



HPE PRIMERA VIRTUAL COPY

Snapshots and copy data management



HPE Primera A630/C630
(2-Node Storage Base)
Base unit 24 drives



HPE Primera A650/C650
(2 / 4-Node Storage Base)
Base unit 24 / 48 drives



HPE Primera A670/C670
(2 / 4-Node Storage Base)
Base unit 24 / 48 drives

CONTENTS

Executive summary.....	4
Introduction.....	4
HPE Primera Virtual Copy basics.....	5
Physical copy.....	5
HPE Primera Virtual Copy Software.....	6
HPE Primera Virtual Copy scalability.....	9
HPE Primera Virtual Copy use cases.....	9
Application development and testing.....	9
Backup.....	10
Recovery.....	12
Offloading processing and decision support.....	12
Data governance and eDiscovery.....	12
HPE Primera Virtual Copy integration.....	12
HPE Primera mySnapshot.....	12
Consistency and autonomic groups.....	13
Scheduling.....	13
HPE Recovery Manager Central software.....	13
HPE RMC Express Protect for flat backup.....	13
HPE Primera Virtual Lock and security.....	14
HPE Primera Remote Copy Software.....	14
HPE Primera Virtual Copy Snapshot architecture.....	15
HPE Primera virtualization overview.....	15
Virtual volumes.....	16
The anatomy of a virtual copy snapshot.....	16
Understanding the HPE Primera COW implementation.....	18
Write operations.....	18
HPE Primera Virtual Copy ROW and HPE Primera thin deduplication.....	19
RW snapshot behavior.....	21
Virtual copy snapshot relationships.....	21
HPE Primera Virtual Copy Snapshot restore and refresh.....	24
How does HPE Primera Virtual Copy snapshot refresh work with updatev?.....	24
Manual methods to refresh snapshots.....	24
Refreshing the snapshot of a base volume.....	25
Refreshing a snapshot of a snapshot.....	26
HPE Primera Virtual Copy Snapshot operations.....	28
Creating consistent group snapshots of virtual volumes.....	32
Promoting a snapshot to a base volume.....	34
Promoting a group of snapshots to their base volumes.....	35
Snapshot refresh commands and procedure.....	35
HPE Primera Virtual Copy Snapshot considerations.....	36
HPE Primera Virtual Copy snapshot refresh considerations.....	37



Technical white paper

HPE Primera Virtual Copy snapshot refresh restrictions and limitations.....	38
Summary.....	38
Glossary	39
Resources and additional links.....	40



EXECUTIVE SUMMARY

HPE Primera storage has brought new levels of storage efficiency, agility, and availability to data centers. With a highly resilient product architecture designed to support massive consolidation, eliminate bottlenecks,¹ and maximize capacity utilization, HPE Primera revolutionized the economics of storage. The platform's highly virtualized operating system (OS) enabled pioneering work in areas such as using data reduction to significantly increase capacity efficiency, reservationless snapshots to simplify and expedite the sharing of data, and flash media to achieve record-setting performance results.

As the use of flash grows and content generation continues to accelerate, the ability to control storage costs while stepping up value creation for the business will continue to require rich, Tier-0 data services such as those offered by HPE Primera. For example, in today's business environments, different processes such as backup, recovery, development, and testing can all be enhanced by the ability to leverage point-in-time (PIT) data copies or snapshots.

Two particularly important trends are driving new demand for snapshot technology. One is copy data management (CDM)—the reduction of storage consumption through elimination of unnecessary duplication of production data. The second is the DevOps culture, which seeks to serve the needs of both development and operations with a single infrastructure, requiring the creation of production copies that do not impact production workloads.

With its highly virtualized operating system and architecture built for delivering IT as a service (ITaaS), HPE Primera excels in these two areas. For example, HPE Primera Virtual Copy software (HPE Primera Virtual Copy)—a reservationless snapshot technology—allows users to simply and efficiently share and protect data from any application.

Because HPE Primera Virtual Copy minimizes the amount of required physical capacity that is actually being consumed by data without degrading performance or availability, customers use this technology to expedite test and development while containing costs. In fact, HPE Primera customers are using HPE Primera Virtual Copy in DevOps environments routinely to create in excess of ten thousand snapshots per day on a single system without incurring negative performance impacts.²

Target audience: This paper is designed to help HPE Primera solution architects, storage administrators, and system administrators who are currently considering or using HPE Primera Virtual Copy understand this technology, its various applications, and its features and functionalities in order to increase storage efficiency using HPE Primera.

INTRODUCTION

An array-based solution, HPE Primera with HPE Primera Virtual Copy eliminates the cost and performance overhead associated with host-based snapshots. The unique performance advantages of HPE Primera ensure that users can leverage PIT snapshots without the performance concerns that impact traditional array-based copies.

For example, the platform's mixed workload support allows an IOPS-intensive production application and a sequential backup process (based on an associated snapshot) to run with negligible impact on each other. The fine-grained virtualization and automatic load-balancing capabilities of the system, for both base and snapshot volumes, are spread over large numbers—as many as hundreds—of drives and controller nodes. This system-wide load balancing eliminates the contention inherent in other architectures, even those that employ wide striping.

HPE Primera Virtual Copy snapshots are also highly flexible. Users can make hundreds of snapshots per base volume, enabling histories of application data to be archived online efficiently. Users also have the option to create read-only (RO) or read-write (RW) snapshots of snapshots. Therefore, organizations can selectively test, use, and change production-quality data safely and easily while protecting production data and other snapshots from accidental or intentional changes. Administrative overhead is dramatically reduced and organizational productivity is greatly improved. The software features a “promote” capability that allows users to take any snapshot associated with a base volume and “promote” it to a base volume, simply and quickly. In this way, a base volume can be rapidly synchronized to any one of its associated snapshots, effectively “rolling back” the base volume to a previous point in time.

¹ See the technical whitepaper [HPE Primera Architecture](#) for details.

² Based on current Hewlett Packard Enterprise telemetry data.



Any RO child snapshot can also be “promoted” to an RW parent snapshot. Furthermore, the user can specify the priority of the copy operation when it is started. This option allows the user to control the overall speed of the snapshot promotion. If this option is not specified, the `promotesv` operation begins with a default priority of medium. High priority indicates that the operation will complete faster. Low priority indicates that the operation will run more slowly than the default priority task. This degree of flexibility is unique to HPE Primera Virtual Copy and gives users more control of the environment in using exactly what fits their needs.

Customers using this technology in their business environments often create several snapshots on a single system. Snapshot “refresh” is an HPE Primera feature that expedites the processes of creating and updating snapshots by significantly reducing space consumption and management overhead of discovering the logical unit number (LUN) and configuring the host application. This makes HPE Primera snapshot technology an efficient way of managing copy data in an HPE Primera environment with an HPE Primera snapshot. This feature eliminates the need to create a new snapshot and instead updates or refreshes an existing snapshot.

HPE Primera systems are centrally administered from a single user interface, the HPE Primera and HPE 3PAR StoreServ Management Console (SSMC). This management and reporting console offers central management of HPE Primera —regardless of model or geographical location. This same console can be used to manage all system functions, including snapshots. In addition to creating HPE Primera Virtual Copy snapshots with the HPE SSMC, users also have the flexibility to employ the scriptable HPE Primera CLI.³

HPE PRIMERA VIRTUAL COPY BASICS

HPE Primera supports two technologies to provide users with copies of HPE Primera base virtual volumes (VVs):

1. Physical copy
2. Virtual Copy snapshots

Each of these technologies allows users to access a PIT copy of a base VV and to read data from or write data to that PIT copy. Both technologies can be leveraged to service customer recovery point objective (RPO) and recovery time objective (RTO) requirements.

Physical copy

A physical copy is a full copy of a volume. A physical copy duplicates all the data from one original base volume to a destination volume (see [FIGURE 1](#)). Any change to either volume causes them to lose synchronization with each other, which can be corrected by resynchronizing the two volumes. The destination volume can be resynchronized with its parent volume by using a saved snapshot so that only the changes since the last copy or resynchronization need to be copied. The user specifies a snapshot of the parent saved during the creation of the physical copy, because it is required for the resynchronization operation. As soon as the resynchronization is complete, a new snapshot for the next resynchronization operation replaces the old snapshot.




FIGURE 1. Physical copies

³ See the [HPE Primera OS 4.2 Command Line Interface Reference Guide](#) and the HPE SSMC Online Help for instructions on performing specific tasks using these interfaces.



Physical copies can be created and managed in groups to reduce the number of management tasks. Users can create a consistent group of physical copies from a list of VVs and group physical copies into autonomic groups that are managed as a single physical copy.

Physical copies can be online or offline:

- With online physical copies, the destination volume is automatically created and can be exported immediately to the host as cloning happens in the background.
- Offline physical copies require a destination volume with a user space size at least as large as the user space of the base volume being copied. Offline physical copies cannot be exported until the data copying is complete.

The physical copy is still space efficient, because the user can choose the destination volume to be a thinly provisioned virtual volume (TPVV) or a TPVV with duplication and compression enabled. The use of TPVV with duplication and compression enabled in the same common provisioning group (CPG) as the parent volume results in significant savings of capacity consumption, because the initial physical copy at creation is an identical copy of the parent volume.

HPE Primera Virtual Copy Software

HPE Primera Virtual Copy is an HPE Primera snapshot implementation used to provide a PIT virtual copy of a VV. It also ensures that the original data can always be obtained if a problem occurs when updating the data on a VV. HPE Primera Virtual Copy implements an efficient variant of copy-on-write (COW) and redirect-on-write (ROW) methodologies. For COW, the HPE Primera OS uses a delayed copy-on-write (DCOW) process that eliminates any performance impact to the host I/O. DCOW is used for snapshots of thinly provisioned volumes without deduplication and compression enabled. DCOW relegates the reading of the original data, updating of the base volume with the new data, and the copying of the original data to a background process after the write update has been acknowledged by the host.

For volumes with deduplication and compression enabled, the HPE Primera OS relies on the ROW method to create snapshots and store the original data for the deduplicated data stored in the Dedupe Store (DDS). COW is used for the rest of the data that are not dedupable and stored in the Dedup Client store (DDC). This process allows HPE Primera to use at most a single write I/O to store the host I/O in an existing block of data of the base volume. Furthermore, the inline dedupe engine will check if the new write is dedupable through the fast and efficient metadata lookup. In case of a match against any data block already in the dedupe store, then only a metadata update will take place, and no I/Os are flushed to the backend.

HPE Primera Virtual Copies are always thin, non-duplicative (only one copy of the change blocks is kept and then referenced), and reservationless.

As a result of the highly efficient way that HPE Primera handles metadata, users can configure thousands of RO and RW snapshots with the flexibility to promote any snapshot without destroying others.⁴ Unlike a physical copy, which duplicates the entire base volume, a virtual copy snapshot initially consumes no space and only increases in size as it records changes to the original volume. This allows an earlier state of the original volume to be re-created by starting with the current state and “rolling back” all the changes that have been made since the HPE Primera Virtual Copy snapshot was created.

Virtual copy snapshots, as shown in [FIGURE 2](#), can be created and managed in groups to reduce the number of management tasks. Users can create a consistent group of virtual copies from a list of VVs and group virtual copies into autonomic groups that are managed as a single virtual copy. A virtual copy snapshot can belong to any CPG, giving users flexibility and efficiency in terms of storage consumption.

⁴ The maximum number of HPE Primera Virtual Copy snapshots that can be created on a system is determined by the system configuration and the HPE Primera OS release. See [TABLE 1](#) for additional details.



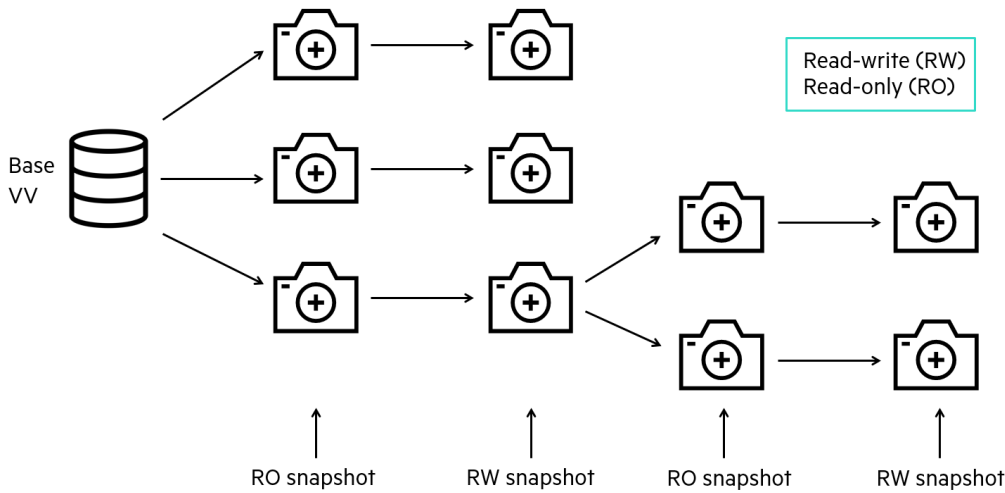


FIGURE 2. Virtual copy snapshots

Virtual copy snapshot space is assigned by setting the Copy CPG policy on the HPE Primera VV. Virtual copy snapshots can be mapped to CPGs different from their base volumes. In many deployments, the snapshot capacity uses the same CPG as the base volume. Another important feature of HPE Primera Virtual Copy is that a single COW is used to maintain the delta changes from the base volume, regardless of the number of snapshots. FIGURE 3 shows the copies of a VV in the HPE SSMC.

Name	System
mgmtdata	A670-3144_20-19-221
srdata	A670-3144_20-19-221
admin	A670-3144_20-19-221
dataset	A670-3144_20-19-221
log_backup	A670-3144_20-19-221
SRM_data1	A670-3144_20-19-221
SRM_data2	A670-3144_20-19-221
SRM_vm1	A670-3144_20-19-221
SRM_vm2	A670-3144_20-19-221
test.vv.1 (2)	A670-3144_20-19-221
test.vv.1.200331101829	A670-3144_20-19-221
test.vv.1.200331140301	A670-3144_20-19-221
testpm	A670-3144_20-19-221

Snapshot ID	Type	Size	Copies	Exports	Timestamp
test.vv.1.200331140301	Snapshot R/W	100 GiB	0 copies	0 exports	31-Mar-2020 02:03:01 PM MDT
test.vv.1.200331101829	Snapshot R/W	100 GiB	0 copies	0 exports	31-Mar-2020 10:18:29 AM MDT
test.vv.1	Base R/W	100 GiB	2 copies	0 exports	31-Mar-2020 09:19:49 AM MDT

FIGURE 3. "Copies" view for an HPE Primera VV in the HPE SSMC



The base VV and a virtual copy can be mapped to different **CPGs**, as presented in [FIGURE 4](#). The **Copy CPG** allocation warning and allocation limit are set independently from the base volume warning and limit. When snapshots are deleted, so are the data and correlating administrative space. HPE Primera Virtual Copy snapshots consume space allocated from the **Copy CPG**. The **Copy CPG** can be based on capacity, performance, and availability.

Create Virtual Volume General ▾ ?

General Advanced options

Name: test.vv1

System: A670-3144_20-19-221 ▾

Domain: <None> ▾

Provisioning: Thin Provisioned ▾

Dedup and Compression: No

CPG: SSD_r6 RAID 6 SSD

Size: 100 GIB ▾

App Volume set: VV_Set

Application: Oracle Database

Protection And Additional Settings

▼ Volume

Number of volumes: 1

Volume set: VV_Set.Snapset optional

Copy CPG: SSD_r6_SRM RAID 6 SSD optional

Comments: optional

► RMC Protection

► Snapshots

► Remote Copy Group

► Export

Changed: Copy CPG to "SSD_r6_SRM"

Create Create + Cancel

FIGURE 4. Setting a "Copy CPG" through the HPE SSMC



HPE Primera Virtual Copy snapshots offer the following benefits:

- **Smart:** Snapshots can be refreshed and promoted, are not interdependent on each other, and can be routinely created or deleted with the scheduler. They can maintain consistent snapshots over a group of volumes. Snapshots can be refreshed with current data on the base VV and can be promoted to other base volumes in addition to the original volume from which the copy originated. More importantly, consistency groups can be maintained for clustered environments such as virtualization hypervisors or Oracle databases when using Oracle Automatic Storage Management (ASM).
- **Easy:** There is no need to create a separate storage pool for snapshots. A snapshot can be created that has a QoS different from its base volume. Simplified management eliminates the need to make decisions up-front about sizing. This alleviates both the problem of capacity waste because of overallocation and snapshot failure resulting from insufficient space.
- **Ready:** HPE Primera Virtual Copy can take multiple snapshots of existing copies and create separate snapshot trees that are not dependent on previous snapshots. HPE Primera Virtual Lock software (HPE Primera Virtual Lock) can also be used to maintain compliance without locking the base volume by simply creating a snapshot and using HPE Primera Virtual Lock against that copy.
- **Efficient:** Reduced space and performance scalability means snapshots require only a single COW to maintain differences from the base volume, regardless of the number of snapshots that are taken against it. HPE Primera Virtual Copy allows you to take hundreds of RO snaps and dozens of RW snaps of any given volume without suffering an additive performance hit.

HPE PRIMERA VIRTUAL COPY SCALABILITY

Because RO and RW copies must alternate (see [FIGURE 2](#)), users can only make an RO copy of an RW volume and an RW copy of an RO volume. For capacity efficiency, expiration dates can be set for virtual copies. The system allows users to produce several hundreds of virtual copies of a base volume.

The maximum values for the HPE Primera OS4.0/4.1/4.2 releases are listed in [TABLE 1](#).

TABLE 1. Maximum snapshots per VV

Maximum	Value
Maximum snapshot per VV: RO	2048
Maximum snapshot per VV: RW	1024
Maximum RW snapshot per RO snapshot	2048
Maximum VVs per consistency group for group snapshot	300

For information regarding the maximum number of VVs and HPE Primera Virtual Copy snapshots that can be created with your specific system configuration, visit the [HPE Single Point of Connectivity Knowledge \(SPOCK\) website](#).

There are also limitations regarding the maximum VVs per system (base VVs + virtual copies) for the HPE Primera family of arrays.

HPE PRIMERA VIRTUAL COPY USE CASES

HPE Primera Virtual Copy provides a foundation for multiple solutions designed to simplify key business processes. This section describes some common use cases for this technology.

Application development and testing

With HPE Primera Virtual Copy, organizations doing development and testing can use multiple RW snapshots to perform parallel development of datasets and parallel testing of new applications. Snapshot data from production sources gives multiple testers their own realistic bases for validating the integrity and performance of new applications. Users can also simultaneously develop datasets by performing different processes on multiple identical snapshots, progressing with the one that best meets their requirements. The option to propagate the results of the test to the base VV or to roll back the changes is immediately available to the user.



Application developers have been using storage snapshot technology for many years, although the way organizations develop software has changed significantly over the last few years. Storage snapshot technology has played an important role in achieving this change. Copies of production data are commonly used in application environments to facilitate testing and development. Using snapshots makes it easy for developers to refresh their copies of data instantly.

As part of CDM, HPE Primera Virtual Copy snapshot technology also supports new trends in data analytics and reporting, which require copies of production data to be created quickly and easily. The snapshot feature can add value to these trends by minimizing the amount of physical capacity required.

The capability of **Snapshot refresh** with `updatevv`, as demonstrated in [FIGURE 5](#), can help reduce management complexity. Snapshots of production volumes can be preconfigured and then refreshed with recent data from the production volume when applications using the snapshots need the latest copy.

The HPE Primera `updatevv` snapshot refresh feature allows application developers to refresh their own snapshots without involving the storage administrator. Super user rights are not required to refresh the snapshot data. This snapshot refresh feature also adds value to CDM, because it helps reduce the number of copies of secondary data within a storage solution.

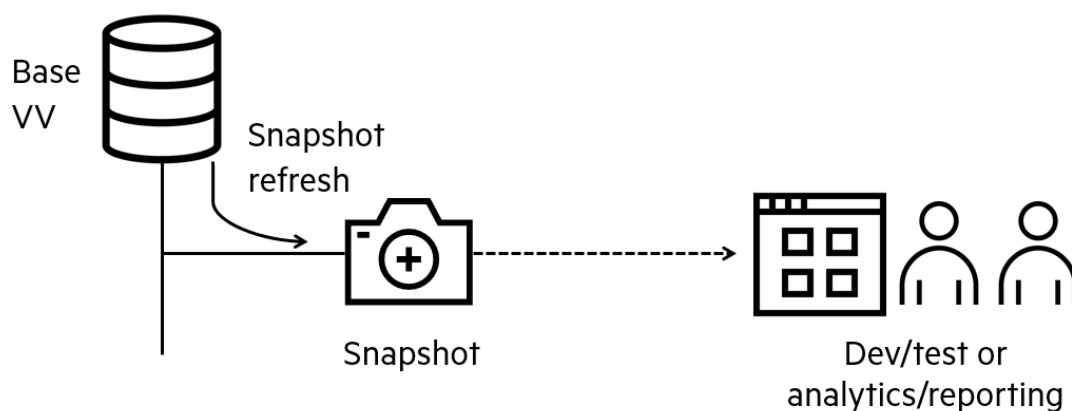


FIGURE 5. "Snapshot refresh" with "Dev/test or analytics/reporting"

Backup

As the requirement for 24x7 availability has become fundamental to business operations, organizations can no longer afford the downtime involved in backing up critical data offline. HPE Primera, with support for simultaneous mixed workloads, is the perfect platform for concurrent backup. With HPE Primera Virtual Copy, users can back up volumes from a snapshot image locally or remotely, allowing business-critical applications to continue in an efficient manner while maintaining performance and consuming only minimal additional capacity. Integration with backup independent software vendors (ISVs) automates the interaction with the storage backup application.

The capability to generate a PIT copy of data while an application is online is vital to creating backups. The ability to use those PIT copies easily and quickly to recover from those backups is a valuable feature that enables efficient backup and recovery for users. Because HPE Primera Virtual Copy snapshots are thin provisioned, if there are one or many snapshots of a base volume, the capacity used by the snapshot does not increase as the number of snapshots increases. [FIGURE 6](#) shows a sample result of performing hourly backups for seven days. The database has an average change rate of 10% per day. Each day, the hourly snapshots consume 100 GB from the Copy CPG space. The space consumed would be the same even if one snapshot was taken daily.



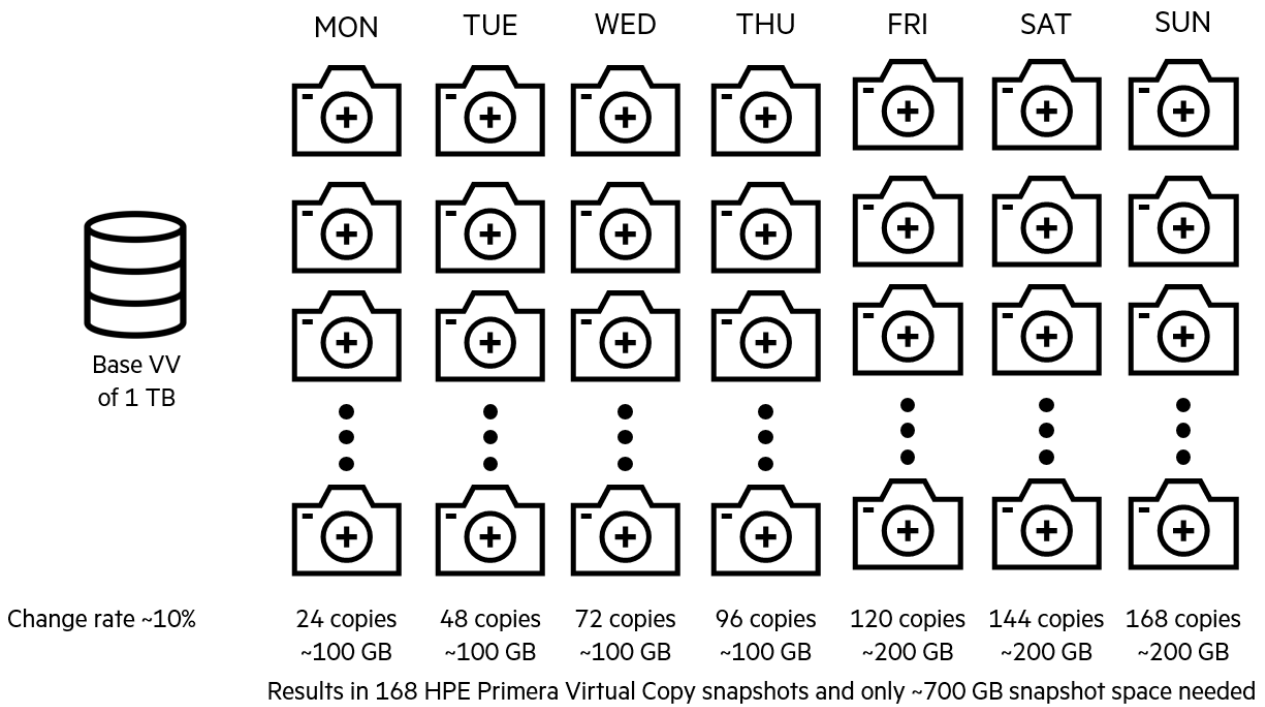


FIGURE 6. Database backup and HPE Primera Virtual Copy snapshot use case

Typically, when a backup application runs a backup job, it creates a snapshot and then reads the data from the snapshot in order to copy it to a backup system. Backup applications then export the snapshot to a backup server and read data from it. With the help of the **Snapshot refresh** feature, as shown in [FIGURE 7](#), backup applications can preconfigure snapshots and perform a snapshot refresh before remounting and reading the data from the snapshot and copying it to the backup device. This can reduce administration overhead to create, export, and discover the snapshot LUNs, thus improving overall backup efficiency.

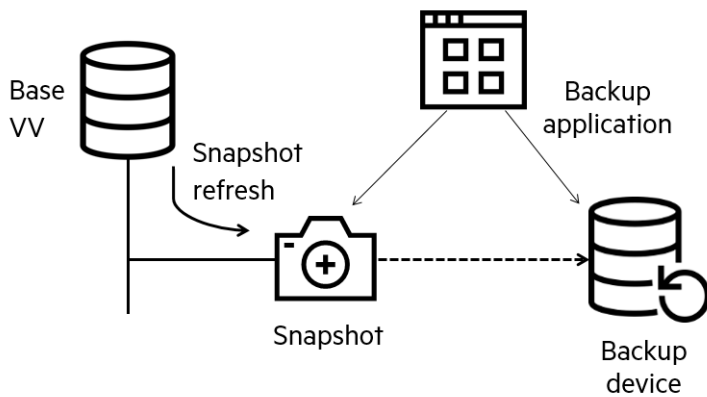


FIGURE 7. "Snapshot refresh" with backup and recovery solution



Administration overhead has the most impact in large-scale environments with many volumes. For example, a backup solution that needs to copy the incremental data of hundreds of volumes to cloud storage in an enterprise environment would face greater challenges regarding creating and managing snapshots than a smaller customer environment. The HPE Primera snapshot refresh feature can be used in this type of scalable solution very efficiently.

Recovery

With HPE Primera Virtual Copy, users can make frequent online copies of production datasets. If an inconsistency or deletion is detected, an administrator can restore from a “clean” copy made just before the inconsistency or deletion occurred. With this safety mechanism, users can minimize the costs of downtime and improve service levels for business users both economically and with reduced opportunity for error.

Depending on the application requirements, recovery can be accomplished in one of two ways:

1. Promote a “clean” snapshot to the base volume.
2. Immediately mount and return to production an instant RW snapshot (taken from the same “clean” snapshot). This can be followed by a later promotion of the RW snapshot to the original base volume when downtime is convenient.

The first method requires a short time to complete the synchronization, but no additional downtime. The second approach allows immediate application recovery, but subsequent scheduled downtime. Either solution can be used across a number of databases and applications, including Oracle, Microsoft SQL Server, Microsoft Exchange Server, Microsoft Hyper-V, and VMware vSphere® (vSphere) environments.

Offloading processing and decision support

By allowing administrators to co-locate production and decision support datasets, HPE Primera Virtual Copy snapshots enable IT departments to create instant snapshots of production datasets for use in extract, transform, and load operations, thereby accelerating data warehouse creation. Afterwards, snapshots of the data warehouse can become multiple sources for parallel creation, speeding data mining and reporting application delivery.

Data governance and eDiscovery

HPE Primera Virtual Copy snapshots can also be used with HPE Primera Virtual Lock to retain tamper-proof storage volumes and volume copies securely. This combination delivers the most efficient and cost-effective approach to data retention for complying with governance and aiding electronic discovery (eDiscovery). With HPE Primera Virtual Copy and HPE Primera Virtual Lock, administrators can create frozen, PIT copies of data by applying a configurable retention period during which copies are protected against deletion—even by a user with the highest-level privileges.

HPE PRIMERA VIRTUAL COPY INTEGRATION

HPE Primera Virtual Copy offers many efficient tools and integration features for efficient snapshot management.

HPE Primera mySnapshot

The HPE Primera mySnapshot utility does not require a separately purchased license. The HPE Primera mySnapshot utility enables safe and easy copy and provisioning access for non-storage professionals such as database administrators, software developers, and test engineers working with systems.

Users can restore their own copies of test data safely and easily, in seconds, without relying on the storage administrator. The HPE Primera mySnapshot utility uses an access control list (ACL) to associate a user with certain administrative permissions and specified storage resources. After these administrative permissions are granted for the specified resources, users can easily replace and restore copies of their own test databases. This enables users who normally only have browse capabilities on the system to update specific snapshots with more recent ones. This results in faster turnaround times for developers who need to have their snapshots refreshed, and alleviates workloads for storage administrators.

The HPE Primera mySnapshot utility can only be configured by using the HPE Primera CLI. See the [HPE Primera OS 4.2 Command Line Interface Reference Guide](#) for instructions and more information about roles and rights when using the HPE Primera mySnapshot utility.



Consistency and autonomic groups

The HPE Primera OS provides a simple, two-step process for the common task of provisioning VVs and exporting them to application servers as virtual logical unit numbers (VLUNs). HPE Primera autonomic group capabilities take autonomic storage management a step further by allowing both hosts and VVs to be combined into “groups” or “sets,” which can then be managed as a single object. Adding an object to an autonomic group applies all previously performed provisioning actions to the new member. For example, when a new host is added to a group, all volumes are automatically exported to that group with no administrative intervention. Similarly, when a new volume is added to a group, that volume is also exported to all hosts in the group—intelligently and without requiring administrator action.

Physical copies and snapshots created with HPE Primera Virtual Copy have several uses for autonomic groups, such as creating a consistent set of snapshots of a group of volumes and simplifying the administration of recurrent snapshots. Autonomic groups may be used when creating an HPE Primera Virtual Copy snapshot to form a consistent group snapshot of all VVs in the set. Each snapshot is created at the same point in time, so related structures in different volumes within the set remain consistent within the snapshot. For example, the use of autonomic groups simplifies the task of performing an online backup of an Oracle database that spans multiple ASM volumes.

Scheduling

Virtual copy snapshots can also be regularly scheduled so that PIT copies of user data can be created and deleted over time. The running of periodic scheduled tasks to create or delete snaps enables flexibility in managing user data, and can automatically remove volumes that have expired. Periodically removing snapshots enables customers to create and remove snapshots at regular intervals. The HPE Primera Scheduler enhances management automation, reduces administration time, and decreases the chance of administrative error by giving users full control over automated virtual copy snapshot creation and deletion.

HPE Recovery Manager Central software

Virtual copies are consistent at the VV level, but not at the host, system, or application level. In other words, virtual copies only preserve the data that was written on the base VV before the virtual copy is created. Virtual copies do not preserve the data that is resident within the application or system buffers, and data is not flushed to disk before the virtual copy is created. Working together, HPE Recovery Manager Central (RMC) software⁵ and HPE Primera Virtual Copy snapshots provide the foundation for rapid online recovery using disk-based copies.

Various HPE RMC features provide application-level consistency by placing databases or virtual machines in “hot standby” mode and then calling HPE Primera Virtual Copy to make a rapid and space-efficient PIT snapshot in Oracle, SQL Server, Exchange, vSphere, or Hyper-V environments. This combination of HPE RMC and HPE Primera Virtual Copy provides not only rapid and space-efficient disk copies but also application-consistent disk-based copies to ensure that all storage-related volumes for an application are available for quick recovery. For example, HPE Recovery Manager Central for Oracle (RMC-O) is a highly efficient solution for creating and managing hundreds of application-consistent, reservationless PIT snapshots of Oracle databases automatically.

HPE RMC Express Protect for flat backup

HPE RMC leverages HPE Primera snapshot technology and HPE StoreOnce Catalyst. This unique combination delivers a converged data protection system that seamlessly integrates a highly robust, flash-optimized, primary storage system and a very fast, scalable, and highly resilient backup system—[HPE StoreOnce Backup](#). The HPE RMC Express Protect feature facilitates the flat backup of HPE Primera snapshot data to HPE StoreOnce. Snapshots backed up on HPE StoreOnce systems are self-contained volumes that can be restored to the original (or different) HPE Primera array in the event of a disaster. This level of data protection cannot be achieved with snapshots alone.

HPE RMC for VMware allows VMware® administrators to create hundreds of nondisruptive, application-consistent virtual machine snapshots, and initiate rapid online recovery directly from within the VMware vCenter Server® (vCenter Server) virtualization management console. HPE RMC for SQL allows SQL Server database administrators (DBAs) to create, schedule, and manage application-consistent snapshots of nonvirtualized SQL databases on an HPE Primera array. The HPE RMC for SQL Express Protect feature enables automatic flat backup of SQL snapshot data from HPE Primera to HPE StoreOnce, independent of backup server software. HPE RMC also provides a programmable interface via Representational State Transfer (REST) Application Programming Interface (API) Software Development Kit (SDK) for plug-in scripting purposes. This enables partners and developers to integrate the snapshot management and backup benefits of HPE RMC with their own business applications.

⁵ HPE RMC is replacing all versions of the HPE Primera Recovery Manager software. HPE RMC currently supports VMware vSphere (6.0, 6.5, and 6.7) and non-virtualized Microsoft SQL Server databases (SQL 2016, 2017, and 2019). For full details, see hpe.com/h20195/V2/GetDocument.aspx?docname=4AA5-5173ENW&cc=us&lc=en.



HPE Primera Virtual Lock and security

Based on the expiration policy setting, the oldest HPE Primera Virtual Copy is removed when the next HPE Primera Virtual Copy is created. HPE Primera Virtual Lock is an optional license that prevents deletion of selected VVs for a specified period.

As discussed in the [HPE Primera Virtual Copy use cases](#) section, users can use HPE Primera Virtual Lock functionality to specify the retention period for each volume or copy of a volume. By locking a VV, a user can prevent intentional or unintentional deletion before the retention period elapses.

HPE Primera Remote Copy Software

HPE Primera Remote Copy software (HPE Primera Remote Copy) provides enterprise and cloud data centers with autonomic replication and Tier-0 disaster recovery technology that allows the protection and sharing of data from any application—simply, efficiently, and affordably.

HPE Primera Remote Copy dramatically reduces the cost of remote data replication and disaster recovery by leveraging thin copy technology and enabling replication with midrange, all-flash, and high-end arrays to deliver unique cost efficiencies for protecting a broader set of applications. The user can set up and test in minutes, reducing the need for professional services. With support for synchronous, asynchronous periodic, multi-site, as well as synchronous long-distance replication, HPE Primera Remote Copy delivers a rich portfolio of replication options for customers to achieve stringent RTOs and fast RPOs.

With HPE Primera Remote Copy, the VVs are related to a volume group, which ensures that the VV data within a volume group maintains write consistency. When a user (or an automated process) creates a PIT snapshot of VVs in a volume group, HPE Primera Remote Copy blocks writes to all volumes in the group. This ensures a consistent PIT virtual copy of the whole volume group. A source RW snapshot can be replicated to a base volume or to another RW snapshot. If a user replicates the source snapshot to a target snapshot, then the target RW snapshot will contain as much data as the source volume from which the RW snapshot was created.

The **Snapshot refresh** feature (see [FIGURE 8](#)), can also be used to host an application that requires a copy of production data at a target site of a disaster recovery solution. In an HPE Primera Remote Copy configuration, target volumes synchronize data source volumes. Because target volumes are write-protected, an RW snapshot can be created and configured in the hosts target site. The snapshot can be refreshed at regular intervals by using the HPE Primera Remote Copy snapshot refresh feature to allow applications to get access to updated data.

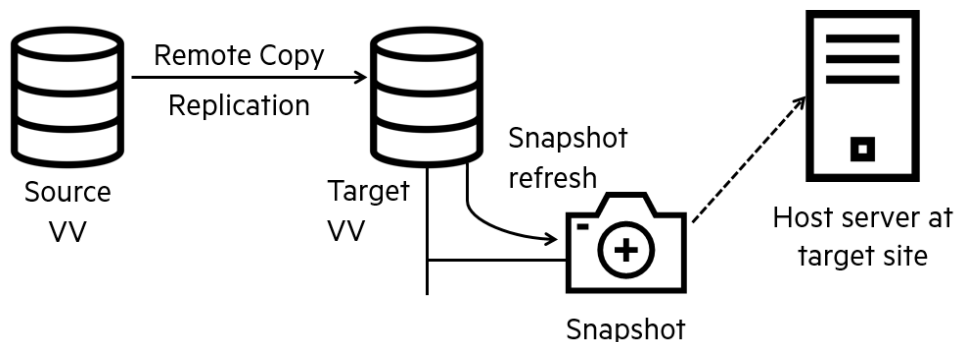


FIGURE 8. "Snapshot refresh" at Remote Copy target site



HPE PRIMERA VIRTUAL COPY SNAPSHOT ARCHITECTURE

To understand the architecture of an HPE Primera Virtual Copy snapshot, it is useful to introduce the key concepts of HPE Primera storage virtualization. This section provides an overview of these concepts, and then discusses product architecture.

HPE Primera virtualization overview

The HPE Primera operating system has a volume manager that provides VV abstraction. It comprises several layers, as shown in [FIGURE 9](#). Each layer is created using physical disks, logical disks, and CPGs.

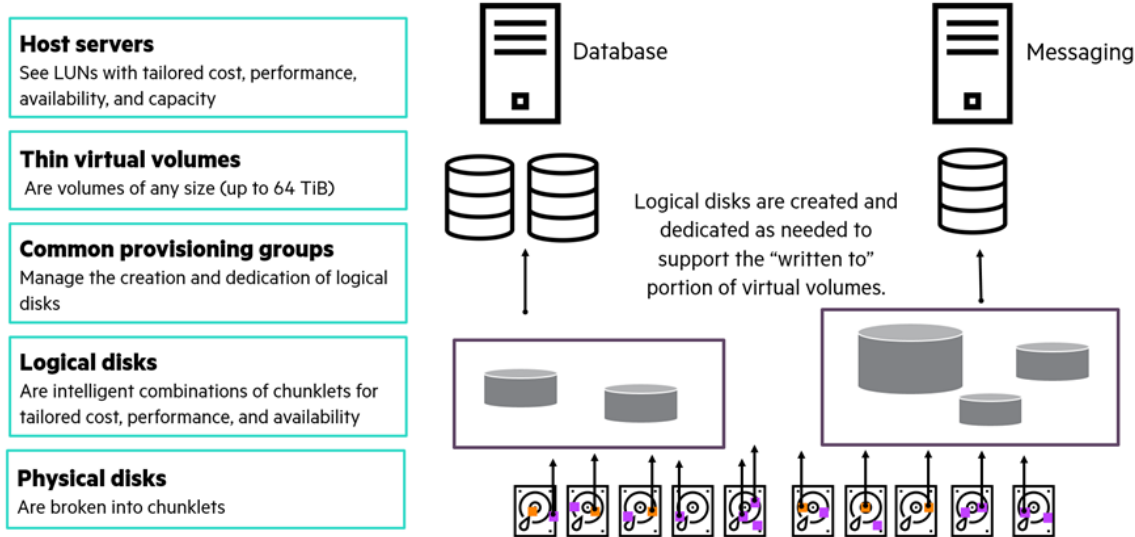


FIGURE 9. Layers of virtual resource abstraction

Physical disks

Every physical disk that is admitted into the system is divided into 1 GB chunklets. A chunklet is the most basic element of data storage in an HPE Primera array. These chunklets form the basis of the RAID sets; depending on the sparing algorithm and system configuration, some chunklets are allocated as spares.

Logical disks

The logical disk layer is where the RAID functionality occurs. Multiple chunklets RAID sets, typically from different physical disks, are striped together to form a logical disk. All chunklets belonging to a given logical disk will be from the same drive type. Logical disks can consist of all Nearline (NL), Fibre Channel (FC), or SSD-type chunklets. There are three types of logical disks:

- User (Usr) logical disks provide user storage space to fully provisioned VVs.
- Shared data (SD) logical disks provide the storage space for snapshots, physical copies, or TPVVs.
- Shared administration (SA) logical disks provide the storage space for snapshot and TPVV administration. They contain the bitmaps pointing to the pages of the shared data logical disk that are in use. The logical disks are divided into “regions,” which are contiguous 128 MiB blocks. The space for the virtual volumes is allocated across these regions.

Common provisioning groups

The next layer that defines the logical disk creation characteristics (such as RAID type, set size, disk type for chunklet selection, total space warning, and limit points) is the CPG. A CPG is a virtual pool of logical disks that allows volumes to share resources and to allocate space on demand. A TPVV created from a CPG will automatically allocate space on demand by mapping new regions from the logical disks associated with the CPG.



Virtual volumes

The top layer is the VV. Volumes draw their resources from CPGs and are the only data layer visible to hosts when they are exported as VLUNs.

The provisioning type classifies a VV as one of the following:

- **TPVV:** Thinly provisioned VV with space for the base volume allocated from the associated user CPG and snapshot space allocated from the associated snapshot CPG (if any).
- **Administrative volumes:** Created by the system for system use only.

TPVVs have three separate data components:

1. **User space** is the area of the volume that corresponds to the logical disk regions in the CPG that are available to the host. The user space contains the user data and is exported as a LUN to the host.
2. **Snapshot space** (or copy space) is the area of the volume that corresponds to logical disk regions in the CPG that contain copies of user data that changed since the previous snapshot. The snapshot space contains the copy data.
3. **Administration space** (or admin space) is the area of the volume that corresponds to logical disk regions in the CPG that track changes to the volume since the previous snapshot was created. The administration space contains pointers to copies of user data in the snapshot space. Administration space is managed by the system and not by the tools users use to manage user and snapshot space.

The anatomy of a virtual copy snapshot

The HPE Primera OS uses page tables to map the VV address space to physical storage pages, which are 16 KiB in size. The process is similar in concept to the virtual memory mapping page tables of an operating system. The logical block address (LBA) of a SCSI RW command is then used as an index into three different page tables to find a particular 16 KiB block. New writes result in a free 16 KiB page being allocated (or multiple 16 KiB pages for larger writes); rewrites of data simply reuse any previously allocated space.

The data from a virtual copy snapshot is stored in shared data (SD) logical disks in the Copy CPG. The bitmap and page table metadata indexes that point to the user data in the SD logical disks are stored in shared administration (SA) logical disks. Because the SA logical disks are critical to the system, they are three-way RAID 1 protected with high availability (HA) cage availability. In addition, the SA metadata indexes are cached in policy memory for best performance and efficiency of the lookups.

A snapshot volume consists of SA and SD areas:

- SA is the management (metadata) space that contains first level/second level (L1/L2) bitmaps. It indicates which SA and SD pages are used with the exception of allocation tables (L1/L2/L3) pointing to the snapshot location within the SD.
- SD contains the actual user data. For a base VV, this is the “old” base volume data that is copied to the SD before the new data is written to the base volume. The SD contains both new and old data.

The snapshot data space grows by the aggregate total of the base volume and has separate policy settings that can limit the growth of the space or establish allocation warnings from a separate policy established by the user. The size of the SD area is determined at time of creation based on the base volume it is associated with (volumes that obtain SA and SD from CPGs can automatically grow these areas). The size of the SA is automatically assigned based on the SD area size. Both areas are divided into 16K pages. By design, the system separates the space assigned to administration functions and the actual user data.

The HPE Primera OS uses bitmaps and tables to manage snapshots efficiently:

- The SD logical disks are made up of 128 MiB regions that contain 8192 16 KiB pages.
- The L1 bitmap indicates whether a given 128 MiB region in SD space is in use or not.
- If a region is in use, then an L2 bitmap indicates if a particular 16 KiB page in that region is in use.



The relationship between regions, pages, and bitmaps is presented in [FIGURE 10](#).

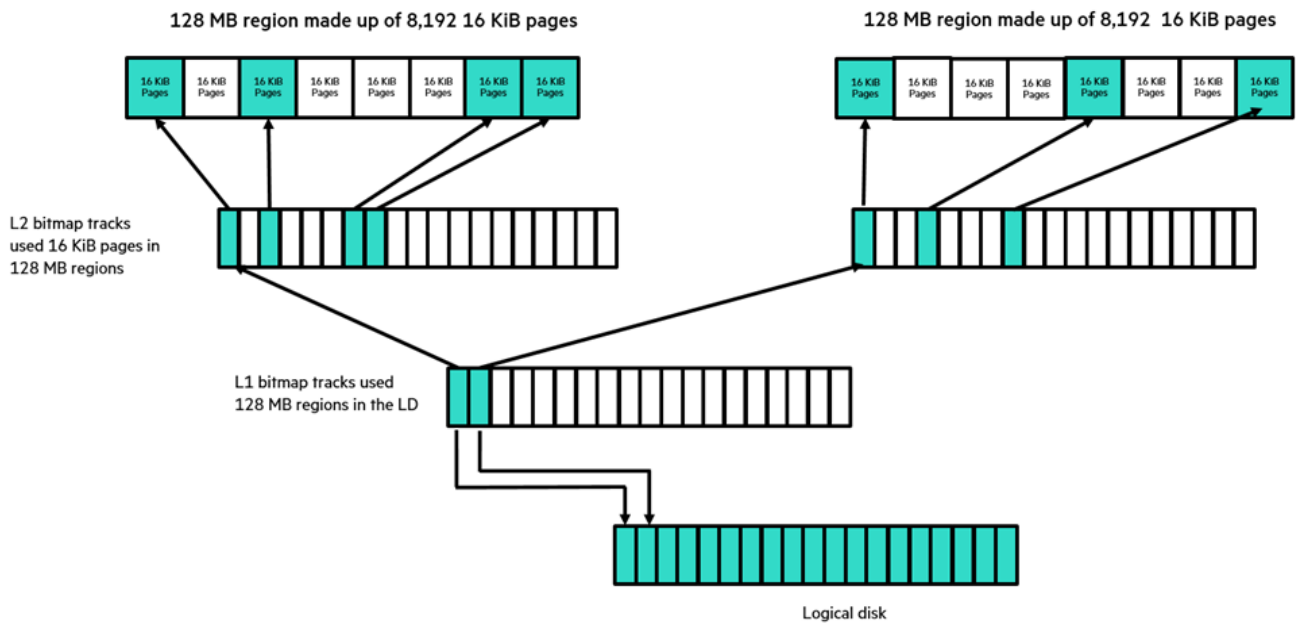


FIGURE 10. Overview of the SD space tracking bitmaps

There are three levels (L1, L2, and L3) of exception tables for each snapshot volume. They indicate the SD location (offset into the SD area) to write or retrieve the snapshot data. The host provides an LBA as part of each operation that indicates the page being accessed. Bits from the LBA are used to access the L1, L2, and L3 exception tables.

The L1 table consists of entries that are either:

- **INVALID:** Indicates that the snapshot range has not been written.
- **Pointer to an L2 table:** Points to the location of the L2 table in the SA area for the snapshot range.

Depending on the operation, the HPE Primera OS either stops processing based on the L1 content or moves to the L2 table to index into it for further processing.

The L2 table contains either:

- **INVALID:** No snapshot written.
- **Location of the L3 table:** Points to the location of the L3 table in the SA for the snapshot range.

The L3 table contains exception entries. The L3 table contains either:

- **INVALID:** No snapshot written.
- **Pointer to one 16 KiB page:** Points to the location within the SD area to which the data is written.

For operations consisting of multiple pages, the HPE Primera OS segments the individual read/requests into 16 KiB pages for processing. It walks through the three exception tables to determine if a snapshot has been taken for the requested location and, if so, accesses the data based on the specific operation requested.



Understanding the HPE Primera COW implementation

HPE Primera Virtual Copy snapshots leverage HPE Primera Thin Provisioning technology. From a performance perspective, they behave like TPVVs. Writes to a base VV initiate a DCOW for TPVVs. This process copies data from the original base VV in a cache memory page (CMP). When it is time to flush data, the flush process aggregates CMPs for a given logical disk into a larger write.

These CMPs might be different for virtual copies or TPVVs. This HPE Primera implementation mitigates overhead or performance impact on the host because of the COW. The COW technique is space efficient, because the storage space required to create a snapshot is minimal to hold only the data that is changing. The data is only valid when the original copy is available.

Write operations

When the host updates a volume for which a snapshot has been created, before updating the volume, the VV layer determines if a COW needs to be performed. The COW must occur before the base volume data is updated so that the original data is preserved for snapshots.

After a write operation by a host, the data is written to land in a node’s memory at a CMP. Before the platform returns an acknowledgement to the host that the write has completed, the platform must have written two copies of the host data. Because an acknowledgement implies that the platform contains multiple copies of the data, the acknowledgement does not occur until this is true. After the data is replicated to another node in the array, the host can be acknowledged for the write operation. At a later point, the flusher pushes the data to disk.

To improve performance from the host’s perspective, the HPE Primera OS uses a DCOW. [FIGURE 11](#) illustrates, at a high level, a base volume where a snapshot has been taken. An RO snapshot is created first, then an RW snapshot. No data is overwritten yet in the base volume, so the base VV and RO as well as RW snapshots all point to the same data blocks.

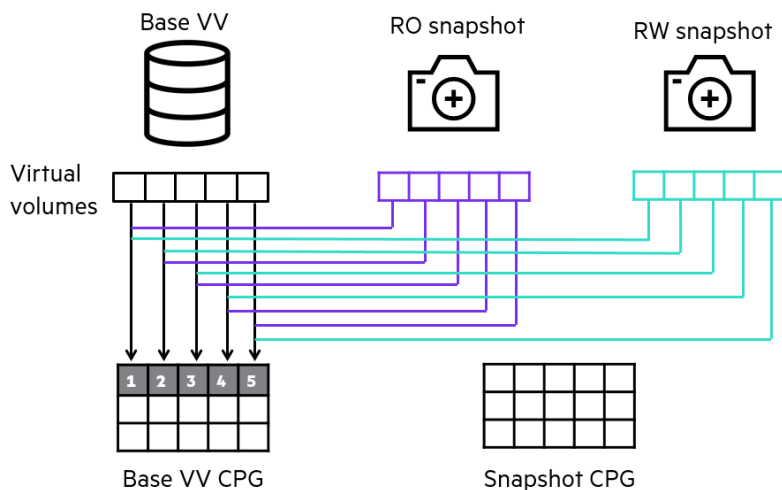


FIGURE 11. Write operation to a base virtual volume (TPVV)

When a write I/O is issued to the source volume of a VV or snapshot, the COW function preserves the data that is to be overwritten. The data is copied to the snapshot space associated with the original VV before the write operation is completed, and a pointer in the administration space points to the copied data.



An overwrite to block 1 with new data 6 results in copying block 1 to the Copy CPG space and modifying the RO and RW snapshot pointers to point to that CPG space, as exhibited in [FIGURE 12](#). A new write to the RW snapshot results in the data being stored in the Copy CPG space. An example is the write of the new data 7, as illustrated in [FIGURE 12](#).

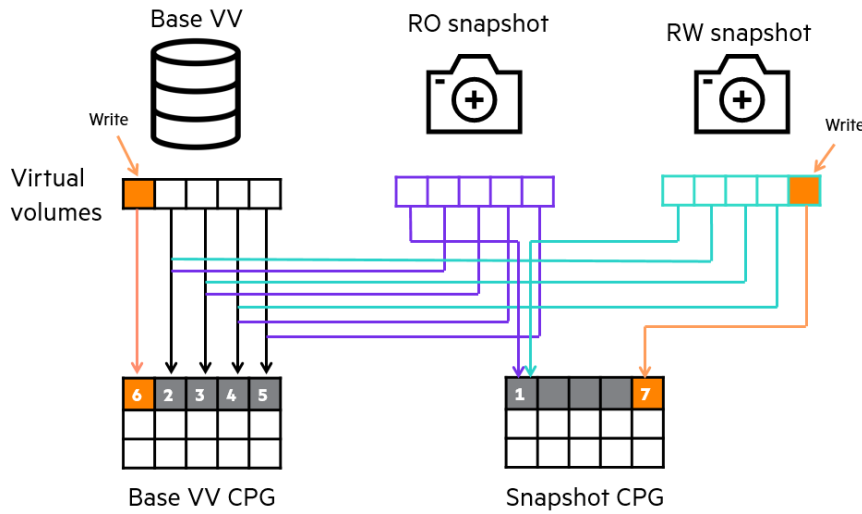


FIGURE 12. Preserving data to be overwritten for TPVV

HPE PRIMERA VIRTUAL COPY ROW AND HPE PRIMERA THIN DEDUPLICATION

Volumes with deduplication enabled can have snapshots. These snapshots can have the same or a different CPG from the user's CPG. Data for snapshots of these volumes resides in the snapshot space. Only duplicated data for snapshots resides in the dedupe store if the snapshots and the base volumes have the same CPG.

Snapshot data and base volume data can point to the same place if they have the same data. When a snapshot is created from a deduplication enabled volume, it shares all data with the base volume. But, if a duplicate page of data is written to the snapshot, it is deduped and both the snapshot and base volume point to the same place in the dedupe store. [FIGURE 13](#) shows the layout for a deduplicated volume when a snapshot is taken. Note that in this case, a single CPG for both user and Copy CPG is being used.

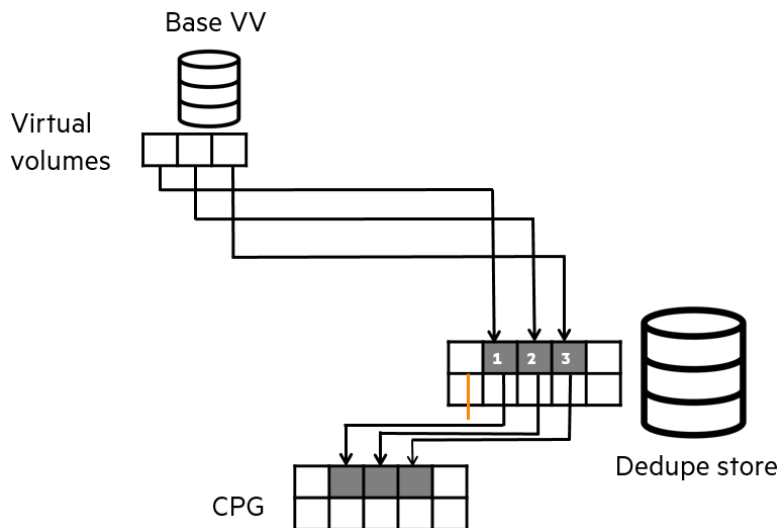


FIGURE 13. Reduplicated volume layout for a snapshot



The following two paragraphs describe how ROW works with deduplicated data stored in DDS.

FIGURE 14 displays at a high level a deduplication enabled volume where a snapshot has been taken. An RO snapshot is created first, and then an RW snapshot is created. No data has been overwritten yet in the base volume, so base VV and RO as well as RW snapshots all point to the same data blocks. When a data block of the source volume of a VV or snapshot is stored in DDS and is written to, the ROW function preserves the data that is to be overwritten. The new write is stored in a new block in the base volume private space (assuming it did not match any previously stored block in the dedupe store so it can be deduped).

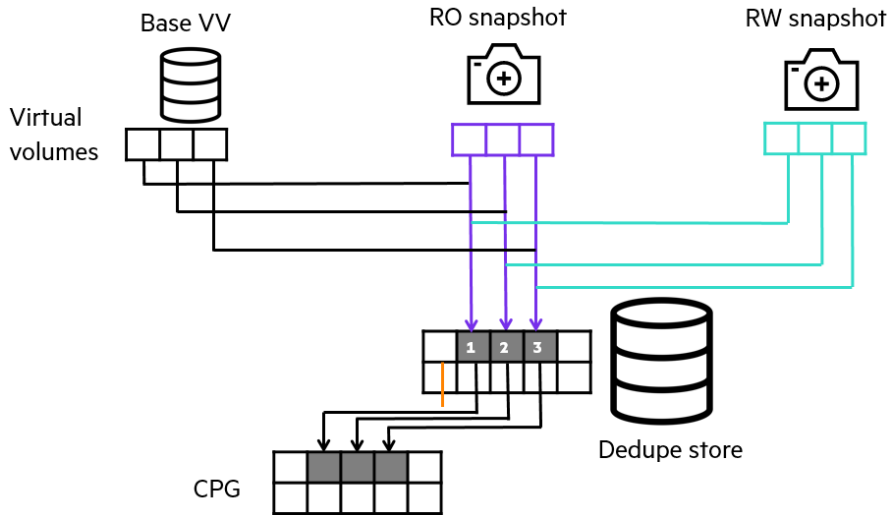


FIGURE 14. Deduplication enabled volume and its snapshot

An overwrite to block 1 with new data 4 results in writing the data block 4 to the base volume and maintaining the RO and RW snapshot pointers that reference the old copy of the data block 1, as presented in FIGURE 15. A new write (block 5) to the RW snapshot simply results in the data being stored in the Copy CPG space.

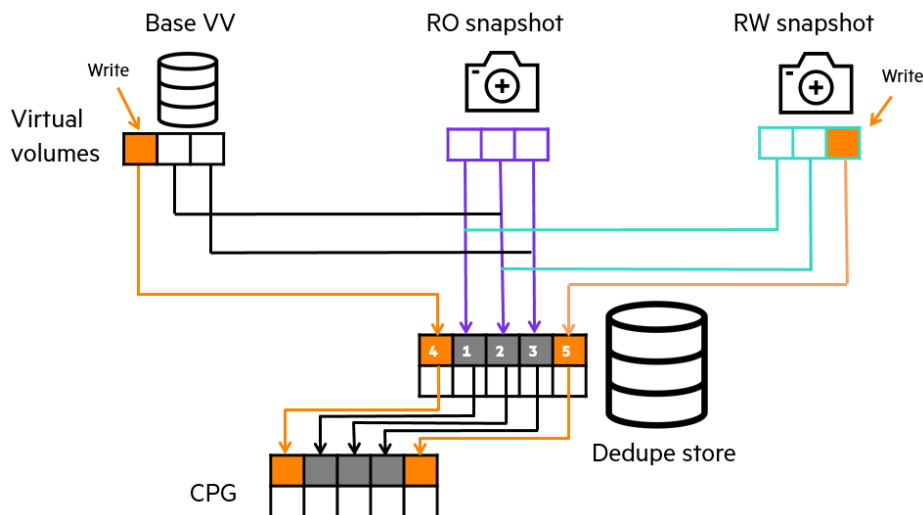


FIGURE 15. Copied data pointer



RW snapshot behavior

The following points summarize RW behavior, depending on the type of snapshot being operated on:

- Write to a RW snapshot if the address was not already written:
 - For a TPVV, perform a COW, update the exception tables, and update the bitmaps.
 - For a deduplication enabled volume, perform a COW for data stored in DDC or perform an ROW for data stored in DDS, update the exception tables, and update the bitmaps.
- Write to a RW snapshot if the address was already written:
 - If the snapshot has already been written at the specified offset, a page in the SD exists for the address. The new data is written on the same SD area (there is no update to the base VV).
- Read from a RW snapshot if the address was not already written. Check the exception tables. If an exception for the given address exists, the data is returned from the SD. Otherwise, examine the VV tree:
 - If a parent snapshot has created an exception for the specified address, return the data from the SD area as pointed out by the exception.
 - If the entire tree is examined without finding an exception, return the data after reading it from the base volume.
- Read from a RW snapshot if the address was already written. Walk through the exception tables to find the address from which to return the data.
- Read from a RO snapshot if the address was not written in the parent after creating the snapshot. Return the data after reading it from the base volume.
- Read from a RO snapshot if the address was written in the parent after creating the snapshot. Return the data from the SD area pointed out by the parent's exception.

Virtual copy snapshot relationships

The relationships between the virtual copies are derived from the base volume and can be represented as a tree. The following rules are enforced by the system when creating a snapshot:

- The tree grows in alternating layers of RW and RO snapshots. You can only make RO copies of RW volumes, and you can only make RW copies of RO volumes.
- **S0** is the first virtual copy made of the base VV.
- **S2** is the most recent virtual copy.
- Each copy tracks changes made to the base VV from its own creation date until the next snapshot is made.
- **S1_0** can be created at any time after S1 is created.

A virtual volume cannot be deleted if a child copy of it exists. For example, **S1** cannot be removed unless **S1_0**, **S1_0_0**, and **S1_0_1** are deleted first.



The example in [FIGURE 16](#) shows two different tree structures: the dashed blue arrows show the **Copy of** relationships, and the black arrows show the parent relationship. For example, **S0** is an **RO** copy of the base VV, and **S1** is the parent of **S0**. The **Copy of** relationship shows that the snapshot was created by copying another VV. The parent relationship refers to the internal organization of the administration space. The parent volume contains information to reconstruct the snapshot represented by the child volume. A parent volume can have a creation date after its child if the parent volume is modified.

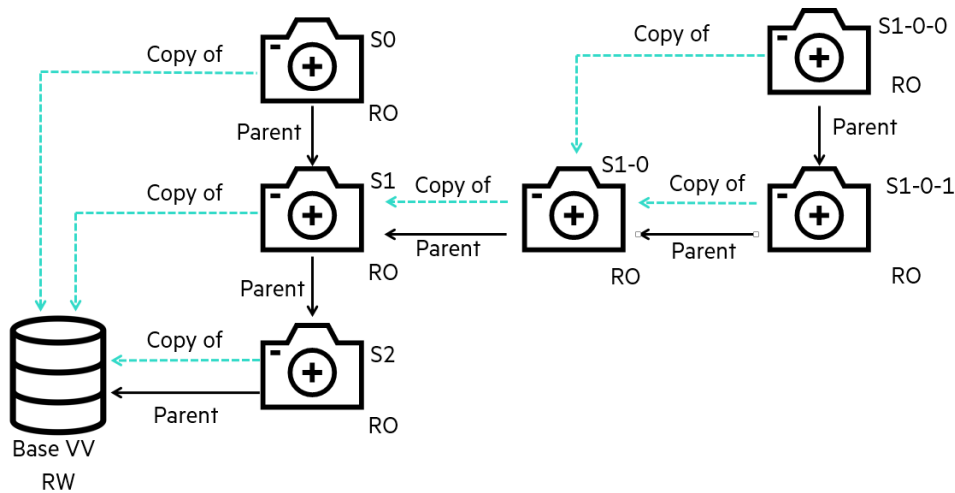


FIGURE 16. HPE Primera Virtual Copy snapshot child and parent relationships

A parent relationship is useful for two reasons:

1. For understanding the performance consequences of virtual copies. The tree representing the parent relationship shows the look-up paths in the administration space that are needed to reconstruct the earlier state of the VV. The farther a virtual copy is from the base volume, the longer it will take to retrieve it. If a snapshot is expected to be kept in use for a long time, consider making a physical copy instead of a virtual copy.
2. For understanding which virtual copies become stale if the administration space is full and the COW data cannot be written. A stale snapshot is one that cannot be completely re-created, because the most recent changes will not be included. The current snapshot and all its children become stale when a write fails. For example, if there is no space to write the COW data when a host writes to **S1_0**, then **S1_0**, **S1_0_1**, and **S1_0_0** become stale.

Users create snapshots using the HPE Primera CLI, HPE SSMC, or the HPE Primera User Interface. They have options during the creation of the snapshot to a particular volume to customize the snapshot volume, including whether it is RO or RW. Base volumes are always RW, but virtual copy snapshots can be RW or RO. The rules that govern the relationships between a base volume and its virtual copy snapshots are based on the difference between RW and RO volumes. RO copies must alternate. Users can make a RO copy of a RW volume, and they can only make a RW copy of a RO volume.

Users can specify whether a snapshot is RO when creating the snapshot (by default, snapshots are RW). However, there are certain rules that are followed regarding whether a snapshot can be written to:

- The first snapshot (closest to the base volume) is always a RO snapshot. If a user tries to create a RW snapshot on the first snapshot associated with a volume, the attempt will not be successful.
- The most recent snapshot will get the latest page for a volume being written. This is because it is quicker to write the most recent snapshot, as it is closest to the base volume. (Conversely, it takes longer to retrieve pages the farther the snapshot is from its base volume). RO snapshots are considered child snapshots of a parent (that is RW-capable).



An HPE Primera implementation of the RO snapshot saves the original data from a run-time allocated snapshot data space before a new write is committed to the base volume or RW snapshot. This way, a read-only snapshot only saves the original data if there are new writes to the base volume or read-write snapshot. [FIGURE 17](#) shows a simple alternating RO and RW snapshot relationship from a base volume.

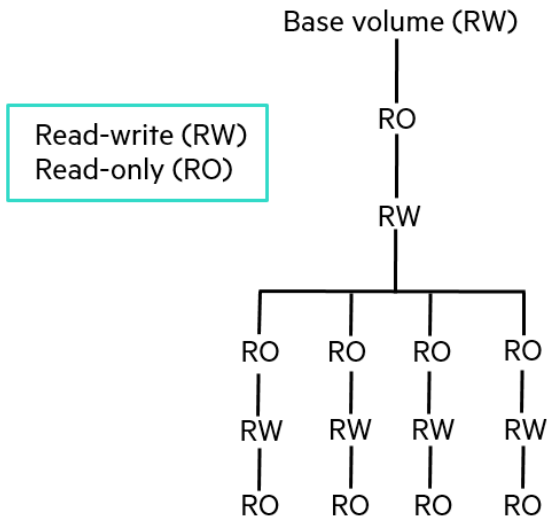


FIGURE 17. Alternating RO and RW HPE Primera Virtual Copy snapshots

Snapshots can also be taken from other snapshots. [FIGURE 18](#) shows a complex relationship between the base volume and the correlating snapshot tree. Notice that the snapshots also have their own tree associated with them. Any one of the virtual copy snapshots can be promoted back to another snapshot in the tree and become a base volume. The default for promotion is back to the base (or parent) volume.

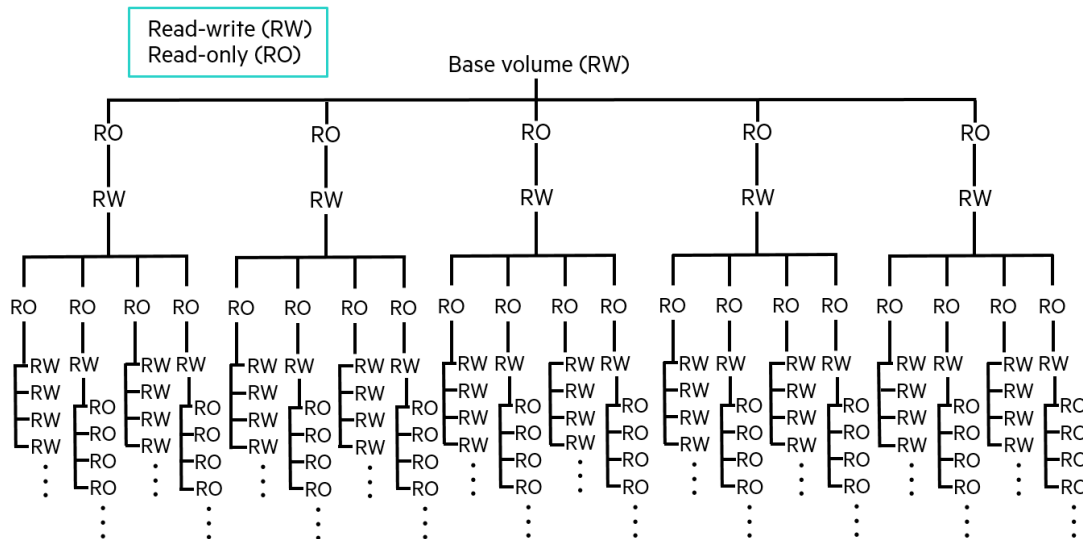


FIGURE 18. Base volume and virtual copy snapshot relationships



Although [FIGURE 18](#) is a complex illustration, some customers do use such a dispersal in virtual copy snapshots for various PIT recoveries as well as test and development environments. Because there is only one COW or ROW needed for this tree, there is little impact on the performance of the base volume.

HPE PRIMERA VIRTUAL COPY SNAPSHOT RESTORE AND REFRESH

HPE Primera Virtual Copy snapshots can be used to restore a parent VV. The VV will be restored to match the PIT data when the snapshot is taken. Such an operation offers the user an option to recover data in case of file deletion or data inconsistency.

The data can be restored using the **promote** task within the HPE SSMC or the HPE Primera CLI. This operation can be done using the online option. Some tasks must be performed from the host operating system for the operation to be successful. The next section shows a sample restore operation using the HPE Primera CLI `promotesv` command.

Another efficient way of managing the data in an HPE Primera environment with HPE Primera Virtual Copy is the ability to refresh the snapshots. Such an option mitigates the need to create a new snapshot and instead updates or refreshes an existing snapshot. The `updatevv` command refreshes a snapshot VV with a new snapshot. The command first verifies if the users have permission to run the `updatevv` command on the specified VVs to protect the user data.

If the snapshot is RW, then its RO parent is replaced by a new RO snapshot of the same name. Then a new RW snapshot is created of the new RO snapshot. However, if the specified snapshot is a RO snapshot, the command only updates the requested snapshots. The next section has a sample refreshing operation using the HPE Primera CLI `updatevv` command.

HOW DOES HPE PRIMERA VIRTUAL COPY SNAPSHOT REFRESH WORK WITH UPDATEVV?

This section describes the steps needed to perform different virtual copy snapshot refreshes. The snapshot refresh uses the HPE Primera CLI `updatevv` command. The virtual copy snapshot refreshes that can be performed are a manual refresh snapshot, refresh the snapshot of a base volume, and refresh a snapshot of a snapshot.

Manual methods to refresh snapshots

As shown in [FIGURE 19](#), in most traditional implementations, refreshing a copy of production data requires executing the following steps at regular intervals:

1. Deconfigure snapshot LUNs from the application.
2. Unexport the previous snapshot.
3. Delete the snapshot.
4. Create a new snapshot.
5. Export the snapshot to the application host.
6. Discover the new LUN.
7. Reconfigure the application to access the LUN.



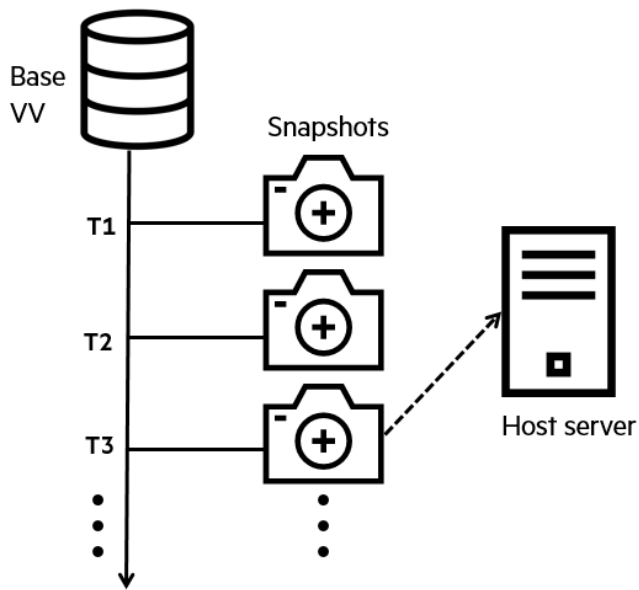


FIGURE 19. Steps for refreshing a copy of production data

The repetitive tasks of deconfiguring application storage, creating a new snapshot, discovering a new LUN, and configuring application storage creates additional overhead for an administrator.

The `updatevv` CLI command refreshes a snapshot with current data on the base virtual volume while retaining the snapshot's LUN information, such as a World Wide Name (WWN) and host access details. Syntax and usages of this CLI command are discussed in the [Snapshot refresh commands and procedure](#) section of this white paper.

The refresh process eliminates the need to create a new snapshot, export it to the host application, and configure the application to access the LUN. This process quiesces the I/O on the parent volume and flushes all data to the drives during its operation.

For efficiency, instead of physically copying the updated data blocks from the parent volume to the target snapshot, the refresh process creates a new snapshot internally with all the existing LUN information. The host recognizes the new snapshot as the old device, but with new data on it after remounting the file system or reopening the block device. The refresh process also changes the creation time of the snapshot.

Snapshots are put in standby mode during the refresh process. In this mode, hosts receive disk resets for all I/O requests. However, there is no impact to the base virtual volumes.

Refreshing the snapshot of a base volume

When using the `updatevv` command to refresh snapshots, it is necessary to create an initial snapshot of the production volume and export it to the application host only one time. On the host, the user would then scan and discover the LUN, and configure the application to access the data on the snapshot.



As shown in [FIGURE 20](#), to refresh the snapshot from a current copy of the production data, follow this simple procedure:

1. Pause the host application I/O, or unmount the file system.
2. Initiate an updatevv refresh, and wait for the operation to complete successfully.
3. Remount the file system, and reinitialize the application to recognize the new, refreshed data.

You can repeat this procedure whenever the snapshot data needs to be refreshed.

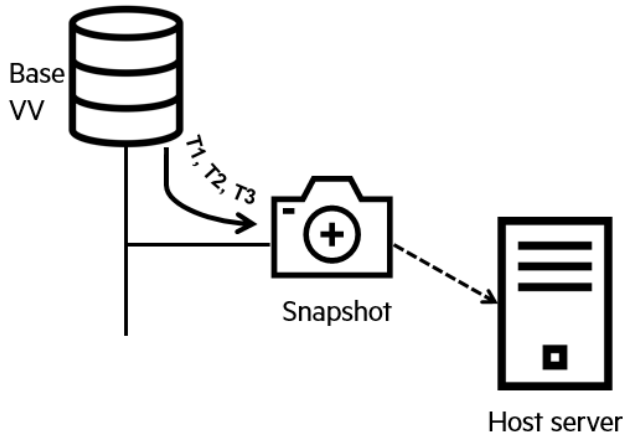


FIGURE 20. HPE Primera snapshot refresh steps

Refreshing a snapshot of a snapshot

If a refresh is performed on a snapshot that is a child of a snapshot, then the refresh process refreshes the data with its immediate parent snapshot. In the example shown in [FIGURE 21](#), **Snap2** data is updated with **Snap1** data.

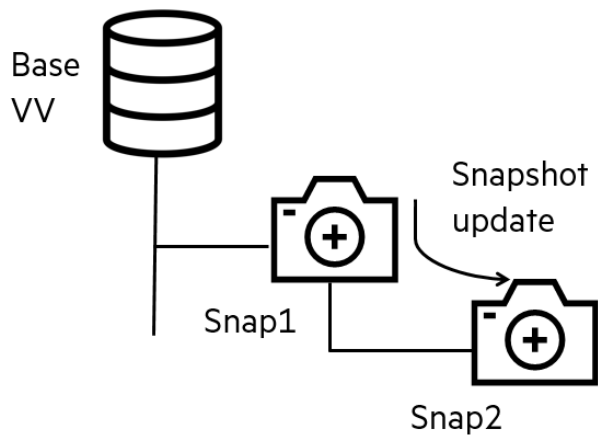


FIGURE 21. Snapshot refresh for a snapshot of a snapshot



If the specified snapshot is a read-write snapshot of a read-only parent snapshot, the snapshot and its parent snapshot can be refreshed with base virtual volume data. As shown in [FIGURE 22](#), data is first refreshed on the **Snap1** snapshot and then on the **Snap2** snapshot.

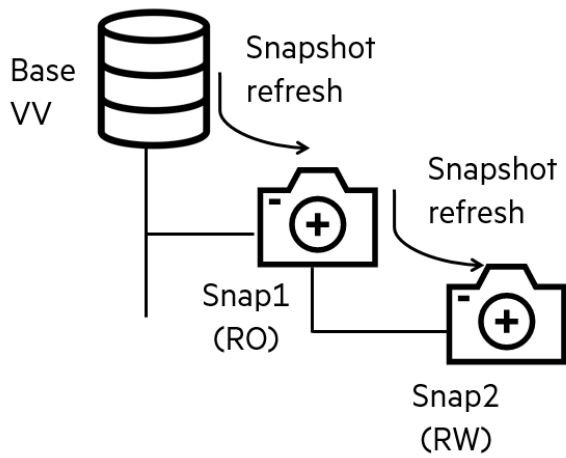


FIGURE 22. A snapshot and its parent snapshot both are refreshed with the base volume

If the read-only parent snapshot has multiple read-write snapshots (see [FIGURE 23](#)), the refresh operation does not refresh the parent snapshot. Instead, a new read-only parent snapshot is created and then the data on the targeted snapshot is refreshed.

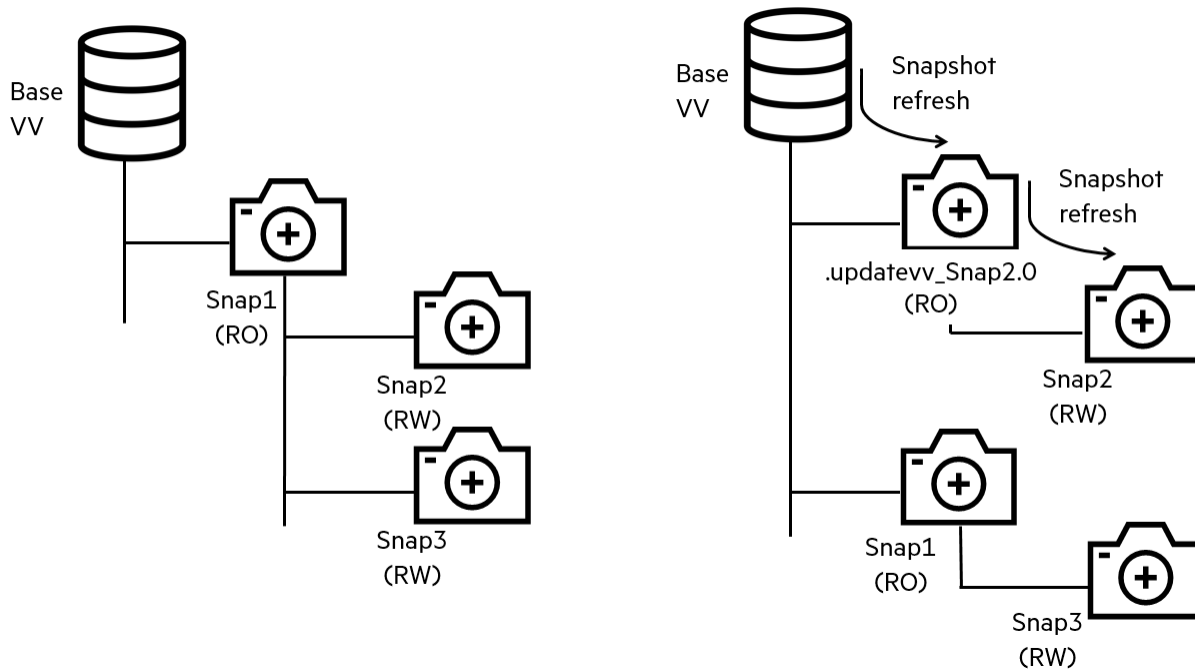


FIGURE 23. Before and after of a snapshot refresh with multiple snapshots and the "-ro" option



HPE PRIMERA VIRTUAL COPY SNAPSHOT OPERATIONS

The HPE SSMC provides an easy way to manage HPE Primera Virtual Copy snapshot operations. It provides an easy-to-read graphical view of virtual copy snapshots. The snapshot tree and parent relationships are viewable in the HPE SSMC and the HPE Primera CLI. FIGURE 24 through FIGURE 27 illustrate the creation of a PIT snapshot copy of a VV. These snapshots can be used to perform tasks such as backing up data on the base volume and allowing multiple copies of a base volume to be modified without affecting the original base volume.

FIGURE 24 shows the **Create Snapshot** screen. This screen allows the user to select the **Access permissions** as well as **Retention time** and **Expiration time**.

Create Snapshot Additional Settings ?

Additional Settings

▼ Snapshot

Access permissions **Read/Write** ▾

Volume ID **Auto ID**

Retention time **Enabled** | 3 | days ▾ ~Apr 3, 2020 10:21:07 AM MDT

The retention time cannot be removed or reduced once it is set. This volume cannot be removed until its retention time expires.

Expiration time **Enabled** | 5 | days ▾ ~Apr 5, 2020 10:21:10 AM MDT

Comments optional

Changed: Expiration time to "5" **Create** **Cancel**

FIGURE 24. Creating a snapshot



FIGURE 25 shows where to select the **Create snapshot** option for a specific virtual volume by using the **Actions** drop-down menu.

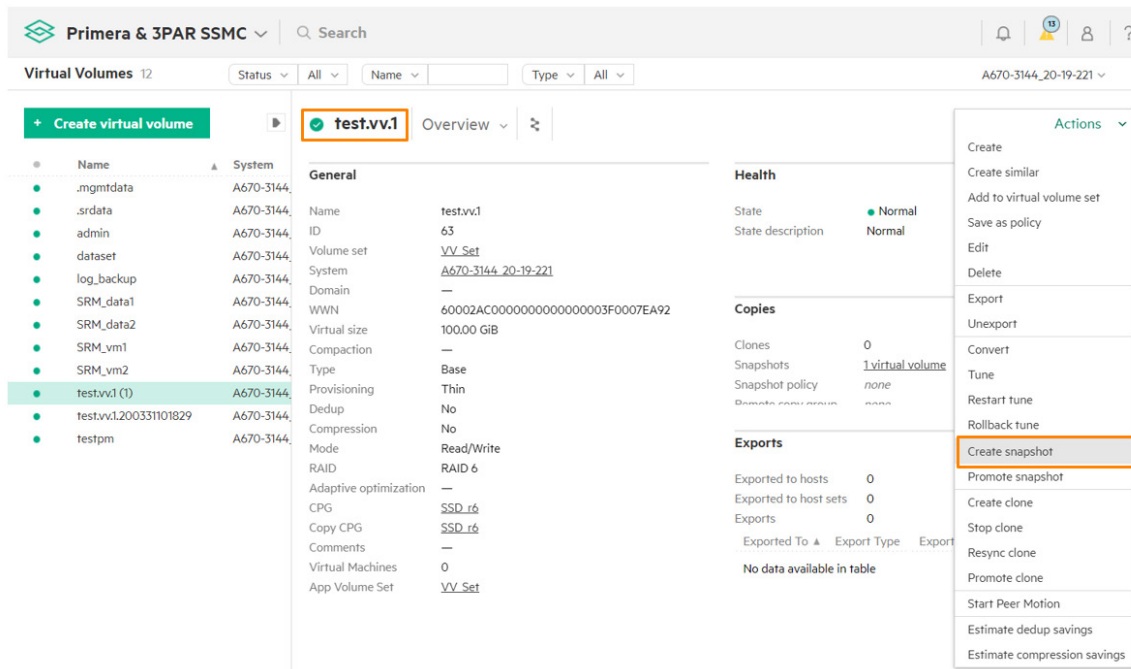


FIGURE 25. Using the HPE SSMC to create a snapshot



FIGURE 26 shows the **Additional Settings** screen from the HPE SSMC for the advanced options of a snapshot.

The screenshot displays the 'Create Snapshot' interface in the HPE SSMC. It is divided into two main sections: 'General' and 'Additional Settings'.
General Section:
- System: A670-3144_20-19-221
- Domain: —
- Parent volume: test.vv.1 (highlighted with an orange box)
- Compression: No
- Dedup: No
- Volume snapshot policy: none
- Create snapshot: Once Assign policy
- Snapshot name pattern: Parent Volume & Time Stamp (dropdown menu)
- Snapshot name: VolumeName.170605141808
- Schedule pattern: Once now (dropdown menu)
- Schedule Time on the array: 03/31/2020 01:50:47 PM MDT
Additional Settings Section:
- Snapshot (expanded)
- Access permissions: Read/Write (dropdown menu, highlighted with an orange box)
- Volume ID: Auto ID
- Retention time: Disabled
- Expiration time: Disabled
- Comments: [Text area] optional
At the bottom right, there are two buttons: 'Create' (green) and 'Cancel' (white with green border).

FIGURE 26. Setting advanced options for creating a snapshot using the HPE SSMC



FIGURE 27 shows a mapping of the virtual volumes and any associated snapshots. This particular example shows only one virtual volume that has a single snapshot.

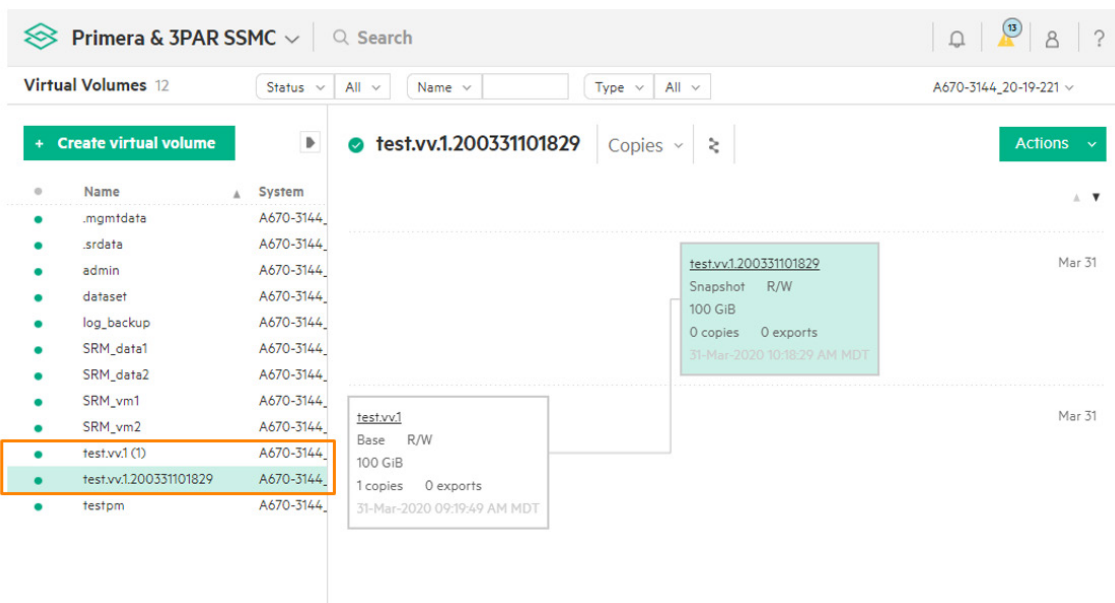


FIGURE 27. Mapping of a snapshot to its parent VV

The HPE Primera CLI can also be used to create a virtual copy snapshot. The following example creates a RO snapshot volume `ha.1-snapro` from volume `ha.1`:

```
cli% createsv -ro ha.1-snapro ha.1
cli% showvv ha.1*
```

Id	Name	Prov	Type	CopyOf	BsId	Rd	-Detailed_State-	---Rsvd(MB)---	Adm	Snp	Usr	VSize
173	ha.1		tpvv	base	---	173	RW		normal	1024	2048	2048
178	ha.1.1601311175705	snp	vcopy	ha.1	173	RW	normal	--	--	--		20480
179	ha.1-snapro	snp	vcopy	ha.1	173	RO	normal	--	--	--		20480
-----								1024	2048	2048	61440	
3	total											



Creating consistent group snapshots of virtual volumes

Consistent group snapshots are all created at the same point in time so that related structures on different volumes in the group remain consistent. In this example, HPE SSMC snapshots are created for the three VVs in a volume set. [FIGURE 28](#) shows where to select the **Create snapshot** option for a specific app volume set by using the **Actions** drop-down menu.

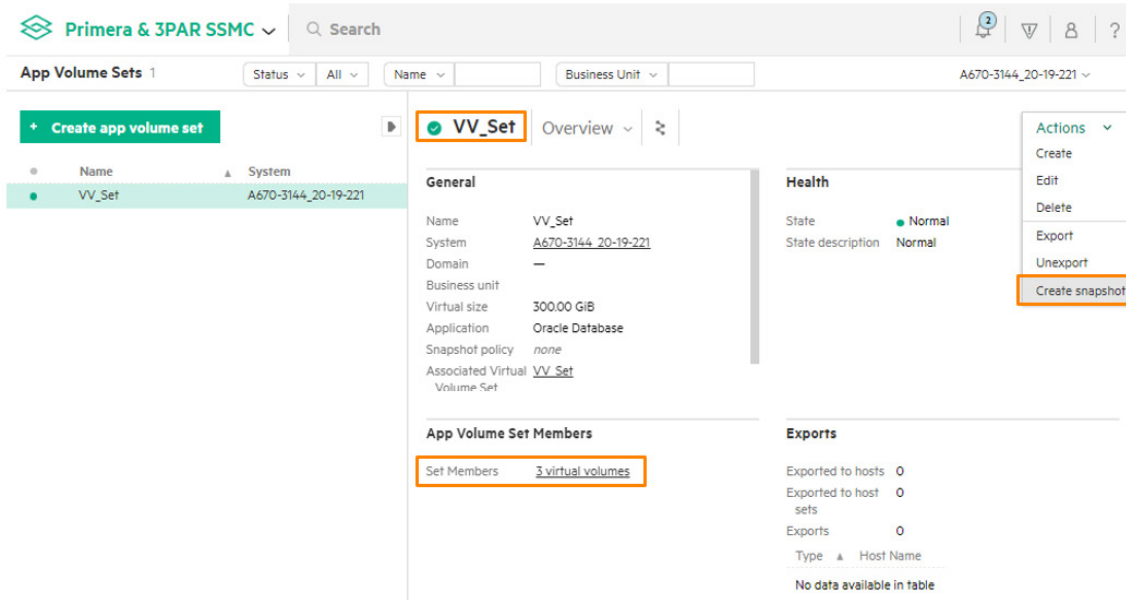


FIGURE 28. Creating a group snapshot for three volumes



FIGURE 29 shows the additional settings screen in the HPE SSMC for the advanced options of a group snapshot.

The screenshot displays the 'Create Snapshot' configuration page in the HPE SSMC. The page is divided into two main sections: 'General' and 'Additional Settings'.
General Section:
- System: A670-3144_20-19-221
- Domain: —
- Parent app volume set: VV_Set (highlighted with an orange box)
- App volume set snapshot policy: none
- Create snapshot: Once, Assign policy
- Snapshot name pattern: Parent Volume & Time Stamp
- Snapshot name: VolumeName.170605141808
- Schedule pattern: Once now
- Schedule Time on the array: 03/31/2020 01:59:17 PM MDT
Additional Settings Section:
- Snapshot (expanded):
- Volume ID: Auto ID
- Retention time: Disabled
- Expiration time: Disabled
- Comments: [Text area] optional
At the bottom right, there are two buttons: 'Create' (green) and 'Cancel' (white with green border).

FIGURE 29. Setting options when creating group snapshots



FIGURE 30 shows the HPE SSMC view of consistency group snapshots.

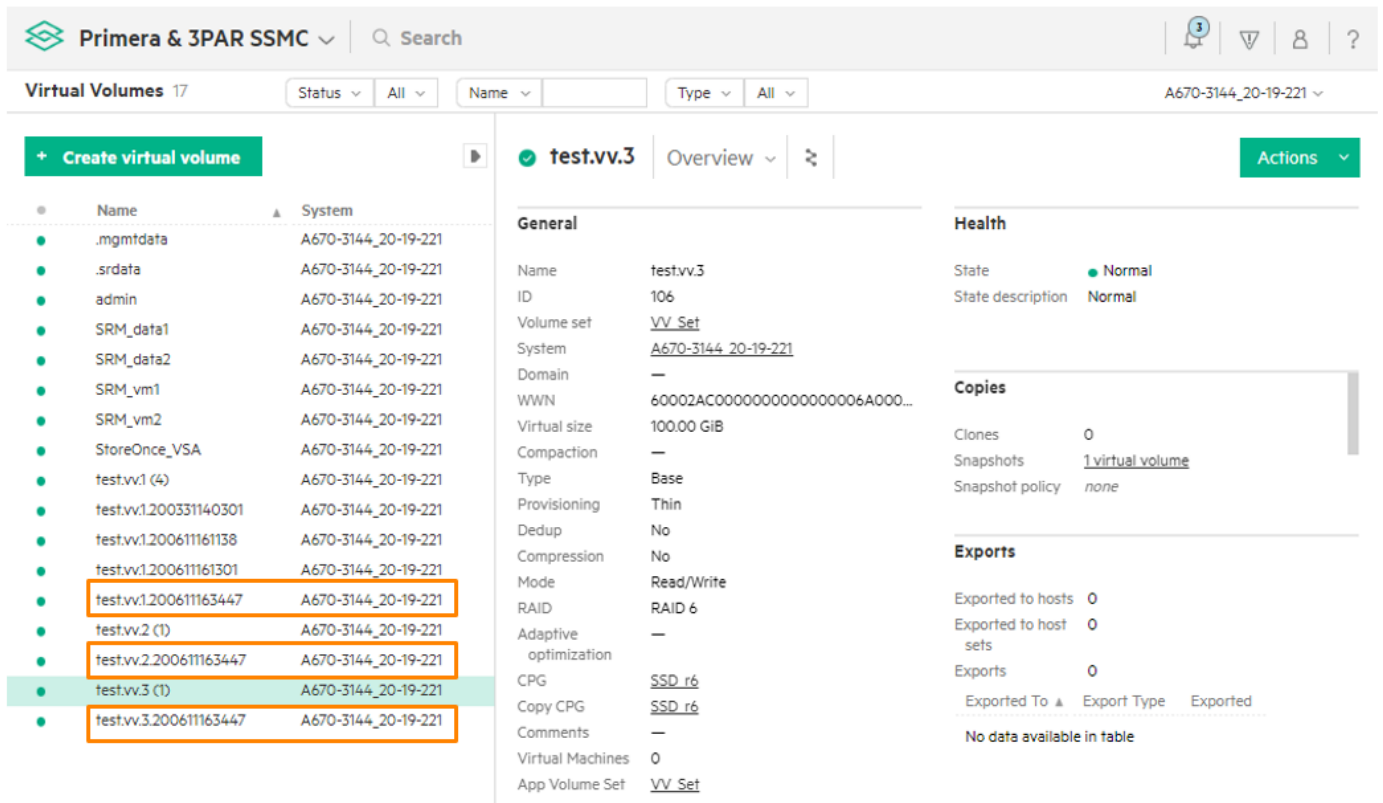


FIGURE 30. View of consistency group snapshots

The HPE Primera CLI `creategroupsv` command is used to create a snapshot of each of the volumes. The snapshot is displayed as:

```
cli% creategroupsv ha.2
ha.3 ha.4 CopyOfVV SnapshotVV
ha.2      ha.2.rw
ha.3      ha.3.rw
ha.4      ha.4.rw
```

Promoting a snapshot to a base volume

The promoting a snapshot to a base volume task copies the differences of a snapshot back to its base volume, allowing the user to revert the base volume to an earlier point in time. If the parent volume is exported, then the use of the `promotesv` command will fail, illustrated as:

```
cli% promotesv
ha.3.rw Error: RW parent (ha.3)
has exports defined
```

In this case, one alternative would be to use the online option to promote a snapshot with the `promotesv` command, shown as:

```
cli% promotesv -online ha.3.rw
Task 3874 has been started to promote virtual copy ha.3.rw
```



If the volume is not exported to a given host, the snapshot can be promoted to a parent volume, exhibited as:

```
cli% promotesv ha.5.160131200257
Task 3879 has been started to promote virtual copy ha.5.160131200257
```

Promoting a group of snapshots to their base volumes

The `promotegroupsv` command copies the differences of snapshots back to their base volumes, allowing you to revert the base volumes to an earlier point in time, presented as:

```
cli% promotegroupsv -online ha.2.160131181036 ha.3.160131181036 ha.4.160131181036
Task 3881 has been started to promote virtual copy ha.2.160131181036 to base
Task 3882 has been started to promote virtual copy ha.3.160131181036 to base
Task 3883 has been started to promote virtual copy ha.4.160131181036 to base
```

Snapshot refresh commands and procedure

This section explains the steps involved and procedure to follow to perform a snapshot refresh operation.

HPE Primera user permissions

HPE Primera users, other than those with “Super” or “Edit” roles, require permission to run the `updatevv` command. The permission can be set using the `setuseracl` command. The following example sets the permission for user `testuser1`, and allows the user to refresh the snapshot of virtual volumes `vv1` and `vv2`:

```
cli% setuseracl testuser1 updatevv vv1 vv2
```

For more information on the `setuseracl` command usage and syntax, see the [HPE Primera OS 4.2 Command Line Interface Reference Guide](#).

Refreshing snapshots

The `updatevv` CLI command is used to perform a refresh operation. In the following example, `testsv` is a snapshot of the virtual volume `testvv`. The command `updatevv` is used to refresh the snapshot with data from its parent volume:

```
cli% showvv -s test*
-----Snp----- -----Usr-----
--(MiB)-- -(% VSize)-- --(MiB)-- -(% VSize)--
  Id Name   Prov Compr Dedup Type Rsvd Used Used Wrn Lim Rsvd Used Used Wrn Lim
65231 testvv tppv No    No    base 1536 780 7.6  0  0 3584 2659 26.0  0  0
65241 testsv snp  NA   NA   vcopy --  *0 *0.0  0  0  --  --  --  --  --
-----
  2 total                1536 780                3584 2659
cli%
cli% updatevv testsv
Updating VVs testsv
select q=quit y=yes n=no: y
cli% showvv -s test*
-----Snp----- -----Usr-----
--(MiB)-- -(% VSize)-- --(MiB)-- -(% VSize)--
  Id Name   Prov Compr Dedup Type Rsvd Used Used Wrn Lim Rsvd Used Used Wrn Lim
65231 testvv tppv No    No    base 1536  0 0.0  0  0 3584 2659 26.0  0  0
65241 testsv snp  NA   NA   vcopy --  *0 *0.0  0  0  --  --  --  --  --
-----
  2 total                1536  0                3584 2659
cli%
```



A group of snapshots can be refreshed as a set. For instance, add the snapshots to a set, and perform the refresh on the set. The following example updates all snapshots in the VVset `snp_set`:

```
cli% updatevv -f set:snp_set
```

For more information on the `updatevv` command usage and syntax, see the [HPE Primera OS 4.2 Command Line Interface Reference Guide](#).

The following `updatevv` command can be used to refresh a snapshot of another snapshot. See [FIGURE 21](#) for the snapshot that is refreshed.

```
cli% updatevv Snap2
Updating VVs Snap2
select q=quit y=yes n=no: y
cli%
```

As shown in [FIGURE 23](#), with the `-ro` option, data is first refreshed on the **Snap1** snapshot and then on to **Snap2**.

```
cli% updatevv -ro Snap2
Updating VVs Snap2
select q=quit y=yes n=no: y
cli%
```

Scheduling a snapshot refresh task

Snapshots of production volumes can be preconfigured, and the snapshot refresh can be initiated later at desired intervals in order to continuously update the snapshot with recent data from a production volume. The following `createsched` command creates a scheduled task named `daily_db_refresh_task` that runs every day at 8 am to refresh the snapshot set `db_snap_set` with its parent volumes:

```
cli% createsched "updatevv -f set:db_snap_set" "0 8 * * *" daily_db_refresh_task
```

For more information on the `createsched` command usage and syntax, see the [HPE Primera OS 4.2 Command Line Interface Reference Guide](#).

HPE PRIMERA VIRTUAL COPY SNAPSHOT CONSIDERATIONS

This section provides certain best practices that can be used when implementing HPE Primera Virtual Copy snapshots.

The time needed to create and remove snapshots impacts the ability to scale the number of snapshots that can be created or deleted. The challenges revolve mainly around the time taken to perform metadata updates during snapshot creation and removal, and the time taken to perform snapshot removals that result from the number of metadata updates. Hewlett Packard Enterprise recommends that the user keep the number of HPE Primera Virtual Copy snapshots created to no more than 300 snapshots per hour.

Consider using scheduling when creating snapshots with a specified retention count to provide a defined set of recovery point granularities over time. Create a schedule for snapshot clean tasks, because this will ensure that when the snapshots are deleted, orphaned blocks are recovered and made available for use by the system.

HPE Primera Virtual Copies can be mapped to other CPGs than their base volumes. This means that they can have different QoS characteristics. For example, the base volume space can be derived from a CPG on SSD disks and the virtual copy space from a CPG on Fibre Channel disks. The base volume space and the virtual copy space can grow independently without impacting each other. Dynamic optimization can tune the base volume space and the virtual copy space independently. Consider using Copy CPG with different characteristics for certain use cases such as development and testing.



HPE Primera Virtual Copy snapshot refresh considerations

For effective usages of the virtual copy snapshot refresh feature, the following best practices are recommended:

Tasks that must be completed on the host are:

- Snapshots need to be exported to hosts only during the initial deployment of the solution. Later, only refresh operations are performed to refresh the data of the snapshot from its parent volume.
- As part of the virtual copy snapshot refresh procedure, Hewlett Packard Enterprise strongly recommends that the host application be stopped from accessing data on the snapshot that is being refreshed. Un-export of the LUN from HPE Primera is not required.

As the file system on the snapshot might have changed due to the refresh, some operating systems require that the file system be remounted.

Performing concurrent virtual copy snapshot refreshes

Snapshot refresh can be initiated on an individual snapshot and on a group of snapshots. Snapshots must be added in a set to refresh them concurrently.

Although there is no hard limit to the maximum number of snapshots that can be refreshed at one time, Hewlett Packard Enterprise strongly recommends a maximum of 300 VVs in a set to minimize the I/O pause that ensures consistency across the set. Avoid scheduling multiple groups at the same time to avoid increasing the I/O pause. The HPE Primera array allows scheduling to the minute.

The `createsched` command can be used to spread out the number of simultaneous snapshots at a time. The following schedule would allow for 3000 (using the max 300 each set) volumes at 4 snapshots a day with a retention of 48 hours. The total number of volumes in this example would be 3000 base VVs + 9*3000 snapshot VVs = 30,000 volumes.

```
createsched "createsv -ro -exp 48h @count@.@@@H@@M@@S@ set:ss_set0000" "0 0,6,12,18 * * *" ss_job0000
createsched "createsv -ro -exp 48h @count@.@@@H@@M@@S@ set:ss_set0001" "1 0,6,12,18 * * *" ss_job0001
createsched "createsv -ro -exp 48h @count@.@@@H@@M@@S@ set:ss_set0002" "2 0,6,12,18 * * *" ss_job0002
createsched "createsv -ro -exp 48h @count@.@@@H@@M@@S@ set:ss_set0003" "3 0,6,12,18 * * *" ss_job0003
createsched "createsv -ro -exp 48h @count@.@@@H@@M@@S@ set:ss_set0004" "4 0,6,12,18 * * *" ss_job0004
createsched "createsv -ro -exp 48h @count@.@@@H@@M@@S@ set:ss_set0100" "0 1,7,13,19 * * *" ss_job0100
createsched "createsv -ro -exp 48h @count@.@@@H@@M@@S@ set:ss_set0200" "0 2,8,14,20 * * *" ss_job0200
createsched "createsv -ro -exp 48h @count@.@@@H@@M@@S@ set:ss_set0300" "0 3,9,15,21 * * *" ss_job0300
createsched "createsv -ro -exp 48h @count@.@@@H@@M@@S@ set:ss_set0400" "0 4,10,16,22 * * *" ss_job0400
createsched "createsv -ro -exp 48h @count@.@@@H@@M@@S@ set:ss_set0500" "0 5,11,17,23 * * *" ss_job0500
```

Plan to prevent I/O timeouts at the parent volume. The snapshot refresh quiesces the I/O on the parent volume and flushes all data to the drives during its operation. I/O timeouts are highly dependent on the intensity of an application workload and the storage system configuration. To avoid I/O timeouts at the host application level, limit the number of concurrent refresh operations to the maximum. Review and configure the appropriate timeouts on the host operating system and application in order to avoid I/O stalls during a snapshot refresh. Hewlett Packard Enterprise also recommends that the refresh process is performed during off-peak hours.

Snapshot CPG considerations

The HPE Primera array allows for data modified after a snapshot has been taken to be written to a different CPG (tier) than the base CPG. If a base volume is on an SSD tier and the snap is on a Fast Class or Nearline drive, a heavy write load to the SSD VV could impact the slower tiers. The same is true for Fast Class base VVs to a Nearline snap space tier. However, the difference in performance is least an order of magnitude greater with SSDs than with spinning disk media. Hewlett Packard Enterprise highly recommends keeping snapshots on the same SSD tier as the SSD base VVs.



HPE PRIMERA VIRTUAL COPY SNAPSHOT REFRESH RESTRICTIONS AND LIMITATIONS

The following restrictions and limitations should be considered when performing a virtual copy snapshot refresh:

- The snapshot refresh feature updates current data on a base volume to a previously created snapshot. Therefore, the snapshot refresh might not be applicable to solutions that involve snapshots scheduled at regular intervals for data recovery in the event of accidental corruption or loss of data.
- A snapshot refresh cannot be performed on a snapshot that has another snapshot. For example, in [FIGURE 21](#), **Snap1** cannot be refreshed as it has another snapshot, **Snap2**.
- Multiple snapshots from the same parent snapshot cannot be refreshed simultaneously.
- Retention time can be set for a snapshot to retain it until a specific period of time. The snapshot refresh operation cannot be performed on a snapshot that has a retention time set.
- In an HPE Primera Remote Copy configuration, snapshots are created internally for replication purposes. Those internal snapshots are not allowed for the snapshot refresh operation. User-created snapshots of HPE Primera Remote Copy volumes are allowed. This operation is also restricted when replication is in progress.
- Volumes that are being migrated by using HPE Peer Motion are not allowed to perform a snapshot refresh.

The HPE SSMC GUI interface does not provide options to perform a virtual copy snapshot refresh.

SUMMARY

As the use of flash arrays grows and data generation continues to accelerate, the need to control storage costs and manage the number of copies of production data also increases. In today's business environment, important functions such as backup, recovery, development, testing, analytics, support, and reporting, can all be enhanced by the ability to leverage point-in-time (PIT) snapshot data copies.

Two important trends are driving this new demand for snapshot technology. One is copy data management (CDM)—the reduction of storage consumption through eliminating unnecessary duplication of production data. The second is a DevOps culture that seeks to serve the needs of both development and operations with a single infrastructure, requiring the creation of production copies that do not impact production workloads.

HPE Primera Virtual Copy snapshot technology can efficiently meet both of these requirements by minimizing the amount of physical capacity that is actually being consumed by data without degrading performance or availability.

The HPE Primera Virtual Copy snapshot refresh and promote features provide efficient ways of managing the data copies required for a storage solution. For example, the HPE Primera `updatevv` command significantly reduces the management overhead at the host for discovering LUNs and reconfiguring applications. Several business organizations' operational needs, such as backup, recovery, development/testing, data analytics, and reporting, can leverage this feature to achieve their business objectives more efficiently.



GLOSSARY

Admin volume	The base volume that is used by the system to store administration data such as the system event log. The admin volume is created as part of the system installation and setup process.
Administrative space	The area of the volume that corresponds to logical disk regions that track changes to the volume since the previous snapshot was created.
Base volume	A TPVV that has been copied.
Child volume:	A VV (virtual or physical copy) created from a parent volume.
Chunklet	On HPE Primera, chunklets are 1 GB.
Copy data	Data that occupies the snapshot data space (HPE Primera Virtual Copy space) on a VV.
Copy space	The area of the volume that corresponds to logical disk regions that contain copies of user data that changed since the previous snapshot.
Logical disk	A collection of chunklets that reside on different physical disks and that are arranged as rows of RAID sets. When you create a CPG, the system creates groups of logical disks and assigns them to the CPG.
Physical copy	A point-in-time copy of an entire virtual volume.
Promote	Can be used for: <ul style="list-style-type: none"> • Physical copies, to break the association between a physical copy and a base volume by changing the physical copy into an independent base volume. • Virtual copies, to copy the changes from a virtual copy.
Resynchronize	To copy changes from one volume in a physical copy pair to the other volume, because the original volume was modified at some point after the physical copy operation took place.
Snapshot	A virtual or physical copy of a VV.
Snapshot administration space	The space on a VV that is used to track changes of the data from the time that a snapshot of a VV was created.
Snapshot refresh/Updatevv	Updates a snapshot with the current data on the base virtual volume.
Source volume	The VV from which a copy is made.
Table of contents (TOC)	The space on a physical disk that contains the internal description of the system. The TOCs on all physical disks in the system contain the same information.
Thinly provisioned virtual volume (TPVV)	A VV that maps to logical disk space associated with a CPG. Therefore, it can grow on demand.



RESOURCES AND ADDITIONAL LINKS

Planning and preparing for the HPE Primera 600

hpe.com/hpesc/public/docDisplay?docId=a00088908en_us

HPE Primera UI 1.0 User Guide

hpe.com/hpesc/public/docDisplay?docLocale=en_US&docId=emr_na-a00088935en_us

HPE Primera OS 4.2 Command Line Interface Reference Guide

hpe.com/hpesc/public/docDisplay?docLocale=en_US&docId=emr_na-a00088929en_us

HPE Primera OS 4.2 Command Line Interface Installation and Reference Guide

hpe.com/hpesc/public/docDisplay?docLocale=en_US&docId=emr_na-a00088937en_us

LEARN MORE AT

hpe.com/storage/primera

Make the right purchase decision.
Contact our presales specialists.



Chat



Email



Call



Get updates

© Copyright 2020 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft, SQL Server, Exchange, and Hyper-V are registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Oracle is a registered trademark of Oracle and/or its affiliates. VMware, VMware vSphere, and VMware vCenter Server are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All third-party marks are property of their respective owners.

a00097607ENW, July 2020