

Frequently Asked Questions about TimeFinder SnapVX on VMAX All Flash Storage Arrays

This document addresses the most commonly asked questions about EMC® TimeFinder® SnapVX™ software in Open System Environments. Many of the questions come from users who are familiar with other snap or clone software products and are trying to draw comparisons between them and the SnapVX software.

Questions and Answers

Session limits

Question: What is the maximum number of sessions allowed in a single array?

Answer: A VMAX® All Flash Array can have up to 64,000 thin devices (TDEVs). Each TDEV can have 256 snapshots. So theoretically, a single array could have over 16 million snapshots.

Practically, the resources in the array, such as memory, backend capacity, and so on, determine the system limit. When an array is ordered it is critical to accurately specify details such as how many snapshots will be created, how many will be linked, what performance is expected, and so on. This information will help ensure that the system is configured with sufficient resources to meet the business requirements.

Snapshot resource usage

Question: Does a targetless snapshot consume any resources?

Answer: Simply creating a targetless snapshot does not consume any system resources. The snapshot resource usage grows as the source device is updated. As new writes come into the source device, the original backend data becomes snapshot deltas that preserve the point in time, which causes the backend usage to grow. Cache and flash resources are also required to store and manage snapshot data and metadata, such as source device, snapshot, timestamp, and so on. Snapshots will begin to fail when these resources are exhausted.

Snap of a snap

Question: Does SnapVX software have an option to create a snap of a snap?

Answer: SnapVX software does not have an explicit concept called *snap of a snap*. However, the software certainly offers functionality to meet business requirements. The key here is to first understand the goal and then use the SnapVX software to meet that goal.

Targetless snapshots preserve only point-in-time data as source tracks are updated. These snapshots contain only differential data, are not full copies, and cannot be accessed by a host. A snapshot can be presented to a host by linking a

target LUN, and multiple targets can be linked to a single snapshot. Therefore, if the goal is to provide another version of the same point-in-time copy, this goal can be accomplished by simply linking another target volume to the existing targetless snapshot. There is no need to duplicate a snapshot to make a point-in-time of the original snapshot available to another LUN.

Example:

```
symsnapvx -sid -sg -name establish
symsnapvx -sid -sg -snapshot_name link -lnsg <target_sg_1>
symsnapvx -sid -sg -snapshot_name link -lnsg <target_sg_2>
```

If a linked target LUN is modified by host writes, for example, during testing or development purposes, and the goal is to create a new copy that contains those changes, a snapshot can be taken of the linked target. Targets can then be linked to that snapshot and presented to a host to allow continuation of the testing, development, or other activity. With SnapVX software, this action is referred to as *cascading* rather than snap of a snap because the linked target, not the targetless snapshot, is replicated. A snapshot can be taken of a linked target after the target is defined. The following diagram is one example of cascading with the SnapVX software.



Many of the reasons for taking a snap of a snap, for example, when the point-in-time data of a snapshot or clone must be preserved before additional development, are unnecessary with SnapVX software because the integrity of the snapshot is always preserved. When data is written to a linked target, the changes are applied only to the target LUN and do not affect the point-in-time data of the snapshot.

Note: When cascading with SnapVX software, the cascade can be broken at any level by unlinking a target after it is defined. Unlinking a target will not affect the validity of the data on any upstream or downstream snapshot or linked target.

Question: What does Defined mean?

Answer: When a target is linked, the defining process scans the entire target device. The defining scan changes the pointers of each track on the target to share the appropriate backend data for the point-in-time snapshot, which will typically be a combination of source tracks and snapshot changes.

The defining process and shared allocations allow the target to access the backend data directly and without redirection. This capability allows SnapVX to offer unlimited levels of cascading, improved performance of accessing linked targets compared to traditional snapshots and nocopy clones, and access to point-in-time data by nocopy targets after they are unlinked.

The defining scan is a lower priority, background operation that defers to higher priority operations. Therefore, the time for a defining process to complete depends on a number of factors, such as the size of the involved volumes, allocated portion of those volumes, and other activity across the array.

Users do not need to wait for defining to complete before they access data on the linked target. However, whether or not a target is fully defined can have the following implications on other operations:

- If a nocopy linked target is not yet fully defined, it cannot be unlinked, and a snapshot cannot be taken of the linked target.
- If a full-copy linked target is not yet fully copied, it cannot be unlinked, and a snapshot cannot be taken of the linked target.

Question: Can a snapshot be taken of multiple Storage Groups at once?

Answer: The short answer is no. Only a single Storage Group can be specified in a single operation. However, several methods are available that allow you to take a snapshot of devices across multiple Storage Groups.

The Unisphere[®] for VMAX Job List Group operation can be used to schedule snapshot jobs to be run within the same time frame. However, this option is not advised if application consistency is required across the snapshots for the multiple Storage Groups. Each snapshot job is performed independently. So, even though each Storage Group snapshot is consistent, consistency across the snapshots is not guaranteed.

The snapshots for the Storage Groups are also likely to have different timestamps. However, if the operation spans multiple operations, even a dependent, write-consistent snapshot can have slightly different timestamps when it is viewed across the entire set of devices.

If consistency across the snapshots is required, users can choose from the following options:

- Implement Cascaded Storage Groups and select the option to take a snapshot of the Parent Storage Group, the individual Child Storage Groups, or both.
- Create a separate Storage Group that contains the devices to be snapped.
 - Devices can be in multiple Storage Groups as long as only one of the Storage Groups is assigned to a Storage Resource Pool (SRP) or displays in a Masking View.
- Create a Consistency Group, Device Groups, or Device File that contain the devices to be snapped.

Question: A snapshot of a Storage-Group was taken. Is it possible to issue a restore on only a subset of the devices?

Answer: Yes. This result can be achieved through some of the previously described methods, like creating a separate Storage Group, Device File, or Device Group, or by using the `-devs` option to specify a device range directly in the command line.

Example:

Some databases consist of archive and log volumes. A user might want to take snapshots of all, but if a restore is required, the user might not want to restore the log volumes. Users can choose to implement cascaded Storage Groups and create

separate child Storage Groups for archive and log volumes. A snapshot across the parent Storage Group will take a snapshot of all child Storage Groups, but the user can issue a restore to a specified child Storage Group if needed.

Copy process and data availability

Question: After linking or relinking a full-copy linked target or restoring a snapshot to the source LUN, must I wait for the copy process to complete before accessing data?

Answer: No, the data is available immediately after the link, relink, or restore operational command completes. Any data that is accessed on the target before it is copied will be accessed through the source device. TimeFinder has provided this functionality for many years, long before the SnapVX software was available.

Question: Why does the copy process seem stuck or take a while to begin?

Answer: The value that the Solutions Enabler and Unisphere report is actually the sum of three underlying values, which is capped at the size of the device and will not decrease until the sum is less than the size of the device. A common observation is that progress is being made across all the values; however, the user interface does not reflect the progress until the sum is less than the size of the volume.

More detailed information can be found in Appendix C of the *Dell EMC HYPERMAX OS TimeFinder Local Replication Technical Note*:

<https://www.emc.com/collateral/technical-documentation/h13697-emc-vmx3-local-replication.pdf>

Host Unmount and Rescan

Question: Do I need to unmount the target before the link and relink operations?

Answer: Yes. The target will be toggled to the not ready state and then back to the ready state during the link and relink operations and will become momentarily unavailable. So the target should be unmounted before issuing either the link or relink command and then remounted after the operation completes. This action also ensures that the host operating system does not cache any file system data.

Question: Do I need to rescan the target after the link and relink operations?

Answer: There is no EMC TimeFinder requirement to rescan after a link or unlink operation. Typically, the rescan requirement is dependent on the specific host operating system or if devices were added to the set, for example, the Storage Group.

Incremental Operations

Question: What is the process to incrementally update a snapshot?

Answer: There is no process to incrementally update a SnapVX targetless snapshot. Create a new snapshot and then terminate the old snapshot, or use the Automatic Expiration Date capability when you create snapshots to eliminate the need to manually terminate snapshots.

Example:

```
symsnapvx -sid -sg -name establish -ttl <delta | absolute>
```

Question: How can a linked target be incrementally updated from a new snapshot?

Answer: This update is done through the relink command, which performs an incremental/differential operation. In the case of full-copy linked targets, only differences between the snapshots are copied to the linked target. You can use the Automatic Expiration Date capability to eliminate the need to manually terminate snapshots that are no longer needed.

Example:

Initial snapshot created and a target linked:

```
symsnapvx -sid -sg -name establish -ttl
symsnapvx -sid -sg -snapshot_name link -lmsg
```

At every update iteration, a new snapshot is created, and target relinked:

```
symsnapvx -sid -sg -name establish -ttl
symsnapvx -sid -sg -snapshot_name relink -lmsg
```

Note: If the same snapshot name is used for each snapshot, the most recent snapshot will be generation 0, which might need to be specified with the relink command.

Question: When relinking a linked target that is also a Symmetrix Remote Data Facility (SRDF) R1 device, is the subsequent SRDF copy full or incremental?

Answer: Incremental. Because the relink operation is differential, only changes will be propagated across the SRDF link. The linked target must be in full-copy mode and the `-copy` parameter must be specified with both the link and relink commands.

Note: During the RDF copy, the data on the secondary (R2) site is not in a dependent-write, consistent state, so if you need to always have a recovery PiT at the R2 site, you should take a snapshot of the R2 before you link and relink the R1.

Linked Targets and Backend Capacity Usage

Question: Will the data on a nocopy linked target be returned to the SRP after the target is unlinked?

Answer: No, unlinking will not cause data to be deallocated on the target, and free capacity will not be returned to the SRP.

Question: How can linked target capacity be returned to the SRP?

Answer: After an unlink process, a target can be freed by using the `symdev` command or through Unisphere by right-clicking the device and selecting Start Free, the same way that any other TDEV would be deallocated.

Nocopy targets will also be deallocated during link and relink operations. This applies to any device being used as a nocopy target regardless if it was previously a target or not.

Snapshot Retention Periods

Question: Are there any guidelines about snapshot retention periods?

Answer: Snapshot retention periods depend on business requirements. You should consider when a snapshot is too old and out-of-date from which to restore. After a snapshot has aged to a point where restore is no longer a reasonable option, you should probably terminate it. However, the business might have regulatory restrictions that require snapshots to be retained for specific periods of time.

It is important to understand that as snapshots age, they typically have larger delta sets, and so they use more cache and disk resources. To free up resources, whenever possible, terminate snapshots after they are not needed.

The Automatic Expiration Date (Time-to-Live) functionality in SnapVX is recommended to effectively manage resources. Expiration dates can be set when the snapshot is created, as well as set, modified, or removed on existing snapshots.

Snapshots are automatically terminated without user intervention when the expiration date is reached, as long as no targets are linked to the snapshot or the snapshot does not have a restore session. If targets are linked, the snapshot terminates when the targets are unlinked. If the snapshot is in a restore operation, the snapshot terminates when the restore session is terminated.

Note: A background process searches for and terminates expired snapshots. Expired snapshots are terminated after this background process finds the snapshot, which might not occur immediately. This process typically takes a few minutes to find an expired snapshot. During this time, it is common to see snapshots reported as being `EXPIRED` in Solutions Enabler and Unisphere for VMAX outputs.

Terminating Snapshots

Question: Is there any impact on the host when snapshots are terminated?

Answer: No, resources are freed in the background with no host impact.

Snapshot termination and the subsequent cleanup of snapshot deltas, such as deallocating and freeing the SRP and cache or Replication Data Pointer (RDP) resources, are background tasks that are designed to defer to higher-priority tasks such as host I/O.

The time for the cleanup process to complete depends on the size of the snapshots and the overall activity in the array. Users need to plan accordingly; specifically, they need to allow enough time between terminating large snapshots, terminating a large number of snapshots, or both and creating new snapshots.

Note: The background cleanup process also needs to complete before a device can be deleted or expanded.

Emulation modes (Clone, Mirror, and VP Snap)

Question: Do the previous TimeFinder modes work with HYPERMAX OS?

Answer: Solutions Enabler 8 and HYPERMAX OS support TimeFinder Clone, Mirror, and VP Snap. The commands are the same as in the previous generation of the VMAX family. However, HYPERMAX OS emulates these sessions on top of SnapVX. Therefore, some aspects of the behavior, such as copy rate and copy process monitoring, will be somewhat different. This difference is described in the *Dell EMC HYPERMAX OS TimeFinder Local Replication Technical Note*:

<https://www.emc.com/collateral/technical-documentation/h13697-emc-vmax3-local-replication.pdf>

The emulation modes allow users to continue to use their existing processes and scripts on a new VMAX All Flash array while they learn about the SnapVX software. However, the emulation modes cannot take advantage of the key benefits of SnapVX software, such as storage group operations, automatic expiration dates, and increased scalability and versatility. Also, future investment will be dedicated to the SnapVX product. Feature enhancements and interoperability with other features might be incorporated into the SnapVX product but not into the emulation modes. Users who opt to stay with the emulation modes must understand the implications of this choice. Transitioning to SnapVX software sooner might be easier than doing so later.

References

Question: Where can I find more information on SnapVX?

Answer: Dell EMC HYPERMAX OS TimeFinder Local Replication Technical Note:
<https://www.emc.com/collateral/technical-documentation/h13697-emc-vmax3-local-replication.pdf>

VMAX All Flash – SnapVX Lightboard:
<https://www.youtube.com/watch?v=hXrk87ruCxo>

https://youtu.be/ma6ouQApn_4
