

Lab Insight Report

Validation of Cohesity Accelerated Recovery from Ransomware

**By Krista Macomber, Senior Analyst
and Randy Kerns, Senior Strategist**

January 2021



Evaluator Group

Enabling you to make the best technology decisions

Overview

Ransomware attacks continue to grow in their frequency, severity, sophistication, and costs to organizations. Already a highly imminent threat to address, ransomware has been exacerbated in its severity by the COVID-19 pandemic. Cybercriminals are seeking to capitalize on new security vulnerabilities resulting from the shift to remote working, as well as the global need for information and treatment (e.g., launching coronavirus-related phishing attacks and targeting healthcare and medical research organizations).

Against this backdrop, organizations increasingly understand that it is not a matter of “if” they will experience a ransomware attack, but “when.” An effective data backup and recovery strategy is crucial to deterring hackers from planting ransomware, avoiding having to pay ransoms, and to avoiding data loss and lengthy downtime (resulting in expensive [costs](#), including governmental fines and a loss of credibility, on top of ransom fees that may have been paid). This is especially true as ransomware variants have evolved to focus on infiltrating the backup environment. Effective data protection also determines how quickly an organization can get back up and running again following a ransomware attack. Even if a customer chooses to pay the ransom, typically they are simply dumped a text file with encryption keys that each have to be manually typed into various system consoles until they are unlocked, without copy and paste functionality. This is a laborious and time-consuming process.

Ransomware protection is not easy, however. Bad actors get creative when it comes to obtaining access to production infrastructure and sabotaging backups. Especially as the amount of data being generated and protected grows, it becomes increasingly challenging to have comprehensive visibility and know exactly when the ransomware was triggered. Lastly, but certainly not least importantly, there is the question of recoverability; systems must get back up and running as quickly as possible.

This paper is intended to review the capabilities of Cohesity in mitigating ransomware threats by means of protecting the enterprise’s data and production and protection infrastructure alike from ransomware, detecting that an attack is occurring, and facilitating rapid recovery. First, however, we will outline key attributes of a ransomware protection strategy as well as Cohesity’s product security strategy as a whole.

Background for Ransomware Protection

As alluded to in the “Overview” section of this report, the core pillars of an effective ransomware strategy include data protection, attack prevention, attack detection, damage control response, and rapid recovery.

Oftentimes, ransomware presents as an automated virus that infiltrates the victim’s IT environment and does damage (namely encrypting files) without human intervention. Protecting the backup system from falling victim to these attacks is a requirement for avoiding detrimental impacts, such as paying ransoms,

of these ransomware viruses. The repository holding the backup data must be immutable from malware attacks. The backup system must not allow the backup catalog to be damaged by a virus attack. Converged backup solutions with a self-contained operating system and a self-contained snapshot catalog in an immutable file system reduces the attack surface.

Automated ransomware variants might be able to disable backups, but they typically have no knowledge of – and will conduct encryption regardless of the status of – the backup environment. There is a second class of ransomware that is arguably more worrisome – especially for mid-to-large enterprises – in terms of its potential impact. This class of ransomware is hacker-operated. The hacker first infiltrates the victim's production systems – typically the network – through means that have nothing to do with ransomware. Once bad actors have intruded into the production environment, they must compromise backup systems (e.g., through phishing or insider attacks) in order to monetize the significant effort that it typically takes to infiltrate the production environment. If they know that the backup environment is sound, there is no sense in encrypting data because the victim will be able to simply recover from their backups as opposed to paying the ransom demanded. As a result, an immutable backup solution might deter cyber criminals from attacking the production environment.

An essential component of protecting the backup environment from a ransomware attack is controlling access to the system through user authentication such as two-factor authentication (2FA) and identity and access management (IAM). These tools help to ensure that only authorized individuals can access the environment. Prevention also includes applying immutability protection and write once read many (WORM) designation – which, when used in conjunction with one another, inhibit data from being modified for a period of time specified by the customer. This helps to ensure that, if a bad actor finds a way to access the environment despite the previously mentioned identity control measures, they cannot manipulate the data.

There are also two different types of detection that are important in ransomware protection. The first is detecting vulnerabilities in the production IT environment by scanning backup copies of production VMs. The second is identifying that a ransomware attack has occurred by analyzing patterns of changes in backup data.

If an attack has occurred on the production IT environment, customers must be able to control the amount of damage that can be caused, and they must be able to recover as quickly as possible. Role-based access control (RBAC) and restricting who can access the cluster and data as well as who can control things such as immutability, play a role here. The Active Directory (AD) environment is also important because this controls access to the applications – which ransomware attackers sometimes manipulate to pressure the organization to pay their ransom.

If the backup solution can detect the vulnerabilities that hackers can use to gain access to production systems, such vulnerabilities may be closed before a hacker even gains access to the production systems, thereby preventing not only ransomware attacks but also extortionware that exfiltrates data. In

those attacks, the hacker monetizes their access to the victim's data not by demanding a ransom to give back what they took away, but instead by threatening to disclose the data, or by selling the data – or both. Therefore, detection of vulnerabilities is also a way of better protecting the production data itself.

Also critical to mitigating the impact of a ransomware attack is the ability to mitigate both data loss and downtime. Organizations need to understand what can be recovered, and how quickly it can be recovered. Capabilities like instant restoration of NAS services, live mounting, and granular and mass restores factor in here. The ability to spot check to verify that recovery points are clean and to the point of time that is required is also important.

Cohesity Strategy Overview

For its part, Cohesity's ransomware protection strategy is layered. At its foundation, Cohesity helps customers to mitigate ransomware. Cohesity has an immutable file system, encryption, internal WORM capabilities, and it works with WORM-enabled public cloud retention storage. Additionally, Cohesity can test for predictable recoverability. Also important for preparation is its documented ability to adhere to compliance regulations, such as its Federal Information Processing Standards (FIPS) validation.

For controlling and restricting access to a Cohesity backup environment, Cohesity supports 2FA, RBAC and whitelisting. It also has secure API access and a secure app ecosystem (for further explanation about Cohesity's app ecosystem, please see the "Attack Prevention" section of this report).

Cohesity can help customers detect that an attack has occurred by analyzing the backup data to look for anomalous changes made to the production environment, and by providing visibility into affected objects and their sources. AD granular comparisons aid in detecting changes hackers may have made to the AD environment; "before and after" comparisons also help the process of restricting and later restoring permissions during recovery from an attack. Proactive monitoring includes antivirus capabilities on data dumped to Cohesity via NAS protocols, auditing, and vulnerability scanning.

In the event that an attack does occur, Cohesity has designed its solution to respond quickly and efficiently. As explained in more detail later in the paper, Cohesity leverages parallel backups and recoveries for faster performance, and its architecture leverages metadata pointers as opposed to requiring data copies to be rehydrated – further accelerating the recovery process. It also provides a ML-based recommendation of clean, recoverable snapshots.

To reduce recovery time, Cohesity provides instant access, live mount and instant mass restore capabilities.

Cohesity Lab Validation

To verify Cohesity's effectiveness when it comes to ransomware protection, Evaluator Group monitored a number of detailed demonstrations performed by Cohesity. The version of Cohesity's software tested was 6.5.1c.

As previously discussed, preventing a bad actor's ability to access data is an important first step to take in ransomware prevention. Cohesity demonstrated its ability to help reduce the vulnerability of source VMs, as well as its ability to prevent damage to backups. It also demonstrated its various capabilities to accelerate recovery.

Summary of Validated Cohesity Capabilities

Data Protection

- MFA, defined user roles and RBAC
- Immutability and WORM
- Different encryption keys per tenant, and ability to customize the hardware console password
- One-time use token for support
- Two-person concurrence for root access by the customer and for support

Attack Prevention

- Scanning snapshots of VMs against known vulnerabilities

Attack Detection

- Identification of anomalous behavior and the last backup prior to the behavior
- Identification of boot issues

Damage Control Response

- Rapid parallel backup
- Global search
- Analysis and recovery options for AD environment

Accelerated Recovery

- Rapid parallel recovery
- Production-grade resiliency for running production on the backup data
- Instant NAS access
- Live mount
- Instant mass restore

Data Protection

Lab Validation: User Authentication

- Checked Cohesity's integration with Okta for 2FA to protect the backup admin and cluster admin accounts.
- Observed definition of the cluster manager role through the Cohesity UI.
- Verified the enablement of a support user. This login was disabled by default; the admin user needed to log in and enable the support account, and the password was then set by the trusted second party. As an additional safeguard, the admin subsequently requires this support user-designated password in order to change the password on this account.

The cornerstone of an effective ransomware protection strategy is protecting access to the backup infrastructure as well as to backup copies and snapshots. For its part, Cohesity employs user authentication and permission control to prevent bad actors from damaging backups through hacking or guessing a password. Immutability and WORM designations are used to prevent hackers and other malicious individuals from modifying or deleting data, should they be able to access it. Different encryption keys can be applied to different data, and the hardware password can be customized. Additionally, to prevent a single person on their own from being able to possibly damage the Cohesity cluster, two-person concurrence is required for root bash access.

From a user authentication perspective, Cohesity integrates with IAM products that use the Security Assertion Markup Language 2.0 (SAML 2.0), such as Okta, to authorize individuals attempting to log in to Cohesity.

Additionally, Cohesity leverages RBAC to manage individuals' visibility into and control over the Cohesity cluster. The idea is that each user should have only the visibility and control needed.

The predefined "Admin" role has full access to make all changes that are allowed on the Cohesity cluster; they cannot take actions such as deleting data with a WORM designation and they cannot access the Linux systems that the cluster is running on. This "super-user" admin assigns each user their role and manages the unique permissions that are associated with each role. This role is different than that of the cluster manager, who performs cluster maintenance such as software upgrades, but does not do backups.

Cohesity offers pre-defined delegated admin roles, and the customer can also create their own roles. As such, only the customer holds the permissions to the various individuals that can access their Cohesity system – Cohesity does not. For security purposes, Cohesity recommends that customers vault or otherwise heavily protect the password to the general ("super-user") admin account, like they would for an AD admin account.

Two other roles are important when it comes to ransomware protection. A “Data Security” role, typically managed by a security officer, sets up and manages WORM policies (including Cohesity’s DataLock feature) and legal holds. Cohesity’s immutability capabilities are examined in more detail in the next section of this paper.

Additionally, a privileged support user helps with support needs. The role of the support user when interfacing with Cohesity for support and for preventing any one individual from modifying the Linux underneath the Cohesity OS is discussed in more detail later in this paper.

Along a similar vein, Cohesity has multi-tenancy, so different tenants can have different encryption keys. This way, if one tenant leaves, their encryption key leaves with them and the cluster does not need to be destroyed to ensure that their data cannot be accessed.

Lab Validation: Immutability and WORM

- Verified that backup data was written to the file system and could not then be over-written.
- Proved that, once applied, the WORM policy could not be removed, and the retention period could not be shortened – even by the Data Security officer.
- Checked that the backup and general (“super-user”) admin roles could not delete WORM backups.
- Checked that retention periods for previous backups could be extended by the Data Security officer, but WORM data could not be shortened – even by the “super-user” admin. Those two roles encompass all the things that any user of the Cohesity software can do, so no user of the Cohesity software can delete WORM data before it has expired.
- Witnessed that the factory reset command cannot be used through the CLI if there is WORM data on the cluster.
- Checked that users’ permissions were retained during the cloning process.
- Observed that Live Mount and Instant Access could make changes to backup data presented by the cluster, but the original backed up data was unaffected; subsequent Live Mounts and Instant Access showed that the original data remained in place.

Cohesity’s immutability is based on its SpanFS™ file system, which uses B+ Tree metadata structures and spans across nodes in the cluster.

As opposed to overwriting data, SpanFS™ writes new and updated data in free space and then places metadata pointers to the new data, leveraging a distributed redirect-on-write mechanism. The metadata pointers are used to clone and present the data back out, leaving the original, immutable data block unchanged.

From an immutability perspective, this means that the original, immutable data point is not altered if it needs to be restored or otherwise accessed – for example, during a Live Mount or Instant Access, which are detailed in the “Recovery” section of this paper.

WORM designations and legal holds are controlled by the Data Security officer; backup and cluster admins cannot delete WORM backups. This is frequently referred to as “compliance mode,” put in place to prevent a rogue admin or a bad actor who has infiltrated the admin account, from inflicting damage.

Additionally, the storage container cannot be deleted or reset by the

cluster or backup admins if there are unexpired locked snapshots. Sudo privileges – which grants Linux root privileges – would need to be enabled for a user (specifically the support user, which is explained in the “Two-Person Concurrence for Linux Root-Level Access” section of this report) to be able to sabotage the cluster through more manual means. This is not a new threat, nor is it one unique to Cohesity. We

explore more about safeguards in the form of two-person concurrence that are applied to control this threat.

Lab Validation: Glacier Archival with WORM Retention

- Monitored Cohesity backing up two VMs, each approximately 142 GB (one being a clone of the other), with a change rate of approximately 2%. Each VM consisted of a CentOS VM with a 32GB eager-zeroed disk with VMware Tools and Cohesity tools and 101 GB of data created by VDBench set for 2x dedupability and 2x compressibility. The backup required 64.7 GB of space in Glacier, for approximately 80% data reduction in data written.
- Noted the WORM designation was set through the customer's AWS account as opposed to through Cohesity.
- Observed archiving to Glacier with WORM retention set for three days.
- Observed that AWS WORM lock would not allow the object to be deleted by the AWS account owner until the end of the retention period designated in AWS. This "delete" attempt action placed a "delete marker" on the object in AWS, but the data remained in AWS because its retention period had not yet expired.

As an additional safeguard, Cohesity supports WORM-enabled public cloud storage for backup purposes. In this case, the test involved AWS Glacier Archive storage with WORM retention (Glacier Vault Lock). If data has been marked in AWS as immutable, deleting it from the Cohesity cluster will not delete the data in Amazon. Cohesity explained that the AWS WORM lock also makes the metadata immutable, so ideally objects should not be deleted until Cohesity can send another full backup. For example, set retention in Amazon for 91 days and send a full every 90 days. This also can potentially avoid costly up-tiering and egress fees because the customer is not pulling data back from AWS and because the data does not have to be reprocessed.

Lab Validation: Two-Person Concurrency for Linux Root-Level Access

- Noted that, when remote support was disabled, Cohesity support could not reach the Linux servers or the cluster software.
- Observed having the support Linux account user briefly use the admin account to set the support Linux account password, prior to logging out and having the original admin change the admin password. Observed that only someone knowing the support password can change the support password. As a result, no one person has the passwords to both accounts.
- Observed that the admin account cannot reach the Linux layer at all, and can disable the root access for the Linux support account. Consequently, before root access can be used, first the admin must enable root access and second the support Linux user must log in to the Linux to use that access.

Having some means of Linux root access is necessary because support personnel, or customers working with support, could require access to the server bash shell with root access in order to fix a bug or other abnormal operation. Requiring two-person concurrence for root access is important because with the power to fix things that are broken comes the power to break things that are not – potentially creating an opening for an attack.

Cohesity demonstrated to Evaluator Group that two-person concurrence is required to manipulate the servers that its software is running on, as well as the software and data files, to prevent a bad actor from gaining access.

The Linux account called “support” provides bash shell access, but not root-level access unless sudo rights for this support user have been enabled by the UI Admin account. Leaving support sudo disabled allows for diagnosis but not for taking action. The support Linux account can be accessed in two ways: by Cohesity support personnel when the

customer so authorizes for a limited period of time, or by the customer locally using SSH to connect to the cluster or its nodes, logging in as the support user using a password. That password is separate from the password that is required for cluster administration and is also in the customer’s control without any back door for Cohesity employees.

In the event that a customer requires remote support from Cohesity, the local admin must first enable support access and then provide Cohesity access to the cluster through a one-time-use support token. This token ensures that only Cohesity is granted access to the credentials needed to access the cluster. By using the support token, Cohesity does not need to know the password that the customer has selected for the Linux support account. A remote tunnel server that is controlled by Cohesity is enabled, to which Cohesity users must connect to be able to reach the customer’s cluster.

If sudo is disabled, whomever is using the Linux support account cannot make changes to the Linux OS files or the Cohesity software and data files. If changes are required, the admin account must enable sudo rights for the support account.

In other words, if the support user is granted sudo access by the admin user, the support user has Linux root privileges and can do anything that can be done from Linux, including making changes to the Linux system and wiping out the Cohesity program and its data. As a safeguard against one individual being able to do so by acting independently, the person who knows the admin password must change sudo access from disabled to enabled, and the person who knows the support user password must log in to execute any commands.

The Intelligent Platform Management Interface (IPMI) is another way that users might access the Cohesity cluster. IPMI facilitates remote management and monitoring of systems and is relevant in the context of ransomware because it can be used to recover in a lockout solution. Lockouts can occur, for example if a rogue admin or a hacker locks other users out of the cluster and has disabled sudo access so that the cluster cannot be unlocked by the support account. The customer can use IPMI to recover either by gaining access to the hardware console or by loading an alternate boot image.

Cohesity sets the hardware console password, but the sole ability to change the console password lies with the customer. To eliminate a back door to accessing the system console through the IPMI, the customer can change the password. Not even a Cohesity insider would then be able to access the console. Of course, it is paramount that the customer does not lose this password. It is also advisable that the customer implement a two-person protocol whereby whomever has the console password does not have access to the physical hardware or to the IPMI except in very unusual situations such as password recovery.

As an additional layer of protection against a single person from having the ability to destroy the Cohesity system, customers can set a BIOS password to be required to load a different boot image. IPMI can be used to change what location (image) the server boots from, thus facilitating access. The person with the BIOS password should not be able to use the IPMI or physically connect to the IPMI without someone else taking action to give them network or physical access.

Attack Prevention

Lab Validation: Prevention of Damage to Production Systems

- Observed setup of CyberScan via the Cohesity Marketplace.
- Verified functions of CyberScan app, notably scanning snapshots of VMs and checking those snapshots against a list of known vulnerabilities that is kept updated by Tenable.io.
- Established that CyberScan could automatically identify that multiple snapshots were taken of a VM in the same day, and that the user could select which snapshot to be scanned.
- Confirmed CyberScan's vulnerability status of the snapshots (low, medium and critical) in the Cohesity UI, and that a Tenable report – which can be used for documentation to, and for assessment by, security teams – was automatically generated and could be downloaded directly from the Cohesity UI.
- Checked that information regarding vulnerabilities was also integrated into the recovery view in the Cohesity UI, so that the backup admin could decide which recovery point to use.

Attack prevention consists of preventing an infection or intrusion, and removing incentive for an intruder to plant ransomware. This is important because, in recent years, with most enterprise data being backed up, the largest and most costly ransomware attacks are orchestrated by hackers who sabotage the backup system to force payments of ransoms. The previously mentioned methods of protecting the backup creates a disincentive for the hacker to unleash ransomware if production systems are compromised. Additionally, consistent scanning of backup copies can prevent the ransomware attack from occurring in the first place through preventing a successful infiltration of production systems.

Cohesity's CyberScan app is used to identify vulnerabilities of source VMs. For example, CyberScan can uncover if a server needs to be patched and is therefore capable of exposing the environment to hackers if it were to be recovered.

It's important to note that scanning production systems impacts performance; as a result, customers tend to scan production environments infrequently. CyberScan scans the backup snapshots or copied within the Cohesity runtime environment in lieu of the production systems themselves. This allows customers to conduct more frequent scans (even daily scans) without negatively impacting production workloads. As a result,

customers could have earlier warning if production systems have become vulnerable to newly discovered modes of attack.

By way of background, Cohesity developed CyberScan jointly with Tenable, allowing the Tenable.io cloud-based vulnerability management platform to communicate with the Cohesity cluster. Like all of the Cohesity Marketplace apps, CyberScan runs in containers on the Cohesity platform. Essentially, this approach facilitates a secure environment in which the customer chooses what can run.

CyberScan directly accesses data sitting on the Cohesity cluster; data does not need to be transported to another system to be processed. CyberScan accesses the Cohesity cluster through an API and cannot bypass the Cohesity cluster security mechanisms. For added security, Cohesity performs checksums as app software packages are being installed to ensure compatibility and to verify that the software package has not been altered or tampered with.

Ransomware Attack Detection

Lab Validation: Ransomware Attack Detection

- Cohesity created a directory of approximately 7,000 files of a variety of types including pictures, XML and text – for a total of 1 GB of unique data. That included uncompressed data like PNGs as well as compressible data like text. Additionally, Cohesity applied deduplication.
- Cohesity created a script that randomly grabbed files and populated them into a directory, it took a backup, made changes to the data, and then took another backup. Cohesity then applied a PowerShell encryption module that simulates a ransomware attack by running the original data copies through an encryption module and then making sure that the original copies are not recoverable.
- Observed in the Cohesity UI Cohesity's identification of the anomalous encryption behavior, including files deleted and modified/changed as well as Anomaly Scores over time.

If a ransomware attack occurs, Cohesity has a number of capabilities that help customers uncover that an attack has taken place. Cohesity assesses backup data for patterns that are suggestive of typical crypto ransomware attacks. It also can potentially uncover issues with VM bootability, which also could signal an attack.

Cohesity looks at the effect that ransomware is having on the data that has been backed up by Cohesity, as opposed to trying to chase every executable file or service. Endpoint detection products already evaluate executables, and executables are continuously being updated and evaluated by bad actors to get past these products. While endpoint protection products ideally can detect in real time when an attack is occurring in real time on the host, Cohesity facilitates identification and recovery once an attack has started encrypting the production IT environment and the changes have been backed up.

The Cohesity software looks for “entropy changes,” defined as randomness in terms of which data blocks have changed, how random the changes are, and how much the ability to dedupe and compress the data has changed. For example, Cohesity can identify if only a part of a file system has been changed, or if changes have occurred across the entire file system. It also can detect if data is no longer reducible (i.e., the data has been scrambled due to ransomware-inflicted encryption). Based on this entropy assessment, Cohesity assigns to each VM an “Anomaly Strength” score that can be tracked over time. If an entire VM is encrypted, the Anomaly Strength score will be higher than if only a component of it is encrypted.

The backup job metadata is forwarded into Cohesity's cloud-based, SaaS-delivered management platform for processing. The CyberScan information previously mentioned is also sent to Cohesity. Cohesity proactively tracks over time and charts for visualization of the VMs' anomalous occurrences. It also allows admins to send push notifications of security alerts, with the Anomaly Strength score typically providing a threshold to determine what is worth evaluating by the customer. Customers can leverage this insight not only to identify ransomware, but also to make decisions about recovery points. This may be a collaborative process between the admin overseeing data protection processes and the VMware admin who understands the content of the VM. The VMware admins can be granted access privileges to manage their own searches and restores.

The training module is maintained by Cohesity in the Cohesity network. End clients do not need to be updated because the training module is delivered through the cloud.

Also relevant to discuss is the ability to detect attacks on VM bootability – which would be detected if someone monitoring the backups were to notice reports of CyberScan being unable to scan a VM. Cohesity sometimes skips scans because they are a lower priority than completing backups and restores. It does not currently screen the causes of scans not running and send alerts out if the problem is on the VM side. However, many scan failures could provide a warning of ransomware-related issues if the admin were to identify that the scans weren't being skipped by Cohesity.

Damage Control Response

Lab Validation: Rapid Parallel Backup

- Cohesity performed an on-demand backup of a 442GB TPC-C database created by HammerDB.
- Monitored the backup being broken up to occur in parallel across a four-node cluster. Specifically, because the log file was large (almost one-third the size of the database itself), Cohesity split the log file and the database was split into three chunks of approximately 128GB. Logging of jobs confirmed that all nodes worked in parallel.
- Noted a compression ratio of approximately 4.4x. The log files and database itself all totaled approximately 518 GB, but after compression totaled only approximately 118 GB in Cohesity.

Lab Validation: Global Search

- Observed the ability to search globally, across the entire Cohesity environment, to identify which VMs an offending file (e.g., “ransomware.exe”) exists on.

Once the customer has identified that an attack has occurred, they must take quick action to limit the spread of damage. Important considerations for damage control include:

- Quickly making a copy of the data before it’s overwritten by a restore operation.
- Searching for known bad files across all VMs.
- Managing, locking down and subsequently reopening their AD environment.

An initial step to take for damage control is to back up the server’s final state prior to recovery, for forensic analysis purposes and to preserve any data that might be un-damaged. Cohesity conducts parallel backups, which allows the backup process to be completed more quickly – as a result, allowing customers to move to recovery and getting business systems back online after a ransomware attack more quickly. In addition, Cohesity facilitates protection of the AD environment – creating an additional barrier to bad actors accessing and damaging applications.

Cohesity’s parallel backups are granular

(in the demonstration to Evaluator Group, breaking up the backup process to occur across four nodes simultaneously). This allows Cohesity to support environments with very large files – including SQL and Oracle database environments – efficiently. The Cohesity software sizes each stream automatically; performance tuning could be done by Cohesity if needed.

Lab Validation: Cohesity and Active Directory

- Verified that some AD accounts were changed, and that access was then restored to the rightful AD admin.
- Observed granular comparison of the live AD environment to snapshots of the AD environment, to identify damages, through the Cohesity user interface.
- Witnessed the deletion of users from the AD environment – as well as their subsequent restoration from the AD Recycle Bin via Cohesity.

Protecting the AD environment also helps to minimize potential damage from a ransomware attack because it facilitates access to applications. Hackers will often take over portions of the AD environment to damage applications in the hope that it will drive the customer to pay the ransom.

Cohesity directly queries AD for its live state, and a copy of the AD environment that can be queried is spun up in the Cohesity cluster.

Cohesity will highlight what changed – for example, if a hacker inserted an account in the domain admins. Group policies can be protected directly from the Cohesity UI, allowing customers to bypass the typically convoluted process required to protect them. This is important because enterprises will often conduct two backups of the AD environment; one by the backup admin and one by AD admin. Typically, the AD admin will back AD up as an app, and the backup admin will back AD up as another VM. Cohesity has the granularity and RBAC required

to address these use cases.

Also through Cohesity, the admin can granularly restore AD objects and other specific affected portions of the environment. A user object can be recovered into a disabled state so that an admin can perform additional checks or changes to the object before enabling it. For example, if someone wrote a script that deleted half of the roster of AD users, the admin could merge the users that aren't showing up back in. Bulk overrides are also possible. Additionally, Cohesity supports merging of backed up object attributes with live object attributes. Permissions and changes to settings are tracked during the backup and recovery process.

Cohesity leverages the AD recycling bin – if it has been enabled – to restore AD objects. What this means is that the security identifiers (SID) of AD objects are preserved when the object is restored, simplifying recovery and preserving group memberships.

The Cohesity cluster admin oversees RBAC for the AD environment and facilitates the AD admin's access to conduct backups, compare snapshots, and perform recoveries. Cohesity also can compartmentalize the restore of AD objects and GPOs. In a large-scale AD environment, only select AD admins may have access to the entire AD, with a subset of AD admins having access to specific AD objects and group

policies. By allowing a different username and password for the restore, Cohesity allows this subset of AD admins to recover only what they have been granted AD access to.

Accelerated Recovery

Lab Validation: Parallel Database Recovery

- Verified Cohesity's parallel restore capabilities for the same 442GB database used to demonstrate parallel backups.
- For the restore, Cohesity broke this database into 8 streams, and the log file into two streams – all of which could be processed in parallel.

Cohesity's architecture can accelerate recoveries for customers – important when trying to get systems back online as soon as possible following a ransomware attack. Specifically, Cohesity offers:

- Parallel database recovery
- Inherent resiliency to support production use of backup infrastructure
- Instant NAS access
- Live mounting
- Instant mass restore

Traditional incremental backup and recovery approaches are time consuming because all incrementals need to be rehydrated before the recovered image is available to the customer.

Even if the chain is completed on the back end, the copy rehydration still needs to occur.

Cohesity's SnapTree-based approach of leveraging metadata pointers means that fully hydrated snapshots can always be accessed. An added benefit is that customers can quickly inspect the snapshot and identify that it is the correct one to be used for the restore. Furthermore, production applications or systems do not need to be brought down for scanning – Cohesity spins up a copy of the backup data and scans that. This approach also makes it easier to re-scan the environment after remediation has occurred. Also notable for accelerating recoveries is Cohesity's ability to conduct parallel restores.

Recoveries – including instant recoveries – and individual file spot checks can be controlled locally or via cloud-delivered software running at Cohesity. This facilitates global searching as well as browsing (for example, by a point in time) for file-based restores.

In the case of a ransomware attack, Cohesity would uncover if a large batch of files have been deleted and replaced with encrypted files. The cloud-delivered global UI shows a sample of the changes that were made, to assist in confirming whether or not unusual changes were caused by ransomware.

Individual files and entire systems can be recovered from what Cohesity has identified as the last clean snapshot or from a custom recovery point. Recoveries can occur to another location, which is typically desired in the event of ransomware because customers are unlikely to trust the VM that has been attacked, until forensics can occur.

Lab Validation: Production-Grade Resiliency

- Observed the powering down of the node that a Linux VDBench VM and a Windows SQL VM were connected to – and that another node then picked up the virtual IP address of the powered-off node.
- Noted the period of unresponsiveness of the VMs before they became responsive again (in this instance, approximately one minute) – as well as that they did not reboot
- Observed that access to a video file shared using SMB is not interrupted when the node sharing the file is powered off without warning to simulate a failure; the video continued playing without interruption.

While in failover during the recovery process, Cohesity is running critical applications in lieu of the production system. In order to sufficiently do so, Cohesity was designed with a number of fault tolerance measures.

The Cohesity architecture is strictly consistent at the block level; if the system is writing data to one node, that node doesn't acknowledge back and confirm that it has the data until another node also has a copy of the data. This ensures that, if a node dies, production data is not lost. Cohesity also supports erasure coding. Customers can configure the resiliency levels (e.g., one node and one disk, one node and two disks, etc.). Additionally, as demonstrated to Evaluator Group, production VMs do not crash if any Cohesity nodes fail – even if it is the one the VMs were connecting to.

Lab Validation: Instant NAS Access and Live Mounting

- Monitored instant NAS access based on a backup of a raw NetApp SMB NAS device. Providing NAS access took less than a second after making the selections in the GUI. Observed that user permissions on the files were retained.
- Checked the live mount of a Windows VM directly from the Cohesity system, leveraging directory structures that were backed up.

Instant access is key in ransomware recovery because it facilitates immediate access to business-critical departmental shares. If an insider has compromised the source filer, the instant access capability can accelerate the recovery. For its part, Cohesity can facilitate instant access for NAS services and VM backups. In the event that a NAS system itself is compromised during a ransomware attack, Cohesity can facilitate instant access to the entire NAS service function – as opposed to just the data – without requiring the original NAS system be restored, or that backups be restored to the original NAS system.

Instant access is logged for chain of custody (this includes updating of the last access date) – a requirement for many customers that Evaluator Group speaks with. In the Evaluator Group evaluation of instant access to a NetApp filer, NetApp SnapLock properties were also captured. Restores are vendor agnostic – data can be restored to another vendor’s array – because Cohesity doesn’t rely on NDMP or proprietary file scanning like NetApp SnapDiff. Cohesity accesses

through the native protocol, and from there reads and writes like any other client.

Performance for instant NAS access is selected by the customer based on quality of service. Irrespective of how large the source NAS volume is, the access will take the same amount of time because Cohesity is creating a clone of the file system.

Live mounting allows directory structures that have been backed up to be accessed, scanned, and validated. Single files and entire directories can be restored to a specific point in time, to the original location or to a different VM. To do so, Cohesity requires integration into the host to ensure that the disks are online and to check for the file system on the disk. The Cohesity agent can be automatically deployed, or an existing one can be deployed. Cohesity will alert the user if there is not an agent present. Changes to the live mounted data do not impact original immutable backup copies.

Lab Validation: Instant Mass Restore

- Observed that 100 Linux VMs were created and backed up. They were recovered and powered on via Cohesity Instant Mass Restore in less than three minutes.
- Observed that individual VMs could be browsed and searched for to be recovered, that entire drives could be recovered, and that a protection group (for example, called “Instant Mass Restore”) could be recovered.

Arguably most impressive, Cohesity offers an Instant Mass Restore capability that can accelerate the investigation and recovery processes. In the evaluation, Cohesity restored and powered on 100 VMs in less than three minutes with locked down networking. They could then be inspected and reconnected into the production environment. Unsophisticated ransomware attacks may leave most backup solutions unharmed, but the cost to the business is still high if the recovery process is lengthy. Instant Mass Restore and Cohesity’s other rapid recovery features discussed in this paper have the potential to make ransomware attacks substantially less costly to the business.

Summary

With the current rise in frequency and severity of ransomware attacks, effective ransomware protection is critical for organizations across industries. Cohesity covers the core tenets of ransomware protection, including reducing the risk of production systems being compromised, preventing backup systems from being compromised and thus removing the incentive to make an encryption attack on production systems, identifying that an attack has occurred, mitigating the spread of damage, and helping to get the business back online as quickly as possible following an attack. Core immutability and access control requirements are addressed, and Cohesity offers a graphical view into – and easy access to details about – vulnerabilities in the environment and anomalous changes that could signify a ransomware attack. Cohesity was architected to accelerate backups and recoveries at scale as much as possible, helping to reduce downtime and data loss and to facilitate forensic analysis of affected machines. Cohesity even has a number of features to simplify the comprehensive protection of AD – which is notable because it facilitates access to critical applications. After a series of lab validations with Cohesity, Evaluator Group concludes that Cohesity is an effective and differentiated solution suitable for enterprises struggling with the challenging task of protecting their environment and business against ransomware.

About Evaluator Group

Evaluator Group Inc., an Information management and data storage analyst firm, has been covering systems for over 20 years. Executives and IT Managers rely upon us to help make informed decisions to architect and purchase systems supporting their data management objectives. We surpass the current technology landscape by defining requirements and providing an in-depth knowledge of the products as well as the intricacies that dictate long-term successful strategies.

Copyright 2021 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group, Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, consequential, or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.