

VMware Virtual Machine Protection Using Transparent Snapshots

Abstract

This white paper provides insights on how to protect and restore VMware virtual machines using Transparent Snapshots, available with Dell EMC PowerProtect Data Manager.

September 2021

Revisions

Date	Description
September 2021	Initial release

Acknowledgments

Author: Debyeet Bagchi

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [9/10/2021] [Technical White Paper] [H18884]

Table of contents

- Revisions.....2
- Acknowledgments.....2
- Table of contents3
- Executive summary.....4
- 1 Introduction.....5
- 2 Transparent Snapshots architecture6
- 3 Integration with PowerProtect Data Manager8
- 4 Transparent Snapshots life cycle10
 - 4.1 Full sync operation10
 - 4.2 Transparent Snapshots creation11
 - 4.3 Delta sync operation12
 - 4.4 Snapshot retire operation13
- 5 Restoring virtual machines15
- 6 Performance test results17
- A Technical support and resources19
 - A.1 Related resources19

Executive summary

This white paper provides insights on how to protect and restore VMware virtual machines using Dell EMC Transparent Snapshots, available with PowerProtect Data Manager. This document details the architecture and life cycle of Transparent Snapshots and describes how this new capability integrates with PowerProtect Data Manager. It also includes an overview of the process for restoring VMs, and presents test results that show the performance benefits of this solution.

1 Introduction

The VMware virtual machine backup process transfers or exports data from a virtual machine (VM) within a VMware environment to a secondary protection storage system. This system, such as a Dell EMC PowerProtect appliance, can be on a primary or secondary site or on the cloud. A backup engine or software such as PowerProtect Data Manager manages this process. PowerProtect Data Manager can perform data- and copy-management operations on the backup copies and ensure that all data is cataloged properly. This function makes available a consistent VM copy as part of a restore requirement in a disaster scenario.

PowerProtect Data Manager can protect VMware virtual machines in a reliable and efficient manner using VMware API for Data Protection (VADP) snapshots (see Figure 1). These VADP snapshots are reliable, proven, and certified by VMware, and can be used as part of backup operations.

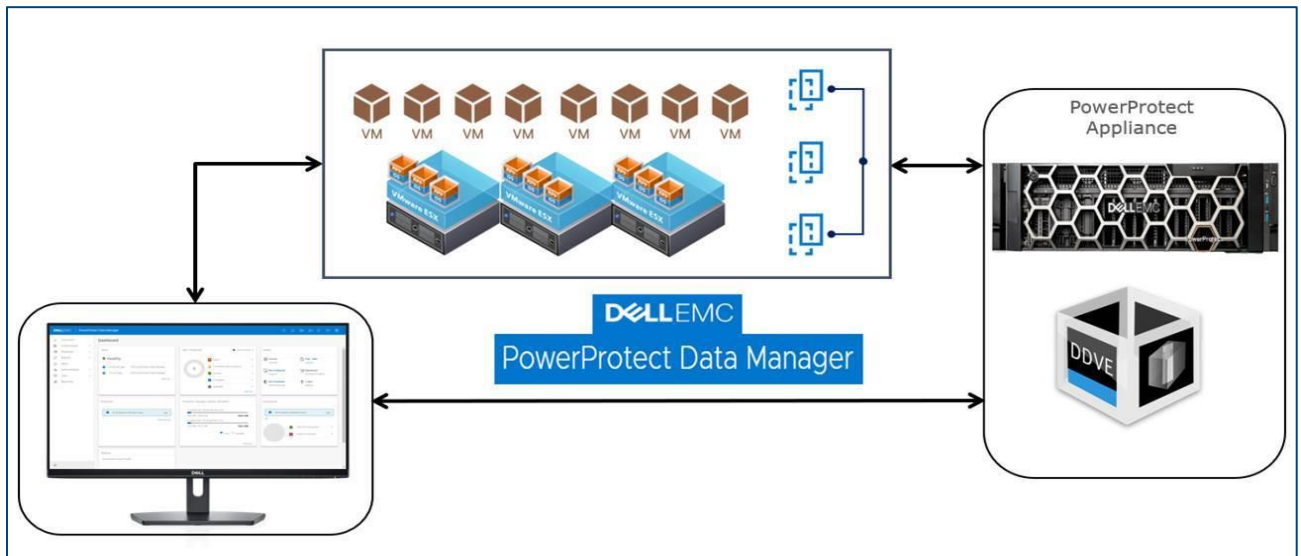


Figure 1 PowerProtect Data Manager for VMware

However, the VADP snapshot process pauses the execution of the VM and allows in-flight disk I/Os to complete. This action may increase the read and write latency and affect the snapshot and VM ecosystem life cycle. When analyzing the life cycle of VADP snapshot, the snapshot entry and exit points inflict a penalty on a VM. After producing a snapshot of a VM disk file, requiring the VM to be stunned, a snapshot of the VM disk file is ingested and deltas must be consolidated into the base disk. When you create a snapshot of a high-transactional application, like a database, adverse effects include lengthy backup windows and application timeouts when the stun to ingest and to consolidate workflow is not efficiently managed.

Addressing these issues requires a solution which can deliver not only backup and restore capabilities, but also an alternative way to reduce the adverse effects of the VM stun operation.

2 Transparent Snapshots architecture

As represented in Figure 2, PowerProtect Data Manager uses Transparent Snapshots by using a variant of the vSphere API for I/O (VAI/O) Filtering framework. In version 7.0U3, VMware added this framework which deploys components in the VMware ESXi infrastructure to create consistent VM backup copies. It then uses an internal component (TSDM) that is created through the PowerProtect Data Manager VIB to write the copies to the protection storage (PowerProtect appliance). During a restore, these same components use the backup copies.

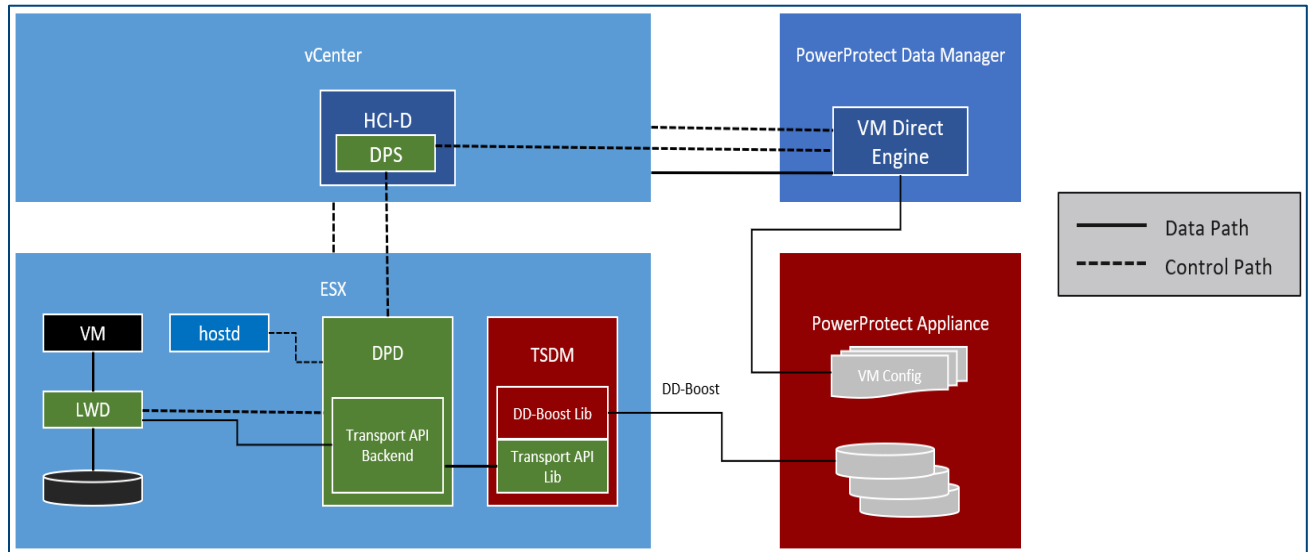


Figure 2 Transparent Snapshots architecture

The following points describe the components to provide you with an understanding of the control and data path as represented in Figure 2:

- PowerProtect Data Manager assumes the role of an orchestrator where it identifies the VM assets in the VMware environment and provides scheduling capabilities.
- PowerProtect Data Manager uses VM Direct Engine to communicate with the VMware vCenter level APIs provided by VMware. The VM Direct Engine connects to a component provided by VMware called Data Protection Service (DPS) which is a vCenter level service and has two key responsibilities in various scenarios:
 - Creates and tracks the progress of the vCenter level tasks that are visible to the end users such as sync, restore, and snapshot operations
 - Responsible for locating the relevant ESXi host on which the operation (backup or restore) is to be performed, based on the placement of the VM asset to be protected
- On each ESXi host, there is a Data Protection Daemon (DPD) provided by VMware that facilitates the protection-related APIs and workflows from VMware.
- The DPD also communicates to the Dell EMC Transparent Snapshot Data Mover (TSDM) component, which is responsible for the VM-backup data movement.
- The backup and restore processes transfer the Transparent Snapshots respectively to and from the PowerProtect Appliance.
- TSDM also consists of the PowerProtect Appliance SDK (DDBoost Library) which helps the framework access the storage units on the PowerProtect Appliance. It also helps write and read data from those storage units.

Note: PowerProtect Data Manager manages the TSDM component by using the VIB (VMware Certified) from Dell Technologies. This component is installed dynamically as part of the integration of PowerProtect Data Manager that requires protection of VMs using Transparent Snapshots. The DPD component is part of the VMware ESXi 7.0 U3 image.

3 Integration with PowerProtect Data Manager

This section examines the steps to integrate PowerProtect Data Manager within the VMware infrastructure, including deploying all necessary components and enabling VM protection using Transparent Snapshots.

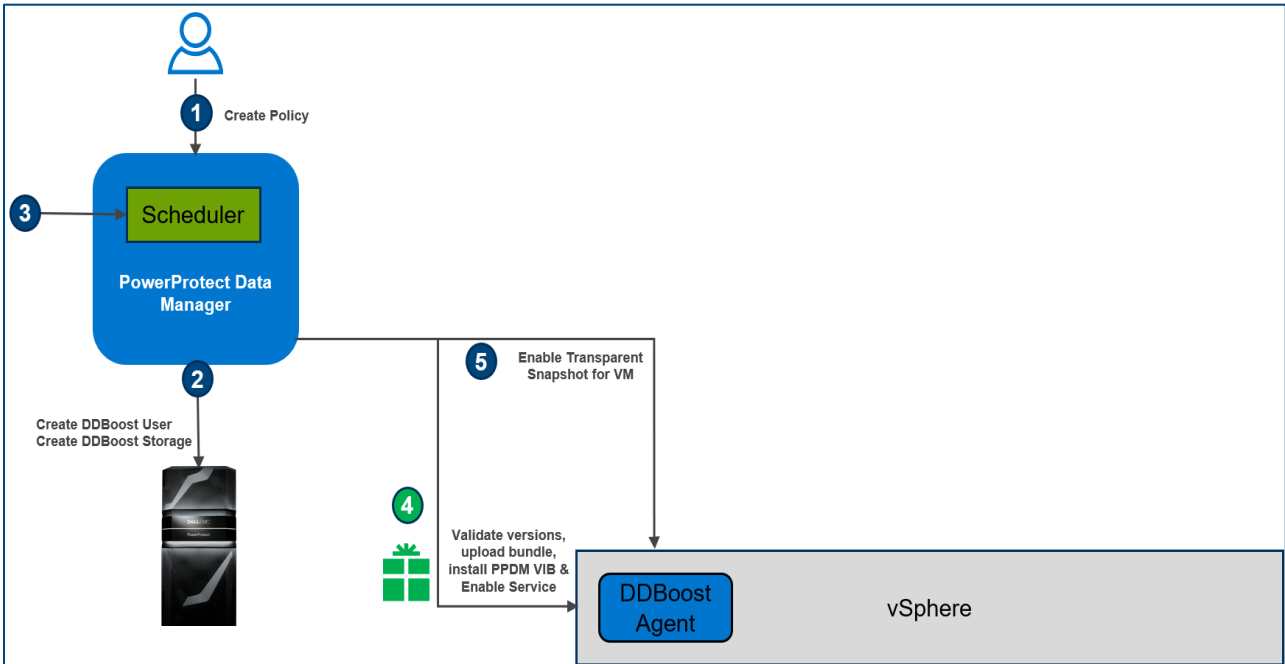


Figure 3 Integration with PowerProtect Data Manager

1. When the VMware VMs have been discovered as PowerProtect Data Manager assets, the next step is to create a protection policy and add the VM assets for protection.

Note: The criteria for policies to be eligible for Transparent Snapshots are: use Crash Consistent Type and Performance Optimized Mode, and disable Guest File System Quiesce and Swap Files Exclusion options.

2. PowerProtect Data Manager creates the Storage unit in the PowerProtect appliance for storing backups.
3. According to the schedule defined in the Protection Policy, the scheduler activates.
4. As these steps occur, the PowerProtect Data Manager VM Direct Engine initiates API calls to the DPS component. It then validates the ESXi version (7.0U3 and above), uploads and installs the PowerProtect Data Manager VIB, and enables the service. The LWD filter is then attached to each VM disk (internally using the DPS and DPD components). In this step, the TSDM component is created but remains idle (running but not used) because there is no data movement. You can see the VIB file installed on the ESXi host that houses the VM being protected (Figure 4).

Summary						
VIB Name	: tsdm					
VIB Operation	: vib_install					
vCenter	: VC					
Hosts Managed	: esxi03.tme.local					
		<input type="checkbox"/>	DEL_bootbank_tsdm_19.9.0-10EM.703.0.0.17990185.v...	4,201.92 ...	07/01/2021, 3:03:22 PM	File [esxi03_DS1] DEL_bootbank
		<input type="checkbox"/>	metadata.zip	3.52 KB	07/01/2021, 3:03:22 PM	File [esxi03_DS1] metadata.zip

Figure 4 The VIB file installed

5. When the DPS components return an acknowledgment of success, PowerProtect Data Manager marks the VM to be protected by Transparent Snapshots.

Note: Aside from the DEL_bootbank VIB file, Figure 4 shows a metadata.zip file that contains information related to the VIB, such as dependencies on the host, system requirements, summary, and version. The files are in the Datastore Files section of the ESXi host. This VIB installation is also shown in the PowerProtect Data Manager Policy Config Job summary.

4 Transparent Snapshots life cycle

Transparent Snapshots provide for a simple, fast, and efficient VM backup that is outlined using the following life cycle:

- Monitor: Track delta changes in memory
- Process: Transfer delta changes directly to protection storage
- Release: Remove delta table and any temporary data blocks

To provide a better understanding of the VM backup process, this section describes the synchronization and data transfer process which consists of four major steps:

- Full sync
- Transparent Snapshot creation
- Delta sync
- Snapshot retire

4.1 Full sync operation

1. During a full sync operation (see Figure 5), PowerProtect Data Manager issues a full sync request that includes all required parameters such as VM information and disk inclusion and exclusion details.
2. The DPS locates the relevant ESXi host, and the operation is transferred to the DPD.
3. The DPD syncs with the TSDM component, uses the VMware Transport API Library, and makes the TSDM aware that a full sync should be performed on the specific asset.
4. The TSDM first uses the DPD Transport APIs to read and query the allocated areas of the disks. After resolving the allocated areas, the TSDM then starts to read the data.
5. The TSDM also uses the DDBoost library to establish a connection to the PowerProtect appliance. Empty files are created in the secondary storage (each file corresponding to the flat VMDK file of the VM asset), and the data transfer begins. Eventually, all allocated areas are transferred and written to the PowerProtect appliance.
6. When the full sync operation is complete, the TSDM sends an acknowledgment to the DPD. The DPS marks the task as complete.

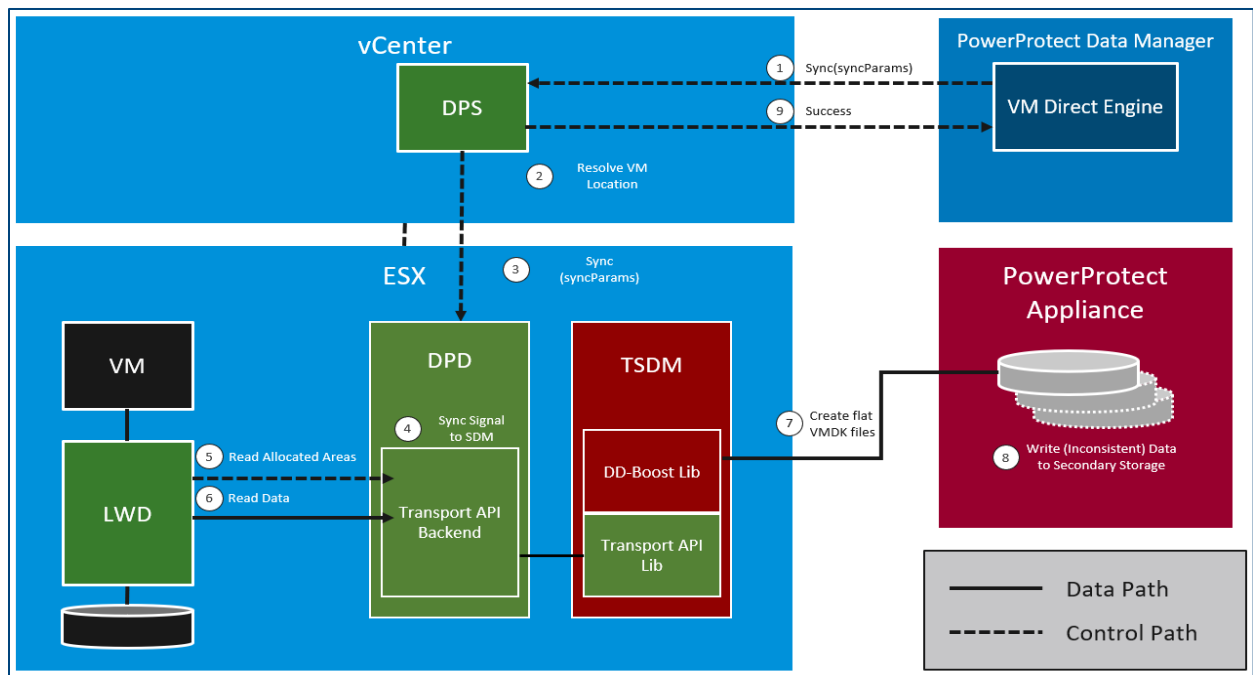


Figure 5 Full sync operation

Note: Because files created during the full sync flow are inconsistent, we cannot use them to restore our VM to a consistent point in time. That is because during the full sync flow, the VM still serves I/Os, therefore data in disks might change during the full sync operation itself. This means that a delta sync operation needs to be performed right after the full sync operation, which creates a consistent point-in-time copy that can be used later for a restore operation.

4.2 Transparent Snapshots creation

1. After the full sync is complete, PowerProtect Data Manager issues a snapshot creation operation (see Figure 6).
2. The DPS API is invoked against the vCenter.
3. The API calls are passed to the DPD component in ESXi after the location is resolved. In this step, the TSDM has no active role because no data transfer occurs.
4. The DPD communicates with all relevant LWD filters for the VM asset, and the snapshot is persisted to the Snapshot Extent Store (SES).

Note: The Snapshot Extent Store (SES) is dynamically created during snapshot creation and is deleted when the snapshot is retired. The SES stores the bitmaps that correspond to the data in the disk at that time. The SES uses thin allocated space across the overall data store space and does not affect the specific VM quota.

5. Using the SES, the LWD Filter takes the bitmap in memory and saves it. Because all of it is bitmap based, creating the LWD snapshot or transparent snapshot is fast, which reduces the read/write latency.
6. After the bitmap is persisted, the LWD Filter can start tracking changes on the disk again, using a new bitmap.
7. The snapshot operation is marked as finished. PowerProtect Data Manager can use the VM Direct Engine to access the vCenter level task completion and get the snapshot UUID that was created.

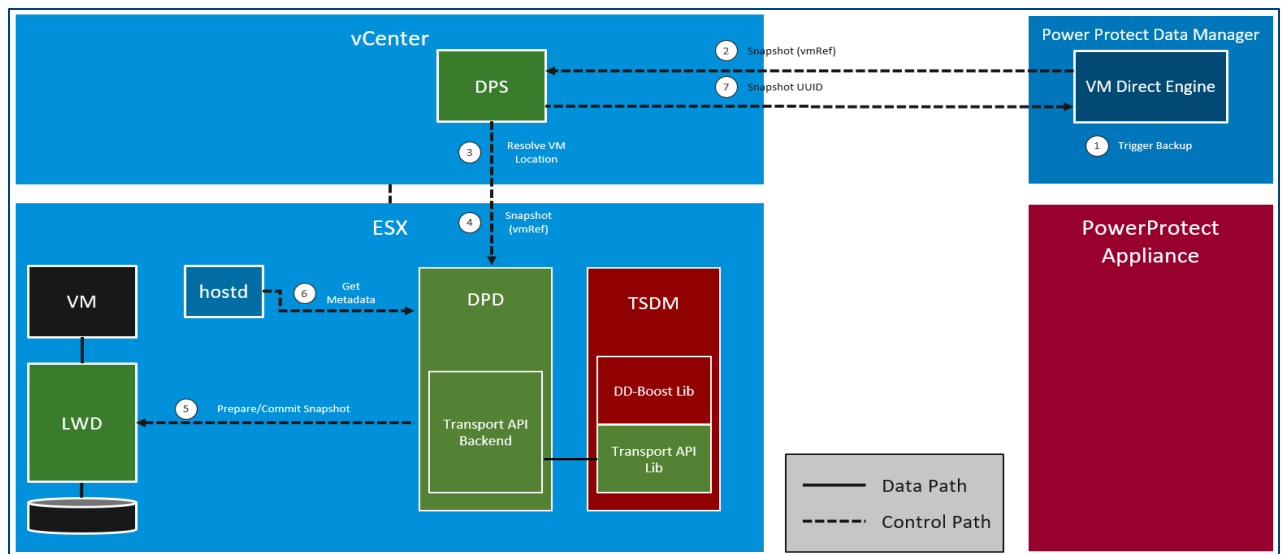


Figure 6 Transparent Snapshot creation

4.3 Delta sync operation

1. After the transparent snapshot creation operation is complete, you can initiate a delta sync operation or an incremental operation (see Figure 7). PowerProtect Data Manager issues an API to the vCenter for a delta sync operation and provides the snapshot UUID that was created in the previous step.
2. When the DPD component is resolved using the DPS, the DPD signals the TSDM to start a delta sync flow.
3. The TSDM uses the Transport API Library to query and track the changed areas that the Transparent Snapshot bitmap represents.
4. Internally, the DPD notifies the TSDM that the delta sync operation has begun.
5. For each changed area, the data is read from the disk.

Note: The delta sync operation uses a Fast Copy overwrite approach. In this approach, the previous point-in-time files are first fast copied. The fast-copied files are partly overwritten with the incremental data. Only the delta or changes that are represented by the currently synced snapshot are copied. These changes are copied in the Snapshot Extent Store (SES).

6. The changed data or delta is read from the disks to create a consistent data flow.
7. The read data is now written to protection storage using the transport API and DDBoost Library.
8. The changed data write to protection or secondary storage is now complete.
9. When all data has moved to protection storage, the TSDM sends an acknowledgment to the DPD that the operation is complete. The vCenter level task is marked as complete by the DPS.
10. The last step is the metadata write to protection storage. In this step, the VM metadata (such as VMX Files, Manifest, and the last TSDM Snapshot information) is transferred using the VM Direct Engine VMware APIs.

Note: From this point in time, the files on the PowerProtect appliance are crash consistent and can be used for recovery. For a full backup, both full sync and delta sync are performed. However, for an incremental backup, only delta sync is performed. A full sync can back up four VM disks in parallel; a delta sync can back up 10 VM disks in parallel.

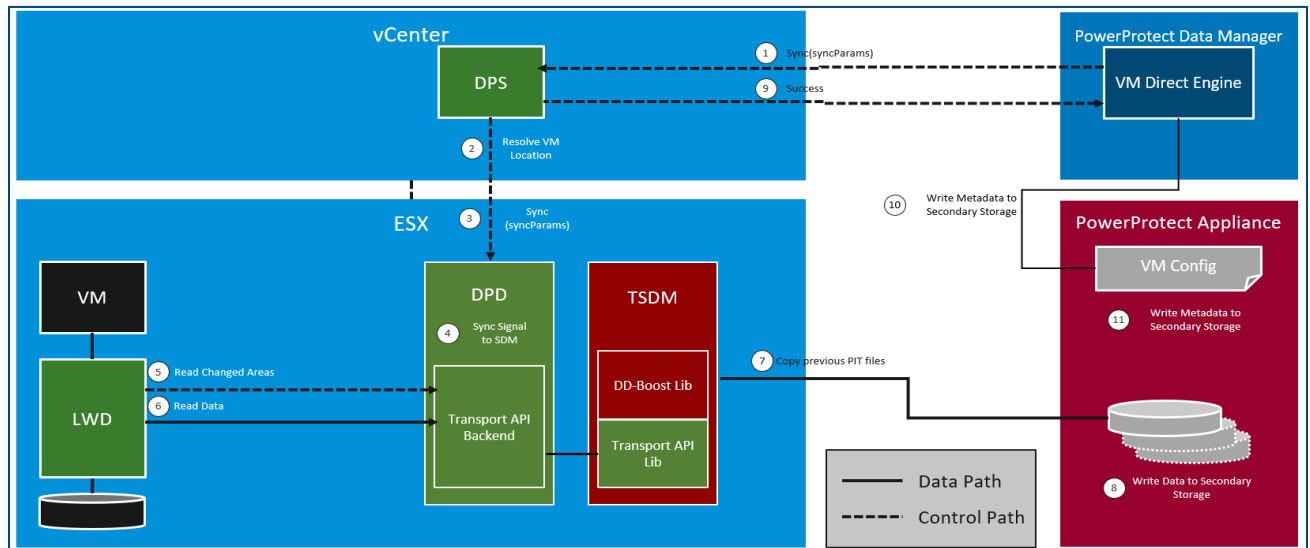


Figure 7 Delta sync operation

4.4 Snapshot retire operation

1. After the delta sync operation is complete, PowerProtect Data Manager ensures that the snapshot that was created before is deleted (see Figure 8). PowerProtect Data Manager now uses the DPS to invoke the retire snapshot API to the DPD.
2. The DPD relays this information to all the relevant LWD Filters to delete all the bitmaps and copy-on-write data residue left from the snapshot creation and delta sync stages.
3. The DPD then sends an acknowledgment of the snapshot-retire success to the DPS and PowerProtect Data Manager.
4. PowerProtect Data Manager records the backup copy set information in the PowerProtect Data Manager Catalog and informs the search node (if any) to start gathering metadata for indexing.

Asset	Stat...	Er...	Size	Data ...	Redu...	Summary
Linux-02_NoStun	✓ Succ...	--	17.2 GB	3.8 GB	88.9 %	Description : Backup has SUCCEEDED. Total VMs backed up: 1 ProxyHostName : localhost TransportModeUsed : SDM ParallelismUsed : 1 BackupLevel : Full Overall : CompressedSize : 418.5 MB AverageThroughput : 39.5 MB/s

Note the following:

- Transparent Snapshots are supported on VMFS, NFS, and vSAN Datastores. VMware vSphere Virtual Volumes (vVols) are not supported currently.
- Encrypted VMDKs are not supported with Transparent Snapshots.
- vRDM and pRDM are not supported with Transparent Snapshots.

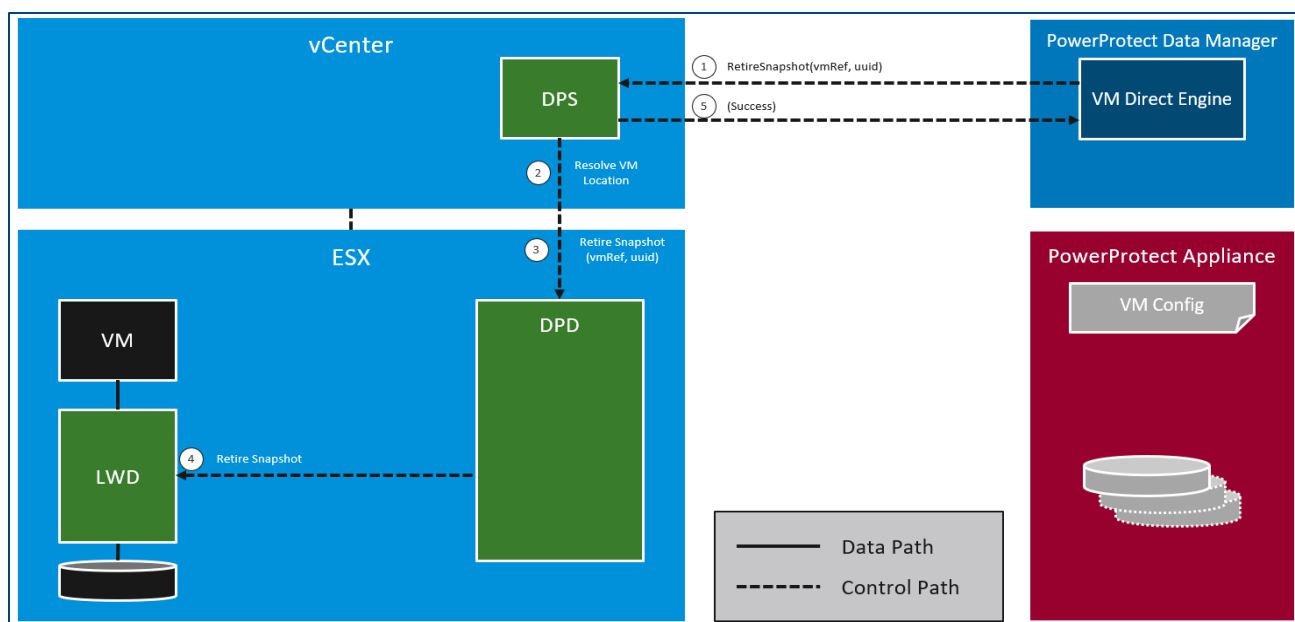


Figure 8 Snapshot retire operation

5 Restoring virtual machines

1. Before invoking any APIs from PowerProtect Data Manager to vCenter, PowerProtect Data Manager takes care of the configuration of the VM. For example, if a disk should be added or removed because the VM changed after the snapshot was taken, the disk might be added. For this reason, the reconfiguration part is performed beforehand.
2. A metadata only snapshot is then taken using the same snapshot creation flow described in the section Transparent Snapshots creation.
3. PowerProtect Data Manager invokes the restore operation, directing which VM should be restored to what point in time.
4. The VM is checked to determine if it is in a powered off state. If it is not, then the VM is powered off.
5. The DPS then locates the relevant ESXi host, and the DPD communicates with the TSDM to initiate a restore operation.
6. For a restore workflow, the strategy is first to reserve all the areas of the disks of the VM that should be reverted to the previous point in time, to:
 - Minimize the data we transfer from the protection storage back to the disk, and
 - To identify which parts of each disk have changed since the point in time to which we are trying to revert.
7. To resolve the changes (described above), the restore workflow leverages two Get Diff APIs, namely Get LWD Diff API and the DD Get Diff API.
8. The DPD resolves the Get LWD Diff API. This difference includes the changes that the VM has made to the disks since the last point in time that was previously synced or backed up to the PowerProtect appliance. The LWD Diff uses the metadata only snapshot, taken before the VM was powered off, to get the details on what was never written to the PowerProtect appliance.
9. The DD Get Diff API provides the delta details between the last sync point-in-time and the one to which we want to revert. It also merges the delta with the delta returned from the previous step. This step provides the complete set of extents that can now be read from the PowerProtect appliance.
10. From the point-in-time copy to which the user wants to revert, data is read and then finally written on top of the VMDKs.
11. After all the data movement is complete, the TSDM sends an acknowledgment to the DPD that the restore process has been completed.
12. DPS marks the task completed and sends an acknowledgment to PowerProtect Data Manager for catalog update.
13. The VM can now be powered on and should have been successfully reverted to the previous point in time.
14. The metadata only snapshot can now be retired, in the same manner as after every delta sync operation.

Note the following:

- Multiple streams are opened on the TSDM (when ESXi receives the request) to achieve restore parallelism. You can achieve a higher level of parallelism if the VMs are spread across multiple ESXi hosts.
- Restore in parallel supports up to eight disks of a VM using Transparent Snapshots.
- Only Restore to Original and New are supported with PowerProtect Data Manager 19.9 release using Transparent Snapshots. Although instant access and File Level Restore are not supported at this time with Transparent Snapshots, if tried, they would default to a restore workflow using VM Proxies (vProxy).

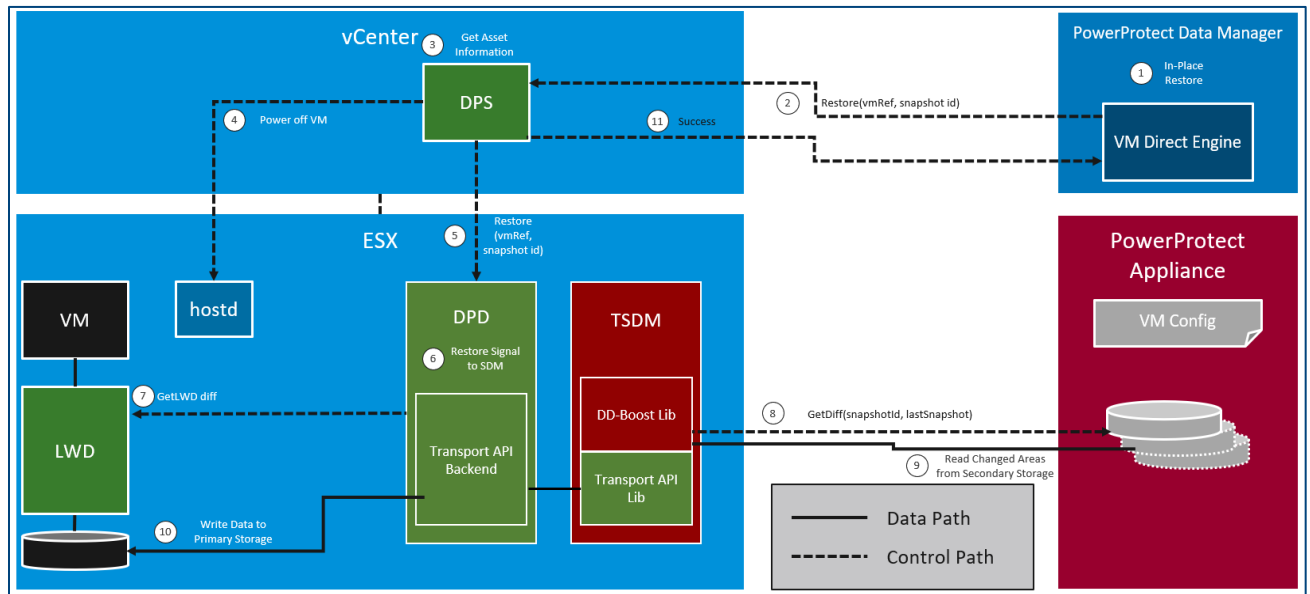


Figure 9 Restoring virtual machines

6 Performance test results

This section compares the I/O characterization between VMware API for Data Protection (VADP) and Transparent Snapshots using PowerProtect Data Manager 19.9. We infer from these results that with Transparent Snapshots, you overcome the penalties of the write and the read latencies, reducing VM latency by up to five times and providing up to five times faster backups.

Disclaimer: These results compare PowerProtect Data Manager 19.9 with Transparent Snapshots backup performance to PowerProtect Data Manager with VADP backup performance. The results are based on Dell Technologies internal testing in August 2021.




Figure 10 Backup performance comparison

	Transparent Snapshot	VADP
Test 1	Effective IOPS during Sync: AVR - 10K Sync transfer rate: AVR - 550 MB/s Latency: During Sync: AVR - 1 ms No Sync: AVR - 0.5 ms	Effective IOPS during Sync: AVR - 10K Sync transfer rate: AVR - 110 MB/s Latency: During Sync: AVR - Read: 1 ms / Write: 2.2 ms No Sync: AVR - 0.5 ms
Test 2	Effective IOPS during Sync: AVR - 10K Sync transfer rate: AVR - 580 MB/s Latency: During Sync: AVR - 1 ms No Sync: AVR - 0.5 ms	Effective IOPS during Sync: AVR - 10K Sync transfer rate: AVR - 110 MB/s Latency: During Sync: AVR - Read: 1 ms / Write: 2 ms No Sync: AVR - 0.5 ms
Test 3	Effective IOPS during Sync: AVR - 10K Sync transfer rate: AVR - 800 MB/s Latency: During Sync: AVR - 1.5 ms No Sync: AVR - 0.5 ms	Effective IOPS during Sync: AVR - 10K Sync transfer rate: AVR - 173 MB/s with 12 DD Streams Latency: During Sync: AVR - Read: 1.1 ms / Write: 2.5 ms No Sync: AVR - 0.5 ms

Figure 11 Backup performance comparison test details

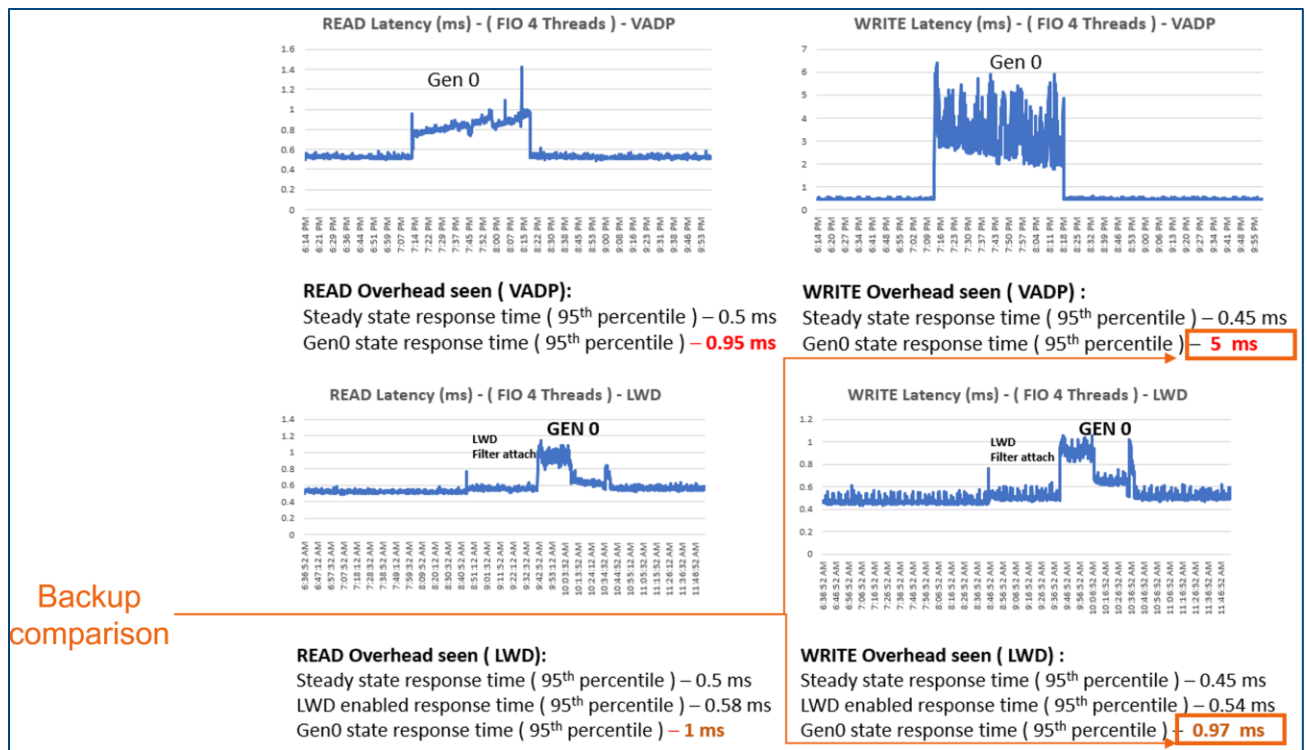


Figure 12 VM latency (read and write) performance comparison

A Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

The [Dell Technologies Info Hub](#) provides expertise to ensure customer success with Dell EMC data protection products.

A.1 Related resources

Additional resources:

- [PowerProtect Data Manager Interactive Demo](#)
- [VMware vSphere APIs for I/O Filtering \(VAI/O\)](#)