Hewlett Packard
Enterprise

Technical white paper

Check if the document is available
in the language of your choice.

RANSOMWARE DATA RECOVERY ARCHITECTURES

The ability to recover your data if you are the victim of a ransomware attack

CONTENTS

INTRODUCTION

Ransomware attacks have been on the rise in recent years. They are extremely insidious in nature and have been foremost in many CEOs' and CIOs' minds lately. Ransomware is malware that prevents an enterprise from accessing and using its data, usually by encrypting the data in an inconspicuous manner so the victim does not know the attack is occurring. Generally a ransom is demanded to provide the key necessary to decrypt the data and make it accessible again, hence the name "Ransomware." The ransom demanded can be in the thousands if not hundreds of thousands or even the millions of dollars. The ability to recover from a ransomware attack in the enterprise is becoming a requirement higher in value than the ability to recover from a data center disaster like a power outage, flood, fire or other natural disaster. This is due to the fact that a ransomware attack can infect and render not only the production site data but also the disaster recovery (DR) site data and even backup data unusable, making recovery impossible unless special precautions are taken to protect the data. This means that a typical disaster tolerant or disaster recovery solution is not sufficient to protect data from and allow recovery from a ransomware attack.

The means hackers use to execute a ransomware attack and the holds they put on data to make it inaccessible are not only cunning but are also constantly evolving. The attackers go to great lengths to hide the fact that an attack has occurred so that they can spend weeks or even months working their way deeper into the enterprise compromising data before announcing the attack and demanding a ransom. This means the tools and processes used to detect a ransomware attack must be run regularly, consistently, and be constantly kept up to date. It is beyond the scope of this document to address the tools, processes, and methods used to detect that a ransomware attack is occurring. This paper focuses on solutions that provide a means for restoring data following a ransomware attack. Hewlett Packard Enterprise recommends that any time a new tool, process, or set of virus definitions is deployed or updated that the user run the new tool or latest set of virus definitions on the last copy of data that was saved at each protected data layer in the solution. This step ensures no infection is present that was not detected by the older version of the tools and virus definitions.

OVERVIEW

Special measures must be taken to protect data to be able to provide the potential for recovery following a ransomware attack. This paper explores various solution architectures that can be deployed to offer the potential for data recovery should the enterprise suffer a ransomware attack. For recovery to be achieved, a clean un-infected copy of the data must have been saved before any attack occurs. That data in turn must be protected so it cannot be attacked and compromised. Because the length of time an attack may go undetected is non-deterministic, the amount of time data must be archived to provide a recoverable data store is also non-deterministic. Hewlett Packard Enterprise recommends at least a six-month repository of data be saved for recovery if needed. Because of this, Hewlett Packard Enterprise recommends a very long-term archive storage layer like backup to cloud and tape as the very last line of defense.

The solutions architectures in this paper look a lot like standard DR solution architectures (they use many of the same tools and processes), but a solution that can provide recoverable data in the event of a ransomware attack is deployed, managed, and treated completely differently than a DR solution. Unlike a traditional DR solution, the intent is not to provide a guaranteed recovery point objective (RPO) or recovery time objective (RTO) but to try and ensure a non-infected recovery point is available should an attack occur. The corollary to this is that a solution designed to recover from a ransomware attack can be used quite successfully to recover from a DR scenario although with more complicated recovery processes than a traditional DR solution requires and with a larger RTO than a traditional DR solution can generally provide.

This paper uses HPE Alletra 9000 in its sample architectures, but everything covered in this paper is also applicable to HPE Primera arrays.

THE FOUNDATION OF A GOOD RANSOMWARE PROTECTION SOLUTION IS SECURITY

At the heart of any good ransomware solution is a secure environment in the enterprise. Having a solid line of defense against all malware, what is referred to as a "perimeter defense," is paramount to preventing a ransomware attack from occurring in the first place.

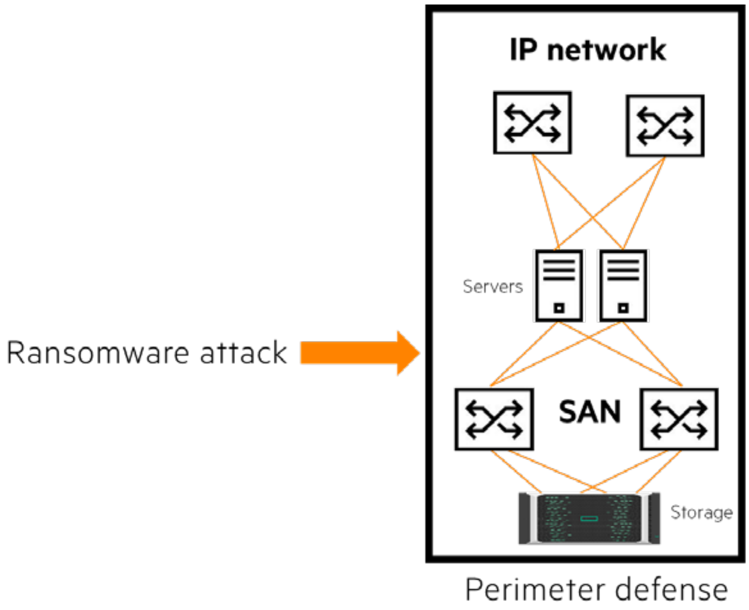

FIGURE 1. The best ransomware attack is the one that is stopped before it occurs

A strong perimeter defense protects the servers, storage, networking, and SAN infrastructure in the enterprise from attack. It encompasses but is not limited to the hardware infrastructure in the enterprise. It ensures all appropriate software updates are applied in the enterprise, secure passwords are in place, users are properly trained so they do not inadvertently open the door to an attack vector, and that the most current malware detection software is in place and is run on a regular basis. It also includes customer security measures such as requiring personnel training on protecting the enterprise, dual authentication or dual authorization (when available) for data destructive administrative tasks, checking space usage on the storage arrays to see if deduped and compressed data is being encrypted and hence is using more capacity than would normally be expected, and checking if backups are consuming more space than expected. Protecting from a ransomware attack is not a one-time operation that you put into place and then ignore. It requires constant vigilance because the threats are constantly changing and evolving.

THE 3–2–1 RULE FOR DATA PROTECTION

The best way to ensure you can recover from a ransomware attack is to have a solid data protection strategy in place. When thinking about protection from ransomware skimping on your data protection strategy is not a good idea. A best practice data protection rule that can help effectively mitigate the threat of data loss and malware attacks is the 3-2-1 rule. The 3-2-1 rule specifies that you should:

3. Have at least three copies of the data—the primary data and two copies.

2. Store the copies on two different types of media—tape, disk, or cloud.

1. Keep one backup copy offsite—either on tape or in the cloud—in the event of local hazards or infections within the network.

Hewlett Packard Enterprise takes the 3-2-1 rule a step further and recommends that in addition to an offline copy of data in the cloud, an offline copy to tape with an air-gap always be created. This is what Hewlett Packard Enterprise calls the "3-2-1-1 rule." A tape removed from the tape library cannot possibly be affected by a malware attack. Following the 3-2-1-1 rule means you should always have an available and usable backup of your data and systems. In a world where ransomware can instantly take you offline, it is a vital precaution.

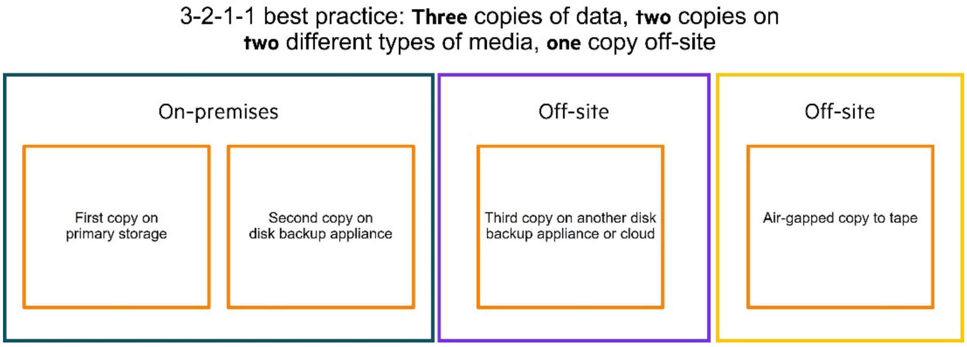
3-2-1-1 best practice: Three copies of data, two copies on two different types of media, one copy off-site
On-premises
First copy on primary storage
Second copy on disk backup appliance
Off-site
Third copy on another disk backup appliance or cloud
Off-site
Air-gapped copy to tape

FIGURE 2. 3-2-1-1 best practice rule for data protection

CLEAN ROOMS

The ransomware data recovery architectures presented in this paper use the concept of clean rooms for protecting data. A clean room contains an immutable and locked copy of the data being protected. It cannot be modified, and it cannot be removed for a predetermined time interval. Hewlett Packard Enterprise refers to this concept as a virtual lock on HPE Alletra 9000 arrays.

After an immutable copy of data is put into a clean room, it is then inspected for infection. If it is determined to be clean, it is then locked so it cannot be deleted. It is then available to be moved to the next level of clean room and can be used for recovery if necessary. Each successive level of clean room is more secure, and therefore less likely to be affected by a malware attack than the previous clean room. Any time a new ransomware detection tool, process, or new ransomware virus definition is deployed, Hewlett Packard Enterprise recommends that it be run against the last copy of data that was saved in each clean room to check for intrusion that was not detected by the older version of the tools.

How frequently new copies of data are created and moved into a clean room and how long the data is kept before aging out are the functions of a number of factors. Like a disaster-tolerant solution, the choices are most often driven by cost. The smaller the RTO desired from the solution, the more expensive the solution becomes. Data needs to be stored for a longer period of time in expensive clean rooms closer to the production environment to help deliver a smaller RTO. Ransomware attacks are insidious and might not actually be detected or announced until quite some time after the attack has initially occurred. The farther a clean room is from the production volumes, the safer the data is but the longer it will take to restore the production volumes back to a usable state if that clean room data must be used, ultimately increasing the RTO.

It can become very expensive to keep enough clean room data near the production array to ensure a small RTO. Figure 3 provides some perspective on the timeline of a ransomware attack.

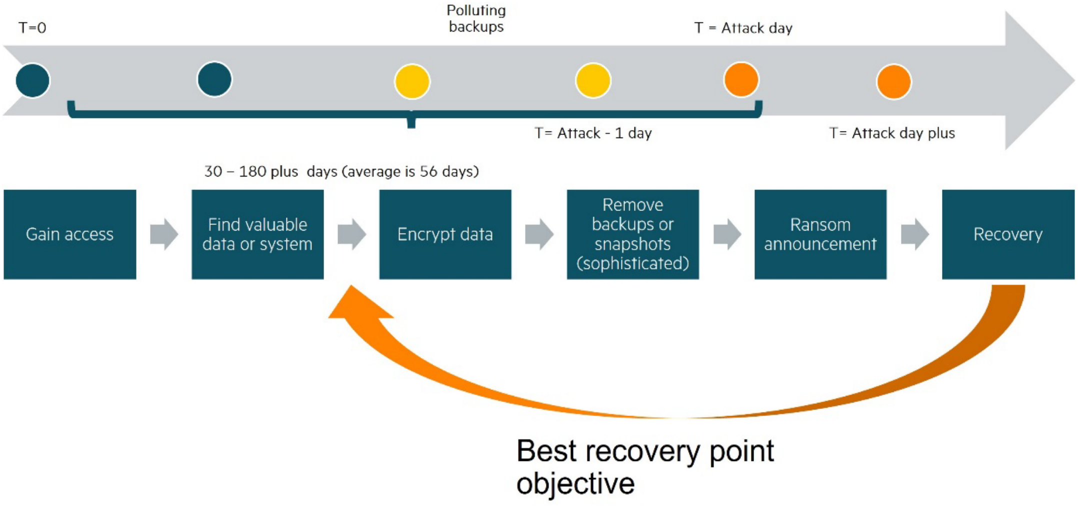
T=0
Polluting backups
T = Attack day
T= Attack - 1 day
T= Attack day plus
30 – 180 plus days (average is 56 days)
Gain access
Find valuable data or system
Encrypt data
Remove backups or snapshots (sophisticated)
Ransom announcement
Recovery
Best recovery point objective

FIGURE 3. Ransomware attack timeline

Data gets into a clean room through one of three means:

1. Immutable (read-only) snapshots that have a retention time assigned to them
2. Immutable backups, including having an immutable backup catalog
3. Array replication (after being replicated to a target array, an immutable copy of the data is created via read-only snapshots on the target array with a retention time)

The next sections take a deeper look at each of the different types of clean room.

Immutable, read-only, snapshots with retention time

The first level of clean room is created on the production enterprise array and consists of read-only snapshots with a defined retention time (virtual lock) applied to them. Because these snapshots are created as read-only they cannot be written to, and because they have a retention time applied to them, they cannot be deleted until the retention time expires. After they are created, they are inspected for infection and if found to be clean, they are then locked. They are then available to provide data to the next level of clean room. Keep in mind that a volume with a retention time applied to it cannot be removed from the array under any circumstances until the retention time expires, so Hewlett Packard Enterprise suggests setting retention times for a reasonable duration that are regularly extended.

IMPORTANT
A volume (BaseVV or snapshot) with a retention time applied to it cannot, under any circumstances, be removed from the array until the retention time expires. This includes production volumes with snapshots that have a retention time applied to them.
Hewlett Packard Enterprise recommends that retention times of a reasonable duration be applied to snapshots and that those retention times be extended when they expire.

NOTE
If the array runs out of space, snapshots on the array will go stale and will not be available to use for recovery. Monitor space usage on the array to ensure it does not run out of space. Unusually high space consumption on the array is a sign of a possible ransomware attack that is encrypting data on the array.

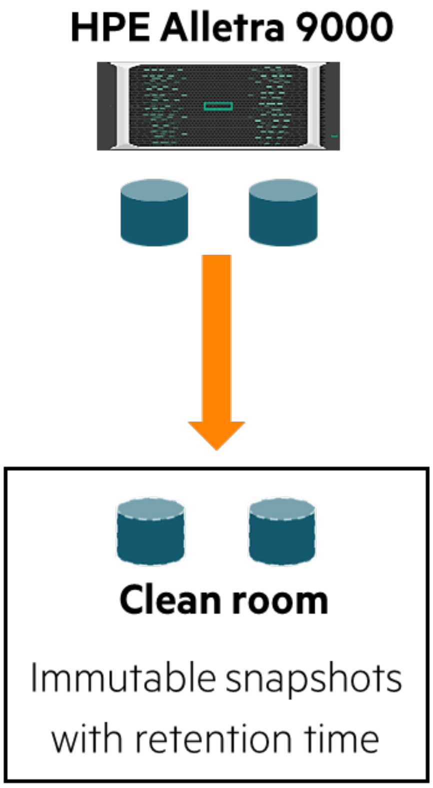

FIGURE 4. Basic snapshot clean room

Backup clean rooms

The next level of clean room is provided with immutable and offline backups. These backups can land on a public or private cloud if desired, but Hewlett Packard Enterprise highly recommends that off-line tape also be used. Landing the data on a public or private cloud provides a copy of the data that can be used for relatively quick recovery if necessary. Even when immutable backup to cloud is used, Hewlett Packard Enterprise recommends also creating an offline backup to tape. An offline air-gapped tape copy of the data is very safe, and it is difficult for it to be compromised by malware. This is why Hewlett Packard Enterprise highly recommends an offline tape copy of the data for the highest level of protection.

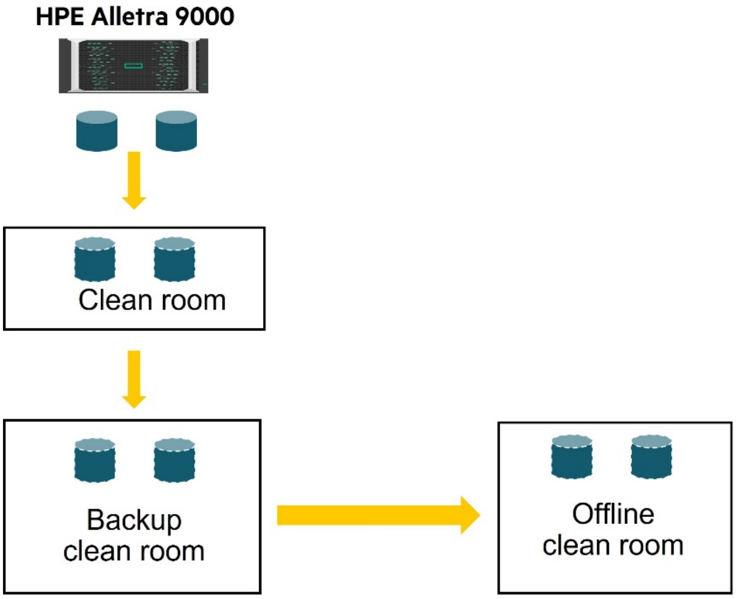

FIGURE 5. Backup clean room and offline tape clean room

Array replication

Array replication can be leveraged to provide an extra level of clean room protection that can be used to provide data to the backup clean room. The replication target array is not used for production, test, or development—it is only a clean room replication target. You can think of it as a vault that no production servers should have access to. It needs to be hardened and restricted to protect it from a direct ransomware attack.

The only servers that should have access to the clean room vault array are the servers used to check data for infection and servers used for backup. The backup servers should only be given access to snapshots of the replicated volumes on the array, not to the replication target volumes themselves. All manner of precaution should be taken to isolate any server connected to the clean room vault array from all networks and attack vectors that can potentially infect the production servers. Direct connect replication links are recommended versus switched links. Replication target volumes on the clean room vault array should never, under any circumstance, be exported to any production servers. The clean room vault array is not intended as and should not be used as a DR failover target. Exporting these volumes to production servers exposes them to the very ransomware attack vectors you are trying to protect from.

Periodically production data is replicated to the clean room vault array. After it is there, it is then inspected for infection and if determined to be clean an immutable copy is placed in the vault clean room, is locked, and can then be used to feed a backup clean room. Hewlett Packard Enterprise recommends that backup clean rooms use the clean room vault array as opposed to using the production array directly as the backup source because a vault array provides an extra level of protection for the data.

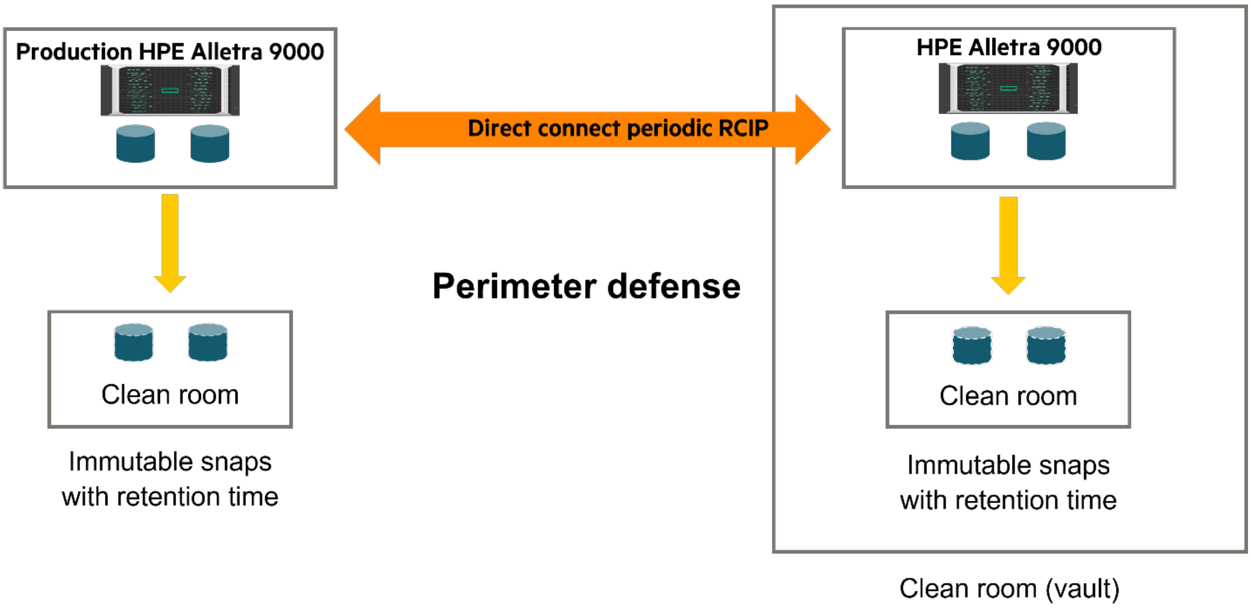
Production HPE Alletra 9000
Direct connect periodic RCIP
HPE Alletra 9000
Perimeter defense
Clean room
Immutable snaps with retention time
Clean room
Immutable snaps with retention time
Clean room (vault)

FIGURE 6. Array replication clean room vault

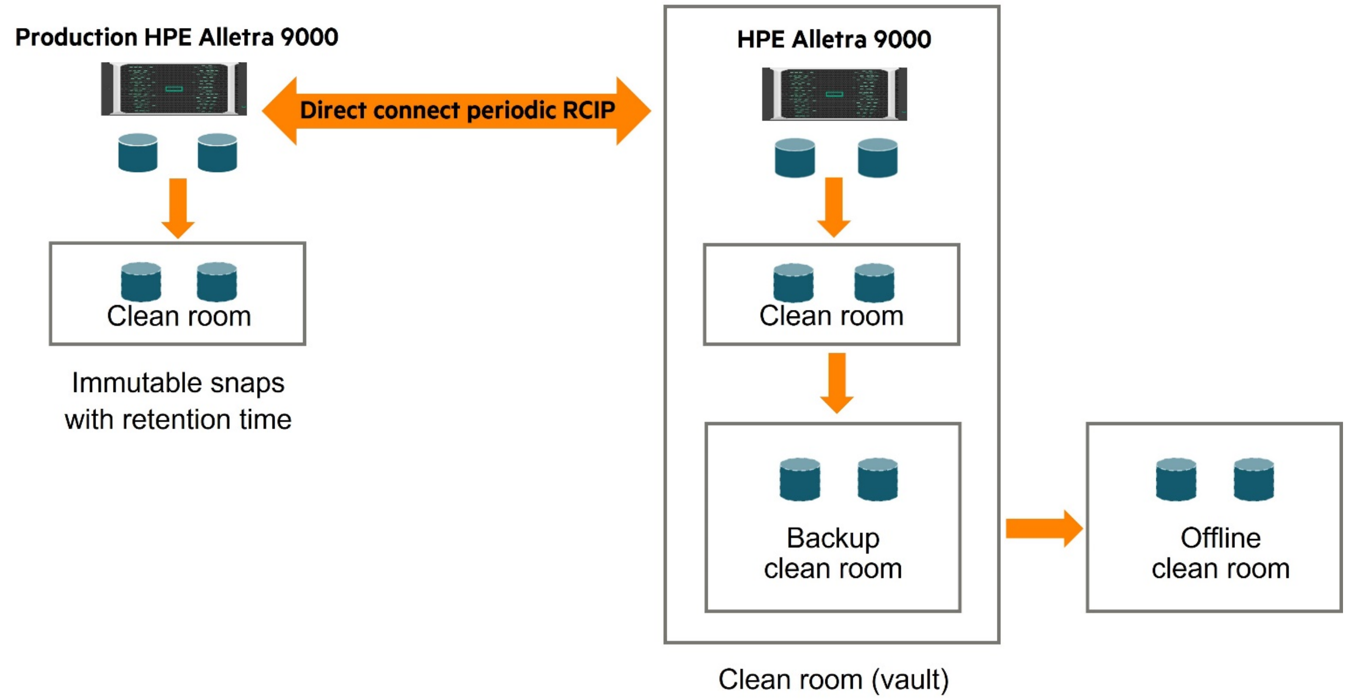
Production HPE Alletra 9000
Direct connect periodic RCIP
HPE Alletra 9000
Clean room
Immutable snaps with retention time
Clean room
Backup clean room
Offline clean room
Clean room (vault)

FIGURE 7. Backup clean room on the clean room from the vault array

Backup clean rooms

The backup clean rooms referenced in these architectures are based on HPE StoreOnce appliance. The use of HPE StoreOnce is not required but is highly recommended for the backup control and space savings it can provide.

Backups should be taken from data that resides in at least one level of clean room removed from the production data and that has been inspected for infection to the greatest extent possible before being backed up.

Immutable backups

The backup data (backup store) and backup catalog must be stored on immutable storage to ensure they cannot be deleted or overwritten during a ransomware attack. The backup catalog created by the backup software of choice must be saved on immutable storage so that it cannot be deleted or otherwise compromised, which would render the backups unusable. A backup to tape can be used to restore backups even if the backup store or backup catalog has been compromised. A backup to tape where the tape cartridge has been removed from the tape library, known as “air-gapped,” is the highest level of backup protection so including backup to tape in addition to backup to cloud is highly recommended.

NOTE
Regardless of where the backup store and catalog reside, Hewlett Packard Enterprise recommends tape also be used as an offline backup target for further protection from a possible ransomware attack. A backup to tape can be used to recover even if the backup catalog or backup store are compromised.

NOTE
An air-gapped backup to tape where the tape cartridge is removed from the tape library is highly recommended even when data is backed up to private cloud, public cloud, or a local storage device.

Data Immutability with HPE StoreOnce

HPE StoreOnce can provide immutability for Catalyst backup stores. A Data Immutability flag is set on a Catalyst backup store so that the content of the backup store cannot be deleted or overwritten by the HPE StoreOnce Catalyst client until the flag expires. You can also set HPE StoreOnce immutability via a Data Immutability flag on a Cloud Bank store.

The Data Immutability flag must be set manually via the HPE StoreOnce GUI interface (it cannot be set via the Catalyst Client API) and needs to be of a long duration because a ransomware attack might not become evident until many months after it has occurred. Exactly how long the Data Immutability flag should be set on a backup store is non-deterministic; it must be longer than the time period from when a ransomware attack occurs to the time it is discovered and there is no way to tell how long this will be.

Note that immutability set by HPE StoreOnce is not the same as the S3 Object Lock on a data store in the cloud when public or private cloud are used with Cloud Bank. There is currently no way to set S3 Object Lock on a cloud data store used by HPE StoreOnce Cloud Bank Services, and, in fact, HPE StoreOnce Cloud Bank services cannot use a cloud data store with S3 Object Lock applied to it.

NOTE
The data immutability offered HPE StoreOnce does not use S3 Object Lock in a private or public cloud used for Cloud Bank storage. The data immutability is set by and controlled at the HPE StoreOnce appliance.

Today a Data Immutability flag applied to a backup store via HPE StoreOnce can be removed by anyone with credentials that gives them access to the HPE StoreOnce appliance and the backup store. This mean a disgruntled employee in the data center or a bad actor who has compromised the credentials necessary to access the StoreOnce appliance could, if they so choose, delete backup stores.

RANSOMWARE DATA RECOVERY ARCHITECTURES

Data Recovery Architecture #1

The first data recovery architecture is a very basic architecture that is actually the foundation for all other architectures discussed in this paper. Hewlett Packard Enterprise does not recommend this architecture because it does not adhere to the best practice 3-2-1 rule for data protection. However, for a customer with a limited budget who cannot afford to deploy one of the other architectures, it is better than doing nothing.

NOTE
Hewlett Packard Enterprise does not recommend data recovery architecture #1. It does not conform to the 3-2-1 best practice rules for data protection.

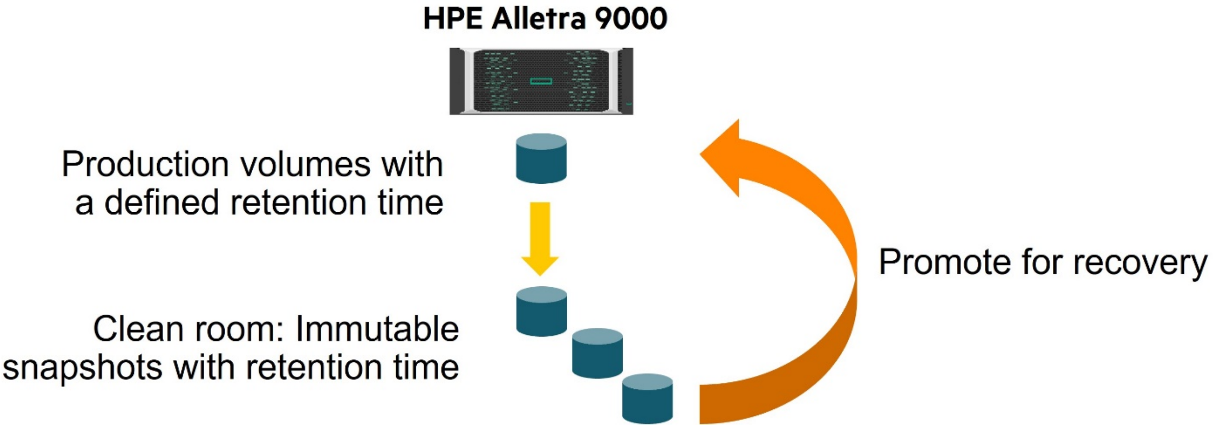
HPE Alletra 9000
Production volumes with a defined retention time
Clean room: Immutable snapshots with retention time
Promote for recovery

FIGURE 8. Data Recovery Architecture #1—Basic snapshot clean room

The clean room in this architecture contains a set of immutable (read-only) snapshots that are created from the production data volumes. After these snapshots are created, they are inspected for infection and if found to be clean, a retention time (virtual lock) is then applied to them so they cannot be deleted by a malware attack. Hewlett Packard Enterprise recommends that the production volumes also have a defined retention time applied to them so they cannot be deleted by any malware. Consideration must be given to the retention time placed on the clean room snapshots. A volume with a retention time applied to it cannot, under any circumstance, be removed from the system until the retention time expires. Setting retention times of extremely long duration can result in a situation where a volume that is desired to be removed or required to be removed cannot be removed in a timely manner. The retention times chosen should be of short to moderate duration and can be extended if necessary. This provides the flexibility to remove volumes and snapshots if the need arises. If a ransomware attack occurs, the newest set of clean, uninfected clean room snapshots can be promoted to recover the production volumes back to a state before the malware attack occurred.

How frequently a new set of clean room snaps are created on the production array and how long they are maintained is a function of several things:

- **The RPO desired by the customer.** RPO is a definition of the amount of data that can be lost if a malware attack occurs and is defined in an amount of time. The smaller the desired RPO, the more frequently a new set of snapshots must be placed in the clean room. There is no guarantee that the desired RPO will be achieved; the malware attack may have occurred long in the past, rendering the desired RPO unachievable.
- **System configuration.** There are system limits on the total number of volumes supported and the number of snapshots allowed per VV. Also, snapshots utilize capacity on the array as the volumes they are created from are written to. These limits and how long a set of snapshots is retained may limit how frequently a new set of snapshots can be placed into a clean room.
- **RTO desired.** RTO is a measure of how long it takes to recover from an attack and to get back up and running. Before a set of snapshots in the clean room can be used for recovery, they must be inspected and be declared clear of infection. Although frequent clean room snapshots may help provide a small RPO, if the clean room is being populated with snapshots too frequently, it might not be possible for the antivirus software and other processes to inspect and certify they are clean quickly enough to keep up with a desired small RTO. For example, consider daily clean room snapshots that are taken on the production array. If a malware attack is detected 30 days from today, there are 30 days' worth of clean room snapshots to potentially check. If the last good clean set of snapshots was taken 15 days ago, then 15 sets of snapshots must be checked before the first good clean set is discovered. However, if clean room snapshots are taken four times per day, then 60 sets of snapshots must be checked before the first clean one is found. This situation affects the solution's overall RTO.
- **Duration.** A ransomware attack can be months in the making, so it is not realistic to assume snapshot clean room data can be retained long enough on the production array to guarantee detection of a ransomware attack before data starts aging out of the clean room. Because the duration of time from when a ransomware attack occurs until it is detected ultimately is non-deterministic and can be weeks to months, recovery will most likely have to be addressed by a backup clean room that provides a long-term, inexpensive, offline copy of the data created from the clean room snapshots. These are covered in Data Recovery Architectures 3 through 6.

Data Recovery Architecture #2

This data recovery architecture extends the base architecture #1 to include the ability to protect data for virtualized environments via Zerto Continuous Data Protection (CDP) and Zerto backup to an immutable long-term repository (LTR). The architecture only works for virtualized environments running VMware® and Microsoft Hyper-V.

This architecture does not fit the description of a backup clean room as defined in this paper. A backup clean room should be separated from the production data and servers by at least one and preferably two levels of clean room (immutable snapshots on the production array or immutable snapshots on a vault array). With Zerto backup, the production VMware vCenter® has a Zerto Virtual Manager (ZVM) VM and a Zerto Virtual Replicator (VRA) VM installed on it. This protects user-selected VMs with a continuous data protection (CDP) journaling mechanism and backs up data via HPE StoreOnce to an LTR backup store in a public or private cloud as seen in Figure 9. The fact that the ZVM and VRA VMs run on the same Hyper-V or VMware ESXi™ or servers as the virtual machines running production workloads means the Zerto infrastructure is exposed to the same ransomware attack vectors as the production VMs being protected from ransomware.

Because the Zerto backup solution archives data directly from the production VMs in the solution rather than from clean room snapshots that have been inspected for infection, Zerto provides no means for inspecting data for infection before Zero archives to its LTR. The only way of detecting infection is by inspecting the immutable snapshots of the production data created on the production HPE Alletra 9000. Also, Zerto backup does not support backup to tape so an air-gapped backup to tape, as is recommended by Hewlett Packard Enterprise, cannot be created if the Zerto backup solution alone is chosen. If the Zerto solution is your solution of choice, Hewlett Packard Enterprise recommends also deploying a separate backup solution that leverages the clean room snapshots to provide backup to tape.

Figure 9 shows a Zerto cluster with a production site and a DR site. The solution can be deployed in this manner if traditional DR is desired in addition to protection from ransomware. In the case of a site disaster, failover can occur to the target site. To save cost and if only ransomware protection is desired (no DR failover), the solution can be deployed with just the source site performing both the CDP and the archive to the Zerto LTR.

Configuring CDP is required for Zerto backup to an LTR store. It should be noted that with Zerto, CDP can be configured to hold data for as long as 30 days for DR recovery. The amount of storage required in the CDP store can become quite large depending on the number of days' worth of changes you have chosen to store, size of the production data, and change rate on the production data. Because most ransomware attacks have a duration much greater than 30 days from the time of the attack until the time of the announcement, Hewlett Packard Enterprise recommends configuring minimal CDP protection unless you are also using the solution to provide DR protection.

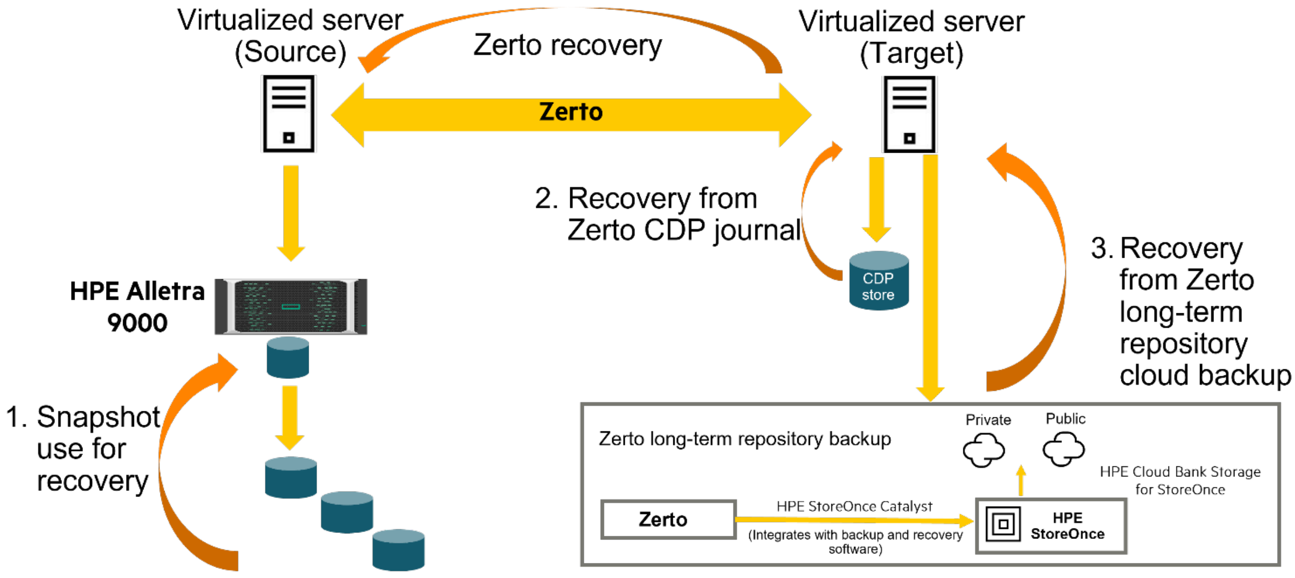
Virtualized server (Source)
Zerto recovery
Virtualized server (Target)
Zerto
2. Recovery from Zerto CDP journal
CDP store
3. Recovery from Zerto long-term repository cloud backup
HPE Alletra 9000
1. Snapshot use for recovery
Zerto long-term repository backup
Private
Public
HPE Cloud Bank Storage for StoreOnce
Zerto
HPE StoreOnce Catalyst
(Integrates with backup and recovery software)
HPE StoreOnce

FIGURE 9. Data Recovery Architecture #2—Zerto Long Term Repository Backup

In this solution, if a ransomware attack occurs, recovery should be attempted first from the snapshot clean room on the production array if the data is held there longer than it is in the CDP store. If no clean data is found there, recovery is attempted from the CDP store. If clean data does not reside in the CDP store, recovery is then attempted from the Zerto LTR backup.

Data Recovery Architecture #3

This solution extends the base architecture #1 to include a backup clean room that can provide onsite, cloud, and offline tape backup. This solution does adhere to the best practice 3-2-1 rule for data protection and the HPE 3-2-1-1 rule if an air-gapped tape backup is included.

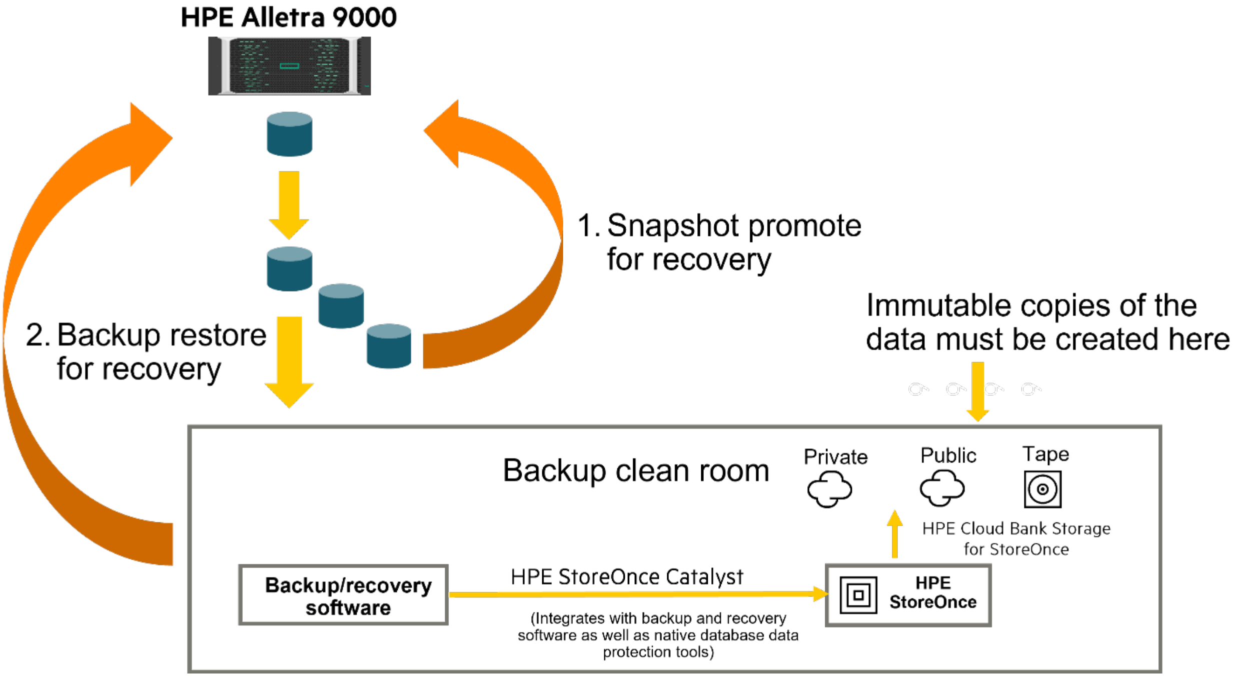
HPE Alletra 9000
1. Snapshot promote for recovery
2. Backup restore for recovery
Immutable copies of the data must be created here
Backup clean room
Private
Public
Tape
HPE Cloud Bank Storage for StoreOnce
Backup/recovery software
HPE StoreOnce Catalyst
(Integrates with backup and recovery software as well as native database data protection tools)
HPE StoreOnce

FIGURE 10. Data Recovery Architecture #3—Backup clean room with offline tape

In this solution, data from the production array clean room is used to feed the backup clean room. If the backup software in use requires read/write copies of the data, then read/write snapshots of the production array clean room snapshots should be created and used by the backup software. It is best if the backup solution does not use the production VVs, but instead uses the production array clean room snapshots that have been inspected and declared free from infection. The backup can land wherever the customer wants, including a private or a public cloud. However, best practice would be to ensure at least one copy lands on tape. If all possible, the tape should be removed from any tape library (becoming an air-gapped copy of the data).

If a ransomware attack occurs, snapshots in the production array clean room are checked for infection. If a clean copy of the data is present, then that copy of the data is promoted to the production volumes. If no clean copy of the data exists in the production array clean room, then successive backups must be restored until a clean copy of the data is found. Having to restore data from the backup clean room can have a significant effect on the RPO of the solution.

A drawback to this architecture is that because both the production servers and the backup servers are connected to and are using the production array, they are subject to common attack vectors. The backup servers should be isolated as much as possible from the networks that the production servers are using (both IP and SAN) to make it more difficult for a single attack vector to affect both the production and backup servers. In essence, they should be given a separate physical perimeter defense to the extent possible even though they share a common array.

Data Recovery Architecture #4

This data recovery architecture uses periodic asynchronous Remote Copy to place a point-in-time copy of data from the production array onto a separate clean room vault array. The vault array should not be used as a common DR site where production workloads are failed over and run on the replication target array. Doing so opens the vault up to a possible ransomware attack. The vault array is intended solely as a clean room to hold immutable copies of data to be used for ransomware recovery and for use as a location for a backup clean room (see Data Recovery Architecture #5 for details). The vault array and related backup and virus scan servers should be isolated as much as possible from the production environment with their own SAN and IP networks and switches. The intent of the architecture is that the vault array should be more secure than the production array by limiting outside access to it. The only servers that should be connected to the vault array in this architecture are the servers used to run virus scan to check for infection and in the case of architecture #5, the backup servers.

Data moves from the production array to the vault array via Asynchronous Periodic Remote Copy. Hewlett Packard Enterprise recommends that the transport between the two arrays be direct-connect Remote Copy IP (RCIP) or Remote Copy Fibre Channel (RCFC) via SAN switches. (Direct connect RCIP is preferred.) If direct connect RCIP cannot be used, then any switches used to connect the production and vault arrays together to provide an RCIP transport should be isolated as much as possible from the corporate network to prevent them from being used as an attack vector. The clean room on the vault array contains immutable snaps of the data with retention times applied to them.

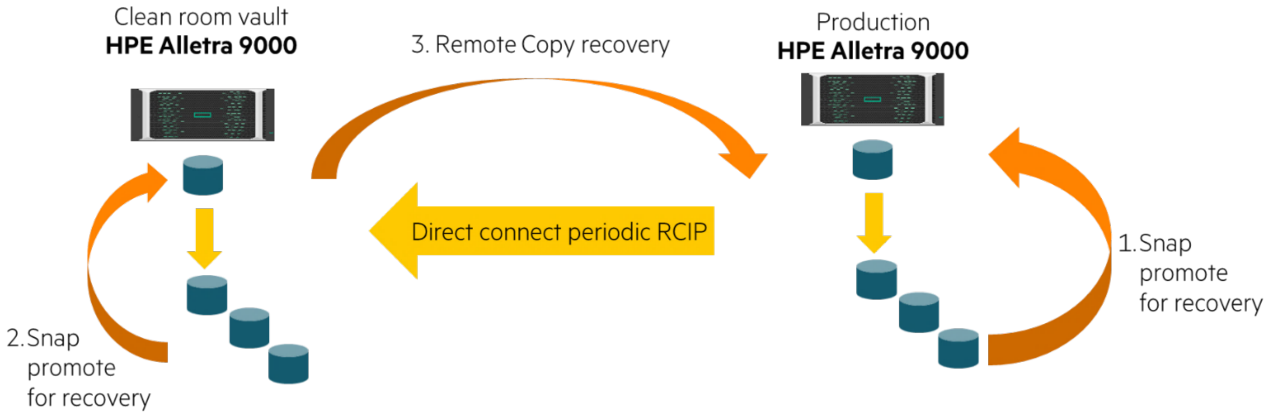
Clean room vault
HPE Alletra 9000
3. Remote Copy recovery
Production
HPE Alletra 9000
Direct connect periodic RCIP
1. Snap promote for recovery
2. Snap promote for recovery

FIGURE 11. Data Recovery Architecture #4—Replicated Clean room Vault array

NOTE
Best practice is that direct connect RCIP be used as the Remote Copy transport between the arrays in the Remote Copy configuration but RCFC with SAN switches can be used as the transport between arrays if desired.

Hewlett Packard Enterprise recommends that periodic asynchronous Remote Copy not be configured for regular delta updates of the vault array. Instead, updates to the vault array should be controlled by setting the Remote Copy group delta resync interval to zero and updating the vault array via the `syncrcopy` CLI command on a planned schedule or by requesting coordinated snapshots of the volumes in the RC group via the `createsv -rcopy` command (see Figure 12). After a Remote Copy resync to the vault array has completed a set of immutable snaps are then created of the replication target volumes. The replication target volumes on the vault array are then checked for infection. If none is found, a retention period is applied to the clean room snaps. After this is complete, another resync can occur immediately afterward if desired. The advantage of this process over using a defined delta resync interval for the Remote Copy group is that all clean room data on the vault array is always checked for infection and is ready for use if necessary.

NOTE
Best practice: Hewlett Packard Enterprise recommends that updates to the vault array be controlled via the `syncrcopy` CLI command executed on a controlled schedule rather than allowing Remote Copy to replicate changes based on a delta resync interval.

Using coordinated snapshots to update the vault array

It is possible to create coordinated snapshots between the production array and vault array. Coordinated snapshots are snapshots taken on the Remote Copy primary and secondary arrays that are created at the exact same point in time. These coordinated snapshots created on the arrays are then used as the clean room snapshots (more frequent clean room snapshots can be created on the production array if desired). This is represented in Figure 12.

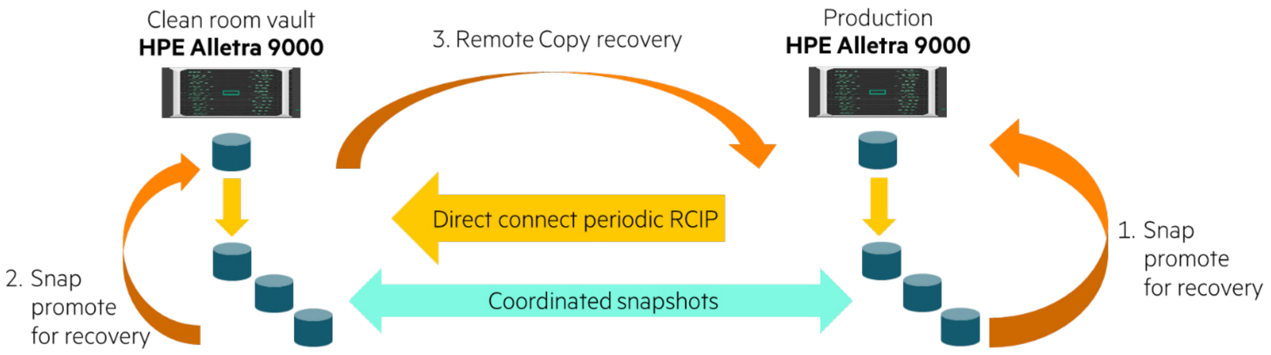

FIGURE 12. Coordinated snapshots to clean room vault array

If the delta resync interval on the Remote Copy group is set to zero, requesting a coordinated snapshot for the Remote Copy group results in a delta resync of the vault array before the coordinate snapshot is taken there. After the vault array is updated via the Remote Copy resync and the coordinated snapshots have been taken, the data on the vault array can be inspected for infection.

NOTE
A cyclic redundancy check (CRC) check can be run on the replication target volumes on the vault array after a delta resync from the primary array has completed with a second CRC check run just before the next resync occurs. When compared, if the two CRCs are different, it is an indication that the data on the vault array has been changed between Remote Copy resyncs (data on the vault array should only change as a result of a Remote Copy resync) and indicates a ransomware attack has affected the vault array.

NOTE
Best practice: Set a delta resync interval of zero on the Remote Copy group and perform controlled resyncs of the Remote Copy group. After the resync is complete, the replication target volumes on the vault array can be checked for infection. If the volumes are found to be clean, immutable snapshots with a retention time can then be created on the array.

If a ransomware attack occurs, snapshots in the production array clean room are checked for infection. If a clean copy of the data is present, then that copy of the data is promoted to the production volumes on the array for recovery. If no clean copy of the data exists in the production array clean room, a Remote Copy failover to the vault array must occur. However, production is not run on the vault array. After the Remote Copy failover, successive clean room snapshots in the vault array clean room are checked for infection. When a clean copy of data is identified in the clean room on the vault array, those snapshots must first be promoted to the volumes on the vault array and then that data must be re-synched back to the production array. After the resync back to the production array has completed, a failover back to the production array must occur to get production back online. Production should never be run on the clean room vault array.

Data Recovery Architecture #5

Architecture #5 is considered the best practice architecture for maximum protection against ransomware and is the data recovery architecture recommended by Hewlett Packard Enterprise. This data recovery architecture moves the backup clean room off the production array and to the vault array discussed in Data Recovery Architecture #4. This architecture makes the backup clean room more secure than having it on the production array.

NOTE
This is the data recovery architecture recommended by Hewlett Packard Enterprise for maximum data protection against a ransomware attack.

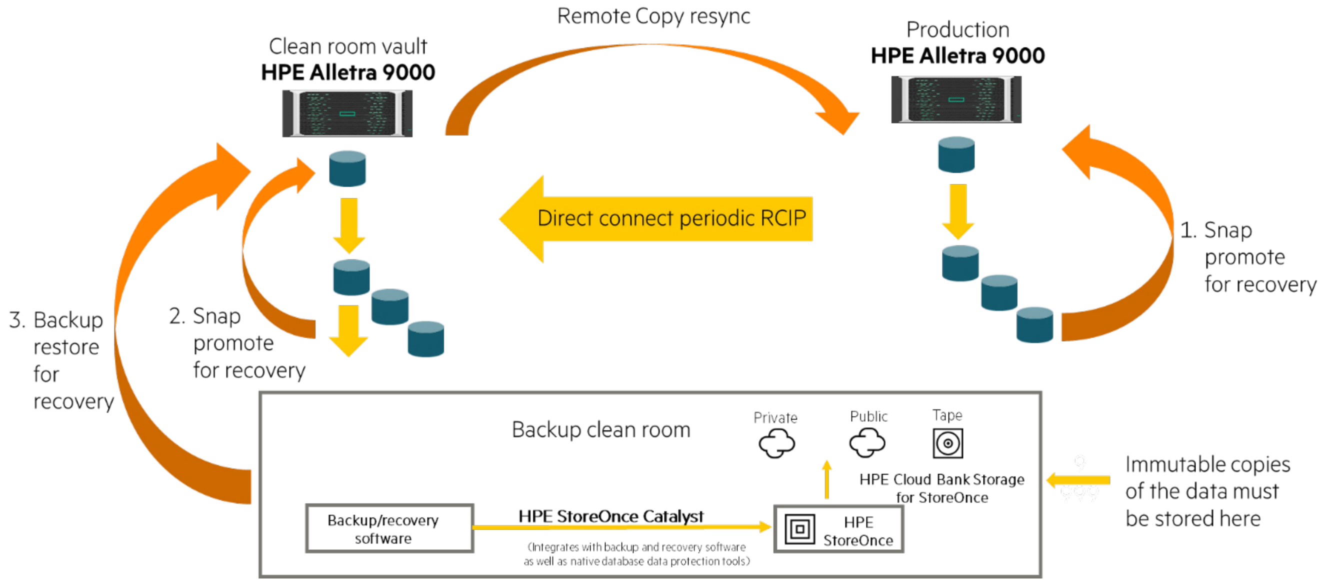
Remote Copy resync
Clean room vault
HPE Alletra 9000
Production
HPE Alletra 9000
Direct connect periodic RCIP
1. Snap promote for recovery
2. Snap promote for recovery
3. Backup restore for recovery
Backup clean room
Private
Public
Tape
HPE Cloud Bank Storage for StoreOnce
Backup/recovery software
HPE StoreOnce Catalyst
(Integrates with backup and recovery software as well as native database data protection tools)
HPE StoreOnce
Immutable copies of the data must be stored here

FIGURE 13. Data Recovery Architecture #5—Backup clean room off of the vault array

If a ransomware attack occurs, snapshots in the production array clean room are checked for infection. If a clean copy of the data is present, then that data is promoted to the production volumes for recovery. If no clean copy of the data exists in the clean room on the production array, a Remote Copy failover to the vault array must occur. Successive clean room snapshots in the vault array clean are checked for infection. If a clean copy of data is identified in the clean room on the vault array, those snapshots must first be promoted to the Remote Copy volumes on the vault array and then that data must be re-synched to the production array. After the data is re-synched, a failover back to the production array occurs to get production back online. If no clean copy of the data is found in the vault array clean room, then successive backups must be restored to the vault array and checked until a good copy of the data is found. After good data is restored from backup, a Remote Copy re-sync back to the production array must occur followed by a failover back to the production array to get production back online. Production should never be run on the clean room vault array.

Data Recovery Architecture #6

Data Recovery Architecture #6 extends Data Recovery Architecture #5 to include an HA/DR option. With this architecture, the production workload is protected from disaster by two production arrays set up in a Peer Persistence configuration with their data being replicated to the vault array in what is known as a three-data center Peer Persistence (3DC-PP) solution or in a synchronous long distance (SLD) solution. Active Peer Persistence is not supported in a 3DC-PP configuration at this time so it cannot be deployed in this architecture if Peer Persistence is used. Unlike architecture #5, coordinated snapshots cannot be created between the production Peer Persistence arrays and the clean room vault array.

NOTE

Active Peer Persistence is not currently supported in this architecture. Check with your HPE representative regarding support for Active Peer Persistence in a 3DC-PP configuration.

NOTE

Currently, the use of coordinated snapshots between the production arrays and the vault array is not supported in this architecture. Check with your HPE representative regarding support of coordinated snapshots to the asynchronous target array in a 3DC-PP or SLD configuration.

Hewlett Packard Enterprise recommends that periodic asynchronous Remote Copy not be configured for regular delta updates of the vault array. Instead, updates to the vault array should be controlled by setting the Remote Copy group delta resync interval to zero and updating the vault array by using the `syncrcopy` CLI command on a planned schedule for the group. After the Remote Copy resync to the vault array has completed, the volumes on the vault array should be checked for infection. If no infection is found, a set of immutable snaps with retention times are then created on the vault array. After this is complete, another resync to the vault array can occur immediately afterward, if desired. The advantage of this is that all clean room data on the vault array is always checked for infection and is ready for use if necessary, resulting in a better RTO.

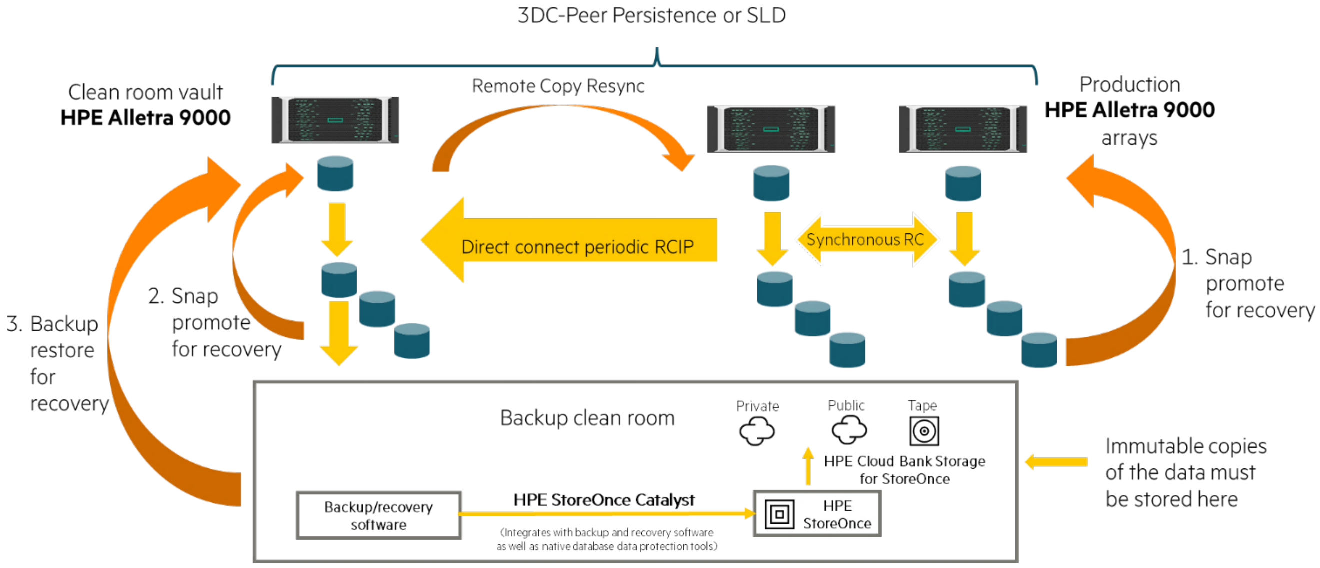

FIGURE 14. Data Recovery Architecture #6—Disaster Recovery Solution

If a ransomware attack occurs, snapshots in the clean rooms on the production arrays are checked for infection. If a clean copy of the data is present, then that copy of the data is promoted to the production volumes. If no clean copy of the data exists in the clean rooms on the production arrays, a Remote Copy failover to the vault array must occur.

After the failover, clean room snapshots in the clean room on the vault array are checked for infection. If a clean copy of data is identified in the clean room on the vault array, those snapshots must first be promoted to the Remote Copy volumes on the vault array and then that data must be re-synched to the production 3DC-PP or SLD arrays. After the data is re-synched, a failover back to the production Peer Persistence arrays occurs to get production back online.

If no clean copy of the data is found in the vault array clean room, successive backups must be restored to the vault array and checked until a good copy of the data is found. After a good backup copy of data is restored to the vault array, then a Remote Copy re-sync back to the production arrays must occur followed by a failover back to the production arrays to get production back online.

Resources, contacts, or additional links

https://www.hpe.com/psnow/doc/c04328820.pdf

https://help.zerto.com/

https://www.zerto.com/solutions/use-cases/security-and-compliance/ransomware/

HPE SPOCK

HPE Alletra 9000: Getting started with data replication using Remote Copy

HPE Alletra 9000: Configuring and managing data replication using Remote Copy

HPE Alletra 9000: Recovering from disaster using Remote Copy

LEARN MORE AT

hpe.com/storage

Make the right purchase decision.
Contact our presales specialists.

Chat Email Call

Get updates

a00119614ENW, Rev. 1

Hewlett Packard
Enterprise