operational needs. By deploying hardened images as the basis of the vSphere infrastructure, an organization can ensure the foundational components of the virtual environment are as secure as possible from the moment of deployment. VMware provides automated tools that allow an organization to measure compliance against both the CIS benchmarks for vSphere and overall NIST SP 800-53 compliance. The deployment of hardened images directly supports the NIST protect function’s information protection processes and procedures category (PR-IP-1).

Updates and patching

One of the key concepts within a defense in depth posture is to establish a regular patching schedule. The NIST Cybersecurity Framework speaks directly to this need for a regular patching schedule in the protect function's maintenance category (PR.MA-1) and the respond function's mitigation category (RS.MI-3). Both functions are foundational for a valid defense in depth posture. VMware produces patches for quality updates as well as for security-related updates on a regular basis. VMware will issue emergency security patches when necessary. It is recommended that an organization develop an update capability that encompasses regularly scheduled patches as well as has an allowance for emergency patching of VMware products. When an organization leverages VMware Cloud Foundation, patching capabilities are integrated to the VMware Cloud Foundation SDDC Manager™. If the organization is not leveraging VMware Cloud Foundation, patching will be performed with VMware Update Manager for vSphere hosts. The latest security advisories issued by VMware can be found at the [VMware Security Advisories page](#). It is recommended that an organization subscribe to the VMware Security Advisories feed.

Limit access

Another key concept within the defense in depth discussion is to ensure that the vSphere environment is protected from unauthorized access. The NIST Cybersecurity Framework speaks directly to this need for securing the foundation of the virtualized environment under the protect function's identity management, authentication, and access control category (PR.AC-1). Both the vSphere hosts and the vCenter Server instance have a firewall capability that can be configured manually or programmatically. This is the foundation of a strong defense in depth posture within the VMware environment. Both vSphere hosts and vCenter must be protected utilizing the built-in firewall. To ensure visibility, vSphere hosts and vCenter should forward their firewall logs to a central logging collection point. This central logging capability directly supports the NIST Cybersecurity Framework, specifically the detect function's security continuous monitoring category (DE.CM-1, DE.CM-7).

Virtual machine protection

VMs must be protected in addition to the vSphere hosts and vCenter. It is important to have granular visibility into the guest OS in the event attackers gain access to the VMs. As documented in the [End-user computing solutions](#) section, VMware Carbon Black Cloud offers this granular level of visibility and advanced prevention mechanisms.

VMware Carbon Black Cloud integrates with vSphere for an agentless experience. The sensor can be deployed by the click of a button in vCenter. In addition to the visibility provided to a security analyst by VMware Carbon Black Cloud, a vSphere administrator has visibility from within vSphere.

Protect with NSX

NSX running in the private cloud provides identification and alignment of workload context to security policies. Characteristics of each workload—such as name, OS, user-defined tags, and more—are exposed to security policy management, allowing for proper deployment of the correct security policies to the correct workloads. End user identity via Active Directory is also configurable as an additional area of identification that can be aligned to security policies.

Purpose-built to protect the data center and powered by machine learning, the NSX Service-defined Firewall™ with advanced threat prevention capabilities offers the industry's highest fidelity insight into advanced threats. The combination of distributed firewalls, network traffic analysis, intrusion detection and prevention, and advanced malware analysis provides comprehensive network detection and response capabilities.

As part of an NSX deployment, NSX Intelligence™ provides a graphical visualization of the data center components, such as groups, VMs, and network traffic flows. NSX Intelligence also creates recommendations for security policies, policy security groups, and services for applications. The recommendations assist with the implementation of micro-segmentation at the application level. This enables an organization to enforce a dynamic security policy by correlating traffic patterns of communication that occur between workloads in the data center.

The NSX Distributed Firewall is a Layer 7 stateful firewall implemented in the data center host hypervisors, providing a firewall for each workload without the need for the deployment of an agent. The firewall policies implemented provide segmentation capabilities for data center workloads. Lateral movement can be reduced or eliminated leveraging the NSX Distributed Firewall. Every packet in/out of every workload is inspected by the firewall, including those that would be sent in the same Layer 2 segment. This allows NSX to block traffic between workloads on the same Layer 2 segment without changing the underlying networking. When micro-segmentation is implemented to block unnecessary inbound and outbound traffic for VMs and the vSphere environment, the possible attack vectors for ransomware are significantly reduced or eliminated.

Detect

If ransomware has infiltrated the private cloud, vSphere, NSX, and VMware Carbon Black Cloud will alert security operations.

Detect with vSphere

vSphere logs and alerts on hundreds of objects directly in the vCenter console. To detect possible ransomware attacks, the following items should be monitored:

- Failed login attempts
- Changes to accounts, including service accounts
- Datastore activity similar to ransomware encryption attacks

These alerts can be configured to notify vSphere administrators within the vCenter console, send alerts to email or systems such as PagerDuty, and be collected by vRealize Log Insight as part of a comprehensive defense in depth security practice.

Refer to the [Detect with VMware Carbon Black Cloud](#) subsection for information on how the integration with VMware Carbon Black Cloud can provide visibility into vulnerabilities in an organization's private cloud.

Detect with NSX

NSX monitors traffic on the network for known threats and potential threats by performing data analysis on traffic patterns. There are many functions available in NSX to detect malicious traffic in the private cloud.

The NSX Distributed Firewall is a component that runs directly on vSphere hosts. VM traffic is analyzed before it leaves the host. In-bound traffic to VMs is also analyzed for malicious content. NSX Distributed IDS/IPS identifies known threats in all east-west traffic using curated signatures based on precise application context.

Potential ransomware is detected by NSX Advanced Threat Analyzer. The VMware Threat Analysis Unit continuously updates NSX Network Detection and Response in real time with threat intelligence, such as active C&C servers, objects with zero-day exploits, toxic websites and malware distribution points, and malware information useful to defend against threats specific to your organization. The patented NSX Advanced Threat Analyzer acts as a network sandbox to deconstruct behaviors engineered into a file or URL to determine if it is malicious. NSX Advanced Threat Analyzer sees all instructions that a program executes, all memory content, and all OS activity. Network traffic analysis utilizes machine learning to detect protocol and traffic anomalies, as well as malware downloads to the compromised machine. Lateral movement from the compromised machine to another would be detected by NSX Distributed IDS/IPS and NTA. Any communication for malicious intent, such as to C&C servers, or for data exfiltration would be detected by NTA. In addition to the historical event features available in NSX, audit and event records are configurable to be logged to external sources based on company policy.

Security operations center (SOC) teams are often overwhelmed by the high volume of low-fidelity alerts generated by their security controls. The unique combination of NTA, IDS/IPS and threat analysis reduces false positives by up to 90 percent and provides unmatched visibility. The result is that NSX Network Detection and Response condenses massive amounts of network data down to a just a handful of high-fidelity alerts. Security analysts can focus on solving valid incidents and proactively protecting the organization rather than responding to false positives.

Detect with VMware Carbon Black Cloud

VMware Carbon Black Cloud provides granular-level visibility into guest operating systems (Windows, macOS and Linux) in an organization's private cloud, which is detailed in the [Detect subsection of the End-user computing solutions](#) section.

The integration with vSphere enables organizations to detect and prioritize response to the most vulnerable systems in their private cloud. The types of data available in vSphere and VMware Carbon Black Cloud in regards to vulnerabilities are:

- Appliance health
- Inventory status
- Affected assets
- Critical product vulnerabilities

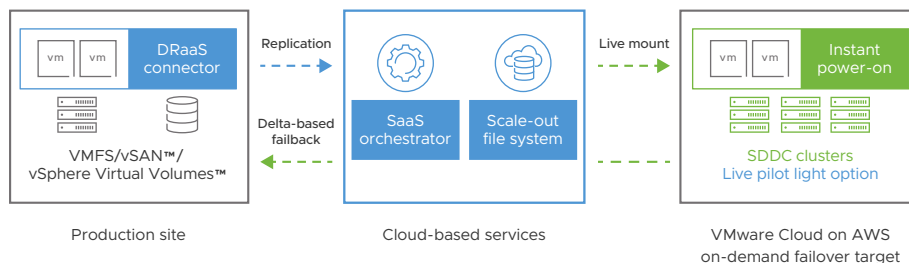
The vulnerabilities reported in both VMware Carbon Black Cloud and vSphere are based on the [Common Vulnerability Scoring System](#) (CVSS) and augmented by the VMware partnership and integration with [Kenna Security](#). Kenna reviews the CVSS to determine if the vulnerability is easily exploitable, if exploit kits have plug-ins for the vulnerability, and if the vulnerability is actively being exploited. With this added context, Kenna can provide a much more accurate score to help focus security, IT and vulnerability management teams on the most exploitable threats in their private cloud. This score is updated every six hours, and vulnerabilities are collected every 24 hours.

Respond

After a possible ransomware event is detected by vSphere, NSX, or VMware Carbon Black Cloud, security operations teams must respond to the event quickly to quarantine the affected systems. If attackers gain access to vSphere in the private cloud, they can encrypt VMs by accessing vSphere hosts through SSH and demand a ransom payment. At this point, the only course of action is to restore VMs from backup. One of the first actions attackers take when deploying ransomware is deleting backups. This step increases the likelihood that the organization must pay the ransom as it would be the only way to regain access to their systems without backups. However, paying the ransom does not guarantee the bad actors will decrypt the machines. Therefore, it is important to implement an air-gapped backup and recovery strategy for all systems in the environment, especially VMs. VMware Cloud Disaster Recovery™ meets these requirements and more to provide a quick and simple way to recover VMs in VMware Cloud on AWS as operations teams secure and rebuild the private cloud environment.

Respond with VMware Cloud Disaster Recovery

VMware Cloud Disaster Recovery enables customers to quickly and cost-effectively enable multiple recovery points for their private cloud workloads. These recovery points are stored on immutable cloud storage through the software-as-a-service (SaaS) orchestrator and scale-out file system, ensuring they cannot be directly accessed or compromised through direct access from the customer’s network. A big challenge with ransomware recovery is determining which backup copy is clean and should be failed over. With instant power-on of VMs, administrators can rapidly inspect dozens of recovery points in a short period of time without copying data, or rehydrate VMs before the VMs are powered on. This capability is enabled via a live Network File System (NFS) datastore mounted by the vSphere hosts in the VMware Cloud on AWS software-defined data center (SDDC) cluster. This allows for rapid recovery of compromised workloads into the pilot light or the newly deployed SDDC instance.



Blue: Steady state; Green: Activated on failover

FIGURE 6: How VMware Cloud Disaster Recovery responds to ransomware.

Respond with NSX

If NSX is deployed in the private cloud, it is possible to respond to ransomware attacks before they affect vSphere hosts and VMs. In alignment with MITRE ATT&CK, NSX classifies malicious activity by attack stage to identify the risk associated with each stage of an attack. It also generates a dynamic intrusion blueprint and a detailed timeline of a threat as it enters and moves laterally across your on-premises and cloud network. These visualizations give the information to quickly understand the scope of the attack and help prioritize response.

Known malicious traffic patterns are detected and blocked by NSX Distributed IDS/IPS. NSX leverages machine learning to analyze the malicious behaviors and malware samples collected from organizations across the VMware global threat intelligence network to automatically create new IDS/IPS signatures and push them out to all NSX sensors at machine scale. Both the Layer 7 distributed firewall and IDS/IPS rules can be configured to either drop or reject traffic automatically based on the configured policy.

Network traffic analysis automatically creates classifiers that recognize malicious network behaviors and previously unknown malware. The detected traffic is evaluated against configured policies to either allow or drop in response. Additional responses can be performed using the API or manual configuration changes to trigger actions, such as security policy changes and quarantine. Integration with third-party products—such as SIEM, security orchestration, automation and response, endpoint protection and firewalls, custom applications, and incident response workflows—assist in creating a defense in depth approach to responding to ransomware in the private cloud. Existing security controls can automatically send unknown objects for analysis and receive actionable threat intelligence in return.

Respond with VMware Carbon Black Cloud

The VMware Carbon Black Cloud features and functionality detailed in the [Respond subsection of the End-user computing solutions](#) section also apply to workloads running in private clouds.

Multi-cloud solutions

Organizations with a multi-cloud strategy must be prepared for ransomware in their private clouds and attacks on their public cloud workloads. The previous section documented many facets to protect the private cloud. In a public cloud environment, the organization trusts the service provider to have taken all necessary steps to deploy capabilities to protect, detect and respond to ransomware. However, public cloud services secured by the service provider can be leveraged in infinite ways by organizations to run modern applications. This section documents VMware solutions that give organizations the insight and tools necessary to prepare for ransomware that might target public cloud workloads and modern applications that run in private and public clouds.

CloudHealth Secure State

For organizations with a public cloud presence, CloudHealth® Secure State™ is a valuable tool in the fight against ransomware. CloudHealth Secure State is a SaaS offering from VMware that gives organizations security insight across public cloud providers. This is particularly relevant if the organization has a multi-cloud strategy. A public cloud security issue appears in the news almost every day due to common problems, such as a misconfigured Amazon S3 bucket, a misconfigured security group, or a misconfigured database. CloudHealth Secure State helps an organization understand their security posture in the public cloud. CloudHealth Secure State provides real-time security insight into how the organization's presence in the public cloud is configured with respect to security.

Protect

There is a concept that a single misconfiguration in the public cloud may not be a problem until there are multiple misconfigurations. Multiple misconfigurations constitute a violation chain. While one link of that chain may not necessarily be an issue by itself, it creates a large security problem when chained together. CloudHealth Secure State identifies and reports on violation chains. CloudHealth Secure State is also contextually aware. It understands users, groups, accounts and the cloud environment. By protecting the environment through proper configuration of cloud services, the likelihood of a successful ransomware attack is greatly decreased.

CloudHealth Secure State provides the ability to create a security baseline policy, verify it, and enforce it across an entire public cloud environment. This addresses PR. IP-1 in the NIST Cybersecurity Framework: A baseline configuration of information technology/industrial control systems is created and maintained, incorporating security principles (e.g., the concept of least functionality).³

Detect

In addition to its built-in detection capabilities, CloudHealth Secure State can integrate with other public cloud threat detection services. CloudHealth Secure State is able to detect a data integrity event that aligns to DE.CM-1 in the NIST Cybersecurity Framework. The encryption performed by ransomware is a data integrity event and must be identified as early in the attack as possible. Ephemeral cloud resources must be monitored to detect security events within minutes of occurring.

CloudHealth Secure State detects environment changes and displays this information in an easy-to-consume manner for security and cloud operations. Security configurations within continuous integration/continuous delivery (CI/CD) pipelines are analyzed. If a CI/CD pipeline is determined to be a potential attack vector, developers and operational teams are notified through integration to systems such as Slack.

CloudHealth Secure State monitors threats by correlating Amazon GuardDuty insights with cloud resource misconfigurations along with relationship context and change activity. CloudHealth Secure State can send cloud resource misconfigurations to vRealize Log Insight, Splunk, and other log aggregation tools for security operations teams to visualize.

Respond

CloudHealth Secure State provides security operations teams with the ability to visualize cloud resource relationships and identify misconfigurations and threats that must be remediated prior to an event occurring. If a ransomware incident occurs in the public cloud, CloudHealth Secure State provides the ability to investigate and resolve identified breaches through reference of security baselines and tracked configuration changes. The security baselines enable developers to monitor and fix security violations in their environments based on permissions provided through role-based access controls. Security operation teams can audit configuration changes and track the resolution of security violations and ransomware incidents.

VMware Carbon Black Cloud

The VMware Carbon Black Cloud features and functionality detailed in the [End-user computing solutions](#) and [Private cloud solutions](#) sections also apply to workloads running in public clouds if they leverage supported operating systems (Windows, macOS, Linux). Whether these are VMs running on vSphere, native Azure, native Amazon Web Services, native Google Cloud Platform, or other clouds, the level of protection, detection and response is the same.

Modern applications

Containerized applications and orchestration platforms such as Kubernetes are growing at a rapid pace. This growth can be attributed in part to the benefits cloud native patterns provide an organization. It does, however, also represent a potential growth in new attack surfaces for threat vectors such as ransomware. While a holistic discussion of a security framework for modern application platforms is beyond the scope of this paper, this section will highlight some of the areas of consideration for a targeted ransomware attack in the context of containers and Kubernetes.

One of the key constructs in the cloud native space is the immutability or declarative approach of applications and platforms. In other words, once a running container or object is deployed, that running object is not updated or changed. If something needs to be updated or it has been compromised, one simply redeploys a complete new instance, destroying the old. This ephemeral quality of cloud native patterns represents a strong reduction in attack surface for a ransomware attack. If malicious code is embedded into or modifies an image, that image can simply be deleted and a new image deployed, therefore destroying the modification.

Protect the build process

Organizations must first focus on the build and certification processes of the images and the components used for cloud native deployment patterns. There are a growing number of examples where code has been embedded through publicly available images of popular software building blocks and then exploited once those applications are deployed on top of those compromised images.

Within the NIST framework previously discussed, Harbor Registry, VMware Tanzu™ Application Service™, and VMware Tanzu Build Service™ offer key capabilities in the areas of image building, deployment, and curation.

Harbor

Harbor is an open source registry that secures artifacts with policies and role-based access control, ensuring images are scanned and free from vulnerabilities. Harbor signs images as trusted. Harbor, a Cloud Native Computing Foundation graduated project, delivers compliance, performance and interoperability to consistently and securely manage artifacts across cloud native compute platforms, such as Kubernetes and Docker.

VMware Tanzu Build Service

Organizations building modern applications should securely automate the code to container pipelines to mitigate the risks of malicious code being embedded into the images. VMware Tanzu Build Service leverages cloud native buildpacks to rebase application images when specialized contractual base images are updated in a registry. Therefore, organizations can resolve certain common vulnerabilities and exposures (CVEs) without a rebuild. The registry (e.g., Harbor, Artifactory) contains the most up-to-date and secure images. In alignment with the NIST controls previously described, Tanzu Build Service automates and secures this aspect of the code to container pipeline with a rich set of role-based access controls to restrict the number of individuals that have access to the build process. This represents a critical reduction in the attack surface for a ransomware threat vector at the build stage.

VMware Tanzu Application Catalog

Organizations often leverage widely available open source software (OSS) as building blocks for their applications. VMware Tanzu Application Catalog provides a mechanism to curate a catalog of these popular OSS building blocks on top of a VMware-provided or organization-provided image. The build pipeline of Tanzu Application Catalog continuously watches upstream projects and rebuilds the artifacts. The artifacts are securely delivered to the organization's registry along with metadata providing critical details about the libraries and binaries in every container image.

Reduce the attack surface after deployment

Protecting running applications from targeted ransomware attacks in the cloud native space covers a multitude of topics. This paper focuses on just a few cloud native examples.

Ransomware threat vectors at their core rely on the ability to gain access to files shared or accessible from multiple locations. Containers and Kubernetes are different from traditional applications because of the separation of an execution layer from the source image. This can be beneficial. For example, read-write access happens at the execution layer and not the source image. The source image is essentially read only, so any restart/redeploy of the image removes the encrypted image. However, access to shared storage becomes a critical decision area. Historically, virtualization leveraged a shared storage model, such as NFS. vSphere vMotion®, vSphere High Availability and other vSphere features require shared storage to function. Containers take a different approach for a storage access model. For example, a ReadWriteMany (RWX) model increases the attack surface for ransomware, whereas ReadOnlyMany (ROX) significantly reduces it. These benefits affect choices in cloud native application architecture and design.

For example, there is a large data lake that holds information for a machine learning or analytics application. If the microservice or process accessing the data has read/write access, then that service is a very attractive target for a ransomware attack. If the process is compromised, the entire data set could be encrypted. However, if the service directly accessing the data only has read access and sends queries to other processes or services for analytics and processing, the large data set has a much lower attack surface from a ransomware threat vector.

This example illustrates the importance of application design and architecture when moving to cloud native patterns. VMware Tanzu Labs™ goes beyond product and works with organizations to examine their existing application catalog or provide guidance on greenfield application efforts. This analysis accelerates an organization's ability to transform into these new arenas and do so in a secure, repeatable manner.

The role of service mesh

Decomposing a monolithic application into microservices results in a distributed system. Application and platform teams must manage different aspects of communication between many discrete services, including the following activities:

- Establishing and maintaining operational visibility into the state of the services
- Connecting, routing, load balancing, and securing communications across distributed microservices
- Reducing latency between the services in the service chain, which can ripple across the entire application and affect the user experience
- Troubleshooting and identifying the root cause of problems in an application composed of many different services written in different programming languages

These challenges represent potential increases in the attack surface for ransomware attacks. One of the emerging tools to help secure communication and define policies to help protect these architectures is a service mesh.

A service mesh is an abstraction layer on top of a microservices application that provides the following capabilities:

- Service-to-service communication, including service discovery and encryption
- Observability through monitoring and tracing
- Resiliency through circuit breaking and retries
- Traffic management with routing and load balancing
- Security with authorization and encryption

Before service mesh, client libraries and API gateways were used to address some of the issues introduced by a microservices architecture. These solutions, however, have challenges of their own. A service mesh solves some of the challenges introduced by distributed microservices by abstracting necessary functions—such as service discovery, connection encryption, error and failure handling, and latency detection and response—to a separate entity called a proxy. The proxy sits in front of each microservice, and all inbound and outbound communications flow through it. The proxy provides the functions previously noted and metrics for observability purposes.

Tanzu Service Mesh, an enterprise-class service mesh, solves the challenges associated with a distributed microservices application by extending service mesh services outside Kubernetes clusters and providing a unified operational layer across heterogeneous platforms and technologies, including virtual machines and other service meshes. Most service mesh implementations focus on services alone, but this limited approach ignores users and data. Users leverage services to access data. If the application runs in a multi-cluster or multi-cloud environment, communication and access between users, services, and data should be managed centrally, without needing to manage the underlying physical infrastructure.

VMware Tanzu Service Mesh™, built on VMware NSX, elevates the service mesh from the physical boundaries and limitations of a single cluster and a single cloud. With Tanzu Service Mesh, organizations can control, measure, secure, scale and operate applications, regardless of where the components are deployed across multiple clusters or clouds. These capabilities contribute to the reduction of the attack surface against ransomware threat vectors when combined with a holistic cybersecurity framework. Tanzu Service Mesh positions organizations to move into the cloud native space with greater confidence against these threats.

Protect

Tanzu Service Mesh protects from ransomware attacks in a number of ways:

- Access policies – Tanzu Service Mesh has a policy framework that controls access between services, regardless of where the services run. For example, one service could be running in AWS and another on premises. But, the policy for which services can access others is not defined by the physical location of these services. The policy is defined by the abstracted service definition within the global namespace. An application owner can run the service in one cloud today and in a different cloud tomorrow, and the same security policy applies. Any attack is containerized by controlling the access between services, which is very similar to the micro-segmentation policies with the NSX Distributed Firewall on Layer 7. The application is protected by Tanzu Service Mesh in Layers 4 through 7.
- Policy score – The Tanzu Service Mesh policy framework can accept a security score from VMware Carbon Black Cloud as well as third-party systems, such as Sysdig. If a specific service has been compromised, VMware Carbon Black Cloud or the third-party system detects the compromise and injects a score to the Tanzu Service Mesh policy framework that will stop all traffic to that service. This can be very useful against ransomware attacks as the policy Tanzu Service Mesh creates is dynamic and responds to a changing situation.
- End-to-end encryption – While not directly protecting from a ransomware attack, it does provide a layer of security. Tanzu Service Mesh provides end-to-end encryption between services with mutual transport layer security (mTLS). These services can run in any cloud and in any Kubernetes cluster. By enabling mTLS, eavesdropping on services communication is not possible. This can prevent insights any attacker needs to advance a ransomware attack.
- NSX Advanced Load Balancer™ integration – Tanzu Service Mesh integrates with NSX Advanced Load Balancer. This integration allows Tanzu Service Mesh to expose an application through the global server load balancing capabilities of NSX Advanced Load Balancer. The WAF capabilities of NSX Advanced Load Balancer are enabled on northbound communication into the Tanzu Service Mesh global namespace if NSX Advanced Load Balancer is used.

Detect

Tanzu Service Mesh deploys an Envoy sidecar to each service that is part of a global namespace in Tanzu Service Mesh. The sidecars log every communication and send the information to Tanzu Service Mesh. This audit information can identify an attack before the bad actors gain a foothold in the organization's cloud. This is available for all global namespaces across Kubernetes clusters and multiple clouds.

Respond

One of the key aspects in responding to a ransomware attack is to repave the affected system. This can be accomplished by restoring from backup or a complete rebuild. Deployment and configuration of modern applications are automated. By abstracting the infrastructure components, modern applications are quickly redeployed to any cloud. Tanzu Service Mesh abstracts the application from the infrastructure configuration. If a service leveraging Tanzu Service Mesh is attacked, the service is redeployed from an automated pipeline. Tanzu Service Mesh will rediscover the service and apply secure communication and security policies automatically, regardless of deployment location. Tanzu Service Mesh enables organizations to repave the application without worrying about how the service will operate within the application framework because communication and security functions are automatically applied once it is re-added to the global namespace.

Summary

Ransomware is a serious threat to all organizations across all industries. VMware provides many capabilities to protect organizations from ransomware attacks. If organizations are infiltrated, VMware technologies enable security operations to protect, detect and respond to these threats. Organizations must focus on deploying end-user solutions—such as Workspace ONE, Horizon, and VMware Carbon Black Cloud—to stop ransomware from entering the environment. vSphere, NSX, VMware Carbon Black Cloud, CloudHealth Secure State, VMware Tanzu, and vRealize Suite provide organizations with robust capabilities to protect against, detect and respond to ransomware in private and public clouds. These solutions must be deployed and configured as part of a larger defense in depth security model. Once VMware solutions are deployed and configured as documented in this paper, organizations will significantly reduce attack vectors and possess the tools and processes to protect, detect and respond quickly to ransomware events.

Authors

Amanda Blevins, Mandy Botsko-Wilson, Niran Even-Chen, Brian Heili, Keith Luck, Dale McKay, Jon Nelson, Adam Osterholt, Scottie Ray, Jeff Whitman, James Murray

Citations

This white paper contains hyperlinks to non-VMware websites that are created and maintained by third parties who are solely responsible for the content on such websites.

1. NIST. "[Framework for Improving Critical Infrastructure Cybersecurity](#)." April 16, 2018.
2. NIST. "[NIST SP 1800-11 – Data Integrity: Recovering from Ransomware and Other Destructive Events](#)." Timothy McBride et al. September 2020.
3. NIST. "[NIST SP 1800-25 – Data Integrity: Identifying and Protecting Assets Against Ransomware and Other Destructive Events](#)." Jennifer Cawthra et al. December 2020.
4. CIS. "[Cybersecurity Spotlight – Defense in Depth \(DiD\)](#)."

