

# Dell EMC PowerMax and VMAX All Flash: TimeFinder SnapVX Local Replication

## Abstract

This document describes Dell EMC™ TimeFinder™ with PowerMaxOS 5978, HYPERMAX OS 5977, Solutions Enabler 9, Dell EMC Unisphere™ 9, and REST API 9.

November 2020

## Revisions

Date	Description
September 2019	Content and template update
September 2020	Updates for PowerMaxOS Q3 2020 release
October 2020	Moved session limits table from Section 4 to Section 1
November 2020	Minor command syntax correction

## Acknowledgments

Author: Michael Bresnahan

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 © 2019–2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [11/16/2020] [Technical White Paper] [H13697.15]

## Table of contents

Revisions.....	2
Acknowledgments .....	2
Table of contents .....	3
Executive summary.....	5
Audience .....	5
1 Introduction .....	6
1.1 TimeFinder session limits .....	6
1.2 Terminology .....	7
2 TimeFinder SnapVX overview .....	8
2.1 Creating snapshots.....	8
2.2 Generation numbers .....	8
2.3 Time-to-live value .....	9
2.4 Host writes to source .....	10
2.5 Redirect-on-write .....	10
2.6 Replication cache .....	11
2.7 Shared allocations .....	15
2.8 Linking target volumes to snapshots .....	16
2.9 Relinking .....	17
2.10 Unlinking .....	17
2.11 Defining state .....	17
2.12 Linking a larger target.....	18
2.13 Source-volume or linked-target SRP relationships.....	18
2.14 Restore operations .....	19
2.15 Cascading snapshots .....	19
2.16 Reserved capacity .....	21
2.17 Interoperability .....	21
2.18 Storage group operations .....	24
2.19 Secure snaps.....	25
2.20 Online device expansion .....	26
2.21 Role Based Access Controls .....	28
3 Implementing TimeFinder SnapVX .....	29

3.1	Creating snapshots.....	29
3.2	Linking, relinking, and unlinking.....	31
3.3	Restoring from snapshots.....	34
4	PowerMaxOS Q3 2020 release enhancements.....	35
4.1	Snapshot policies.....	35
4.2	Cloud Mobility for Dell EMC PowerMax .....	37
4.3	SnapSet and SnapSet ID .....	38
4.4	Enhanced snapshot termination options .....	38
4.5	Snapshot SRP utilization alert .....	41
5	TimeFinder legacy modes: Mirror, Clone, and VP Snap.....	42
5.1	Interoperability between native SnapVX and legacy TimeFinder modes.....	42
5.2	Snapshots created by zDP .....	42
5.3	TimeFinder VP Snap .....	43
5.4	TimeFinder/Snap deprecation .....	43
6	Conclusion.....	44
A	TimeFinder SnapVX state table .....	45
B	Geometry compatibility mode.....	46
B.1	Open Replicator for Symmetrix (ORS) .....	46
B.2	TimeFinder SnapVX .....	47
B.3	TimeFinder Clone and VP Snap.....	47
B.4	TimeFinder/Mirror .....	47
C	Monitoring the copy process .....	48
D	Nocopy Linked Target behavior before HYPERMAX OS 5977.810.184 .....	49
E	Scheduling snapshots with Unisphere for PowerMax.....	50
E.1	Scheduling snapshots with Unisphere for PowerMax 9.1 and earlier versions .....	50
F	System resource planning for capacity and performance.....	56
F.1	Capacity planning and monitoring .....	56
F.2	Snapshot capacity monitoring: .....	57
F.2.1	Logical snapshot capacity .....	57
F.2.2	Physical snapshot capacity .....	60
F.3	Performance considerations.....	62
F.4	Dell EMC ProtectPoint.....	63
G	Technical support and resources .....	66
G.1	Related resources.....	66

## Executive summary

Dell EMC™ TimeFinder™ software delivers point-in-time copies of volumes that can be used for backups, decision support, data warehouse refreshes, recovery from logical corruption, or any other process that requires parallel access to production data.

Previous TimeFinder offerings each had their own characteristics and ideal use cases. These offerings also have several similarities, the main one being that each requires a target volume to retain snapshot or clone data.

TimeFinder in HYPERMAX OS 5977 introduced TimeFinder SnapVX which combines the best aspects of the previous TimeFinder offerings, adds some ease-of-use features, and increases scalability.

SnapVX provides very low impact snapshots and clones. SnapVX supports up to 1024 snapshots per source volume, which are tracked as versions with less overhead and simple relationship tracking. Users can assign names to identify their snapshots and can set automatic expiration dates on each snapshot.

SnapVX provides the ability to manage consistent point-in-time copies for Storage Groups (SG) with a single operation. Up to 1024 target volumes can be linked per source volume, providing read/write access as pointers or full-copy clones.

Solutions Enabler also provides compatibility modes for users who rely on their TimeFinder/Mirror, Clone, or VP Snap command scripts. This allows users to leverage their existing scripts while learning how to take advantage of the features of SnapVX.

This document describes TimeFinder features for business continuity and implementation guidelines, including restrictions and limitations for this product. The features discussed are valid for PowerMaxOS 5978 and HYPERMAX OS 5977.

## Audience

This document is intended for storage administrators, database administrators, and technologists who have an interest in understanding the concepts surrounding Local Replication in PowerMax, [VMAX All Flash](#), and VMAX3™ family storage arrays.

# 1 Introduction

Local replication with SnapVX starts out as efficiently as possible by creating a **snapshot**, a pointer-based structure that preserves a point-in-time view of a source volume. Snapshots do not require target volumes, share back-end allocations with the source volume and other snapshots of the source volume, and only consume additional space when the source volume is changed. A single source volume can have up to 1024 snapshots.

Each snapshot has a user-defined name and can optionally have an expiration date, both of which can be modified at any time. Management interfaces enable the user to take a snapshot of an entire SG with a single command.

A point-in-time snapshot can be accessed by *linking* it to a host accessible volume referred to as a target. The target volumes are standard Thin LUNs. Up to 1024 target volumes can be linked to the snapshots of a single source volume. This limit can be achieved either by linking all 1024 target volumes to the same snapshot from the source volume, or by linking multiple target volumes to multiple snapshots from the same source volume. However, a target volume may only be linked to a single snapshot at a time.

By default, targets are linked in a nocopy mode, but can be linked in a copy mode to create full-copy clones. Snapshots can be cascaded from linked targets, and targets can be linked to snapshots of linked targets. There is no limit to the levels of cascading and the cascade can be broken at any point, with limitations that are discussed later.

## 1.1 TimeFinder session limits

Table 1 lists the per-source-volume session limit of the various TimeFinder modes. Each SnapVX source volume can also have up to 1024 SnapVX linked targets.

Table 1 TimeFinder session limits

TimeFinder mode	Maximum sessions per source volume
Total sessions, combination of all modes	1024
Snapshot policies SnapVX snapshots	1024
Data Protector for z Systems (zDP) SnapVX snapshots	1024
Solutions Enabler SnapVX snapshots	256
Unisphere for PowerMax SnapVX snapshots	256
Dell EMC Mainframe Enablers SnapVX snapshots	256
AppSync™ SnapVX snapshots	256
PowerProtect Storage Direct SnapVX snapshots	256
RecoverPoint SnapVX snapshots	256
TimeFinder/VP Snap	32

TimeFinder/Clone	8
TimeFinder/Mirror (BCV)	8

Each SnapVX source volume can also have up to 1024 SnapVX linked targets. There is no hard limit on the number of sessions in a single array. The total sessions within an array depend on the configured and available resources. Include snapshot use in the overall sizing and planning of new systems to ensure that each array is configured with required resources. After deployment, monitor system resources as usage grows. These best practice recommendations are applicable to all features and applications on a storage array and are not specific to snapshots.

## 1.2 Terminology

The following are explanations of terms that are commonly used throughout the paper:

**Storage Resource Pool (SRP):** A collection of data pools which provide physical storage for thin devices. SRPs are managed by Fully Automated Storage Tiering (FAST).

**Source Volume:** A LUN that has SnapVX snapshots from it, either with or without linked targets.

**Snapshot:** A preserved point-in-time image of a source volume. A snapshot uses pointers in cache to indicate which version of a track is applicable to the specific point-in-time; either a track that resides on the source volume, or a snapshot delta.

**Snapshot Delta:** A point-in-time version of a source volume track that was preserved during a host write to a source volume that had an active snapshot.

**Linked Target Volume:** A LUN that is linked to a SnapVX snapshot to make the point-in-time of the snapshot accessible to a host. Linked targets can have one of two modes:

- **Nocopy Mode:** Does not copy data to the linked target volume but still makes the point-in-time accessible using pointers to the snapshot. Beginning with HYPERMAX OS 5977.810.184, the point-in-time image is still available after the target is unlinked<sup>1</sup>.
- **Copy Mode:** Copies all relevant tracks from the snapshot's point-in-time to the linked target volume to create a complete copy of the point-in-time that will remain available after the target is unlinked.

See appendix D for changes to the nocopy target functionality introduced in HYPERMAX OS 5977.810.184.

**Unisphere:** Unisphere is an HTML5 web-based application that allows you to configure and manage PowerMax, VMAX™ All Flash, and VMAX storage systems. The term Unisphere incorporates "Unisphere for PowerMax" for the management of PowerMax and All Flash storage systems running PowerMaxOS 5978, and "Unisphere for VMAX" for the management of VMAX All Flash and VMAX storage systems running HYPERMAX OS 5977 and Enginuity OS 5876.

## 2 TimeFinder SnapVX overview

### 2.1 Creating snapshots

TimeFinder SnapVX allows snapshots to be taken without the use of a target volume. The snapshots are made as efficient as possible by sharing point-in-time tracks, which are called snapshot deltas. The snapshot deltas are stored directly in the Storage Resource Pool (SRP) where the source volume resides. Each snapshot references the snapshot deltas appropriate for the specific point-in-time image.

SnapVX Snapshots are created with the `symsnapvx establish` command. As part of the establish command the user provides a name for the snapshot. The name has up to 32 characters and can contain letters, digits, dashes, and underscores. The name is not case-sensitive.

Figure 1 shows a volume with three targetless snapshots.




Figure 1 Production volume with three targetless snapshots

### 2.2 Generation numbers

It is possible to have multiple snapshots with the same name for the same set of source volumes, and generation numbers identify the different snapshots. It is important to understand how generation numbers are applied when deciding to either reuse snapshot names or to use unique names for each snapshot. Generation numbers are relative to existing snapshots at the time of creating or displaying snapshots, and they change as snapshots are created and terminated.

When a snapshot name is reused, the new session becomes Generation 0 and all previous sessions with the same name have their generation numbers incremented by one. Similarly, if a more recent snapshot is



terminated, all older snapshots have their generation numbers decreased by one. The most recent snapshot always has a generation number of 0. This is advantageous because after creating a new snapshot, the user knows that it is Generation 0 and does not have to search for the specific snapshot.

For example, if a user creates a snapshot called `hourly_backup` at 1PM, the snapshot will be Generation 0. Then at 2PM they create another snapshot and reuse the name `hourly_backup`. At this point, the 1PM snapshot will become Generation 1 and the 2PM snapshot will be Generation 0. Similarly, at this point if the 2PM snapshot (Generation 0) was terminated, the 1PM snapshot will again become Generation 0.

Reusing snapshot names and employing generation numbers can be very useful. A user could maintain a set number of generations for a specific snapshot. Each day the user creates a new snapshot and terminates the oldest, using the same commands each day. This makes scripting simple.

However, the user must recognize that the generation numbers are dynamic; creating or terminating snapshots may cause generation numbers of other snapshots to change. This could be problematic in situations where multiple users create snapshots on the same set of volumes, especially in situations where a single volume is in multiple groups (Storage Groups, Device Groups, Composite Groups, and so on). As with any other feature or option, the user must make the best choice for their environment.

## 2.3 Time-to-live value

The user can also define a Time-to-live (TTL) value on each snapshot, which acts as an expiration date. Snapshots automatically terminate when the TTL expires, assuming the snapshot does not have any linked targets or is in a restore session. If a snapshot has linked targets and the time-to-live has expired, the snapshot terminates when the last target is unlinked or the restore session is terminated.

TTL values can either be set to a specific date (`-absolute`) or a specified number of days (`-delta`). When a user sets a TTL value to a specific date, Solutions Enabler internally uses a delta that is calculated from the specified date and the current host time. This internal conversion eliminates the need for users to synchronize clocks or calculate times between servers and the storage array. Time-to-Live values can be specified when the snapshot is established and can also be set or changed on existing snapshots.

Expiration can be specified by days or hours beginning with Solutions Enabler 8.4, Unisphere 8.4, and REST API 8.4. Expiration is available in one day increments in earlier releases.

Hours can be specified with both the `-absolute` and the `-delta` options. The maximum value allowed is 23 hours. Specifying hours is optional when setting the expiration date and can be omitted. Scripts that were created with previous versions of Solutions Enabler and do not specify hours continue to function after upgrading to Solutions Enabler 8.4.

- Examples of the CLI format of the `-ttl` option are:
- Expire snapshot in 1 day: `-ttl -delta 1`
- Expire snapshot in 1 day and 12 hours: `-ttl -delta 1:12`
- Expire snapshot in 12 hours: `-ttl -delta 0:12`
- Expire snapshot on March 1, 2018 at 1pm: `-ttl -absolute 03/1/2018:13`

Figure 2 is an example of the Unisphere wizard:

Figure 2 Unisphere 9.0 Create Snapshot Wizard

## 2.4 Host writes to source

Host writes to source volumes create snapshot deltas in the SRP. Snapshot deltas are the original point-in-time version of each track that has been updated since a snapshot was established. Multiple snapshots can share a single snapshot delta if all are active at the time a source track was updated. So, if a source track that is being updated needs to be preserved by multiple snapshots, only one snapshot delta is needed. This sharing allows the snapshots to be as space efficient as possible.

The output of command `symsnapvx list -detail` and other commands report the number of deltas for each snapshot. The `symsnapvx list -detail` output also has a “Non-Shared Tracks” field that provides the number of tracks uniquely allocated to each snapshot. This information is valuable because it can be used to determine which snapshots are using the most non-shared space and are most beneficial to terminate in the event of needing to free some space in the SRP. Only non-shared space is freed from the SRP when a snapshot is terminated.

## 2.5 Redirect-on-write

SnapVX introduces redirect-on-write (ROW) technology to TimeFinder. When a source track is written to and the original data needs to be preserved for a snapshot or snapshots, the new write is accepted and

asynchronously written to a new location in the Storage Resource Pool (SRP). The source volume now points to the new data while the snapshot or snapshots continue to point to the original data (the snapshot delta) in its original location.

Redirect-On-Write is illustrated in the following figures. Figure 3 shows the source volume and snapshot both pointing to the same location in the pool before the track is updated.




Figure 3 Source track before update

Figure 4 shows the source pointing to the new location in the pool after the host write, while the snapshot continues to point to the original location.




Figure 4 Source track after update and snapshot delta

Even though ROW provides many benefits, there are situations where the traditional Asynchronous Copy on First Write (ACOFW) mechanism that has been used by TimeFinder for many years is the ideal option. For example, to prevent source data from being redirected to a less than ideal performance or compression tier.

Therefore, both ROW and ACOFW are available in PowerMaxOS and HYPERMAX OS. The exact method that is selected for an operation is completely transparent to the user. As new data is written, the array will choose the appropriate mechanism to ensure source data resides in a location for ideal performance while also keeping snapshot data as efficient as possible on backend capacity.

## 2.6 Replication cache

A portion of the metadata in cache is dedicated for Replication Data Pointers (RDP) which keep track of the snapshot deltas in the SRP. This portion of metadata is called Replication Cache. If Replication Cache is exhausted, snapshots may begin to fail.

Replication Cache is used by SnapVX snapshots and VP Snap sessions. Replication Cache usage increases as SnapVX and VP Snap source devices are written to, as there is more point-in-time data to manage. SnapVX Linked Targets, Clone, and Mirror sessions do not influence Replication Cache usage.

It is important to understand that Replication Cache resources are not immediately released upon snapshot termination. The process may take some time as it is designed as a background process so to not take away processing resources from other higher priority operations within the array. Therefore, current Replication

Cache usage should be checked before creating new snapshots immediately after terminating existing snapshots.

All systems follow the same algorithm to determine the percentage of metadata to dedicate to Replication Cache. Systems that have very high snapshot usage may be candidates to have the Replication Cache portion of metadata increased. However, this is not a typical use case. In most systems, the standard algorithm provides enough Replication Cache for the environment. Contact your local Dell EMC Service Representative if you feel your environment requires increased Replication Cache.

The PowerMaxOS 5978 Q3 2019 SR contains enhancements that greatly improve replication cache efficiency.

## 2.6.1 Monitoring replication cache usage

Solutions Enabler 8.2 and Unisphere 8.2 introduced tools to monitor Replication Cache Usage. These tools are only available when used with systems running the HYPERMAX OS 5977.810.184 or later. The information described below is also available in REST API beginning with version 8.3.

The `symcfg list -v` output has a field that shows the current Replication Cache usage, as shown in the following:

```
Symmetrix ID: 00019680XYZ (Local)
Time Zone   : Eastern Standard Time
Product Model      : VMAX100K
Symmetrix ID       : 00019680XYZ
Microcode Version (Number) : 5977 (17590000)
-----< TRUNCATED >-----
Max # of DA Write Pending Slots : N/A
Max # of Device Write Pending Slots : 85622
Replication Usage (Percent)      : 12
```

Solutions Enabler provides an alert ID 1222 to report when Replication Cache usage has exceeded the specified thresholds. The default threshold values are the same as the default values for other Solutions Enabler thresholds alerts. The user also can specify other values. However, the alerts are not enabled by default, and need to be enabled by the user. The default threshold values are as follows:

- FATAL severity = 100%
- CRITICAL severity >= 80%
- MAJOR severity >= 70%
- MINOR severity >= 65%
- WARNING severity >= 60%

Unisphere 9.0 reports Replication Cache usage on Efficiency tab of the Performance Dashboard:



Figure 5 Replication Cache Usage

Alerts are also available in Unisphere. The System Alerts have the following threshold values and are enabled by default:

- FATAL severity = 100%
- CRITICAL severity  $\geq 80\%$
- Warning severity  $\geq 60\%$

Create						
Name	Warning	Critical	Fatal	Custom	Enabled	Notifications
Storage Resource Pool Utilization	60%	80%	100%		<input checked="" type="checkbox"/>	NONE
Frontend Meta Data Usage	60%	80%	100%		<input checked="" type="checkbox"/>	NONE
Local Replication Utilization	60%	80%	100%		<input checked="" type="checkbox"/>	NONE
Backend Meta Data Usage	60%	80%	100%		<input checked="" type="checkbox"/>	NONE
System Meta Data Utilization	60%	80%	100%		<input checked="" type="checkbox"/>	NONE

Figure 6 Unisphere 9.0 Threshold Alerts – Standard Alerts

Users can choose to create custom alerts with their own threshold values:

### Create Symmetrix Threshold Alert

Category  
Local Replication Utilization ▼

Warning

60 % default

Critical

80 % default

Fatal

100 % default

Figure 7 Unisphere 9.0 Custom Threshold Alerts

As stated earlier, snapshots may begin to fail if Replication Cache is exhausted. Solutions Enabler and Unisphere snapshot outputs include a flag to indicate if a snapshot is in a failed state.

Solutions Enabler 9.0 and Unisphere for PowerMax introduced an additional flag to indicate the reason for the failure, which includes RDP resources. The following shows an example of the `symsnapvx` structure to list only failed snapshots.

The following is an example of failed snapshots viewed from Unisphere for PowerMax 9.0. Note that the snapshots in the Unisphere for PowerMax example failed due to SRP resources being exhausted, which typically indicates the used capacity of the SRP reached the Reserved Capacity threshold.

```
C:\>symsnapvx -sid XYZ list -failed
Symmetrix ID      : 000197800XYZ (Microcode Version: 5978)
-----
Sym              Num  Flags
Dev  Snapshot Name      Gens FLRG TSEB Last Snapshot Timestamp
-----
0003B Snap1          1 R...   Wed Dec 20 08:30:34 2017
0003C Snap1          1 R...   Wed Dec 20 08:30:34 2017
0003D Snap1          1 R...   Wed Dec 20 08:30:34 2017
0003E Snap1          1 R...   Wed Dec 20 08:30:34 2017
Flags:
(F)ailed : X = General Failure, . = No Failure
          : S = SRP Failure, R = RDP Failure, M = Mixed Failure
(L)ink    : X = Link Exists, . = No Link Exists
(R)estore : X = Restore Active, . = No Restore Active
(G)CM     : X = GCM, . = Non-GCM
(T)ype    : Z = zDP snapshot, . = normal snapshot
(S)ecured : X = Secured, . = Not Secured
(E)xpanded: X = Source Device Expanded, . = Source Device Not Expanded
(B)ackground: X = Background define in progress, . = No Background define
```

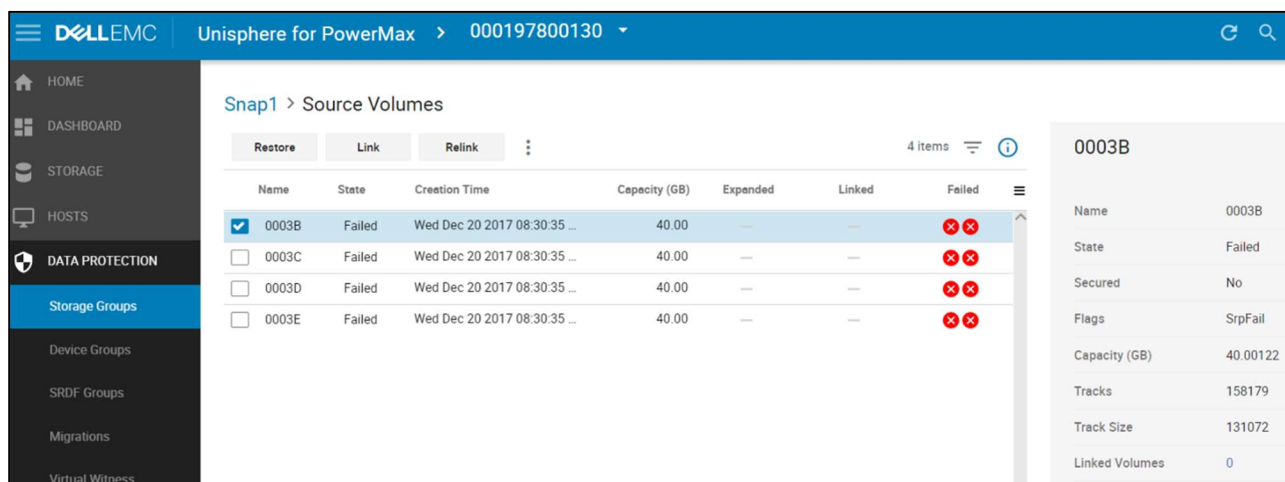


Figure 8 Failed snapshots viewed from Unisphere 9.0

## 2.7 Shared allocations

VP Snap in Enginuity 5876 introduced back-end allocation sharing between VP Snap target volumes; multiple snapshots from a source volume can share point-in-time versions of a track. A back-end allocation could be shared by up to 16 VP Snaps. Allowing multiple targets to reference the same shared copy provides cost-effective space savings. Shared track values on thin pools could be seen with Solutions Enabler and Unisphere.

SnapVX in HYPERMAX OS 5977 greatly enhanced this shared allocation capability. A snapshot delta can be shared between snapshots and shared with target volumes, and a source allocation can be shared with target volumes. And there is no limit to the number of volumes that can share an allocation.

Sharing of a snapshot delta between snapshots is like that of VP Snap in Enginuity 5876; when a source write arrives and the original track needs to be preserved as a snapshot delta for multiple snapshots, only a single snapshot delta needs creating.

However, nocopy-linked targets begin to share tracks with the source volume or share snapshot deltas for the specific point-in-time as part of the define process, which is discussed on page 12. The targets are added to sharing at the time of link, independent of source track updates.

If a source track that is being shared with a target (or multiple targets) is updated, the existing track is preserved as a snapshot delta and continues to be shared with the target or targets.

Writes to a target are applied only to the specific target. If the write is to a track that is being shared then the track for the specific target is split from the shared group, and any other snapshots or targets continue to share the original shared track.

Reporting of shared track counts on pools or SRPs by Solutions Enabler or Unisphere is not applicable in HYPERMAX OS 5977. However, some Solutions Enabler displays report shared track counts on pools for backwards compatibility with older arrays.

Solutions Enabler instances and Unisphere report the non-shared track count on each snapshot. The non-shared count is the most important value because non-shared snapshot deltas are discarded when the snapshot is terminated, which frees space in the SRP.

---

**Note:** See appendix D for changes to the nocopy target functionality introduced in HYPERMAX OS 5977.810.184.

---



## 2.8 Linking target volumes to snapshots

A target volume can be linked to a snapshot when a user wants to present a point-in-time copy to a host. Multiple targets can be linked to a single snapshot. Writing to a linked target does not affect the point-in-time of the snapshot; should the data on the linked target be corrupted by the user activity the target can be relinked to the original snapshot. Figure 9 depicts a source volume with three snapshots, and a target linked to the third snapshot.

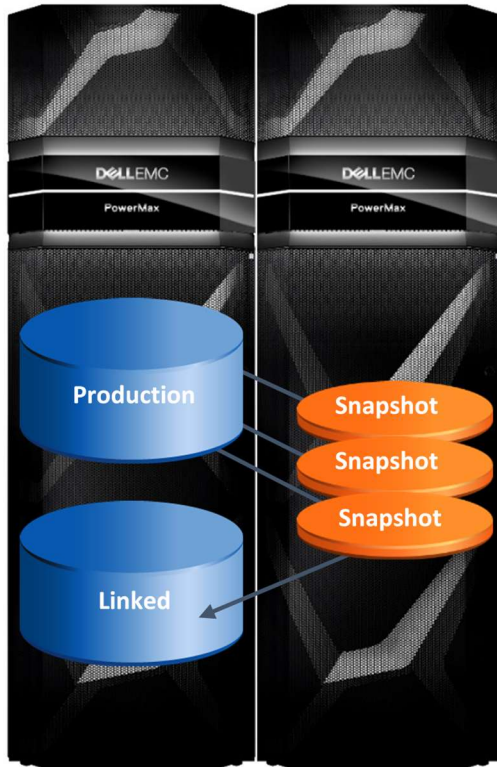


Figure 9 Production volume with three targetless snapshots and one Linked Target

Targets can be linked in either copy mode which duplicates data to create separate, full-volume copies, or nocopy mode which creates space-saving copies that only allocate space during to host writes to the target.

Host writes to a source volume that has a snapshot with a target linked in nocopy mode only preserve point-in-time tracks as snapshot deltas, and do not cause any tracks to be copied to the linked target. The linked targets access the preserved point-in-time by sharing the snapshot delta allocation.

Snapshots that have linked targets cannot be terminated. The user needs to unlink the target first and then terminate the snapshot.

Linked targets can be SRDF R1 volumes with the following requirements:

- Beginning with the PowerMaxOS Q3 2020 release, an SRDF R1 can be linked in nocopy mode.
- Prior code releases require an SRDF R1 to be linked in full copy mode.



## 2.9 Relinking

The `relink` command provides the capability to perform incremental refreshes of linked targets. The `relink` command can also be used to link a target to a different snapshot of the same source volume. The `relink` command can be issued to both copy and nocopy mode linked targets. This functionality provides an easy way for users to check different points-in-time if they are unsure which one is the best for them to access.

---

**Note:** Just like the `link` command, `relink` is in nocopy mode by default unless the user specifies `-copy`, regardless of the mode of the original link.

---

## 2.10 Unlinking

The `unlink` command breaks the relationship between a snapshot and a linked target. Copy mode linked targets can be unlinked once copying is complete and retain a full, usable point-in-time copy of the source volume.

When a nocopy linked target is unlinked, the operation does not deallocate any data on the target. This includes data that was copied during the link operation as well as data that was directly written to the target. Therefore, an unlink operation does not return any free capacity to the SRP.

If the user wants to deallocate the data on the target or targets after unlinking, they can do so by issuing the `symdev free -all` command to the target volume.

---

**Note:** This is a powerful command that completely deallocates all data on the volumes that it is issued against. Take extreme care when using this command. As a safety mechanism, the `free -all` command requires the device to be not ready or unmapped.

---

A target volume that has cascaded snapshots can only be unlinked once the target has completely copied. In the case of a nocopy linked target that has a cascaded snapshot or snapshots, the user needs to `set mode` to copy, then unlink after copy has completed. Cascaded snapshots are discussed further on page 14.

---

**Note:** Before HYPERMAX OS 5977.810.184, a nocopy linked target that has a cascaded snapshot or snapshots cannot be unlinked. The user needs to `set mode` to copy, then unlink after copy has completed. This is no longer required as of the HYPERMAX OS 5977.810.184.

---

## 2.11 Defining state

SnapVX introduces the concept of a “defined” state for linked target tracks. Defining may aid performance of host access to target tracks that have not yet been copied to target by presenting data directly from the SRP and eliminating the need for a redirect to source or snapshot.

The point-in-time of a snapshot is immediately available through the target volume when it is linked. However, its track tables still point to its old track locations and all tracks are considered “undefined.” Shortly after linking, a background define process changes the pointers of each track to reflect the location of the appropriate track version for the specific point-in-time. The defining process is what creates shared allocations between the target and tracks on source or snapshot deltas.

Figure 10 illustrates the defining process. The target is linked to snapshot 2. The first two tracks on the linked target volume have been defined and point directly to the data in the SRP. The last two tracks on the linked target volume have not yet been defined and are still pointing to the snapshot.

**Note:** An undefined track could point to a track on source rather than snapshot if that is the appropriate version of the track for the specific point-in-time represented by that snapshot.

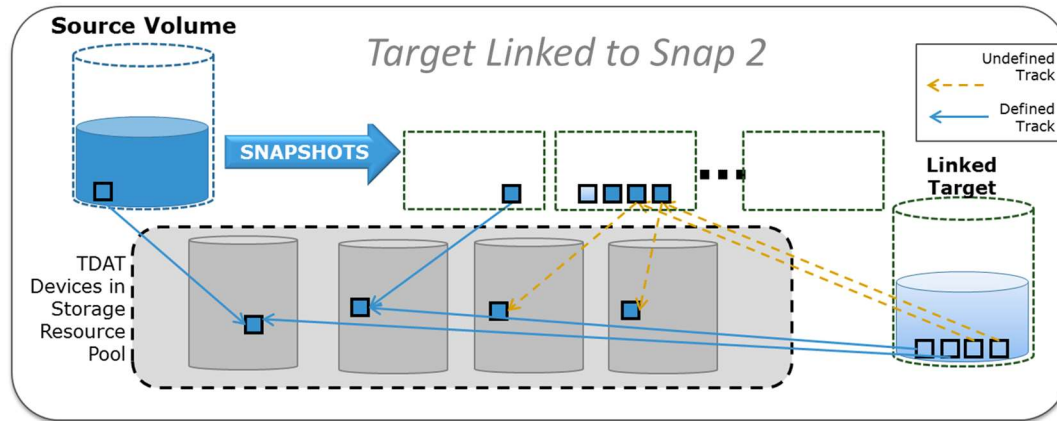


Figure 10 Defined and undefined track states

Defining applies to both copy mode and nocopy mode linked targets. In the case of copy mode linked targets, the defining process and copy process run concurrently on the linked target but each track will be defined before copying.

The point-in-time of the snapshot is immediately available when the target is linked. The user does not need to wait for the target volume to be defined before accessing the point-in-time. A write to an undefined target track will invoke the defining process for the track. A read to an undefined target track will be redirected to the appropriate point-in-time version of the track using either the snapshot or the source volume.

The user can check for the defined state of a linked target with the following command:

```
symsnapvx -sid verify -sg -snapshot_name -linked -defined
```

Beginning with PowerMaxOS 5978 and Solutions Enabler 9.0 and Unisphere 9.0, users can take a new snapshot of a source volume that has another snapshot restore in progress.

## 2.12 Linking a larger target

A target volume that is larger than the source volume can be linked to the snapshot in either copy or nocopy mode. This functionality is enabled by default. There is an environment variable `SYMCLI_SNAPVX_LARGER_TGT` that can be set to prevent larger targets from being used.

## 2.13 Source-volume or linked-target SRP relationships

Source volumes and linked targets can be associated either to the same or separate SRPs. Single SRP systems are recommended because they are easier to manage, and performance may be better when spreading workloads across all available resources in the array. Also, the Reserved Capacity setting in HYPERMAX OS (described on Page 15) allows the user to limit the amount of storage in the SRP that the

replicas can consume, eliminating the need to physically separate production volumes and their replicas (snaps and clones).

---

**Note:** SRP count and configuration need to be specified during the system ordering process.

---

Snapshot deltas will always be stored in the source volume's SRP. There is no advantage to having nocopy linked targets in a separate SRP because writes to a source volume will only create snapshot deltas and will not initiate a copy to target.

If opting for a separate SRP for copy mode linked targets, the target SRP must have enough disks to support the expected capacity and performance requirements of the target volumes.

A single snapshot can have both copy and nocopy linked targets. All linked targets from a single snapshot do not need to be in the same SRP. Source volumes and linked targets (copy and nocopy modes) can be moved between SRPs while TimeFinder sessions are active; FAST will move allocations accordingly.

## 2.14 Restore operations

Snapshots can be restored directly to the source volume or volumes using the `symsnapvx restore` command or using Unisphere. The `symsnapvx terminate -restored` command terminates the restored session only and does not break the source-to-snapshot relationship. Restoring from a snapshot to source does not affect the point-in-time of any other snapshots or linked targets.

Restoring from a linked target back to a source is functionally possible although it is not done with a restore command. Instead, the user would take a snapshot of a linked target, which essentially creates a cascaded scenario, discussed further in the following section. The user could then use the original source volume as a linked target to the cascaded snapshot.

## 2.15 Cascading snapshots

Cascading refers to taking a snapshot of a linked target and linking a target to a snapshot of a linked target. All tracks on a target volume must be defined before a snapshot can be taken of the target volume.

Figure 11 illustrates one example of cascaded SnapVX sessions.

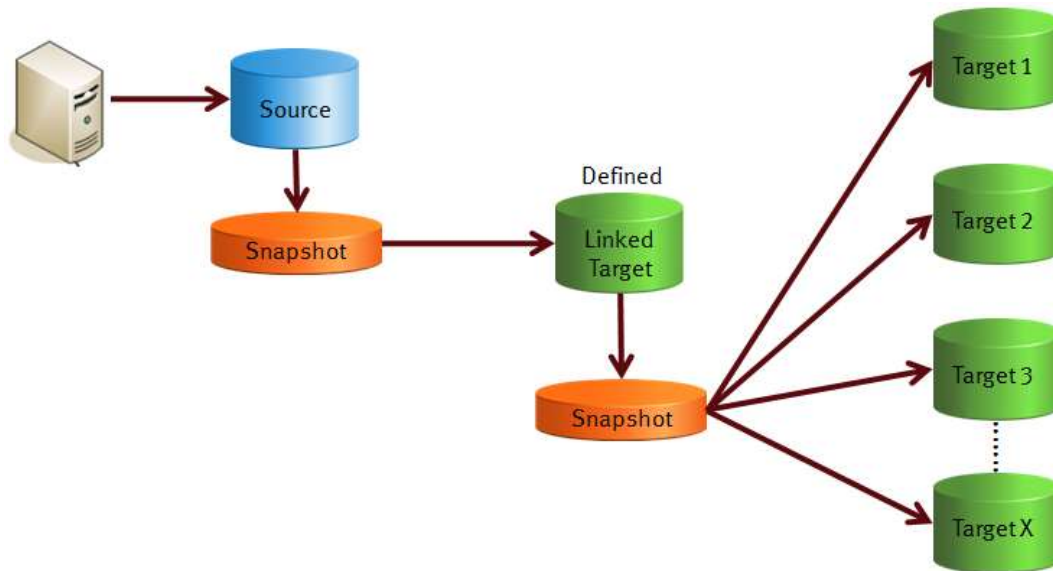


Figure 11 Cascading snapshots and linked targets

The cascading capabilities of SnapVX are very robust. However, several traditional use cases for cascading clone or snap sessions no longer exist with SnapVX.

Since point-in-time data for SnapVX snapshots is stored in the SRP and cannot be directly accessed by the host, the point-in-time of the snapshot is always pristine, and there is no need to make a “Gold Copy” clone to cascade from. Even if a target is linked to the snapshot and written to, the point-in-time of the snapshot does not change. This reduces the number of full volume copies that are required in most use cases making SnapVX very space efficient.

Furthermore, the defining mechanism allows reads of tracks that do not reside on the target volume to be serviced directly from the SRP and not redirected to the source volume as would be the case with traditional snaps or clones. This not only improves performance of the read but also reduces the possibility of target reads affecting performance of the source volume.

Cascaded SnapVX sessions are supported with the following considerations:

- There is no limit on the number of cascaded hops that are allowed.
- Beginning with PowerMaxOS 5978, Solutions Enabler 9.0 and Unisphere 9.0 users can take a new snapshot of a linked target that has the define process still in progress. Earlier code versions only allowed a snapshot of a linked target that is in either the `defined` or `copied` state.
- A cascaded snapshot can only restore to a linked target that is in copy mode and has fully copied.
- The cascaded chain can be broken at any point by unlinking the target from the snapshot with the following rules:
  - A snapshot that has a linked target cannot be terminated.
  - Prior to HYPERMAX OS 5977.810.184 a linked target that has a cascaded snapshot must be in copy mode and fully copied before it can be unlinked. This is not required with HYPERMAX OS 5977.810.184 or later.

## 2.16 Reserved capacity

Reserved Capacity is a percentage of the SRP that can only be consumed by new host writes. Existing and new snapshots are affected if the Free Capacity of the SRP falls below the set Reserved Capacity. The following commands view and change the values:

- `symcfg list -srp -v`
- `symconfigure -sid <sid> -cmd "set srp <srp_name> resv_cap=<value>," commit`

Valid values for Reserved Capacity percentage are from 1% to 80%, and NONE.

For example, if the reserved capacity on an SRP is set to 10% and the SRP becomes 90% full of volume allocations, snapshot deltas, and SRDF Delta Set Extension allocations, only new host writes can consume the last free 10% of the SRP.

When only Reserved Capacity is available in an SRP, a snapshot fails at the next attempt to create a new snapshot delta, and the snapshot needs to be terminated. VP Snap and Clone Nocopy sessions also fail at the next write that requires a source track to be copied to target.

Copy to targets stops and cannot complete. This includes SnapVX linked targets, Clone copy mode sessions, and TimeFinder/Mirror BCVs. Copy resumes if free space is made available in the SRP, or if the Reserved Capacity value is lowered enough.

The Reserved Capacity affects the TimeFinder copy process only. Host reads and writes to the targets are still possible.

VAAI XCOPY and ODX copy operations also obey the Reserved Capacity of an SRP. In these cases, the operation can continue by falling back to host-level copy which appears as host writes to the array rather than copy operations.

## 2.17 Interoperability

TimeFinder SnapVX is interoperable with other HYPERMAX OS features under the following conditions:

### 2.17.1 SRDF

- Snapshots can be taken from R1 and R2 volumes in all SRDF modes (SRDF/S, SRDF/A, and SRDF/Metro).
- Write-pacing is not required when taking snapshots from SRDF/A R2 volumes.
- Snapshots cannot be taken of Storage Groups that contain a mix of SRDF and non-SRDF devices.
- Linked targets can be SRDF source volumes (SRDF R1) with the following requirements:
  - Beginning with the PowerMaxOS Q3 2020 release, an SRDF R1 can be linked in nocopy mode.
  - Prior code releases require an SRDF R1 to be linked in full copy mode.

### 2.17.2 Service levels on PowerMaxOS 5978

- Service Level settings on source or target devices will have no effect on TimeFinder operations.
- Likewise, TimeFinder operations will have no effect on the Service Level of source or target devices.

### 2.17.3 FAST on HYPERMAX OS 5977

- Allocations owned by a source volume will be managed by the Service Level Objective of the source volumes SG.
- Allocations owned by a target volume will be managed by the Service Level Objective of the target volumes SG.
- Snapshot deltas will be managed by the Optimized Service Level Objective.
- Read-miss I/O to target tracks that are owned by the source volume will increment the FAST metrics for the source volume, not the target volume, and may contribute to the extent being promoted.
- Source and target volumes can be moved across SRPs while TimeFinder sessions are active.

### 2.17.4 Zero space reclaim

Zero space reclaim is supported on SnapVX source and target volumes.

### 2.17.5 Persistent allocations

Persistent Allocations are tracks that are unaffected by a standard reclaim operation. The user can mark and unmark tracks as persistent with the Solutions Enabler `symdev` or `symmsg` commands, and with Unisphere.

- The considerations for Persistent Allocations with SnapVX are as follows:
- Snapshots can be taken of a volume that has Persistent Allocations. However, restores from a snapshot are blocked if the source volume has any Persistent Allocations. The user needs to remove the Persistent Allocations from the entire volume to allow the restore operation. Persistent Allocations can be reset on the source volume once the restore session is terminated.
- Volumes with persistent allocations cannot be used as linked targets. Likewise, persistent allocations cannot be set on linked targets.

### 2.17.6 Non-Disruptive Migration

Non-Disruptive Migration (NDM) was introduced in the HYPERMAX OS Q3 2016 Service Release. For more in-depth information, see the *VMAX Non-Disruptive Migration Configuration and Best Practices Technical Notes*, available on [dell EMC.com](http://dell EMC.com).

- Existing TimeFinder sessions from an NDM device can remain in place during the migration.
- New replication from an NDM device can be configured during a specific point in the migration process.
- An NDM device, source or target, cannot be the target of a TimeFinder session.
- The TimeFinder session cannot restore data back to the NDM device at any time during the migration session.

### 2.17.7 Data reduction on PowerMax arrays

The Data Reduction feature in PowerMaxOS 5978 on PowerMax arrays consists of deduplication and compression. For more information about Data Reduction, see the *Data Reduction with Dell EMC PowerMax* white paper available on [dell EMC.com](http://dell EMC.com).

#### Nocopy Sessions (SnapVX, VP Snap):

- Compressed source data remains compressed when becoming snapshot data.
- Uncompressed source data may become compressed as it becomes snapshot data.
- Uncompressed snapshot data may be compressed as it becomes less active.

- Snapshot data is available for deduplication.
- Read activity through a nocopy linked target to uncompressed snapshot data may prevent the data from being compressed.
- Read activity through a nocopy linked target to compressed snapshot data may cause the snapshot data to be decompressed.
- Enabling data reduction on a nocopy linked target will only effect on data owned by the linked target.

#### **Copy Sessions (SnapVX Full Copy Linked Targets, Clone, Mirror):**

- The data reduction settings for both the source and the target affect copy sessions.
- When enabled on the SG for both the source and full copy targets, data is not duplicated to the target volume during the copy operation. This has the additional benefit of allowing a SnapVX full copy linked target to also be an SRDF R1 without consuming additional backend capacity.
- When data reduction is enabled on the source only, the data is decompressed before being allocated to the target.
- When data reduction is enabled on the target only, the data is compressed and possibly deduplicated before being allocated to the target.
- It is not recommended to change the data reduction settings in between differential operations (that is, disabling before each differential operation and then enabling again after the copy completes) as this causes data to go through needless compression, decompression, or deduplication cycles.

### **2.17.8 Inline compression on VMAX All Flash arrays**

Inline Compression was introduced in the HYPERMAX OS 5977.952.892 Q3 2016 Service Release. For more in-depth information about compression, see the *Data Reduction with Dell EMC PowerMax* white paper available on [dell EMC.com](http://dell EMC.com).

#### **Nocopy Sessions (SnapVX, VP Snap):**

- Compressed source data remains compressed when becoming snapshot data.
- Uncompressed source data may become compressed as it becomes snapshot data.
- Uncompressed snapshot data may be compressed as it becomes less active.
- Read activity through a nocopy linked target to uncompressed snapshot data may prevent the data from being compressed.
- Read activity through a nocopy linked target to compressed snapshot data may cause the snapshot data to be decompressed.
- The compression setting of a nocopy linked target will only effect on data owned by the linked target.

#### **Copy Sessions (SnapVX Full Copy Linked Targets, Clone, Mirror):**

- The compression settings for both the source and the target affect copy sessions.
- When compression is enabled on both the source and the target, the data is decompressed before the copy and then compressed again to allocate for target.
- When compression is enabled on the source only, the data is decompressed before being allocated to the target.
- When compression is enabled on the target only, the data is compressed before being allocated to the target.
- Copy times may vary due to decompression and compression of the data.



- It is not recommended to change the compression settings in between differential operations (that is, disabling compression before each differential operation and then enabling again after the copy completes) as this causes data to go through needless compression or decompression cycles.

## 2.18 Storage group operations

SnapVX can act on entire SGs. This reduces complexity by eliminating the need for users to create separate device files, device groups, or composite groups for replication. Those options are still available in Solutions Enabler. The SG option is the only option for SnapVX in Unisphere. This document focuses on the SG option in the following implementation section.

When creating snapshots with the `symsnapvx establish` command, the `-sg` switch is used to specify the source SG, and one snapshot of the same name is created for each volume in the SG. If for some reason a snapshot cannot be created for one of the volumes in the SG, the entire operation fails, and no snapshots are created.

When linking target volumes to a snapshot with the `symsnapvx link` command, the `-lnsg` switch is used to specify the entire target SG to link. Like the `establish` command, if any volume in the SG cannot be linked for some reason, the entire operation fails, and no volumes are linked. Unisphere also has an option that creates a new SG and appropriate volumes for the user.

In the case of cascaded SG, a `symsnapvx` command issued to a parent SG is applied to all child SGs.

---

**Note:** Solutions Enabler continues support of Composite Groups (`-cg`), Device Groups (`-g`) and Device Files (`-f`) options. Composite Groups provide the ability to take consistent snapshots across a set of devices that span multiple arrays.

---

SnapVX also introduces the `-devs` option to specify device range within the command line.

### 2.18.1 Expanding storage groups with active snapshots

SnapVX is designed to replicate entire applications that are configured on volumes contained in an entire SG. It is possible to add volumes to an SG that has active snapshots. The new volumes are included the next time a snapshot is taken. However, the volumes are not included in previously existing snapshots.

If an SG is restored from a snapshot that was taken before the new volumes were added to the SG, the SG is taken back to a point-in-time when the new volumes did not exist. Therefore, these new volumes in the SG should not be presented to the application and are set to not ready. The volumes remain not ready after the restored session is terminated and the user must decide the best course of action to reintroduce the new volumes to the SG. The volumes automatically become ready again if the user decides to restore from snapshots that include these volumes.

Similar behavior applies for linked target SGs that have been expanded. If the target SG is relinked to an older generation of the snapshot that was created before the additional volumes were added, the extra volumes in the target SG are set to not ready and unlinked from the newer snap. The user must decide the best course of action to remove or reintroduce the extra volumes. The volumes remain not ready if the unlinked. But if the SG is relinked to a generation of the snapshot that contains all volumes in the SG, the volumes are automatically be made ready and included in the session.



## 2.19 Secure snaps

Secure snaps are an optional setting for SnapVX targetless snapshots that prevent a user from deleting snapshots accidentally or intentionally. Secure snaps are available beginning with the HYPERMAX OS 5977.1125.1125 Q2 2017 Release using Solutions Enabler 8.4, Unisphere 8.4, and REST API 8.4.

Secure snaps allow the user to set a retention period on snapshots. No user can terminate the secure snap during the retention period. Attempts to terminate a secure snap prior to the retention date are rejected. When the retention time is reached the snapshot is automatically terminated.

Users may extend a retention period in situations where a snapshot is needed longer than originally planned. However, reducing a retention period is not allowed. A traditional snapshot may be converted to a secure snap, but a secure snap may not be converted to a traditional snapshot. All SnapVX operations and rules for traditional snapshots regarding restores, linked target operations, and automatic expiration also apply to secure snaps. If a secure snap has any linked targets or restored sessions when the expiration date is reached the snapshot does not terminate until the targets are unlinked or the restored session is terminated.

When implementing secure snaps, a user should determine how many snapshots on an array must be secure. Users should consider using secure snaps only on certain critical volumes, or only on a subset of the snapshots to capture points-in-time that are critical to the business, and for how long the secure snaps should be retained. As always, proper planning and system sizing is crucial, no matter the types or number of snapshots that will exist in an environment.

The CLI format of the retention period is the same as that of the traditional expiration date. The traditional expiration date and secure retention period cannot be used together on a single snapshot.

The following is the Solutions Enabler command structure along with an example of setting a snapshot to secure for one day and 12 hours:

```
symsnapvx -sg <sg> -name <snapshot_name> establish --secure <-delta | absolute>
```

```
symsnapvx -sg prod_sg -name daily_8am_snap establish --secure -delta 1:12
```

Figure 12 shows how to create secure snaps with the Unisphere snapshot wizard. A snapshot cannot be set as secure if the standard expiration date has already been set. Likewise, the standard expiration date cannot be set on a snapshot if secure snapshot has already been selected.

**Advanced Options**

☐ Enable Force Flag ☒ Enable Secure Snaps

**Secure snapshots may only be terminated after they expire or by customer-authorized Dell EMC Support**

Days: 1 Hours: 0

CANCEL OK

Figure 12 Creating secure snaps in Unisphere 9.0

### 2.19.1 Secure snap behavior and use case comparison to traditional snapshots

Secure snaps are intended for use in environments where protecting specific point-in-time images are considered the highest priority and therefore exhibit several behavioral differences compared to traditional snapshots. For example, if the system is reaching Replication Cache or SRP resource limits the user cannot terminate secure snaps to help remediate the situation as can be done with traditional snapshots.

Secure snaps are also unique in the handling of host writes and snapshots when SRP or Replication Cache limits are reached. When Reserved Capacity of an SRP is reached and a host write to a source volume requires a new allocation in the SRP, an existing traditional snapshot is placed into a failed state. However, secure snaps can grow past the Reserved Capacity limit of the SRP. New secure snaps cannot be created once Reserved Capacity is reached.

Preserving the secure snaps takes priority over host writes to those source volumes in the event the entire SRP runs out of available capacity. This is the intent of the feature. For example, if there is an intentional attack or a runaway application on a system that consuming all available capacity, preserving the secure snaps allows the user to restore from the secure snaps once the situation is resolved.

Note: Secure snapshots may only be terminated after they expire or by customer-authorized Dell Support. See the knowledge base article 498316 for more information: <https://support.emc.com/kb/498316>

## 2.20 Online device expansion

Expansion of a TDEV with active SnapVX sessions is available beginning with PowerMaxOS 5978, Solutions Enabler 9.0, Unisphere for PowerMax 9.0, and REST API 9.0. Both source and target devices can be expanded, and the functionality supports FBA and CKD devices.

SnapVX source devices may be expanded while maintaining the point-in-time image of existing targetless snapshots. Solutions Enabler and Unisphere outputs have fields to report if a device has been expanded since the snapshot was taken, and the device size at time of snapshot creation. The fields are shown in the following output example and in Figure 13.

The following is the output of the `symsnapvx list -detail` command in Solutions Enabler 9.0.

```
$ symsnapvx -sid <sid> -sg <sg_name> list -detail
```

```
Storage Group (SG) Name : ODE_Prod_SG
SG's Symmetrix ID      : 000197800XYZ (Microcode Version: 5978)
```

```
-----
                Snapshot Total
Sym      Flags      Dev Size Deltas Non-Shared
Dev Snapshot Name Gen FLRG TSEB Snapshot Timestamp (Tracks) (Tracks) (Tracks) Expiration Date
-----
3B Snap1  0 .... ..X. Fri Nov 17 12:05:01 2017 327690 452 0 Fri Nov 17 14:05:01 2017
3C Snap1  0 .... ..X. Fri Nov 17 12:05:01 2017 327690 869 0 Fri Nov 17 14:05:01 2017
3D Snap1  0 .... ..X. Fri Nov 17 12:05:01 2017 327690 784 0 Fri Nov 17 14:05:01 2017
3E Snap1  0 .... ..X. Fri Nov 17 12:05:01 2017 327690 142 0 Fri Nov 17 14:05:01 2017
-----
                2247      0
```

Flags:

```
(F)ailed : X = General Failure, . = No Failure
          : S = SRP Failure, R = RDP Failure
(L)ink   : X = Link Exists, . = No Link Exists
(R)estore : X = Restore Active, . = No Restore Active
(G)CM    : X = GCM, . = Non-GCM
(T)ype   : Z = zDP snapshot, . = normal snapshot
(S)ecured : X = Secured, . = Not Secured
(E)xpanded : X = Source Device Expanded, . = Source Device Not Expanded
(B)ackground: X = Background define in progress, . = No Background define
```

DCB_Snaps > Source Volumes						
<div> <div>Restore</div> <div>Link</div> <div>Relink</div> <div></div> </div>						
Name	State	Creation Time	Secured	Linked	Expanded	Capacity (GB)
<input type="checkbox"/> 0003B	Established	Tue Oct 31 2017 14:12:...	—	—	—	40.00
<input type="checkbox"/> 0003C	Established	Tue Oct 31 2017 14:12:...	—	—	—	40.00
<input type="checkbox"/> 0003D	Established	Tue Oct 31 2017 14:12:...	—	—	—	40.00
<input type="checkbox"/> 0003E	Established	Tue Oct 31 2017 14:12:...	—	—	—	40.00

Figure 13 SnapVX source volume display in Unisphere 9.0

## 2.20.1 Considerations for linked targets

- The snapshot device size must be considered during link operations.
- Target devices must be equal or larger capacity of the snapshot device size, not the size of the source device itself.
- Snapshot device size must also be considered during relink operations.
  - For example, if a snapshot has an existing linked target of the same capacity, and then the source device is expanded, then a new snapshot is taken, attempt to relink to the new snapshot will be blocked.
  - However, the target device can be expanded without being unlinked.

- Fully defined linked target devices in both copy and nocopy modes may be expanded.
- Restore from a snapshot that was created before a source device is expanded is allowed.
  - Source TDEV size is not reduced.
  - Data on the source TDEV that is beyond the size of the snapshot is deallocated during the restore operation.

Device expansion is not supported with the following session types:

- Legacy TimeFinder Modes
- Devices in active Non-disruptive Migration (NDM) sessions
- Open Replicator Sessions (ORS)
- Parallel Clone
- VMware® vSphere® Virtual Volumes™ (vVols) with Replication
- FlashCopy or Remote Pair FlashCopy
- Devices involved in ProtectPoint sessions
- Internal and external RecoverPoint devices

## 2.21 Role Based Access Controls

Role Based Access Controls in Solutions Enabler 9.0, Unisphere 9.0, and REST API 9.0 introduced roles which provide the ability to restrict certain TimeFinder operations for users. Roles and permissions as related to TimeFinder are as follows:

- Local Replication:
  - Required to create, modify, and terminate TimeFinder sessions from a source device.
- Device Manager:
  - Required to use a device as a TimeFinder target, and to change copy mode of existing session.
  - Require *with* Local Replication role to restore a TimeFinder session.
- Storage Admin:
  - Required to create and extend retention time for secure snapshots, as well as converting a traditional snapshot to a Secure Snapshot.

More information about Role Based Access Controls is available in the Dell EMC document *Role Based Access Controls (RBAC) Technical Overview & Enhancements*.

## 3 Implementing TimeFinder SnapVX

### 3.1 Creating snapshots

The following details creating snapshots with Solutions Enabler:

```
$ symsnapvx -sid XYZ -nop -sg Prod establish -name ProdSnap -ttl -delta 2
```

```
Establish operation execution is in progress for the storage group Prod.  
Please wait...
```

```
  Polling for Establish.....Started.  
  Polling for Establish.....Done.
```

```
Establish operation successfully executed for the storage group Prod
```

Figure 14 shows creating a snapshot called ProdSnap and keeping for 2 days.

Figure 14 Creating snapshot with Unisphere

The following example shows there are 4 generations of snap from the SG snap source:

```
$ symsnapvx -sid XYZ list -sg Prod
```

```
Storage Group (SG) Name : Prod
SG's Symmetrix ID      : 000197200XYZ (Microcode Version: 5977)
```

```
-----
Sym          Num Flgs
Dev  Snapshot Name      Gens FLRG Last Snapshot Timestamp
-----
00040 ProdSnap          4 .... Mon Oct 20 11:00:00 2014
00041 ProdSnap          4 .... Mon Oct 20 11:00:00 2014
00042 ProdSnap          4 .... Mon Oct 20 11:00:00 2014
00043 ProdSnap          4 .... Mon Oct 20 11:00:00 2014
```

Flgs:

```
(F)ailed : X = Failed, . = No Failure
(L)ink   : X = Link Exists, . = No Link Exists
(R)estore : X = Restore Active, . = No Restore Active
(G)CM    : X = GCM, . = Non-GCM
```

The following shows a more detailed view giving more information including the timestamp of each snapshot and expiration date if one was set on creation.

```
$ symsnapvx -sid XYZ list -sg Prod -detail
```

```
Storage Group (SG) Name : Prod
SG's Symmetrix ID      : 000197200XYZ (Microcode Version: 5977)
```

```
-----
Sym          Flgs          Total
Dev  Snapshot Name      Gen FLRG Snapshot Timestamp (Tracks) (Tracks) Expiration Date
-----
00040 ProdSnap          0 .... Mon Sep 29 11:00:00 2014 26371 25848 NA
    ProdSnap          1 .... Mon Sep 29 10:00:00 2014 26389 25866 NA
    ProdSnap          2 .... Mon Sep 29 09:00:00 2014 26389 5120 Thu Oct 23 10:13:55 2014
    ProdSnap          3 .... Mon Sep 29 08:00:00 2014 26389 5120 Wed Oct 22 10:11:58 2014
00041 ProdSnap          0 .... Mon Sep 29 11:00:00 2014 26379 25886 NA
    ProdSnap          1 .... Mon Sep 29 10:00:00 2014 26393 25900 NA
    ProdSnap          2 .... Mon Sep 29 09:00:00 2014 26393 5086 Thu Oct 23 10:13:55 2014
    ProdSnap          3 .... Mon Sep 29 08:00:00 2014 26393 5086 Wed Oct 22 10:11:58 2014
00042 ProdSnap          0 .... Mon Sep 29 11:00:00 2014 26377 25833 NA
    ProdSnap          1 .... Mon Sep 29 10:00:00 2014 26388 25844 NA
    ProdSnap          2 .... Mon Sep 29 09:00:00 2014 26388 5199 Thu Oct 23 10:13:55 2014
    ProdSnap          3 .... Mon Sep 29 08:00:00 2014 26388 5199 Wed Oct 22 10:11:58 2014
00043 ProdSnap          0 .... Mon Sep 29 11:00:00 2014 26377 25857 NA
    ProdSnap          1 .... Mon Sep 29 10:00:00 2014 26386 25866 NA
    ProdSnap          2 .... Mon Sep 29 09:00:00 2014 26386 5137 Thu Oct 23 10:13:55 2014
    ProdSnap          3 .... Mon Sep 29 08:00:00 2014 26386 5137 Wed Oct 22 10:11:58 2014
-----
                                422172 247984
```

Flgs:

```
(F)ailed : X = Failed, . = No Failure
(L)ink   : X = Link Exists, . = No Link Exists
(R)estore : X = Restore Active, . = No Restore Active
(G)CM    : X = GCM, . = Non-GCM
```

**Note:** See appendix B for information about the GCM Flag.

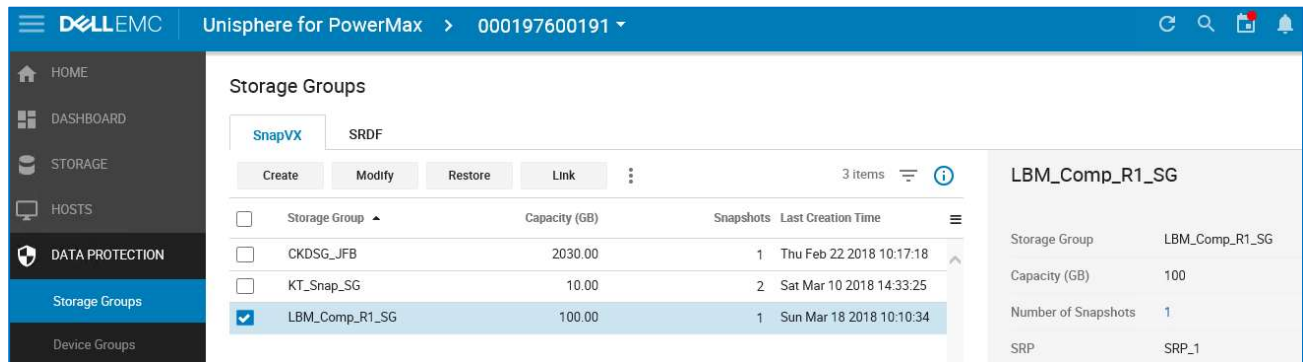


Figure 15 Viewing snapshots with Unisphere 9.0

## 3.2 Linking, relinking, and unlinking

To access a point in time copy, a host mapped target volume is linked to the snapshot data. The links may be created in Copy mode for a permanent copy of the target volume, or in Nocopy mode for temporary use.

Copy mode links create full-volume clones of the data. Nocopy mode links are space saving snapshots that only consume space when host data is written directly to the target volume. Both modes can be in either the same or separate SRP as the source volume.

The following shows the CLI commands for linking to a snapshot, note that a generation number was specified. If no generation is specified, the most recent snapshot for the snapshot name provided will automatically be selected.

```
$ symsnapvx -sid XYZ -sg Prod -snapshot_name ProdSnap -gen 2 link -lnsg NonProd
```

```
Execute Link operation for Storage Group Prod (y/[n]) ? y
```

```
Link operation execution is in progress for the storage group Prod. Please wait...
```

```
  Polling for Link.....Started.
```

```
  Polling for Link.....Done.
```

```
Link operation successfully executed for the storage group Prod
```

**Note:** See appendix C for information about monitoring the copy process to copy mode linked targets.

When linking snapshots using Unisphere, the user is presented with existing SGs to select for the targets. Unisphere can also create a new SG with the appropriate volumes and perform the link operation from one simple wizard, as seen in Figure 16.

**Link Snapshot | LBM\_Comp\_R1\_SG**

Select Snapshot

Enter the name of the snapshot to filter...

✓ Scheduled\_Hourly\_Snaps

Creation Date \*

Sun Mar 18 2018 10:10:34 GMT-0400 (0)

Select Target Storage Group

☒ New ☐ Existing

New Storage Group Name \*

LBM\_Comp\_R1\_LNK\_SG\_001

Advanced Options

? CANCEL ADD TO JOB LIST ▾

Figure 16 Linking target with Unisphere 9.0

Should the user need to transition to a different snapshot copy using the same set of source and target volumes a relink operation can be performed linking to a different snapshot or a different generation of the same snapshot. Users should unmount any target volumes from the mount host and remount after relinking to ensure that the data the host sees is correct.



The following shows the relink process with Solutions Enabler and Figure 17 shows the relink process in Unisphere.

```
$ symsnapvx -sid XYZ -sg Prod -snapshot_name ProdSnap -gen 0 relink -lnsg NonProd
```

```
Execute Relink operation for Storage Group Prod (y/[n]) ? y
```

```
Relink operation execution is in progress for the storage group Prod. Please wait...
```

```
  Polling for Relink.....Started.
```

```
  Polling for Relink.....Done.
```

```
Relink operation successfully executed for the storage group Prod
```

Figure 17 Relinking target with Unisphere 9.0

At any time, the target volume can be unlinked from the snap copy, if the link was created with the copy mode and state transitioned to a copied state the target volumes is an independent copy of the source.

Independent copies that are no longer required can be reused as linked targets in the future providing there are no snapshots cascaded off these volumes. If the space used by linked copies needs to be returned to the SRP the user can run the `symdev free -all` command to return this space to the system for global use.

---

**Note:** It is important to understand that the `symdev free -all` command provides the user with the ability to wipe a volume of all data. The command should be used with caution and only by knowledgeable administrators.

---

### 3.3 Restoring from snapshots

It is possible to restore from any point in time snapshot that is available in the array. The following shows how to restore from a snapshot to the production volume using the CLI commands. It is recommended to unmount the production volume before issuing the restore operation and remount after to ensure that the host sees the correct data. Users can determine which snapshot to use by verifying the data through a linked target with the link or relink commands.

```
$ symsnapvx -sid XYZ -sg Prod -snapshot_name ProdSnap -gen 2 restore
```

```
Execute Restore operation for Storage Group Prod (y/[n]) ? y
```

```
Restore operation execution is in progress for the storage group Prod. Please wait...
```

```
  Polling for Restore.....Started.
```

```
  Polling for Restore.....Done.
```

```
Restore operation successfully executed for the storage group Prod
```

That the example above shows the command-line option restoring from generation 2 of a snapshot to the source volume.

## 4 PowerMaxOS Q3 2020 release enhancements

The following PowerMaxOS Q3 2020 release introduced many enhancements to SnapVX, including the following:

**Snapshot policies:** Automated, compliance-based snapshot protection for applications.

**Increased session limits:** Snapshot policies support up to 1024 snapshots per source volume.

**Cloud snapshots:** Cloud Mobility for Dell EMC PowerMax provides open-systems snapshot movement to and from private and public clouds (Dell EMC ECS™, AWS, Microsoft Azure) for recovering data back to the array and consumed it directly in AWS.

**SnapSet:** A set of consistent snapshots that are taken together. For example, when a snapshot is taken of an SG that contains 10 devices, the resulting SnapSet consists of 10 consistent snapshots.

**SnapSet ID:** A system-generated number that is applied to consistent snapshots across the volumes in a SnapSet. When a snapshot is taken of an SG, all resulting snapshots have the same SnapSet ID. The SnapSet ID is an absolute value that does not adjust as other snapshots are created or terminated.

**Bulk terminate:** A Unisphere term for functionality that allows termination of many snapshots of an SG in a single operation.

**Snapshot SRP Utilization Alert:** Unisphere alert that triggers on preset threshold values for overall snapshot usage across the entire SRP and for designated SGs.

**Interoperability of nocopy linked and SRDF:** Support for using an SRDF R1 as a nocopy linked target. Prior code releases require an SRDF R1 to be linked in full copy mode.

### 4.1 Snapshot policies

Snapshot policies provide automated scheduling of SnapVX snapshots using a highly available and flexible policy engine that runs internally on the storage array. Snapshot policies can be managed through Dell EMC Unisphere™ for PowerMax, REST API, and Solutions Enabler.

Snapshot policies can be customized with rules that specify when to take snapshots, how many snapshots to take, and how long to keep each snapshot. Compliance requirements can also be specified to send alerts if the rules of a policy are not being met. Applications can be protected by multiple policies with differing schedules and retention parameters according to the requirements of the business. Each policy can protect many applications, even protecting a mix of open systems and mainframe applications.

Snapshot policies provide reliable protection for applications in an automated fashion that requires little to no maintenance by the business. Administrators can manually take snapshots of applications that are protected by snapshot policies to satisfy on-demand requirements.

For more information, see [Dell EMC PowerMax and VMAX All Flash: Snapshot Policies Best Practices](#).

### Create Snapshot Policy

Properties

Name:

Location: ☒ Local ☐ Cloud


Type: ☐ Secure


Recovery Point Objective (RPO)

Create a snapshot: By The Minute every 15 mins at :00

Keep 2 Hour(s) Total 8 Snapshots (Max 1024)

Compliance

Show as  If fewer than 7 (88%) snapshots are created

 If fewer than 6 (75%) snapshots are created


 CANCEL CREATE

Figure 18 Snapshot Policy Parameters

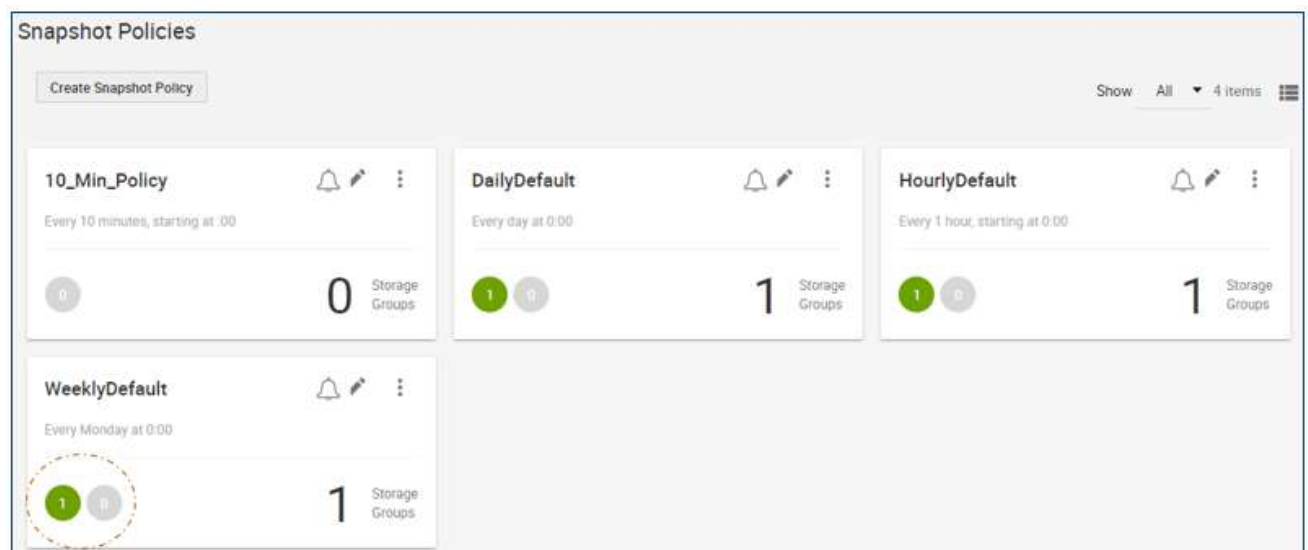


Figure 19 Snapshot Policies dashboard

## 4.2 Cloud Mobility for Dell EMC PowerMax

Cloud Mobility for Dell EMC PowerMax provides application-level protection for block devices through snapshot shipping to private and public clouds (Dell EMC ECS™, Amazon Web Services (AWS), and Microsoft® Azure®). Data that is moved to the cloud can be used to restore data to the array and can be consumed directly in AWS.

Cloud Mobility offers the following benefits:

- Reduces array capacity usage
- Data repurposing: Data analytics, reporting
- Space-efficient snapshot shipping: Full initial copy and differential updates
- Cloud-system-independent compression and encryption
- Separation from PowerMax data services
- In-flight and Data at Rest Encryption (on cloud provider)

Cloud Mobility snapshot policies can be managed through Unisphere for PowerMax and REST API, and the policies can only be created with Embedded Management for Unisphere. Cloud Mobility policy parameters include:

- Cloud provider selection
- RPO options:
  - Daily or weekly
- Keep snapshots for:
  - 3 days
  - 1 week
  - 1, 3, or 6 months
  - 1 to 14 years

**Create Snapshot Policy**

**Properties**

Name: AWS-1year-weekly

Location: ☐ Local ☒ Cloud AWS-Virginia-rename [\(Change\)](#) ⓘ

**Recovery Point Objective (RPO)**

Create a snapshot: Weekly on Tuesday at 00:00

Keep For: 1 Year

**Compliance**

Show as: ⚠ if fewer than 51 (98%) snapshots are created

✖ if fewer than 50 (96%) snapshots are created

ⓘ CANCEL CREATE

Figure 20 Creating a Cloud Mobility policy

For more information, see the Cloud Mobility for Dell EMC PowerMax white paper.

## 4.3 SnapSet and SnapSet ID

A SnapSet is a set of consistent snapshots that are taken together across a group of volumes. For example, when a snapshot is taken of an SG that contains 10 devices, the resulting SnapSet consists of 10 consistent snapshots.

Beginning with the PowerMaxOS Q3 2020 release, snapshots are assigned a SnapSet ID. The SnapSet ID is an absolute value that remains the same regardless of creation or deletion of other snapshot generations. When a snapshot is taken across a SnapSet, the snapshots that are created on the individual TDEVs in the SnapSet have the same SnapSet ID. SnapSet IDs are assigned to all snapshot types.

Figure 21 shows the SnapSet ID of a snapshot.

Storage Groups > LBM\_Test\_SG

DETAILS VOLUMES PERFORMANCE DATA PROTECTION

SNAPSHOTS CLOUD SNAPSHOTS SNAPSHOT POLICIES SRDF

Create Modify Restore Link

Snapshot	Creation Time	Linked	Restored	Expired	Secured	Persistent	Snapshot type
<input checked="" type="checkbox"/> HourlyDefault	Wed Mar 04 2020 12:00:17	---	---	---	---	---	Policy
<input type="checkbox"/> HourlyDefault	Wed Mar 04 2020 11:00:23	---	---	---	---	---	Policy
<input type="checkbox"/> HourlyDefault	Wed Mar 04 2020 10:00:20	---	---	---	---	---	Policy
<input type="checkbox"/> HourlyDefault	Wed Mar 04 2020 09:00:26	---	---	---	---	---	Policy
<input type="checkbox"/> HourlyDefault	Wed Mar 04 2020 08:00:22	---	---	---	---	---	Policy
<input type="checkbox"/> HourlyDefault	Wed Mar 04 2020 07:00:18	---	---	---	---	---	Policy
<input type="checkbox"/> HourlyDefault	Wed Mar 04 2020 06:00:24	---	---	---	---	---	Policy

24 items

Name HourlyDefault

Snapset ID 115857421697

Creation Time Wed Mar 04 2020

Expiry Date

State Established

Secured No

Figure 21 SnapSet ID

The generation ID of a snapshot is relative to the number of snapshots at the time the snapshots are viewed. However, the SnapSet ID remains the same throughout the life of the snapshot. Since the introduction of SnapVX, the standard method to control or view details of a snapshot is to use the relative generation number, and this method continues to be supported.

However, using snapshot policies can increase the snapshot frequency and snapshot count. This usage also increases the probability that the relative generation number is changed shortly before the user issues a control command to a snapshot. This scenario would issue the control command to the incorrect snapshot.

Creating and terminating snapshots on subsets of the devices in a group could also cause snapshots of the group to have different generation IDs on the individual snapshots. The SnapSet ID would remain constant.

The best practice is to use the SnapSet ID rather than the generation ID. We do not recommend correlating the SnapSet ID of a snapshot to the generation ID. The generation ID will change the next time another snapshot is created and or terminated.

## 4.4 Enhanced snapshot termination options

The PowerMaxOS Q3 2020 release introduces options to terminate multiple snapshots simultaneously. System performance is not affected by snapshot termination because the back-end deltas are deallocated by a low-priority background task after the snapshots have been terminated.

The Bulk Terminate options allow the user to terminate a sequential range of snapshots from a **single SG** with the following options:

- Terminate all snapshots
- Terminate all versions of a specific snapshot
- Terminate oldest versions of a snapshot by specifying which to keep
- Terminate oldest versions of a snapshot by specifying which to terminate

The Bulk Terminate options are not exclusive to policy snapshots and can be used on manual snapshots too. However, snapshots that are linked, in a restored state, or created by Cloud Mobility, are not terminated by the Bulk Terminate options.

The **Bulk Terminate** options are intended to simplify deletion of multiple snapshots. The operations will terminate as many snapshots as possible even if some cannot be terminated. This behavior differs from standard snapshot deletion and other snapshot operations which must be successful across the entire set of LUNs. After a Bulk Terminate operation completes, verify that all intended snapshots were successfully terminated.

Examples of the snapshot terminate options are shown in the following figures.

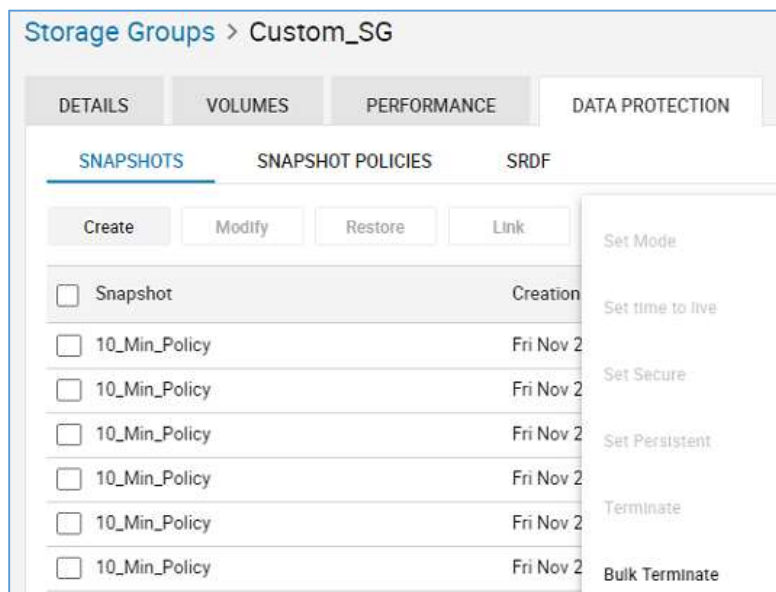


Figure 22 Bulk Terminate

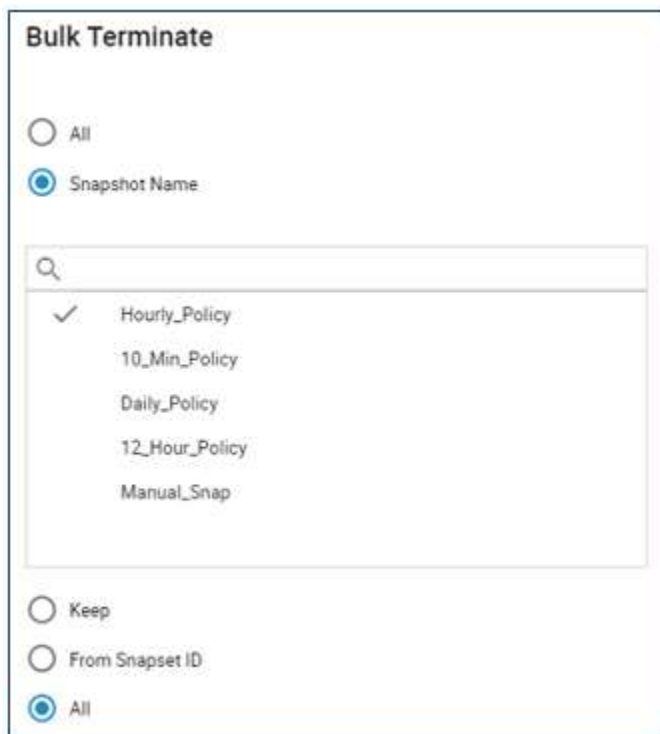


Figure 23 Bulk Terminate

In addition to the Bulk Terminate options, you can use Unisphere to manually select multiple snapshots of an SG and terminate them with a single operation.

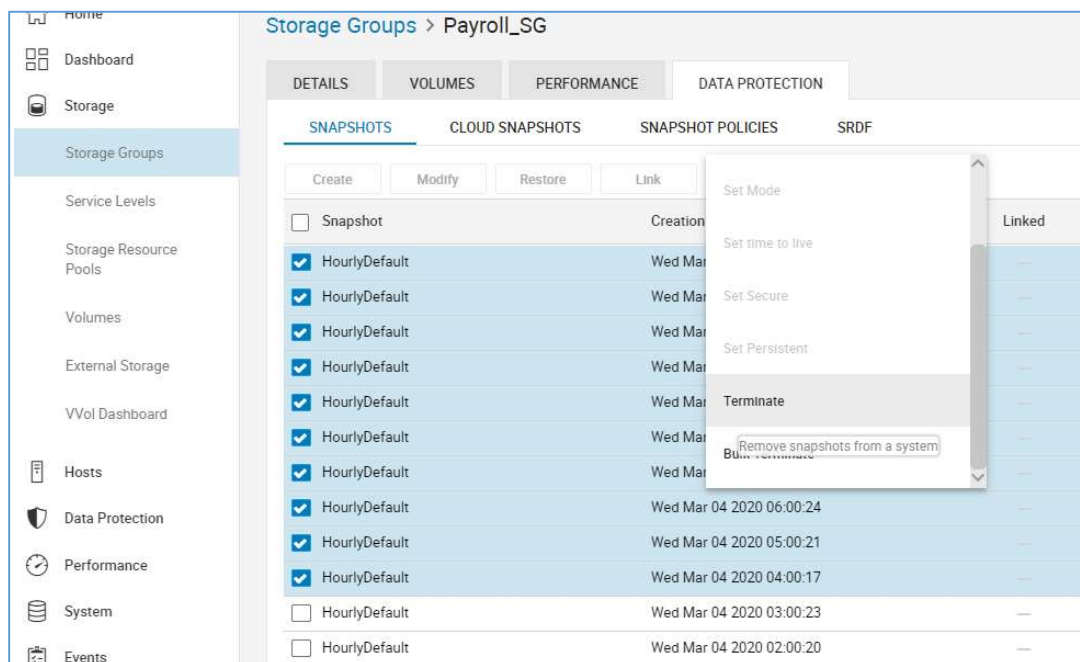


Figure 24 Manually selecting snapshots to terminate



## 4.5 Snapshot SRP utilization alert

Unisphere for PowerMax 9.2 introduced a snapshot SRP utilization alert. The default alert triggers on preset threshold values for overall snapshot usage across the entire SRP. Custom alerts can also be created with customized thresholds for specified SGs.

**Create System Threshold Alert**

Category  
Snapshot SRP Utilization ▼

☒ Storage Group ☐ Wildcard

**Select Instances to Enable**

Available Instances

Name
<input checked="" type="checkbox"/> Local_Rep_SG
<input type="checkbox"/> LBM_KT_Test_SG

Instances to Enable

Name
<input checked="" type="checkbox"/> Name

Warning 60 % default

Critical 80 % default

Fatal 100 % default

Figure 25 SRP utilization alert

## 5 TimeFinder legacy modes: Mirror, Clone, and VP Snap

TimeFinder/Mirror, Clone, and VP Snap are available with PowerMaxOS, HYPERMAX OS, Solutions Enabler 8.0 and later and Unisphere 8.0 and later.

Legacy sessions follow all feature restrictions from Enginuity 5876, including session counts, cascaded hop limits, and requiring source-target pairing. Legacy sessions cannot take advantage of SG operations (`-sg` option). Legacy mode target volumes can be moved between SRPs while sessions are active.

The `symclone` and `symmir` commands do not convert legacy commands to native SnapVX operations. Instead, specific legacy sessions are created which have behavioral characteristics different than native SnapVX. For example, when a Clone or Mirror session is created, a targetless snapshot will not be created. Data will copy directly from source to target without using snapshot deltas.

### 5.1 Interoperability between native SnapVX and legacy TimeFinder modes

The PowerMaxOS 5978 Q3 2019 SR along with Solutions Enabler 9.1, Unisphere 9.1, and REST API 9.1 introduce the ability to create native SnapVX and legacy mode sessions from the same source volume. This can be used either as a short-term solution to maintain existing snaps/clones while transitioning from one mode to another, or as a continuous solution to provide an application with key benefits from the various session types. For example, to create a Clone of a database that is already being protected by regularly scheduled SnapVX snapshots.

Interoperability is not allowed when cascading sessions. For example, a SnapVX snapshot cannot be taken from a Clone target. Interoperability is also limited during restore operations; legacy mode sessions cannot be created if a native SnapVX snapshot is in a restore state, and conversely.

### 5.2 Snapshots created by zDP

Data Protector for z Systems (zDP) is a Dell EMC z/OS-based application that uses TimeFinder SnapVX snapshots to enable rapid recovery from logical data corruption. zDP achieves this by providing multiple, frequent, and consistent point-in-time copies of data in an automated fashion across multiple volumes from which an application level recovery can be conducted.

For more information about zDP see the Data Protector for Z Systems (zDP) Essentials White Paper.

Solutions Enabler and Unisphere can display snapshots that were created by zDP but cannot perform control operations on those snapshots.

PowerMaxOS 5978 Q3 2019 SR increased the zDP snapshot limit from 256 to 1024 snapshots per source volume. Solutions Enabler 9.1, Unisphere 9.1, and REST API 9.1 can report each of those snapshots.

However, versions of management software before 9.1 do not have the ability to report more than 256 snapshots per source volume. If a volume has any zDP snapshots that were created with PowerMaxOS 5978 Q3 2019 SR, a single placeholder snapshot with the name `_Automated_Snapshot` will be displayed for the volume as seen in the following example.

```
C:\Users\Administrator>symsnapvx -sid 191 list -devs 31
```

(Microcode Version: 5978)

Sym	Snapshot Name	Num	Flags	Last Snapshot Timestamp
Dev		Gens	FLRG TS	
00031	_Automated_Snapshot	1	.... ..	Mon Mar 25 14:56:54 2019

Figure 26 Solutions Enabler 8.4 reporting zDP snapshots which were created with PowerMaxOS 5978 Q3 2019 SR

## 5.3 TimeFinder VP Snap

VP Snap control commands are the same as with Solutions Enabler 7 and Enginuity 5876, enabling users to use their existing VMAX VP Snap scripts. However, there are several behavioral differences that need to be understood when preparing volumes and monitoring sessions. The following behaviors apply to VP Snap in PowerMaxOS and HYPERMAX OS:

- All VP Snap target volumes from a specific source do not need to be associated to the same SRP.
- VP Snap target volumes can be moved across SRPs while sessions are active. Non-shared allocations are moved.
- As source tracks are updated, the Protected Track counts do not decrement.
- As source tracks are updated when the session is active, the preserved point-in-time tracks for the VP Snap sessions are stored in the SRP of the source volume, even if the target volume is in a different SRP. Only changed tracks on the target volume are stored in the target SRP.
- Shared allocation counts in Solutions Enabler and Unisphere do not apply due to the way that shared allocations are handled in PowerMaxOS and HYPERMAX OS. These counters always show zero.
- VP Snap target volumes are not deallocated when the sessions are terminated. If the target volume is in a separate SRP from the source volume, any allocations that reside in the source SRP after terminating are relocated to the target SRP by a FAST compliance move.

---

**Note:** See appendix C for information about monitoring the copy process of Mirror and Clone sessions.

---

## 5.4 TimeFinder/Snap deprecation

TimeFinder/Snap (`symsnap`) is not supported on VMAX3 or Solutions Enabler 8.0 and later.

## 6 Conclusion

TimeFinder SnapVX provides enhanced functionality and combines the benefits of the TimeFinder Clone, Mirror, and VP Snap into one easy-to-use software feature.

SnapVX allows the user to create snapshots without the need for a target volume. Snapshots can then be used to link to target volumes in either full-copy or nocopy mode which can then be presented to the host server. SnapVX allows for far greater scalability than previous TimeFinder offerings, with up to 1024 snapshots per source volume and 1024 linked target volumes per source volume. SnapVX also introduces the ability to take snapshots on an SG level and uses advanced redirect-on-write technology.

SnapVX is compatible with many other features. Users can also use legacy TimeFinder/Mirror, Clone, and VP Snap modes which use SnapVX technology in the background, completely transparent to the user.

## A TimeFinder SnapVX state table

Table 2 describes prerequisites, transient states, and final states for all controls:

Table 2 TimeFinder SnapVX state table

Action	Prerequisite	Transient State	Final State
Establish	None	Establish In Progress	Established
Restore	Established  If the source device is a link target, must be fully copied.	Restore In Progress	Restored
Terminate	Established	Terminate In Progress	NA
Terminate with FLAG1_RESTORED	Restored	NA	NA
Setmode Copy	Linked NoCopy	Link Copy In Progress	Link Copied
Setmode NoCopy	Link Copy In Progress or Link Copied	NA	Linked
Set TTL	Established with no links or restores.	NA	NA
Link	Established	NA	Linked
Link with FLAG1_COPY	Established	Link Copy In Progress	Link Copied
Unlink	Linked  Link Copied  Link Defined	NA	NA
Relink	Linked or  Link Copied	NA	Linked
Relink with FLAG1_COPY	Linked or  Link Copied	Link Copy In Progress	Link Copied
Rename	Established	NA	NA

## B Geometry compatibility mode

One of the architectural differences between PowerMaxOS 5978 and HYPERMAX OS 5977 compared to Enginuity 5876 is the doubling of track size, from 64K to 128K. An array running PowerMaxOS 5978 and HYPERMAX OS 5977 cannot create a device that is the same size as a device with an odd number of cylinders on an array running Enginuity 5876. To support full functionality, SRDF R1 and R2 pairs must be same size and TimeFinder source and target pairs must be the same size.

HYPERMAX OS introduced a device attribute, geometry compatible mode (GCM). A device with GCM set is treated as half a cylinder smaller than its true configured size so the device will be treated as the same size as an Enginuity 5876 device with an odd number of cylinders, thus enabling full functionality between PowerMaxOS 5978 and HYPERMAX OS 5977 with Enginuity 5876 for SRDF, TimeFinder, and ORS.

The GCM attribute can only be set for devices on arrays running PowerMaxOS 5978 and HYPERMAX OS 5977. The attribute can be set manually or automatically as part of an operation that creates a local or remote replication relationship.

The `symdev set /unset`, `symdg set/unset`, `symcg set/unset`, and `symsg set/unset` commands have been enhanced with a new option `-gcm` to set and unset GCM for a device or group.

The `symdev show`, `symdev list -v`, `symdg show ld`, `symdg list ld -v`, `sympd show`, and `sympd list -v` commands have been enhanced to report the GCM attribute.

The `symrdf createpair` command has been enhanced to transparently `set/unset` the GCM attribute as part of the create pair operation, as follows:

Set the GCM attribute for a target device that is configured half cylinder larger. The source of the copy can be:

A device on an array running Enginuity 5876 with an odd number of cylinders and capacity that matches the GCM size of the device.

A GCM device on an array running PowerMaxOS 5978 or HYPERMAX OS 5977.

Unset the GCM attribute for a target device that is configured the exact same size as the source of the copy. The source of the copy can be:

A device on an array running Enginuity 5876 with even number of the cylinders and capacity that matches the size of the device on the array running PowerMaxOS 5978 or HYPERMAX OS 5977.

A device on the array running PowerMaxOS 5978 or HYPERMAX OS 5977 without the GCM attribute.

### B.1 Open Replicator for Symmetrix (ORS)

The ORS create operation will be modified to use the GCM size when establishing a push or pull session to the remote target.

## B.2 TimeFinder SnapVX

If a SnapVX snapshot is created from a GCM device, the resulting snapshot is also given the GCM attribute and this is reported as an attribute of the snapshot. A SnapVX link operation changes the GCM attribute on the link target device to match the GCM setting of the snapshot.

## B.3 TimeFinder Clone and VP Snap

Create, full establish, or full restore operations between a GCM device and a non-GCM device result in the device that is the target of the data copy having its GCM attribute changed to match the GCM attribute setting of the device that is the source of the data copy.

When the source and target devices are the same physical size, and the source of a clone session is a GCM device and the target device is not, then the Clone to Larger Target rules are enforced.

## B.4 TimeFinder/Mirror

Full establish and full restore operations between a GCM device and a non-GCM device result in the device that is the target of the data copy having its GCM attribute changed to match the GCM attribute setting of the device that is the source of the data copy.

## C Monitoring the copy process

When monitoring the copy process to a full copy linked target, the number of tracks remaining to copy are determined by a sum of multiple internal counters, capped at the maximum size of the volume. If the sum of these values exceeds the size of the device, it may appear that the copy has not started or is stuck when querying the session from Solutions Enabler or Unisphere, even though the copy has actually begun. For example:

```
$ symsnapvx -sid XYZ list -sg snapsource -linked -detail
```

```
Storage Group (SG Name : SNVX_SG
SG's Symmetrix ID : 000197200XYZ (Microcode Version: 5977)
```

```
-----
Sym          Link Flgs          Remaining Done
Dev  Snapshot Name      Gen Dev  FCMD Snapshot Timestamp      (Tracks)  (%)
-----
00040 hourlysnap        2 0003D ...I Mon Sep 29 14:01:59 2014    18510    0
00041 hourlysnap        2 0003E ...I Mon Sep 29 14:01:59 2014    18510    0
00042 hourlysnap        2 0003F ...I Mon Sep 29 14:01:59 2014    18510    0
00043 hourlysnap        2 00040 ...I Mon Sep 29 14:01:59 2014    18510    0
-----
                                74040
```

Flgs:

```
(F)ailed : F = Force Failed, X = Failed, . = No Failure
(C)opy : I = CopyInProg, C = Copied, D = Copied/Destaged, . = NoCopy Link
(M)odified : X = Modified Target Data, . = Not Modified
(D)efined : X = All Tracks Defined, . = Define in progress
```

Figure 27 Monitoring the Copy Process with Solutions Enabler

The user may be able to notice usage increase on the target volumes or SRPs. This does not slow down the actual copy, just the way that the copy is reported. Once the internal counters drop below a certain threshold the number of tracks reported to the user drops dramatically.

Another effect of this is that incremental relink and restore operations may appear as full operations because the internal counters add up to the full size of the volume. Both behaviors will be more noticeable on larger volumes.

This behavior also applies to monitoring the copy process of TimeFinder/Mirror and Clone sessions.

This behavior is also documented in Dell EMC Knowledgebase article 196700  
<https://support.emc.com/kb/196700>



## D Nocopy Linked Target behavior before HYPERMAX OS 5977.810.184

With HYPERMAX OS 5977.810.184 and later, users can access the data on fully defined nocopy targets after having been unlinked. This functionality is possible through the shared allocations. As discussed earlier, the defining process creates the shared allocations between the target volume and source or snapshot deltas.

When a target is unlinked, the allocation sharing remains in place. In other words, unlinking will not unshare the allocations. Even after unlinked, termination of a snapshot results in the target owning the snapshot delta. And an updated write to the source track results in the target owning the original track. The target also takes ownership of any shared source tracks if the source is deallocated after unlinking.

This enhanced functionality allows the user to continue to access the target data after unlinking in the same way that previously required full copy targets, but without duplicating the entire back-end data from the point-in-time to the target.

There are a few important aspects to understand about this behavior:

- Nocopy target data is valid only if the target is fully defined before unlinking. The following commands help the user determine if the defined process has completed on the target or targets:

```
symsnapvx -<sourcedevs> -snapshot_name list -linked
symsnapvx -<targetdevs> -snapshot_name verify -linked -defined -by_tgt
```

- The unlink operation does not cause any data on the target to be deallocated. This includes data that was copied during the link operation as well as data that was directly written to the target. Therefore, an unlink operation does not return any free capacity to the SRP.
- If the user wants to deallocate the data on the target or targets after unlinking, they can do so by issuing the `symdev free -all` command to the target volume.

---

**Note:** This is a powerful command that will completely deallocate all data on the volumes that it is issued against. Take extreme care when using this command. As a safety mechanism, the `symdev free -all` command requires the device to be not ready or unmapped.

---

- Similarly, after terminating a snapshot, users may not see as much capacity returned available to the SRP as expected because the previous target that was sharing the snapshot deltas has taken ownership of them.
- Also, if that target happens to be in a separate SRP from the source or snapshot, the data is copied across SRPs as the target takes ownership.
- The underlying behavior exists in earlier versions of HYPERMAX OS, but the data verification testing was not completed to qualify the feature until HYPERMAX OS 5977.810.184.

This means that at earlier versions of HYPERMAX OS, users may observe the behavior described in the previous bullets. But the validity of the data on the nocopy targets after unlinking cannot be guaranteed.

- Regarding performance, this feature should be considered to perform as well as accessing linked nocopy targets.

## E Scheduling snapshots with Unisphere for PowerMax

The method for scheduling recurring snapshots that was introduced in Unisphere for VMAX 8.3 is being deprecated in Unisphere for PowerMax 9.2 in favor of Snapshot Policies.

When upgrading to Unisphere for PowerMax 9.2, snapshot schedules that were created with previous versions of Unisphere continue to run as scheduled. However, the snapshots cannot be modified other than being deleted, and new schedules cannot be created using the legacy method.

Do not assign an SG to snapshot policies and legacy snapshot schedules simultaneously. In general, delete legacy schedules and end their snapshots before assigning the SG to a policy.

However, the legacy schedules and snapshots can remain if required to support business cases as snapshot policies are implemented and approach compliance. Use care to ensure that system resources are not exhausted, including ensuring the automatic expiration date is set on the snapshots.

The following information is for environments that are unable to upgrade to the PowerMaxOS Q3 2020 release.

### E.1 Scheduling snapshots with Unisphere for PowerMax 9.1 and earlier versions

This functionality provides the user with the ability to configure a snapshot to be created at specific times and days of the week. This is configured by adding the snapshot task to the Job List in Unisphere.

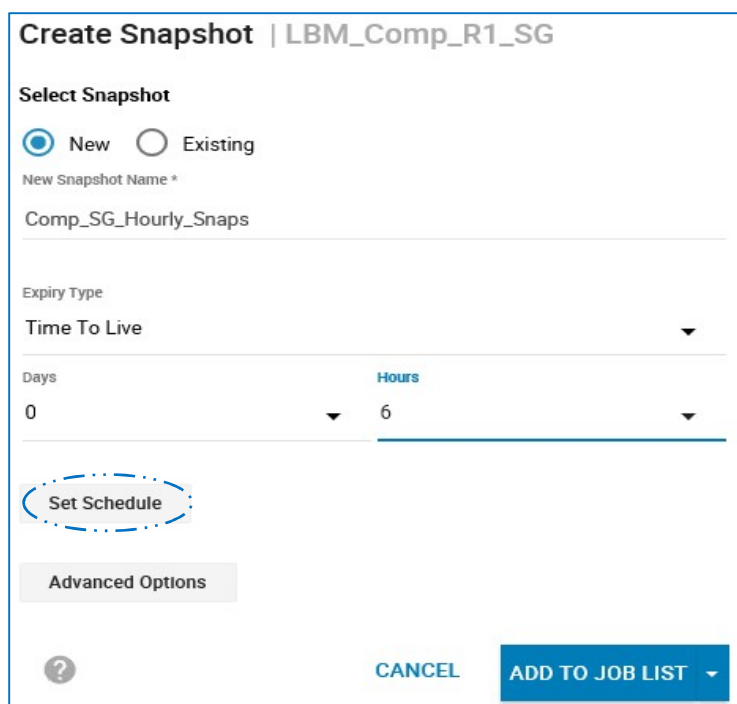
Once the schedule is set up, it runs automatically. To discontinue a scheduled job, delete the entry from the Job List.

Automatic termination of snapshots is performed by the Time-to-Live option as normal, not by the scheduler. Use of the automatic expiration date is strongly recommended for recurring snapshots. The automatic expiration of snapshots that are no longer needed maximizes the efficiency of SnapVX on system resources.

The following examples show snapshot scheduling with Unisphere 9.0.

#### **Process:**

Open the Create Snapshot Wizard from the Protection Dashboard. Specify the snapshot name and expiration days or hours. Then click the Set Schedule button.



**Create Snapshot | LBM\_Comp\_R1\_SG**

**Select Snapshot**

☒ New ☐ Existing

New Snapshot Name \*

Comp\_SG\_Hourly\_Snaps

Expiry Type

Time To Live

Days: 0 Hours: 6

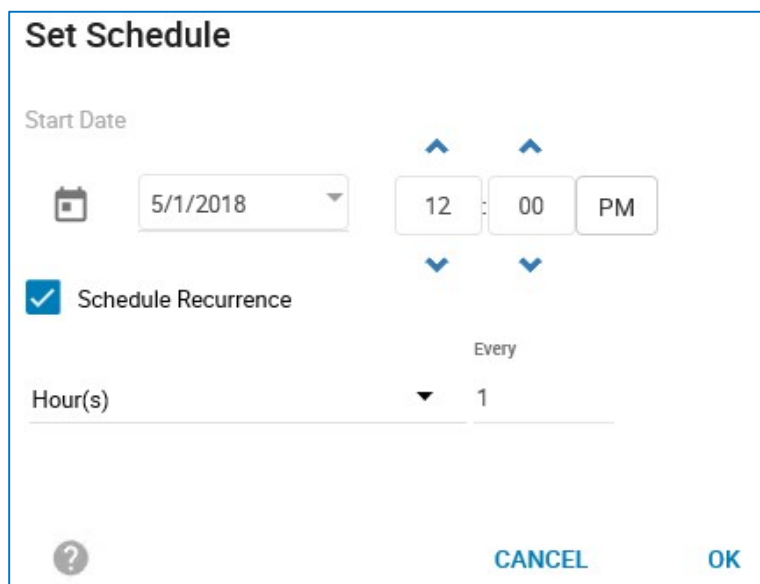
**Set Schedule**

Advanced Options

? CANCEL ADD TO JOB LIST

Figure 28 Unisphere 9.0 Create Snapshot Wizard

In the Set Schedule window, specify the date and time you want the schedule to start and click the Schedule Recurrence.



**Set Schedule**

Start Date

5/1/2018 12:00 PM

☒ Schedule Recurrence

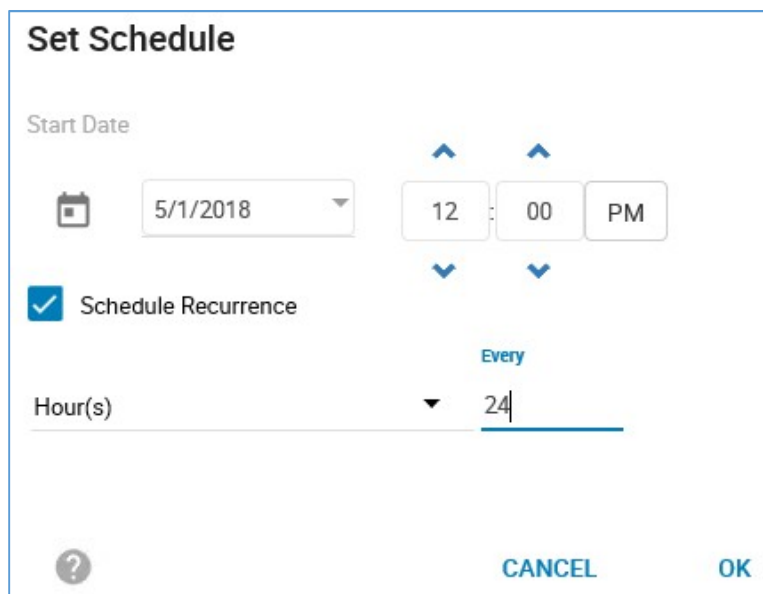
Every

Hour(s): 1

? CANCEL OK

Figure 29 Unisphere 9.0 Set Snapshot Schedule

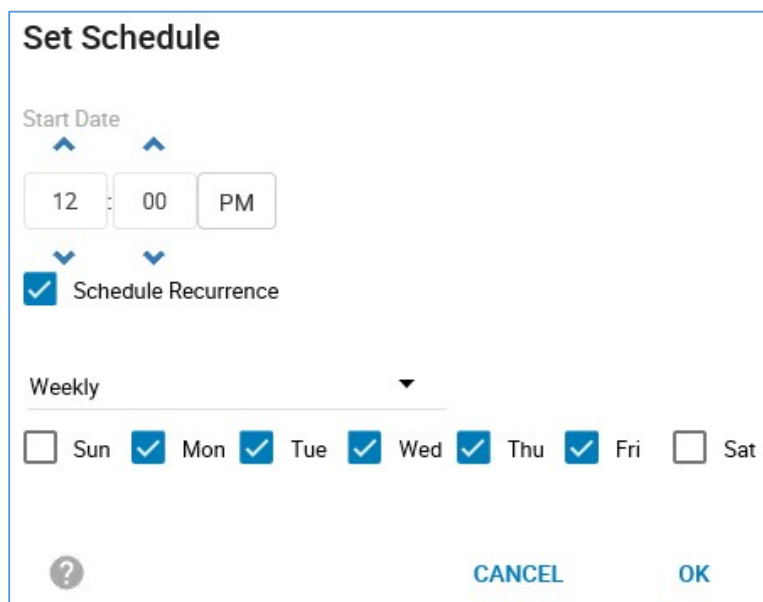
The schedule can be set in either full hours, with the minimum value of 1 hour:



The 'Set Schedule' dialog box in Unisphere 9.0. It features a 'Start Date' section with a calendar icon and a date field set to '5/1/2018'. To the right of the date is a time selector with '12', '00', and 'PM' buttons, flanked by up and down arrows. Below the date, there is a checked checkbox for 'Schedule Recurrence'. Underneath this, the word 'Every' is displayed above a 'Hour(s)' field which contains the value '24'. At the bottom of the dialog are a help icon (question mark), a 'CANCEL' button, and an 'OK' button.

Figure 30 Specifying and hourly schedule in the Unisphere 9.0

Or the schedule can be set for a specific time of day on specific days of the week. Click OK once the schedule is defined.



The 'Set Schedule' dialog box in Unisphere 9.0, configured for a weekly schedule. The 'Start Date' section shows a time of '12:00 PM'. The 'Schedule Recurrence' checkbox is checked. Below it, the 'Weekly' dropdown is selected. A row of checkboxes represents the days of the week: Sun (unchecked), Mon (checked), Tue (checked), Wed (checked), Thu (checked), Fri (checked), and Sat (unchecked). At the bottom are a help icon (question mark), a 'CANCEL' button, and an 'OK' button.

Figure 31 Specifying a daily schedule in the Unisphere 9.0

Click ADD TO JOB LIST in the Create Snapshot Wizard. The job will be added to the Job List and will run at the specified interval.

Figure 32 Unisphere 9.0 Create Snapshot Wizard

The Job List displays basic information about the job.

Job List

ModifyRunScheduleGroup⋮

4 items

<input type="checkbox"/>	Name	Status	Symmetrix ID	Scheduled	Completed	Recurring
<input checked="" type="checkbox"/>	SnapVX create Comp_SG_H...	SCHEDULED	000197600191	Mon Mar 19 2018 12:00:00 GMT...		

Figure 33 Unisphere 9.0 Job List

Double-clicking the job will display additional information.

SnapVX create Comp_SG_Hourly_...	
Status	<span>SCHEDULED (1 task)</span>
Scheduled	Mon Mar 19 2018 12:00:00 GMT-0400
Occurrence	Schedule to Run every Monday, Tuesday, Wednesday, Thursday, Friday
Modified	Sun Mar 18 2018 12:39:01 GMT-0400
Completed	-

Figure 34 Scheduled Snapshot Job Details

The Modify button will allow you to edit the job name and the next run of the job.

**Modify Job | SnapVX create Comp\_SG\_Hourl...**

Name \*

SnapVX create Comp\_SG\_Hourly\_Snaps for Storage Group LBM\_Comp\_R1\_SG

Schedule

3/19/2018 12 : 00 PM ☐ Clear Schedule

Occurrence

Schedule to Run every Monday, Tuesday, Wednesday, Thursday, Friday

Figure 35 Unisphere snapshot scheduler window

#### Notes:

- Be sure to monitor resources used by the snapshots as discussed earlier in this document, such as used or available back-end capacity, Snapshot Deltas and Replication Cache or Replication Data Pointers (RDP), using threshold alerts as applicable.
- As mentioned earlier in the document, proper planning is critical to ensure that the system is configured with enough resources to support your snapshot requirements. As with any other application, as requirements change and grow, resources may need to be added to the system.
- Unscheduled snapshot jobs can be grouped together, and then a single recurring schedule can be applied to that group. This can be helpful in coordinating snapshots for multiple SG.

Be aware that the snapshot for each SG is performed individually and may have slightly differing timestamps. If the same timestamp is required, the user may want to consider creating a separate, larger SG specifically for the snapshots. Unisphere 9.0 provides the user with drag-and-drop capability to choose the order of job execution.

Be sure to remember to set up the day/time schedule after the jobs have been grouped. Previously scheduled snapshot jobs cannot be grouped together.

Job List						
<input type="button" value="Modify"/> <input type="button" value="Run"/> <input type="button" value="Schedule"/> <input type="button" value="Group"/> <span>⋮</span>						
<input type="checkbox"/>	Name	Status	Symmetrix ID	Scheduled	Completed	Recurring
<input checked="" type="checkbox"/>	SnapVX create Comp_Snap_Hourly for Stor...	UNSCHEDULED	000197600191			
<input checked="" type="checkbox"/>	SnapVX create KT_Snap_Hourly for Storag...	UNSCHEDULED	000197600191			

Figure 36 Unisphere job list

**Group Jobs**

Name \*

Hourly Snaps for Comp and KT

Job Order

- ≡ SnapVX create Comp\_Snap\_Hourly for Storage Group LBM\_Comp\_R1\_SG
- ≡ SnapVX create KT\_Snap\_Hourly for Storage Group KT\_Snap\_SG

Schedule

5/1/2018 12 : 00 PM

☒ Schedule Recurrence

Every

Hour(s) 1

< >

? CANCEL OK

Figure 37 Grouping jobs in the Unisphere job list

- If execution of a snapshot at a specific time fails for any reason, an alert is created. The job remains active and attempts to create another snapshot at the next scheduled time.
- The scheduler runs within Unisphere and may be affected if there is a problem with the host where it is installed, or if Unisphere needs to be uninstalled.
- Embedded Management (eManagement) provides high availability. In this case, the schedule continues to run on the secondary instance if the primary fails.

## F System resource planning for capacity and performance

- This section discusses how to determine the effect of snapshot usage on system resources. Determining factors include:
  - Capacity of the applications that will be protected by snapshots
  - Change rate of the applications
    - The amount of source data that is updated during the life of a snapshot
  - Frequency of snapshots
  - Lifetime of the snapshots
  - Number and capacity of linked targets
  - Copy mode of linked targets (full copy or nocopy)
  - Read/write activity to the linked targets

These details help identify the additional capacity required to support the protection requirements of the application when planning a capacity upgrade and planning for new systems. Monitoring resources during normal operation is also critical to ensure available resources are not exhausted.

### F.1 Capacity planning and monitoring

#### Snapshot capacity usage:

A single generation of a snapshot can never exceed the capacity of the source device at the time the snapshot was taken. A single generation of a snapshot could potentially use backend capacity equal to the source device usage at the time the snapshot was taken. However, the efficiency features in PowerMaxOS and HYPERMAX OS, such as thin provisioning, Data Reduction, and shared allocations, greatly reduce the amount of actual snapshot data in the system.

Snapshots can be terminated manually or set to self-expire through the time-to-live (TTL) option. Use of both methods together is the best way to ensure snapshot capacity is released timely. Manually terminating snapshots is recommended if the exact time of termination is an important factor for the environment. Setting an expiration date on every snapshot, including snapshots that will be manually terminated, is the best way to prevent snapshots from being forgotten and continuing to grow and consume system resources. Expiration is especially recommended if any level of automation is being performed to create snapshots, whether via script or Unisphere.

A snapshot that has a target linked cannot be terminated. This applies to manual termination and automatic expiration. Users should query the array for expired snapshots periodically to verify that there are no targets linked preventing termination.

Resources consumed by snapshots that were created by other solutions such as ProtectPoint, RecoverPoint, and AppSync also need to be monitored. Pausing use of those applications without performing proper snapshot cleanup may consume valuable system resources.

When snapshots are terminated, the underlying cleanup process is a lower priority task designed to give priority to other operations in the array, like host I/O for example. This ensures that snapshot termination will not affect performance of the devices involved or performance of the system. The rate at which snapshot capacity will be freed depends available system resources. Therefore, the rate may vary across systems, and may even vary on a single system during different periods of activity.



## F.2 Snapshot capacity monitoring:

Snapshot data is reported in two ways; logical snapshot capacity and physical snapshot capacity.

### F.2.1 Logical snapshot capacity

Logical snapshot capacity is based on delta count for each snapshot and based on the full 128 KB track size. Data that is shared across multiple snapshots will be reported by each snapshot. Therefore, these values cannot be used to determine actual backend snapshot usage. The following Solutions Enabler commands and Unisphere displays report logical snapshot capacity.

The values in the following SnapVX outputs are defined as follows:

**Snapshot Dev Size:** The provisioned size of the source device at the time the snapshot was created.

**Total Deltas:** Data that has been written by the application since the snapshot was created.

**Non-Shared:** Data that is unique to the specific snapshot and not shared by any other snapshots.

The following shows the logical snapshot capacity per snapshot:

```
$ symsnapvx -sid XYZ -sg SNVX_SG list -detail -tb
```

```
Storage Group (SG) Name : SNVX_SG
SG's Symmetrix ID      : 000197600XYZ (Microcode Version: 5978)
```

```
-----
                          Snapshot  Total
Sym          Flags          Dev Size  Deltas Non-Shared
Dev Snapshot Name Gen FLRG TSEB Snapshot Timestamp    (GBs)    (GBs)    (GBs)
-----
0077F hourlysnap    0 .... Wed Sep 05 20:00:00 2018    2.0    0.5    0.0
    hourlysnap    1 .... Wed Sep 05 19:00:00 2018    2.0    0.5    0.0
    hourlysnap    2 .... Wed Sep 05 18:00:00 2018    2.0    0.5    0.1
    hourlysnap    3 .... Wed Sep 05 17:00:00 2018    2.0    0.5    0.1
    hourlysnap    4 .... Wed Sep 05 16:00:00 2018    2.0    0.5    0.3
    hourlysnap    5 .... Wed Sep 05 15:00:00 2018    2.0    0.5    0.5

00780 hourlysnap    0 .... Wed Sep 05 20:00:00 2018    2.0    0.5    0.0
    hourlysnap    1 .... Wed Sep 05 19:00:00 2018    2.0    0.5    0.0
    hourlysnap    2 .... Wed Sep 05 18:00:00 2018    2.0    0.5    0.1
    hourlysnap    3 .... Wed Sep 05 17:00:00 2018    2.0    0.5    0.1
    hourlysnap    4 .... Wed Sep 05 16:00:00 2018    2.0    0.5    0.3
    hourlysnap    5 .... Wed Sep 05 15:00:00 2018    2.0    0.5    0.5

-----
                          6.0    2.0
```

147\_I0Gen > SnapVX Snapshots

Create Modify Restore Link 4 items

<input type="checkbox"/>	Snapshot	Creation Time	Linked	Tracks	Non-Shared ...
<input type="checkbox"/>	James_1_s...	Mon Sep 10 2018 12:01:48 (0)	✓	5464080	5455872
<input type="checkbox"/>	Snap1	Fri Sep 28 2018 09:41:51 (0)	—	5464080	229219
<input type="checkbox"/>	Snap2	Fri Sep 28 2018 11:00:03 (0)	—	5464080	4244974
<input type="checkbox"/>	Snap2	Fri Sep 28 2018 09:44:02 (1)	—	5464080	178339

Figure 38 Logical snapshot capacity per snapshot

The fields in the demand report are defined as follows:

- **Total Snapshot Capacity:** The combined subscribed size of all existing snapshots.
- **Modified %:** Percentage of **Total Snapshot Capacity** written by the application since the snapshots were created.
- **NonShared Snapshot Capacity:** Capacity of the **Modified %** that is not shared by multiple snapshots.
- **Shared Snapshot Capacity:** Capacity of the **Modified %** that is shared by multiple snapshots.

System demand report:

```
$ symcfg -sid XYZ list -demand -detail
```

A R R A Y   C A P A C I T Y   R E P O R T														
	Subscribed Capacity				Snapshot Capacity				Usable Capacity					
	Total	Allc	NonShared	Shared	Total	Mdfy	NonShared	Shared	Total	Used	User	System	Temp	
SymmID	(TB)	(%)	(TB)	(TB)	(TB)	(%)	(TB)	(TB)	(TB)	(%)	(TB)	(TB)	(TB)	
000197600XYZ	192.55	20	37.50	0.65	4.04	28	0.75	0.36	97.65	39	36.82	1.70	0.00	

Verbose system demand report:

```
$ symcfg -sid XYZ list -demand -v
```

```
Symmetrix ID      : 000197600XYZ
Array Usage
User Subscribed (TB) : 192.55
Allocation (%)      : 20
Non Shared (TB)     : 37.50
Shared (TB)         : 0.65
System Subscribed (TB) : 1.70
Snapshot Capacity (TB) : 4.04
Modified (%)        : 28
Non Shared (TB)     : 0.75
Shared (TB)         : 0.36
```

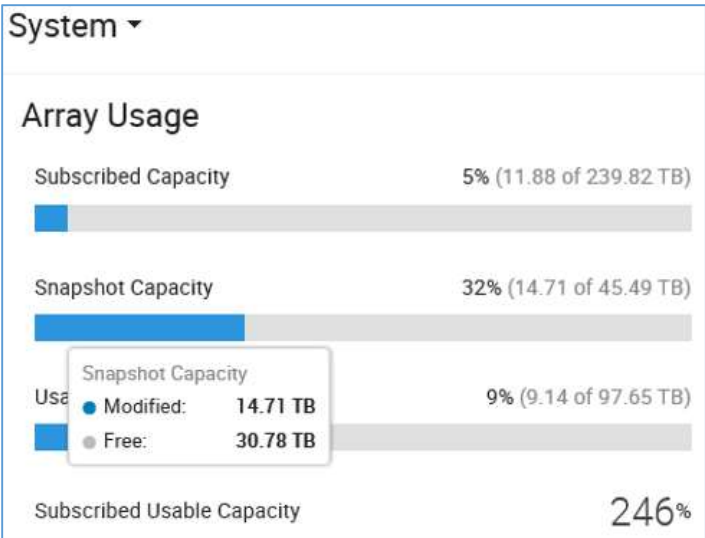


Figure 39 System demand report

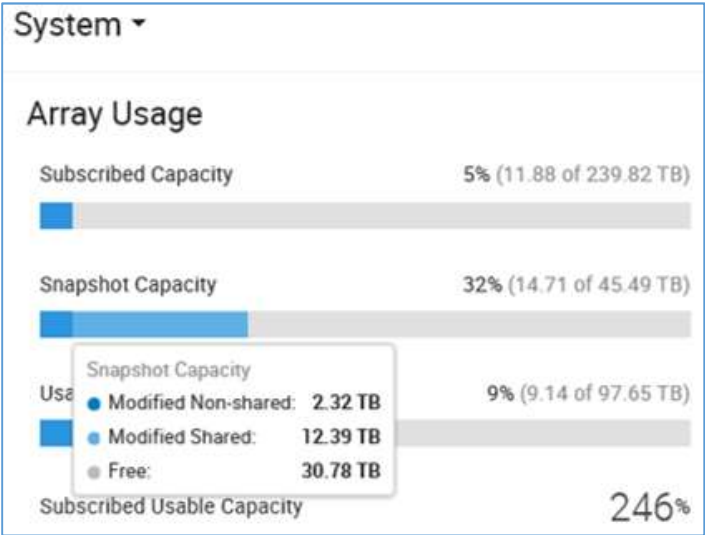


Figure 40 Detailed system demand report

The Efficiency Dashboard in Unisphere provides a graph of snapshot usage as shown in Figure 41:

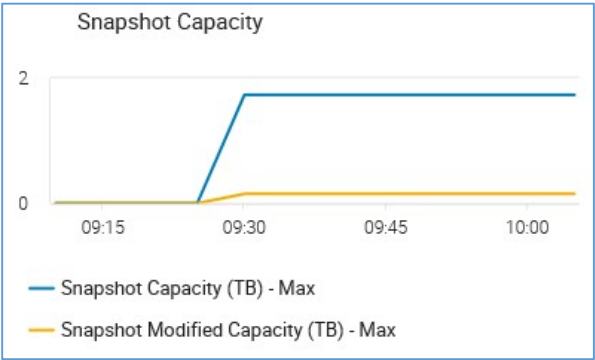


Figure 41 Array Efficiency Dashboard in Unisphere 9.0

## F.2.2 Physical snapshot capacity

Physical snapshot capacity is the actual backend capacity that is snapshot data. Physical snapshot capacity may be significantly less than logical snapshot capacity due to efficiency of features such as shared allocations and Data Reduction.

The storage group demand report displays physical snapshot capacity of each SG in the specified SRP, along with a sum of the combined snapshot usage. The fields are defined as follows:

- **Snapshot Allocated:** Backend capacity that is snapshot data, per SG, based on 128 KB track size.
- **Snapshot Used:** **Snapshot Allocated** after compression has been applied.
- **Snapshot Compression Ratio:** Ratio of Snapshot Allocated to Snapshot Used.

**Note:** The following output only displays Storage Groups that have an explicit Service Level or Workload.

Physical snapshot capacity per SG:

```
$ symcfg -sid XYZ list -demand -srp -type sg -detail -gb
```

STORAGE RESOURCE POOLS

Name : SRP\_1  
Usable Capacity (GB) : 87626.0  
Compression State : Enabled  
Data Reduction Ratio : 2.2:1  
SRDF DSE Allocated (GB) : 0.0

---

SG Name	Subscribed (GB)	Allocated (GB)	Allocated (%)	Snapshot Used (GB)	Snapshot Comp Ratio	Snapshot Allocated (GB)	Snapshot Used (GB)	Snapshot Comp Ratio
Weekly_BU	20.0	8.9	44	8.6	1.0:1	0.0	0.0	2.6:1
Prod_SG	667.0	666.0	99	666.0	-	1878.2	558.6	3.4:1
<not_in_sg>	178127.6	2083.1	1	2083.1	1.0:1	494.3	494.3	1.0:1
Total	178814.6	2757.9	2	2757.7		2372.5	1052.9	

SRP_1 > Storage Group Demand Report							
Export				26 items			
<input type="checkbox"/> Storage Group	Subscription (GB)	Allocated (GB)	Allocated (%)	Used (GB)	Snapshot Allocated (GB)	Snapshot Used (GB)	Snapshot Compression Ratio
<input type="checkbox"/> 147_child_0	50.00	0.00	0%	0.00	0.00	0.00	-
<input type="checkbox"/> 147_IOGen	667.00	666.00	100%	666.00	1,878.16	558.60	3.8:1
<input type="checkbox"/> 147_IOGen_LNK_SG_0...	667.00	666.00	100%	188.80	0.00	0.00	-

Figure 42 Physical snapshot capacity per SG

Physical snapshot capacity per TDEV is available in the `symcfg -sid 191 list -srp -detail -v` command as shown in Figure 48. The output also provides a sum of the combined snapshot usage across the TDEVs in the SRP.

The fields in the output are defined as follows:

- **Snapshot Allocated:** Backend capacity that is snapshot data, per TDEV, based on 128 KB track size.
- **Snapshot Used:** **Snapshot Allocated** after compression has been applied.
- **Snapshot Compression Ratio:** Ratio of Snapshot Allocated to Snapshot Used.

Snapshot usage per TDEV:

		Thin Device			Snapshot		
Sym	Subs	Allocated	Used	Allocated	Used		
Dev	Emul	(GB)	(GB)	(%)	(GB)	(GB)	(GB)
0020D	FBA	10240.0	224.5	2	49.7	1.3	1.2
0020E	FBA	10240.0	214.6	2	41.8	3.6	3.1
0020F	FBA	10240.0	214.5	2	41.8	3.4	2.9
00210	FBA	10240.0	215.6	2	41.8	1.8	1.6
-----							
Total		40960.0	869.2	2	175.1	10.1	8.8

#### Linked target capacity usage aspects:

Linking targets in nocopy mode is the recommended best practice because they provide same performance as full copy linked targets while using less backend resources. Using full copy mode in PowerMax systems with Data Reduction enabled is equal to a nocopy target due to deduplication.

Some use cases require dedicated linked target volumes. For example, when presenting separate copies to multiple users at the same time. Linked targets can also be shared across source volumes in some use cases. For example, if a set of linked targets is used only to verify the validity of snapshot data from a Storage Group, once the data is verified, the targets can be unlinked and then linked to the next Storage Group for the same purpose. Limiting the number of devices in a system whenever possible reduces usage of system resources and simplifies management.

During nocopy link and relink operations, all existing data on the target is returned to the pool. Therefore, it is not recommended to deallocate a target after an unlink operation if the target will be linked again soon. Deallocating the target in this case will not cause any harm, although a device is not available for use while the operation is running. But it is an unnecessary step since the subsequent link operation will deallocate the device anyway.

However, if there are no plans to use the target again soon, then users should consider performing a `free -all` operation on the target to deallocate any exclusive allocations and remove the device from any shared allocations.

---

**Note:** Extreme care must be taken when using the `free -all` command. See page 14 or more information. The `free -all` command is available with the `symdev`, `symmsg`, `symdg`, and `symcg` commands.

---

**Linked target capacity usage monitoring:**

Nocopy linked targets typically have some combination of exclusive and shared data. Shared data is created during the define process. Exclusive data is created by host writes. Operations such as snapshot termination, and deallocation, freeing, or host writes to other targets can cause shared data to become exclusive data.

Solutions Enabler command `symcfg -sid <sid> list -tdev -detail` reports total and exclusive allocations for individual TDEVs. The output includes a flag to indicate the presence of shared tracks.

**Exclusive allocated tracks per TDEV:**

```
$ symcfg -sid XYZ list -tdev -sg SNVX_SG -detail -gb
```

```

      S Y M M E T R I X   T H I N   D E V I C E S
-----
          Pool      Pool Exclusive  Comp
          Bound    Flags  Total Subs  Allocated Allocated Ratio
Sym  Pool Name  ESPT      GBs  (%)   GB  (%)   GBs
-----
0077F -      FS.B    2.00  -    0  0    0.50 1.0:1
      DG2_F_F  -.--      -  -    0.50 25    -  -
00780 -      FS.B    2.00  -    0  0    0.50 1.0:1
      DG2_F_F  -.--      -  -    0.50 25    -  -
00781 -      FS.B    2.00  -    0  0    0.50 1.0:1
      DG2_F_F  -.--      -  -    0.50 25    -  -
00782 -      FS.B    2.00  -    0  0    0.50 1.0:1
      DG2_F_F  -.--      -  -    0.50 25    -  -

Total
Tracks      8.00  -    2.00 25    2.00
Legend:
Flags: (E)mulation : A = AS400, F = FBA, 8 = CKD3380, 9 = CKD3390
      (S)hared Tracks : S = Shared Tracks Present, . = No Shared Tracks
      (P)ersistent Allocs : A = All, S = Some, . = None
      S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,
      D = Deallocating, R = Reclaiming, C = Compressing,
      N = Uncompressing, F = FreeingAll, . = Unbound

```

## F.3 Performance considerations

**Snapshot performance aspects:**

I/O profile and snapshot count are both factors of performance.

A snapshot captures a point-in-time image of a TDEV or TDEVs. The image is preserved track-by-track the first time each track is written after the snapshot was taken. Subsequent updates to a track will not require the preservation process for that snapshot.

SnapVX in PowerMaxOS uses redirect-on-write to preserve data for a snapshot. When an update is sent to data that needs to be preserved, the original data remains in place and becomes snapshot data, while the new data is written to another location in the array. If the update is a full 128 KB track, then the update is

destaged without any other actions. If the update is less than a full track, then the remainder of the track is read into cache and then the updated version is destaged.

#### **Linked target performance aspects:**

It is not necessary to wait for define (nocopy targets) or copy (full copy targets) to complete before accessing the targets. However, there may be some overhead incurred for the first I/O to an undefined target track. This is like accessing a full copy target or clone while copy is in progress. The effect is typically negligible unless there will be heavy I/O to the devices immediately after linking.

Therefore, in some use cases it may be best to wait for the define process to progress before accessing the targets. Even in those use cases, it may not be necessary to wait for the define process to complete. Simply allowing the define process to progress before accessing the targets will reduce or eliminate the chance of accessing undefined target data.

The same rules and considerations apply for snapshot restore operations; Data is available immediately after the restore command completes, there is no need to wait for the copy to complete, and there may be some overhead when accessing tracks that have not completed.

A fully defined nocopy linked target provides the same performance as a fully copied link target. However, some users still choose to use full copy linked targets. The copy rate will dynamically adjust according to system activity to give highest priority to host I/O. Copy times can vary across systems and even within a single system during different periods of activity.

Copy Quality of Service (QoS) allows the user to reduce copy priority below its default value and allow other operations to consume system resources. For more information about QoS, see the *Dell EMC VMAX3 and VMAX All Flash Quality of Service Controls for Multitenant Environments White Paper*.

#### **Benchmark testing:**

Best practice for benchmark testing during SnapVX operations is to use a production-like workload. This will give the most accurate assessment of the performance that can be expected when the application is in production.

The system used in the testing must also be properly sized for the required capacity and snapshot activity, along with any other features that may affect system performance.

Tests using IO levels and change rates beyond that expected in production will not produce realistic results.

## **F.4 Dell EMC ProtectPoint**

Note: Dell EMC ProtectPoint has been renamed to Storage Direct and is included in PowerProtect, Data Protection Suite for Apps, and Data Protection Suite Enterprise Edition software.

Dell EMC ProtectPoint technology is a data protection solution that integrates primary storage, PowerMax, VMAX All Flash and XtremIO, with Data Domain protection storage, to accelerate backup and recovery, help organizations meet application protection SLAs, and minimize backup impact on applications, all while reducing cost and complexity. Using ProtectPoint, every backup is a full backup with the storage and bandwidth efficiency of an incremental backup, which simplifies recovery.

### Configuration requirements and best practices:

The information in this section focuses on PowerMax and VMAX All Flash aspects of ProtectPoint implementations. See the following documents for more information about ProtectPoint and Data Domain requirements:

- Dell EMC ProtectPoint Primary and Protection Storage Configuration Guide
- Dell EMC ProtectPoint Solutions Guide

The minimum configuration requirements for a ProtectPoint implementation are as follows:

- One DX pair on the PowerMax or VMAX
  - Both DX emulations on the same Brick
  - Two DX ports per DX emulation
- Two ports on the Data Domain
- One switch between the PowerMax or VMAX and the Data Domain

The following configuration options will increase performance and availability:

- One DX pair per Brick on the PowerMax or VMAX
  - Two DX ports per DX emulation
- Four ports on the Data Domain
- Two switches between the PowerMax or VMAX and the Data Domain
- Minimum of 30 production LUNs per backup

### Copy throughput:

Copy throughput will be higher when there is a significant change rate and many copy streams available. This applies to backup and rollback processes.

Backups consisting of a minimum of 30 production devices will make optimal use of system resources and maximize the SnapVX throughput internal to the PowerMax or VMAX. Backing up fewer than 30 devices is supported but may not necessarily take less time, regardless of device capacity, because the process will not use as many resources. Backing up only a few devices at a time, as is often seen in proof-of-concept testing, will not maximize use of resources and therefore is not a true assessment of what may be seen in production.

The time for a backup to complete is affected by subscribed capacity of the TDEVs involved, not only allocated capacity. This is because SnapVX cycles through the entire subscribed size of the TDEV to determine which data needs to copy, and this is what is viewed from the SnapVX outputs that display the copy progress. Therefore, backup of a TDEV that is not heavily overprovisioned will likely take less time to complete than backup of a TDEV of the same allocated capacity but with a larger subscribed capacity.

Throughput of an incremental backup may seem less than that of a full backup, although the duration will be shorter. For instance, a differential copy of 5% changed data (in relation to the subscribed capacity) may typically reflect ~40% the throughput of a full copy but would complete nearly 80% faster.

Because SnapVX cycles through the entire provisioned size of the TDEV, the copy rate as viewed from the SnapVX point-of-view may not reflect the copy rate that is seen between the PowerMax or VMAX system and the Data Domain. Only actual data will be transferred from the PowerMax / VMAX to the Data Domain.



Overall system activity influences the ProtectPoint backup and restore processes. Likewise, ProtectPoint backup and restore operations have an influence on the system. If the ProtectPoint backup needs to be slowed down beyond the level determined by default scheduling algorithms, the QoS settings of the encapsulated Data Domain backup devices (vDisks) can be modified. This will extend the ProtectPoint backup time while lessening the effect on overall system performance.

Multiple Data Domains can share DX ports. The DX resources will be shared if backups are run simultaneously.

When possible, schedule backups when the volumes being backed up are under light or moderate write loads. Snapshot overhead is most significant when the snapshot is newly established and nearly all writes need special handling. If the backup window must coincide with a busy time for the system, additional CPU resources may be required to support the increased load and should be considered when the system is being sized.

## G Technical support and resources

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

[Storage and data protection technical white papers and videos](#) provide expertise that helps to ensure customer success with Dell EMC storage and data protection products.

### G.1 Related resources

- Frequently Asked Questions about TimeFinder SnapVX on VMAX All Flash Storage Arrays
- Dell EMC Solutions Enabler TimeFinder SnapVX CLI User Guide
- Dell EMC Solutions Enabler TimeFinder Family CLI User Guide
- Dell EMC PowerMax Family Product Guide
- Dell EMC VMAX All Flash Product Guide
- Data Reduction with Dell EMC PowerMax White Paper
- VMAX Non-Disruptive Migration Configuration and Best Practices Technical Notes
- Role Based Access Controls (RBAC) Technical Overview & Enhancements
- Data Protector for Z Systems (zDP) Essentials
- Dell EMC ProtectPoint Primary and Protection Storage Configuration Guide
- Dell EMC ProtectPoint Solutions Guide
- Dell EMC VMAX3 and VMAX All Flash Quality of Service Controls for Multitenant Environments White Paper