



DELL EMC UNITY: SNAPSHOTS AND THIN CLONES

A Detailed Review

ABSTRACT

This white paper describes Dell EMC™ Unity Snapshots and Thin Clones for Dell EMC Unity storage systems and includes information about their configuration management, and other functionality.

January, 2019

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 © 2017 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [01/19] [White Paper] [H15089.6]

Dell EMC believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

TABLE OF CONTENTS

EXECUTIVE SUMMARY5
Audience5

TERMINOLOGY5

SNAPSHOT OVERVIEW7
Redirect on Write.....7
Consistency Groups7

SNAPSHOT OPERATIONS.....8
Create8
Restore.....9
Copy.....10
Attach10
Share.....13
Refresh.....13
Delete.....14
Replicate14

SNAPSHOT PROPERTIES 15

SNAPSHOT RETENTION POLICY 16
Pool Automatic Deletion17
Expiration18
No Automatic Deletion.....18
Invalidation18
Snapshot Schedules18
Snapshot Shipping19

SNAPSHOT INTEROPERABILITY 20
Thick Storage Objects20
Compression/Data reduction/advanced data deduplication20
Dell EMC Unity File System21
Replication21
Host I/O Limits.....23
Cloud Tiering Appliance (CTA).....24

THIN CLONE OVERVIEW 24

THIN CLONE HIERARCHY 24

THIN CLONE OPERATIONS..... 25

 Create 25

 Refresh..... 26

 Edit Properties..... 27

 Delete 28

THIN CLONE INTEROPERABILITY 28

 Consistency Group..... 28

 Replication 29

CONCLUSION 30

REFERENCES..... 31

EXECUTIVE SUMMARY

As data becomes increasingly important to organizations of all types, these organizations continually strive to find the safest and most effective ways to protect their data. While many methods of data protection exist, one of the simplest and most effective is local snapshots, which allows recovery of data by rolling back to an older point-in-time or copying select data from the snapshot. Snapshots continue to be an essential data protection mechanism used across a wide variety of industries and use cases to preserve the most important mission-critical production data, sometimes in conjunction with other data protection technologies.

Dell EMC Unity storage systems introduce a radically simple and completely unified approach to snapshots. As a unified storage system, Dell EMC Unity offers snapshot capabilities for both block and file storage resources that use the same workflows, operations, and architecture. Redirect-on-write (ROW) technology ensures pool space is used optimally and reduces the management burden by never requiring administrators to designate protection space. Whether snapshots being are created, managed, or deleted, Dell EMC Unity provides a unified look and feel for administrators to manage their local data protection environments from a variety of interfaces including Unisphere GUI, CLI, or RESTful API.

Dell EMC Unity provides full control for administrators to schedule and manage snapshots, both automatically and manually. With various automatic deletion settings, administrators never have to worry about snapshots consuming too much pool space or being retained for an unnecessary amount of time. With flexible management options and the ability to interoperate with most other Dell EMC Unity features, snapshots offer the local data protection administrators can trust to ensure their production data is kept safe at all times..

Dell EMC Unity extends the snapshot capabilities in order to provide Copy Data Management, through the use of the Thin Clones feature. Available beginning in Dell EMC Unity OE version 4.2, Thin Clones use the same underlying pointer-based technology that snapshots use to provide a method for managing multiple copies of Block storage resources. Thin Clones support many data services, which engineers and developers can leverage in their testing development environments. When users create a Thin Clone, despite some differences in the settings, it acts as a regular LUN and is listed on the LUNs page. Similar to snapshots, users can create, manage and destroy Thin Clones through Unisphere, CLI, and the RESTful API.

AUDIENCE

This white paper is intended for Dell EMC customers, partners, and employees who are interested in understanding Dell EMC Unity snapshot and Thin Clone functionality, operations, and management.

TERMINOLOGY

Application Consistency – Ensures all incoming I/O for a given application is quiesced while a snapshot is taken. Application Consistency makes sure that the application can restart from the point in time that a snapshot was taken. This strategy differs from Crash Consistency because I/O is quiesced at the host, instead of being quiesced at the array.

Base LUN – The parent LUN, Consistency Group or VMware VMFS Datastores for a set of derivative snapshots and Thin Clones.

Block Storage Resources – LUNs, Consistency Groups and VMware VMFS Datastores.

Server Message Block (SMB) – An access protocol that allows users to access files and folders from Windows hosts that are located on a network. User authentication is maintained through Active Directory and file access is determined by directory access controls.

Crash Consistent Snapshot – A point-in-time copy of data. A crash-consistent snapshot addresses only the data level; it is unaware of any underlying applications.

File Storage Resources – File Systems (NFS, SMB) and VMware NFS Datastores.

Logical Unit Number (LUN) – A logical unit of block storage that is created in a pool.

Consistency Group – A collection of LUNs that are grouped together. Snapshot operations on a Consistency Group affect all the LUNs contained in the group, providing ease of management and crash consistency if the LUNs are dependent on each other.

Network File System (NFS) – An access protocol that allows users to access files and folders from Linux/UNIX hosts that are located on a network.

Network-Attached Storage (NAS) – File-based storage for a wide range of clients and applications that access storage over network connections. Protocol-specific file systems are located and managed on the storage system, which transfers data to hosts over TCP/IP using either the SMB or NFS file sharing protocol.

Redirect on Write (ROW) – The technology behind Dell EMC Unity snapshots. After a snapshot is taken, new writes to the primary storage are redirected and written to a new location within the storage pool.

Share – A named, mountable instance of file-level storage, that is accessible through a file system or VMware NFS Datastore. Each share is accessible through the protocol (NFS or SMB) that is defined for the file system where it resides.

Snapshot – A point-in-time copy of data that is stored on the storage system.

Storage Processor (SP) – A hardware component that provides the processing resources to perform storage operations such as creating, managing, and monitoring storage resources.

Storage Pool – A group of disk drives that is used for configuring LUNs, Datastores, and File Systems. A drive can be a member of only one pool.

Storage Resource – An addressable and configurable storage instance that is associated with a specific quantity of storage. LUNs, File Systems, and VMware Datastores constitute storage resources.

Thin Clone – A read-write copy of a Block level storage resource (LUN, LUNs within Consistency Group or VMFS Datastores).

Unisphere – A web-based management environment that is used to create storage resources, configure and schedule protection for stored data, and manage and monitor other storage operations.

Unisphere CLI (UEMCLI) – The command line interface for managing Dell EMC Unity storage systems.

SNAPSHOT OVERVIEW

Dell EMC Unity snapshots allow administrators to use the same architecture, operations, and workflows to uniformly manage block and file snapshots. Redirect-on-write technology is used for all storage resources, eliminating the need for a save volume or for copying snapped data away before overwriting data, thereby improving efficiency.

Because snapshots use the same technology for block and file, managing snapshots is very similar between resource types, which all offer the ability to create, delete, copy and restore snapshots from a common interface. Block and file snapshots vary only slightly with regard to making writable snapshots available to hosts: block snapshots are attached and file snapshots are used to create shares. All snapshots can be configured with one of several automatic deletion policies, which ensures that snapshots do not consume an unwanted amount of pool space or are not retained for longer than needed. With configurable pool thresholds and expiration dates, administrators have full control of all snapshots within their storage environment.

While most discussion of snapshot management in this paper refers to the Unisphere GUI, snapshots are also fully supported through the UEMCLI and RESTful API. So whether administrators take manual snapshots through Unisphere, leverage the customizable snapshot schedules, or create advanced data protection scripts, they can fully manage their storage environments using whichever method they prefer. This ability leads to a powerful, flexible foundation for managing data protection regardless of the complexity of the use case or environment.

REDIRECT ON WRITE

The technology used for snapshots of both block and file resources in Dell EMC Unity is called redirect on write (ROW). With redirect on write, new writes to snapped storage resources or their snapshots are redirected to a new location in the same pool, and pointers are updated to point to the new location. This method contrasts to copy-on-write snapshots that, in order to preserve the snapped data, copy blocks to a designated protection space prior to being overwritten. With redirect on write, there is no need to define separate protection space or to copy data to a separate location before writing to snapped blocks because new data is simply written to a new location in the same pool. This technology is illustrated in Figure 1.

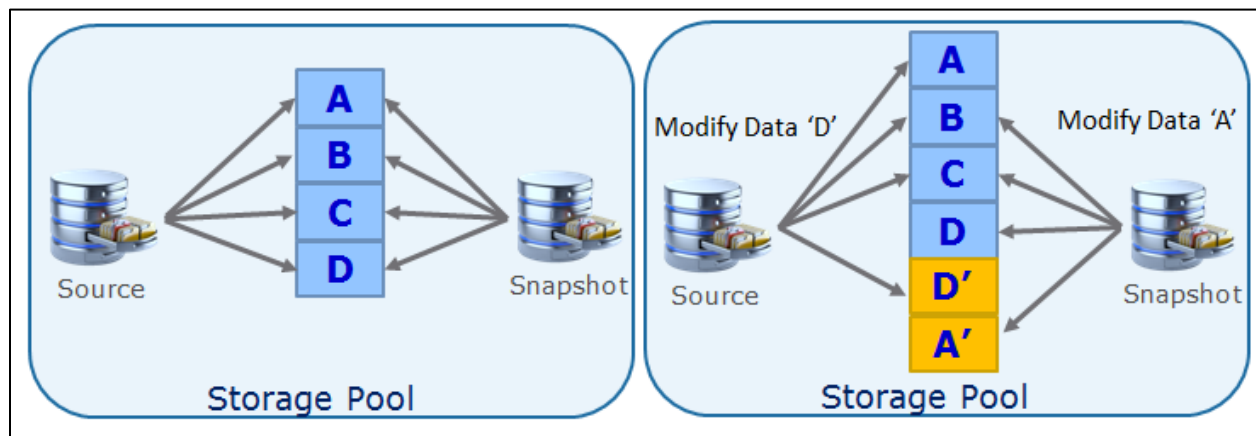


Figure 1. Redirect on Write

In this example, a source LUN contains four blocks of data: A, B, C, and D. A snapshot is taken of this LUN, which also points to A, B, C and D. When a host attempts to overwrite the D block with new data, D', the new data is written to a new location in the same storage pool and the LUN updates its pointer to D'. This is because the data in block D is also associated with the snapshot, so it needs to be preserved rather than overwritten in case the snapshot needs to be restored in the future.

Similarly, if a host has write access to the snapshot, attempting to overwrite block A when writing to the snapshot results in the new data, A', being redirected to a new location in the pool and block A being preserved. This is because block A is still associated with the source LUN, so it cannot be overwritten. However, if the source LUN host later attempts to write new data to block A, it is overwritten without need for redirection because block A is no longer associated with or needed for the snapshot. This functionality also applies to file systems in the same way at a block level.

CONSISTENCY GROUPS

In Dell EMC Unity, there are several types of block resources. Two of these, LUNs and VMware VMFS Datastores, are individual resources that are provisioned one at a time. A third type of block resource, a consistency group, allows administrators to provision multiple LUNs within a crash-consistent group, which allows those LUNs to be protected uniformly at the same point-in-time. Snapshots

are taken at the consistency group level, which automatically protects each individual LUN within the consistency group at a single point-in-time. This achieves crash-consistency across all LUNs in the group, which is useful when the LUNs compose a larger multi-LUN application, which is required to be crash-consistent itself. Consider a database that is composed of several LUNs, which contain both database data and log information that must reflect the same point-in-time across the entire application. Because the data in each LUN is dependent on the others, and the database does not make sense if all do not represent the same point-in-time, this is a good use case for a consistency group. Creating these LUNs as part of a consistency group ensures that all LUNs are always protected and restored to the same point-in-time so the database will function properly after recovery. As an example, Figure 2 shows a database that is composed of several LUNs contained within a consistency group. Two snapshots have been taken, one at 8AM and one at 12PM. Each individual consistency group LUN is also protected at the exact same point-in-time that is associated with the snapshot to ensure crash consistency in the event that either snapshot is restored.




Figure 2. Consistency Group

An important characteristic of consistency groups is that snapshots must cover all consistency group LUNs in order to ensure the entire group can always be restored to a single crash consistent point-in-time. As a result, some operations are not permitted for consistency groups with existing snapshots. In particular:

- A LUN with one or more existing snapshots cannot be added to a consistency group. Any snapshots must be deleted before the LUN can be added to ensure that the snapshot state of all LUNs in the consistency group is consistent.
- No LUNs can be added to a consistency group that has one or more snapshots. Such an addition would create a consistency group that could not be fully restored because the newly added LUN would not be protected at the same point in time as the rest of the consistency group. Therefore, snapshots of the consistency group must be deleted before adding any new LUNs to the group.
- Similarly, no LUNs can be removed from a consistency group with one or more snapshots. This removal would make it impossible to restore the full consistency group back to the snapped point-in-time, because the LUN would no longer be a part of the group. All consistency group snapshots must be deleted prior to removing any LUNs from the LUN group.

SNAPSHOT OPERATIONS

Several operations can be performed on Dell EMC Unity snapshots. These operations are uniform between block and file snapshots, with the exception of making these snapshots accessible to hosts. In Unisphere, block and file snapshots are created and managed in the same manner, and their allocation can be viewed in a single location for each individual pool to help administrators assess which resource snapshots are consuming the most storage pool space. In addition to Unisphere, all operations may also be performed through UEMCLI and the RESTful API.

CREATE

A snapshot creates a new point-in-time view of a block or file resource that is associated with the point-in-time at which the snapshot was taken. Immediately after creation, a snapshot consumes almost no space in the pool because it still shares all of its blocks with the primary block or file resource. However, as new data is written to the parent resource, redirects occur and the snapshot begins to consume pool space that is not associated with the current production version of the parent resource. After a snapshot is created, the administrator can perform other snapshot operations, such as restoring, copying, attaching/detaching, or deleting.

When administrators create a snapshot, they specify several attributes. These include the snapshot Name, Description, and Retention Policy which determines the snapshot's automatic deletion behavior.

When administrators create a file snapshot, they can select an additional option that allows them to create either a Read-only (.ckpt) or Read/ Write (share) type of snapshot. This choice has implications for restoring and accessing file system snapshots. The major difference is essentially Read-only snapshots allow the restoration of individual files through the .ckpt or Previous Versions method in addition to full file system level restore. Read-write snapshots, on the other hand, can be made available to hosts as separate shares. Both snapshot types can be restored at the file system level. Figure 3 shows the File snapshot creation.

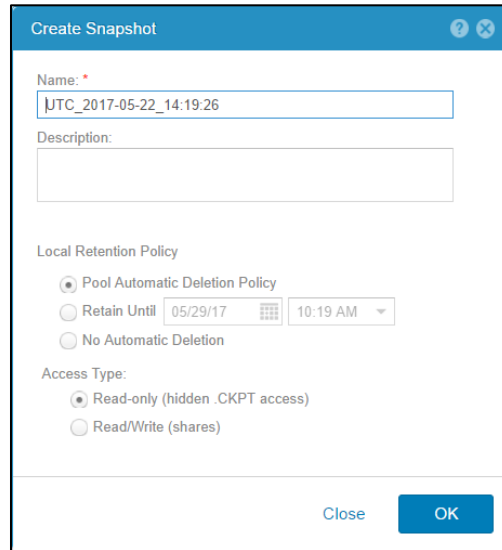


Figure 3. Create Snapshot

RESTORE

Restoring a snapshot reverts its associated parent resource to the point-in-time at which the snapshot was taken. For LUNs, file systems, and VMware Datastores, this is the individual resource. For consistency groups, all LUNs in the group are restored to the single point-in-time that is associated with the snapshot. While restoring the LUNs, the system automatically creates a backup snapshot associated with the current point-in-time which allows administrators to reverse this operation and return the resource to the point-in-time immediately before the snapshot was restored. Before restoring a snapshot, Dell EMC recommends that you quiesce the application that is running on the production host, and then flush the host cache to prevent data corruption during the restore operation.

For Read-only (.ckpt) file snapshots, it is possible to restore individual files from a snapshot without restoring the entire file system. This type of restoration is done through the Windows **Previous Versions** tab on recent versions of Windows, or by appending .ckpt to the share location on UNIX or older versions of Windows. To view the Previous Versions of a file system in Windows, right-click the file system, select **Properties**, and then view the Previous Version tab. Administrators can then copy individual files associated with the snapshot back into the current production file system. This functionality is not available for read-write (share) file system snapshots which are generally taken to provide host access to point-in-time file system data. Figure 4 shows the Previous Versions on Windows host.

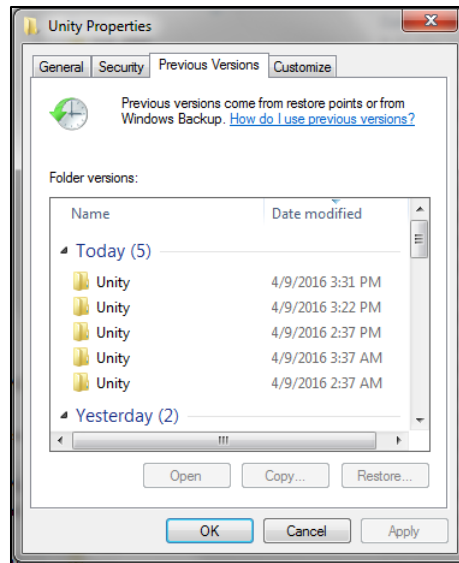


Figure 4. Previous Versions

COPY

The copy operation allows existing snapshots to be duplicated as they are at the time of the copy operation. For example, after administrators makes some changes to a writable snapshot, they might duplicate it to preserve the view of the snapshot at a certain point-in-time. This operation can be thought of as taking a snapshot of a snapshot, and it enables administrators to take hierarchical snapshots that preserve data over multiple iterations of changes. When administrators copy a snapshot, the resulting copy is always created as a Read-Write share eligible snapshot, regardless of the type of snapshot that administrator copied.

ATTACH

Attaching a snapshot applies only to block snapshots. The operation makes the snapshot available for host I/O. A host must have configured Snapshot access or LUN and Snapshot access in order to access the snapshot. When a snapshot is attached, Unisphere provides administrators the option of creating a copy of the snapshot to preserve the current snapshot view for future use. If needed, the copy can be used later to discard any changes that were made to the snapshot while it was attached. Similarly, administrators can detach an attached snapshot in order to remove host access to the snapshot. Figure 5 shows the snapshot attach process.

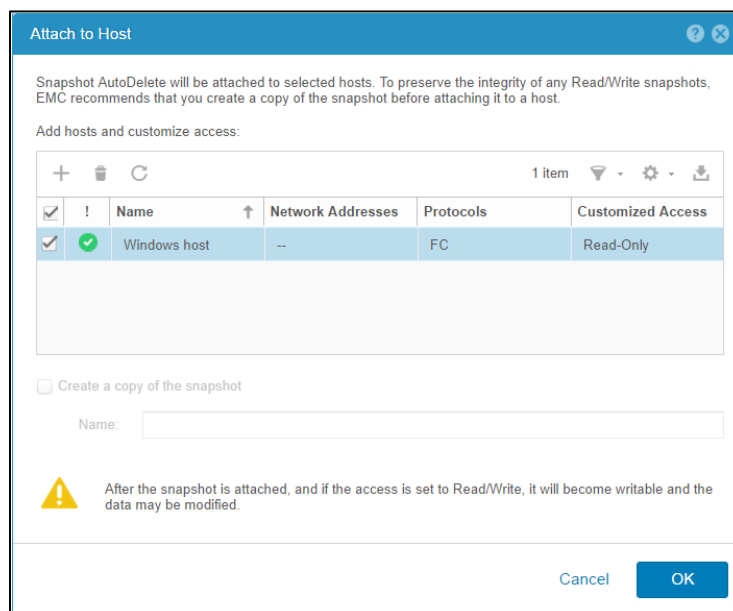


Figure 5. Attaching a Snapshot

In Dell EMC Unity OE version 4.1 and later, multiple snapshots for a particular Block resource can be attached to one or more hosts at the same time – Figure 6 illustrates the multiple snapshot mount. Previously, only a single snapshot could be mounted to one or more

hosts. This change allows snapshots, potentially with different information for a common source, to be mounted to one or more hosts at the same time.




Figure 6. Multiple Snapshots Mount

Also, snapshots can be attached with Read-Only or Read/Write access – shown in Figure 7 shows the **Select Host Access** screen and the **Access Type** option where the Read-Only or Read/Write option is specified. When the Read/Write Access Type is specified, an option is available to copy the snapshot, which preserves the current point-in-time view of the data.

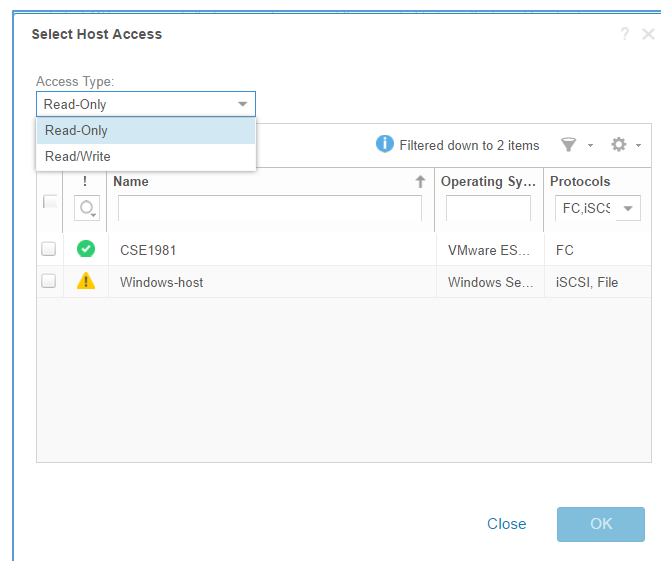


Figure 7. Host Access to Snapshot

Once a snapshot is mounted to a host, the snapshot can be accessed through a mount point as if it is a local storage drive. The snapshot is available through the mount point until the snapshot is detached from the host. Once the snapshot is detached from the host, the mount point is also removed. Reattaching the same or attaching a new snapshot to the host may or may not show up as the same mount point, which is not ideal in cases where a permanent mount point is needed.

Through Unisphere CLI, it is possible to create a permanent mount point, also known as a Snapshot Mount Point, on the host as a placeholder for snapshots of a given LUN. This reserves a mount point for the LUN's snapshot, even if other block storage objects are created and mounted to the same host afterwards. The most common use case for this functionality is scripting. For example, if a user creates a backup script that contains a variable pointing to the Snapshot Mount Point, they do not need to change the variable even when a snapshot is detached and a new snapshot is attached to the same host.

To create a Snapshot Mount Point, use the `uemcli /stor/prov/luns/lun -id <value> set [-snapHosts <value>]` command. The `-snapHosts` option allows the user to specify a comma separated list of hosts to create the Snapshot Mount Point on. When utilizing the `uemcli /prot/snap -id <value> attach -type default` command in combination with the configured snapshot host(s), there is no need to specify the host to mount the snapshot to. This option, unlike the snapshot attach options in Unisphere, only allows a single snapshot to be mounted to each configured host at a time. To utilize the same multiple snapshot attach capabilities as Unisphere, utilize the `uemcli /prot/snap -id <value> attach -type dynamic` command.

Figure 8 below shows an example of a host with a Snapshot Mount Point. This placeholder can be used to attach any snapshot of the LUN to the same mount point. Assuming **LUN1**'s ID is **sv_6**, and **Host**'s ID is **H1**, the following command can be used to create the Snapshot Mount Point on the host:

```
uemcli /stor/prov/luns/lun -id sv_6 set -snapHosts H1
```




Figure 8. Snapshots Hosts Attach

Once you create the Snapshot Mount Point on your host(s), you can attach a snapshot to it. To attach a snapshot to the mount point, use the `uemcli /prot/snap -id <type> attach -type default` command. In the example below, the **Snap1** snapshot is being attached to the host using the mount point. After running the command below, the host now has access to the **Snap1** snapshot for LUN1, as shown in Figure 9.

```
uemcli /prot/snap -id Snap1 attach -type default
```

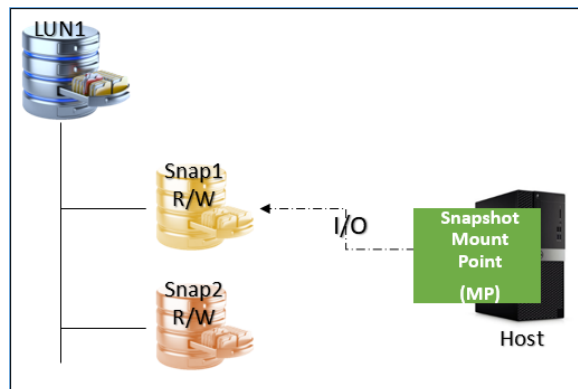


Figure 9. Default Snapshot Mount Point – Mount Snap1

If the **Snap1** snapshot is detached and the **Snap2** snapshot is attached, the same mount point is used, as shown in Figure 10. Note that attached snapshots using this method are only shown via Unisphere CLI. The Unisphere GUI does not show attached snapshots using the “default” attach method.




Figure 10. Default Snapshot Mount Point – Mount Snap2

For more details about how to use Unisphere CLI, please reference *Dell EMC Unity Family Unisphere Command Line Interface User Guide* on *Dell EMC Online Support*.

SHARE

In order to access a writable file snapshot, the user must create a new share for this file system using the snapshot. This operation is accomplished through the same share creation wizard that is used to create regular shares, except Snapshot for File System is selected as the source during configuration. The administrator then provides a share name and path for the snapshot share similar to a regular file system share. Figure 11 shows the wizard page the administrator uses to create the new snapshot share. After the share is created, users can access the point-in-time data that is associated with the snapshot by navigating to the appropriate share.

The screenshot shows the 'Create an SMB Share' wizard. On the left, a sidebar lists 'File System', 'General', 'Advanced', 'Summary', and 'Results'. The 'File System' tab is active. The main area is titled 'Select a source for the new share'. It shows 'File System: *' with a dropdown menu set to 'ProductA'. Below this, there are two radio button options: 'File System "ProductA"' and 'Snapshot for File System "ProductA"'. The second option is selected. Below the radio buttons is a dropdown menu showing 'UTC_2017-05-22_18:30:25'. At the bottom right, there are 'Cancel' and 'Next' buttons.

Figure 11. Creating a Snapshot Share

REFRESH

In Dell EMC Unity OE version 4.1 and later, the snapshot Refresh option is available for Block storage resources level only. This option allows a snapshot to be updated with the current data within the parent LUN. A snapshot refresh operation deletes the current contents of the snapshot, including any changes to the snapshot, and updates the snapshot with the current contents of the source device. Before the refresh operation, the user can run a copy operation to preserve the current contents of the snapshot. If a snapshot is refreshed for a consistency group LUN, the point-in-time image is refreshed for all LUNs in the consistency group. Figure 12 shows from where the user can do **Refresh** of a snapshot.

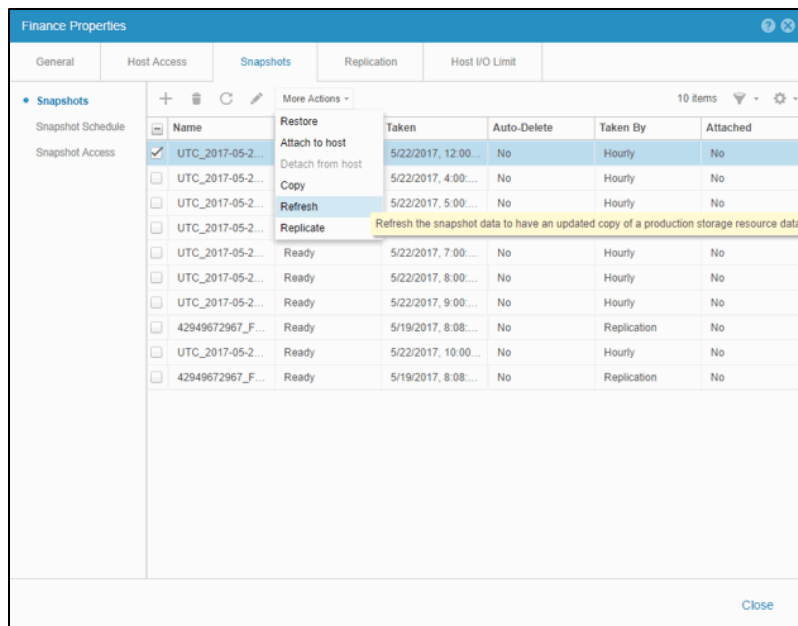


Figure 12. Retention Policies

DELETE

Deleting a snapshot removes the snapshot from the storage system, which might return space to the storage pool, depending on whether the snapshot contains any unique data. For example, a snapshot that was just created would return minimal space to the pool because almost all of its blocks would be shared with its parent resource, and therefore no space can be reclaimed. When a snapshot delete is initiated, the snapshot name is immediately prepended with the word `Destroying_<timestamp>` so the name can be reused immediately. The snapshot might take a short time to delete and continues to count against the snapshot limit until completely deleted. The length of time it takes for the delete operation to complete depends on the amount of unique data that is associated with the snapshot, and the total system load at the time of deletion.

Deleting snapshots might have a performance impact under some circumstances. These include when many snapshots are deleted at the same time, which can occur as a result of the automatic delete functionality. Because of this potential performance impact, Dell EMC recommends to stagger snapshot deletions whenever possible through snapshot scheduling or scripting delays. Performance might also be affected when the user deletes many snapshots on pools with a large number of NL-SAS drives or while FAST VP relocations are also running.

REPLICATE

With Dell EMC Unity OE version 4.2 and later, Asynchronous and Synchronous Replication supports the replication of read-only file and block snapshots to either a local or a remote site along with the source storage resource's data. Both scheduled snapshots and user created snapshots can be replicated. Snapshot replication can be enabled on all resources that support Asynchronous Replication, including: LUNs, Thin Clones, Consistency Groups, File Systems, and VMware VMFS and NFS Datastores. The Replicate option can be used to manually set a snapshot for replication and for it to include in the next RPO synchronization. For more information on snapshot replication, review the white paper titled *Dell EMC Unity: Replication Technologies* on Dell EMC Online Support. Figure 13 shows the Replicate option in the snapshot tab.

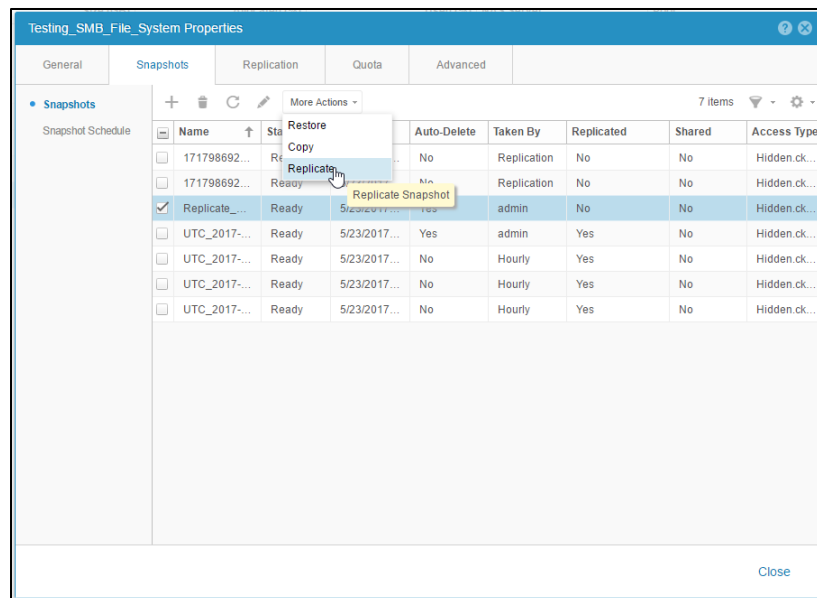


Figure 13. Replicate a Snapshot

SNAPSHOT PROPERTIES

As administrators create, modify, and access snapshots in Dell EMC Unity, these snapshots' properties can be changed to reflect a number of different states. To understand these states, we will first define the properties associated with all snapshots, and then cover the possible states along with the conditions which might contribute to each state. Figure 14 shows some of these properties, some of which are set upon snapshot creation. These can be viewed on the **Snapshots** tab of the storage resource **Properties** page in Unisphere.

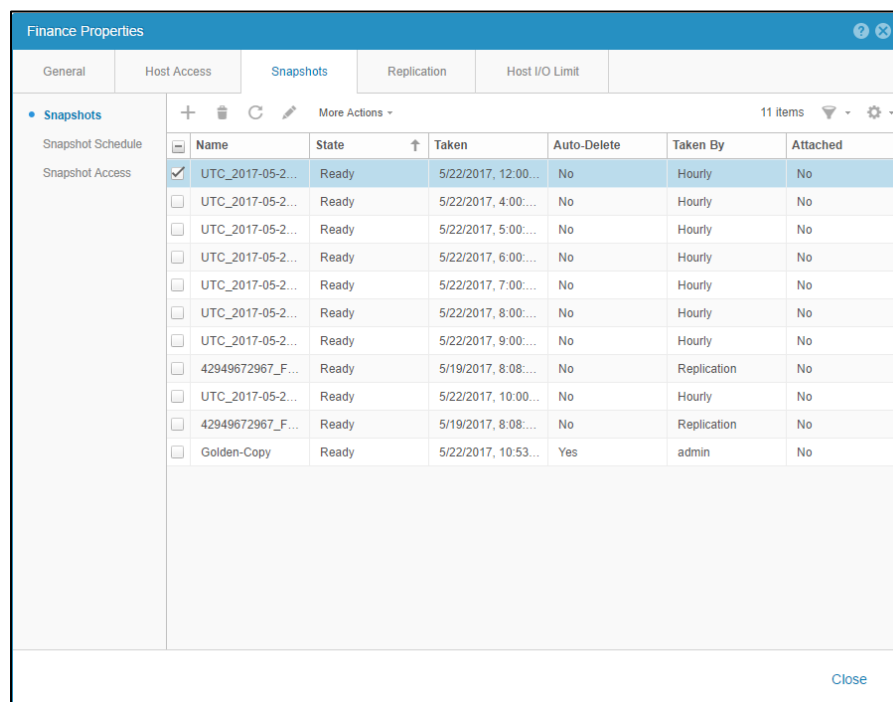


Figure 14. Snapshot Properties

Here are the available snapshot properties:

- **Name** – By default, the snapshot name is the timestamp that the snapshot was taken, in UTC_YYYY-MM-DD_HH:MM:SS format. The default name can be changed to a user defined name during or after snapshot creation. Snapshot names must be unique across all snapshots on the system.

- **State** – This property represents the current status of the snapshot. The possible state values include Ready, Initializing, Offline, and Destroying. There is no Expired state for snapshots which have passed the expiration date. These snapshots will simply be deleted when the next hourly deletion process begins as described in more detail in the Expiration section of this paper.
- **CLI ID** – The identification number of the snapshot, to be used in CLI commands which require a particular snapshot to be referenced.
- **Taken** – This represents the date and time at which the snapshot was taken.
- **Auto-Delete** – This indicates whether the snapshot is configured for pool space threshold based automatic deletion. Possible values are Yes or No. Snapshots configured for expiration based deletion display an auto-delete property of No.
- **Expiration Time** – This property is only applicable for snapshots configured for expiration based deletion, and displays the time at which the snapshot expires. Note that the actual deletion may occur up to one hour after this time, as described later in the Expiration section of this paper.
- **Taken By** – This displays the administrative user that created the snapshot.
- **Source** – This property displays the parent resource of which the snapshot was taken. When taking a first level snapshot or a copy of an unattached/unshared snapshot, the source resource is set to the source LUN or file system. When taking a copy of an attached/shared LUN/file system snapshot, the source is that snapshot. If a source 1st level snapshot is deleted or detached later, the source of the 2nd level snapshot (copy) becomes the source LUN.
- **Modified** – This represents whether or not the snapshot has been changed from its original state. This property defaults to No when a snapshot is initially taken, but is changed to Yes for writable snapshots when they are attached or shared. Read only file system snapshots can never be modified, and as a result, this will always show as No.
- **Last Writable Time** – This displays the last time a snapshot was able to be written to, either because it was attached (block) or shared (file). Snapshots that have never been writable will not display any value.
- **Attached** – This property is applicable to block snapshots only, and indicates whether the snapshot is currently attached. Multiple snapshots for a block storage resource may be attached at the same time.
- **Shared** – This property is applicable to file snapshots only, and indicates whether the snapshot is currently available to clients as a share. The default value is No, and changes to Yes when the read-write file system snapshot is used to create shares. For read-only file system snapshots, this property is always set to No.
- **Async Replicated** – This attribute states the replication state for the snapshot. Possible values can be **No** for a snapshot that is not marked for replication, **Pending** for a snapshot that is marked for replication but is awaiting transfer, **Yes** for a snapshot that has successfully been transferred to the disaster recovery resource, and **Failed** for a snapshot that failed to be replicated. Note here, only Read-Only snapshots can be replicated. Read-Write snapshots cannot be replicated.
- **Sync Replicated** – This attribute states the replication state for the snapshot. Possible values can be **No** for a snapshot that is not marked for replication, **Yes** for a snapshot that has successfully been transferred to the disaster recovery resource, and **Failed** for a snapshot that failed to be replicated. Note here, only Read-Only snapshots can be replicated. Read-Write snapshots cannot be replicated.

SNAPSHOT RETENTION POLICY

Because snapshots leverage the storage pool for protection space, Dell EMC Unity offers several automatic deletion options to control the amount of space that snapshots consume. Automatic deletion can be caused by the amount of space that snapshots are taking from a storage pool or by a snapshot's retention date expiring. Automatic deletion can be disabled individually when taking a snapshot. An automatic deletion method, or retention policy, is a required parameter for each snapshot created, and can be modified after creation. Figure 15 shows the **Retention Policy** of snapshot for a block resource storage.

Create Snapshot

Name: *
JTC_2017-05-22_15:09:19

Description:

Local Retention Policy

☒ Pool Automatic Deletion Policy

☐ Retain Until 05/29/17 11:09 AM

☐ No Automatic Deletion

Close OK

Figure 15. Retention Policies

POOL AUTOMATIC DELETION

When you create a snapshot, the default automatic deletion option is Pool Automatic Deletion Policy, which deletes the oldest snapshots as necessary depending on the amount of space snapshots consume from the pool and/or the total pool space consumed. Thresholds are set on a pool level and only snapshots set to Pool Automatic Deletion Policy on this pool are deleted when the thresholds are crossed.

The following two types of thresholds exist per pool, and they can be enabled for disabled independently of one another. Illustrated on Figure 16:

- Total pool consumption threshold – Enabled by default, start deleting at 95% full until 85% full
- Snapshot pool consumption threshold – Disabled by default, start deleting at 25% consumed until 20% consumed

Flash-Pool:1 Properties

General Drives Usage FAST VP Snapshot Settings

Automatically delete oldest snapshots

☒ Total pool consumption

Start deleting when pool space reaches: * 95 % full

Continue deleting until pool space reaches: * 85 % full

☐ Snapshot pool consumption

Start deleting when pool space reaches: * 25 % full

Continue deleting until pool space reaches: * 20 % full

Automatic Deletion State: Idle

Close Apply

Figure 16. Pool Properties Snapshot Settings

Each threshold has a configurable high and low watermark setting. In the case of the total pool consumption setting, when the total amount of used space in the storage pool passes the total pool consumption high watermark, the oldest snapshots residing on that pool with retention policy set to pool automatic deletion begin being deleted. This continues until the low watermark is passed. The snapshot pool consumption setting works similarly if enabled, except that the total pool space consumed specifically by snapshots must exceed the high watermark in order for automatic deletion to be initiated. Note that any snapshots with time based or disabled automatic deletion settings are not subject to automatic deletion.

EXPIRATION

Snapshots can be configured for time based expiration by specifying a particular time at which the snapshot is automatically deleted. When snapshots expire, they are marked for deletion but might not be deleted immediately. An hourly task is responsible for deleting any expired snapshots, so any expired snapshots might remain on the system for up to an hour. The state of expired snapshots continue to display as Ready until the task begins the actual deletion operation, then the state updates to Deleting. Snapshots that are configured for time based expiration are not eligible for automatic deletion when pool or snapshot space thresholds are crossed. Table 1 has some details about the snapshot retention timeframe:

Table 1. Snapshot Retention

	CLI		GUI		REST	
	Manual	Scheduled	Manual	Scheduled	Manual	Scheduled
LOCAL SNAPSHOTS	1 year	1 year	5 years	4 weeks	100 years	No limit
REMOTE SNAPSHOTS	5 years	255 weeks	5 years	255 weeks	5 years	255 weeks

NO AUTOMATIC DELETION

If no automatic deletion of either pool-based or time-based type is desired, snapshots can be designated with No Automatic Deletion policy. In this case, the snapshot is not affected by pool threshold settings or expiration dates, and it will continue to exist until it is manually deleted or invalidated. Unless the user modifies the individual snapshot properties after creation, the user cannot configure the **No Automatic Deletion** option for scheduled snapshots. This setting prevents a large number of snapshots from being created and incrementally filling the pool without the administrator's knowledge, which could lead to an unnecessarily filling of the storage pool and triggering snapshot invalidation.

INVALIDATION

While the Dell EMC Unity snapshot retention policies aim to reduce the likelihood that an oversubscribed pool becomes entirely full as the result of snapshots, this condition can occur in rare cases when snapshots exist that are not configured for automatic deletion or are retained for too long. Because snapshots configured for time based automatic deletion or no automatic deletion are not subject to automatic deletion on the basis of the pool or snapshot space thresholds, it is possible that an overprovisioned pool can become entirely filled while at the same time, advertising available space to its clients. In this case, either the host I/O must be failed, or snapshots must be invalidated to make room for new writes.

The snapshot invalidation behavior depends on the resource type and configuration. For block, this behavior is different for Thin LUNs and Thick LUNs. For file, this behavior depends on the `-poolFullPolicy` setting, which can be configured using the UEMCLI `/stor/prov/fs` command. This provides the administrator with the option to configure each individual file system to either delete all associated snapshots (including system replication snapshots) or fail host I/O to that resource. The default option is `failWrites`. The behavior for each resource type and configuration is described in Table 2.

Table 2. Invalidation Behavior

TYPE	BEHAVIOR
THICK LUN	Complete host I/O & Invalidate Snapshots
THIN LUN	Fail host I/O & Preserve Snapshots
FILE SYSTEM (DELETEALLSNAPS)	Complete host I/O & Invalidate Snapshots
FILE SYSTEM (FAILWRITES)	Fail host I/O & Preserve Snapshots

In addition to snapshot invalidation, file resources also include low space handling. When the pool's free space goes below the "start" threshold (4GB), the file system and its read/write snapshots are mounted as read-only. If the pool's free space goes above the "stop" threshold (20GB), the file systems that changed to read-only return back to read/write.

SNAPSHOT SCHEDULES

In addition to being taken manually, snapshots can be scheduled to be taken automatically at specific times or frequencies. These schedules can be applied to both file and block resources, and a single schedule can be applied to many individual resources. By

default, three snapshot schedules exist in Unisphere and are titled Default Protection, Protection with shorter retention, and Protection with longer retention, with each taking daily snapshots and retaining them for various lengths of time. When creating custom snapshot schedules, the options are similar to those available for taking a manual snapshot. Two scheduling options are available - Every number of hours and Daily/Weekly, of which one or both can be used. When configuring snapshots for every number of hours, snapshots are scheduled to occur from every 1 to 24 hours, where daily/weekly snapshot scheduling allows administrators to specify certain times on specific days for snapshots to occur. A retention policy must also be set for each selected option, which can be set to pool automatic deletion or time-based retention. Figure 17 shows the Snapshot Schedule page.

DASHBOARD

SYSTEM

System View

Performance

Service

STORAGE

Pools

Block

File

VMware

ACCESS

Snapshot Schedule

4 items

<input type="checkbox"/>	Name	Type	In Use
<input type="checkbox"/>	Protection with shorter retention	System Defined	No
<input type="checkbox"/>	Default Protection	System Defined	No
<input type="checkbox"/>	Protection with longer retention	System Defined	No
<input type="checkbox"/>	NewSchedule	User Defined	Yes

Figure 17. Snapshot Schedules

A new snapshot schedule can be created ahead of time on the Snapshot Schedules page, or when a new storage resource is created from within the creation wizard as an alternative to selecting an existing schedule. When creating a schedule in this manner, the schedule is still added to the system-wide list of snapshot schedules, which allow it to be reused for other storage resources in the future. Snapshot schedules, except for system-defined schedules, can be modified or deleted after creation. However, note that snapshot schedules cannot be deleted while they are associated with a storage resource. Figure 18 shows the creation of the snapshot schedule.

Create Schedule

Name: *

Schedule01

☐ Synchronize snapshot schedule to remote system

Snapshot Frequency

☐ Every 6 hours

Retention Policy

☐ Pool Automatic Deletion Policy

☒ Retain for 7 Hours

☒ Daily / Weekly

☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☒ Sat ☒ Sun

Snapshot time: 12:00 AM

Retention Policy

☐ Pool Automatic Deletion Policy

☒ Retain for 7 Days

Note: Times are displayed in Local Time (UTC -04:00) in 12-hour format.

Cancel

Create

Figure 18. Creating a Snapshot Schedule

SNAPSHOT SHIPPING

Beginning with Dell EMC Unity OE version 4.2, users can enable snapshot replication for scheduled snapshots when users create the replication session. It can also be enabled or disabled at any time throughout the lifetime of the replication session from the **Replication Session** tab of the **Storage Resources Properties** page and then selecting the appropriate option. The scheduled and user snapshots can be replicated with a remote retention policy that is longer than the source snapshot. For example, the retention policy for the source snapshots might be 2 hours, while the retention policy for the replicated snapshot can be set for 1 week. This can be configured through creating replication session and **Snapshot Schedule** setting that can be found on the **Snapshots** tab in **LUN/File System properties**

page. Figure 19 shows the replication session details in Unisphere. Review the white paper titled *Dell EMC Unity: Replication Technologies* on Dell EMC Online Support.

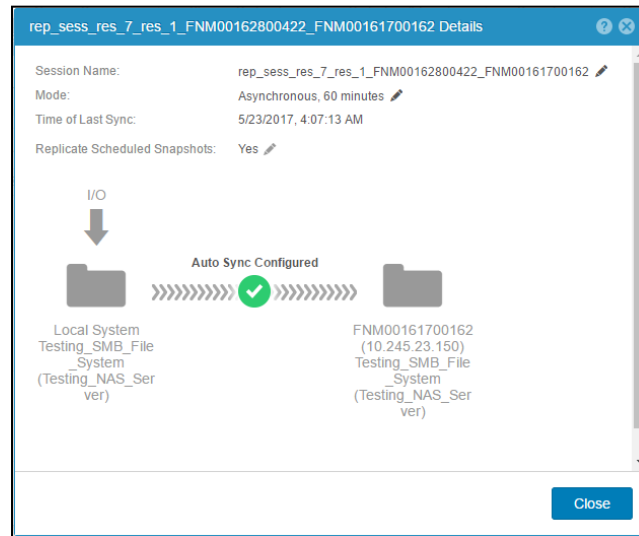


Figure 19. Snapshot Schedules

SNAPSHOT INTEROPERABILITY

Dell EMC Unity snapshots are fully compatible with other data services, which allow administrators to easily protect their critical data regardless of their environment, configuration, or use case. For most features and data services, there is no direct interaction with snapshots. However, several features have interactions with or dependencies on snapshot functionality. Considerations for using these features in conjunction with snapshots are discussed in this section.

THICK STORAGE OBJECTS

Although snapshots can be taken on both Thick and Thin storage objects, Dell EMC recommends only taking snapshots of Thin storage objects. Snapshots are a space-efficient feature that require Thin provisioning; therefore, taking a snapshot of a Thick storage object will convert it to Thin. This conversion process can strand some storage capacity in the pool, making it unusable while the snapshots are in place. There can also be performance impacts from taking a snapshot on a Thick object. If a snapshot needs to be taken on a storage object that was created as Thick, it is recommended to first migrate to a Thin storage object, and then create the snapshot.

COMPRESSION/DATA REDUCTION/ADVANCED DATA DEDUPLICATION

In Dell EMC Unity OE version 4.1 and later, Compression can be enabled on LUNs and VMware VMFS Datastores. In Dell EMC Unity OE version 4.2, Compression support was added for File Systems and VMware NFS Datastores for resources created on the 4.2 release or later. The snapshot feature is fully supported with Data Reduction that was introduced in Dell EMC version 4.3, and benefits from the space savings achieved on the source storage resource. When taking a snapshot of a Data Reduction enabled storage resource, the data on the source may be compressed. The data is left in its Data Reduction state, and the snapshot inherits the savings achieved on the source storage resource. This savings is also calculated and reported as part of the GBs saved Data Reduction information.

When a snapshot is mounted, and the source storage resource has Data Reduction enabled, Data Reduction is utilized on any snapshot I/O. Savings can also be achieved on writes to a snapshot. As write operations are received, if the source storage resource has Data Reduction enabled, snapshot writes are also passed through the Data Reduction Algorithms. This saving is tracked and reported as part of the GBs saved for the source storage resource.

In Dell EMC Unity OE version 4.5, Data Reduction includes an optional feature called Advanced Deduplication, which expands the deduplication capabilities of the Data Reduction algorithm. Advanced Deduplication is included as an optional feature to Data Reduction algorithm for the Dell EMC Unity 450F, 550F, and 650F systems. Note, this only applies on Dynamic Pools.

DELL EMC UNITY FILE SYSTEM

The Dell EMC Unity File System introduces a new 64-bit NAS architecture. This new file system brings many benefits to Dell EMC Unity including enhanced scalability, efficiency, and flexibility. It offers the ability to automatically and manually shrink and extend file systems to adjust to new client capacity demands as well as dynamically optimize storage pool and file system utilization. When shrinking a file system, it is possible for space to be returned to the underlying storage pool to be reused by other storage resources. However, the returned amount depends on many factors, including the presence of snapshots.

When a file system shrinks, space can be returned to the storage pool as long as that space is not associated with an existing snapshot. For example, if a 100GB fully allocated file system shrinks to 80 GB, approximately 20 GB will be returned to the storage pool if there are no existing snapshots of the file system. This is because the 20 GB being shrunk is no longer needed by the file system, or any other resources. However, suppose the same shrink operation was initiated immediately after taking a snapshot of the 100GB file system. In this case, almost no space would be returned to the pool because the 20 GB being shrunk from the file system is associated with the snapshot, and the point-in-time at which the snapshot was taken, and the file system was 100 GB, must be preserved. This process illustrated on Figure 20.

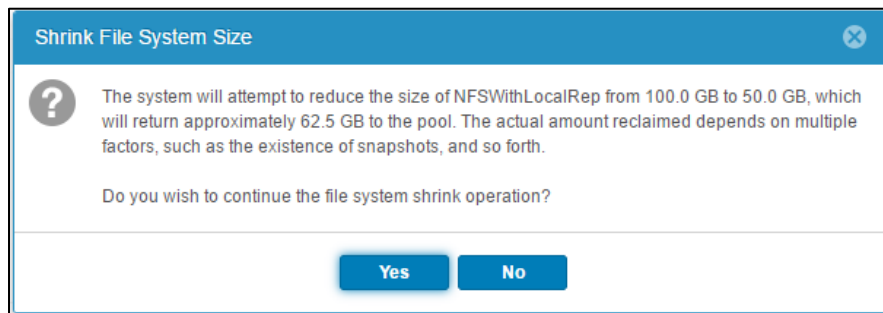


Figure 20. Shrink after taking Snapshot

These are two extreme examples to illustrate the potential impact that the presence of snapshots can have on file system shrink and reclaim operation. Under normal operation, however, the effects will be more moderate. In general, the portion of shrunken space actually reclaimed to the pool in the presence of snapshots depends on the amount of changed data since the most recent snapshot was taken. For example, if a snapshot is taken of a file system, and then the file system is completely overwritten prior to initiating the shrink operation, all of the expected space is returned to the pool because the current production file system no longer shares blocks with any snapshots. In other words, shrunken space can only be reclaimed if it is not being referenced by one or more snapshots.

REPLICATION

Dell EMC Unity offers several options for remotely protecting production data, including native asynchronous unified replication, native synchronous block and file replication, and RecoverPoint support. Asynchronous unified replication, leverages Dell EMC Unity snapshot technology to replicate point-in-time data. As a result considerations when using snapshots should also be applied to native asynchronous replication.

When a native asynchronous replication session is created, two internal system-managed snapshots are created on each of the source and destination resources. Whenever a sync is performed, either automatically due to RPO or manually, one of the snapshots on the source resource is refreshed to reflect the point-in-time when the sync started. This snapshot is then differentially copied to the destination resource, so only new or changed data is transferred. When this transfer is complete, the corresponding snapshot on the destination resource is refreshed to reflect the point-in-time associated with the copy. At this point, these snapshots on the source and destination represent a common base which reflects the point-in-time of the last sync on both the source and destination. The next time a sync begins, the process is repeated using the second pair of snapshots on the source and destination resource, which then becomes the most recent common base. This process continues to ensure that the common base is always preserved until the next sync completes, so the session can always be recovered from either system if necessary. Figure 21 shows the replicated snapshots on the **Snapshots** tab.

Because the system manages these internal replication snapshots, no user operations can be performed on these snapshots, however they are visible in Unisphere. Similarly, replication snapshots will not expire (other than being refreshed) or be automatically deleted based on storage pool or snapshot space used thresholds. However, these snapshots are not exempt from invalidation, in the event that the storage pool becomes completely full. If the storage resource is configured to invalidate snapshots in this scenario, replication snapshots are also invalidated as necessary in order to not fail host I/O. When this happens, the associated replication sessions become unavailable, having lost the common base, and require a full synchronization to the destination upon recovery.

In Dell EMC Unity OE version 4.2 and later, when asynchronous replication is enabled on the storage resource, there is the option for scheduled and manual snapshots to be replicated. The destination can have a remote retention policy that is longer than that of the source resource. This operation can be set during the snapshot creation, replication sessions and while editing existing snapshots. For example, users can set the retention policy for the source snapshots might be 2 hours, while the retention policy for the replicated snapshot be 1 week. For more information, review the white paper titled *Dell EMC Unity: Replication Technologies* on Dell EMC Online Support.

Snapshot Schedule	Name	State	Taken	Auto-Delete	Taken By	Async Replica...	Sync Replicated
Snapshot Access	4294967306...	Ready	11/19/2018, ...	No	Replication	No	No
	4294967306...	Ready	11/19/2018, ...	No	Replication	No	No
	UTC_2018-1...	Ready	11/19/2018, ...	Yes	admin	Yes	No
	Rep_UTC_2...	Ready	11/19/2018, ...	Yes	admin	Yes	No
	UTC_2018-1...	Ready	11/19/2018, ...	No	admin	No	No

Figure 21. Internal Snapshots

With Dell EMC OE version 4.4, Synchronous File Replication, also known as Dell EMC Unity MetroSync is available. Included with the MetroSync, Read-Only Snapshots and a Snapshot Schedule can be replicated to the destination system seamlessly. When a user creates Read-Write snapshots on the source, it will not be replicated to the destination system. When a snapshot gets deleted on the source, it will automatically get deleted on the destination. Creating, modifying or deleting a snapshot on the destination does not affect source snapshots.

Additionally, with MetroSync a Snapshot Schedule can be replicated. By applying a replicated snapshot schedule to an associated resource, that resource automatically applies the same schedule to the destination. When creating a new Snapshot Schedule, the user has the option to select the "Synchronize snapshot schedule to remote system" option. This option is only available when creating a Snapshot Schedule and is only applicable to Synchronous File Replication shoes in Figure 21 below. During the creation of a File System, a user can select the desired Snapshot Schedule to be applied. Once the selected Snapshot Schedule is applied, it cannot be modified on the destination system. However, a user can modify the Snapshot Schedule on the source system, and the changes will be reflected on the destination system.

Figure 22. Snapshot Schedule

There are some limitations regarding the **Restore** operation with snapshots depending on the replication type:

- File Synchronous Replication – You cannot restore a snapshot while a file synchronous replication session is in place. In order to restore a snapshot, you must first delete the replication session, perform the restore, and then re-configure the replication session. This requires a full sync after the replication session is re-established.
- Block Synchronous Replication – You must first pause the block synchronous replication session prior to restoring to a snapshot. Once the restore is complete, you can resume the replication session.
- File and Block Asynchronous Replication – You can restore a snapshot while a replication session is in place and pausing the session is not required.

HOST I/O LIMITS

Host I/O Limits is a Dell EMC Unity quality of service feature that allows administrators to specify I/O and bandwidth maximums for block resources or groups of block resources. Administrator can set Host I/O limits on all writable block resources, including attached snapshots. There is no special consideration when enabling host I/O limits on attached snapshots, other than snapshots must be attached prior to becoming available for setting host I/O limits. When users detach an attached snapshot with a currently applied host I/O limit, the snapshot is disassociated with the host I/O limit and needs to be re-associated with the host I/O limit if the snapshot is later reattached. **Error! Reference source not found.** shows the **Host I/O Limit Properties** page and the **Attached Snapshots**.

Host_IO_Limit_01 Properties

General

Storage Resources

Attached Snapshots

Performance

1 item

<input type="checkbox"/>	LUN Name	Snapshot Name	↑	Consistency Group Name	SP Owner
<input type="checkbox"/>	ThinCloneLUNBase	Snap1Read			SPA

Figure 23. Host I/O Limit on Attached Snapshot

CLOUD TIERING APPLIANCE (CTA)

Beginning with Dell EMC Unity OE version 4.2 and Cloud Tiering Appliance (CTA) version 12, block snapshots can be archived to and restored from the cloud. For more information on Cloud Tiering Appliance (CTA) and how it works with Dell EMC Unity, review the white paper titled *Dell EMC Unity: Cloud Tiering Appliance (CTA)* on Dell EMC Online Support.

THIN CLONE OVERVIEW

A Thin Clone is a read-write copy of a Thin Block level storage resource (LUN, LUNs within Consistency Group, or VMware VMFS Datastores). A Thin Clone shares the underlying data blocks with the parent storage resource. When creating a Thin Clone, no data needs to be copied from the source. Upon creation, the data is immediately available on the Thin Clone. Any changes to the data on the Thin Clone will not affect the Base LUN and vice versa.

There are many operations available for Thin Clones including the ability to create, view, modify, refresh and delete Thin Clones. In addition, the user can configure a number of data services for Thin Clones using Unisphere, Unisphere CLI and RESTful API. For example, the administrator can provide host access, configure a Host I/O limit, replicate snapshots, and set up a snapshot schedule. Thin Clones have the same look and feel as their source LUN, LUN within Consistency group or VMware VMFS Datastore.

Thin Clones use pointer-based technology, which means a Thin Clone does not consume much space from the storage pool, however, additional space is consumed only when new data is written to the base LUN or the Thin Clone. On the contrary, Thin Clones share the space with the base resource, rather than allocate a copy of the source data for itself, which gives a lot of benefits to the user.

There are several use cases for Thin Clones:

- Application Testing – The user can test applications on Thin Clones before putting them in production. This allows development and test personnel to work with real workloads and use all data services associated with production storage resources without interfering with production.
- Parallel processing – Parallel processing applications that span multiple servers can use multiple Thin Clones of a single production data set to achieve results more quickly.
- Online backup – The user can create Thin Clones to maintain hot backup copies of production systems. If there is corruption in the production data set, the user can immediately resume the read-write workload by using the Thin Clones.
- System deployment – The user can use Thin Clones to build and deploy templates for identical or near-identical environments. For example, the user can create a test template that is cloned as needed for predictable testing.

THIN CLONE HIERARCHY

In order to create a Thin Clone, users need to take a snapshot of a Block storage resource. After the snapshot is taken, it can be used to create a Thin Clone. A single LUN/LUNs within a Consistency Group and VMFS Datastores can have up to 16 Thin Clones. Note, when creating a Thin Clone of a Consistency Group, it creates a Thin Clone of all the LUNs that are associated within the Consistency Group at same time. The total maximum number of Thin Clones plus snapshots per LUN, Consistency Group, or VMFS Datastore is equal to 256. For example, if a user creates 250 snapshots, the user can only create 6 Thin Clones from this LUN. Also, a Thin Clone limit is associated with the LUN count per system and a snapshots limit, which is different from model to another. For more information on the system limits, please reference the *Dell EMC Unity Simple Support Matrix*.

Additionally, the user can create a Thin Clone from another Thin Clone's snapshot. For example, if the Base LUN has 10 Thin Clones associated with it, a user can create a Thin Clone from another Thin Clone's snapshot, up to 6 additional Thin Clones. In this case, to create more Thin Clones for that LUN, the user can delete a Thin Clone since the 16 Thin Clone limit had been reached for the particular resource. Figure 24 shows an example hierarchy of Thin Clones.

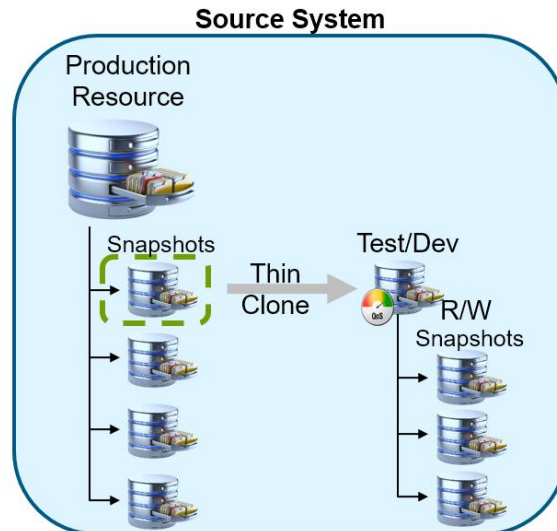


Figure 24. Thin Clone Hierarchy

THIN CLONE OPERATIONS

CREATE

To create a Thin Clone of a Block resource, on the LUNs/Consistency Groups page, select a LUN, Consistency Group, or VMFS Datastore, and then click **More Actions** > **Clone**. Figure 25 illustrated how to create Thin Clone.

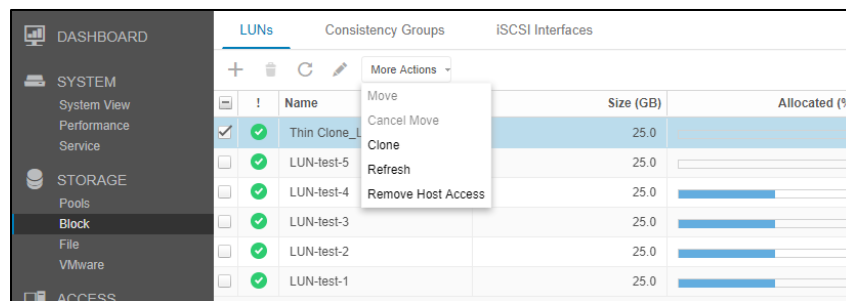


Figure 25. Create Thin Clone

Prior to creating a Thin Clone, users should create a snapshot of the resources. The snapshot can have read-only host access or be an unattached snapshot. To create a Thin Clone from a snapshot that is providing read-write access to a host, users need to create a copy of the unattached snapshot, and then create a Thin Clone. User can then give a read-write access to that snapshot. The Create Thin Clone wizard is simple to use, like any other wizard in Unisphere. There are two options in the first step of creating a Thin Clone: 1) Clone using an existing snapshot, which is usually chosen when the snapshot was created ahead of time. 2) Clone using the latest version of snapshot, which can be selected when no snapshot exists. In the Thin Clone wizard, users can give access to Thin Clone, set up a specific Snapshot Schedule and start a Replication session. Figure 26 shows the Thin Clone wizard.

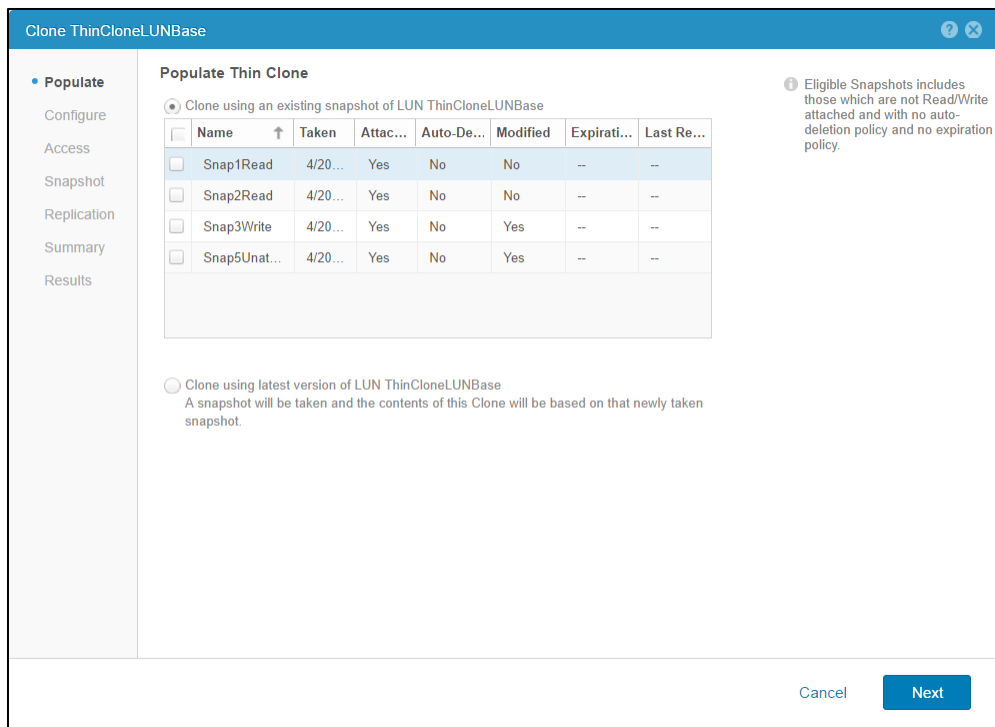


Figure 26. Thin Clone Wizard

REFRESH

Refresh functionality is part of the Thin Clone feature. A refresh process can make a Thin Clone return to a previous image or to the original snapshot image. Users can specify which snapshot they want to refresh to. In order to refresh, a snapshot cannot have either automatic snapshot deletion or a snapshot expiration policy enabled or be providing Read-Write access to a host. When the user attempts to select a snapshot that has a Read-Write access or configured a deletion policy, the GUI will prevent the user from continuing to the next step. Figure 27 shows the **Refresh** Thin Clone wizard.

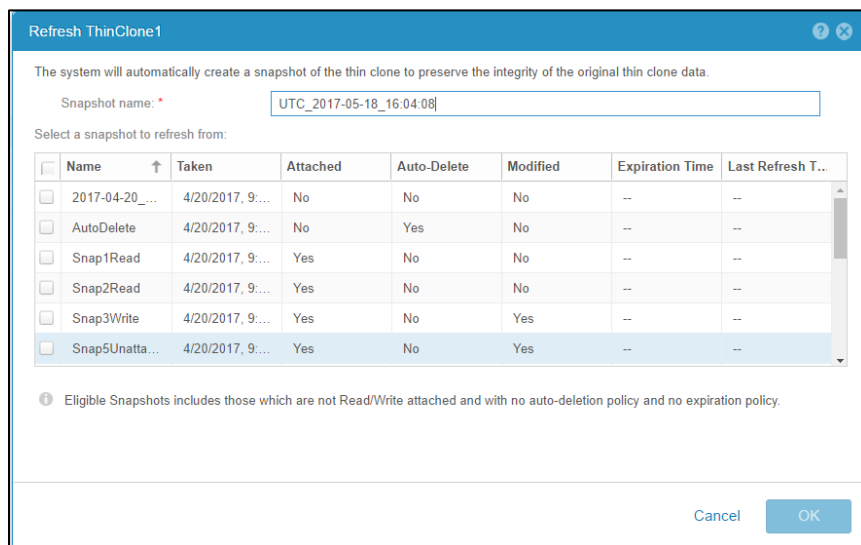


Figure 27. Refresh Thin Clone

In Dell EMC Unity OE version 4.2.1 and later, a LUN can be refreshed by any snapshot created under the Base LUN. This includes snapshots of a Thin Clone taken of the Base LUN. Being able to refresh the Base LUN from any snapshot of a Thin Clone, allows the LUN to go to any image the user wants. One use case is allowing the user to push any changes that have been made in a test environment into production. When refreshing a Thin Clone, the snapshot selected cannot have an auto-delete policy set or Read-Write access. These rules also apply when refreshing a Base LUN. Figure 28 shows the **Refresh** operation for the Base LUN.

Refresh LUNBase

The system will automatically create a snapshot of the storage resource to preserve the integrity of the original thin clone data.

Snapshot name: *

Select a snapshot to refresh from:

<input type="checkbox"/>	Name	Taken	Attached	Auto-Delete	Modified	Expiration T...	Last Refresh...
<input type="checkbox"/>	UTC_2017-09-07_15:54...	9/07/2017, ...	No	No	No	--	--
<input checked="" type="checkbox"/>	ThinClone_Snapshot	9/14/2017, ...	No	No	No	--	--

Eligible Snapshots includes those which are not Read/Write attached and with no auto-deletion policy and no expiration policy.

Figure 28. Refresh a LUN

Table 3 shows different scenarios that accompany the Refresh operation:

Table 3. Refresh Rules Operation

OBJECT TO REFRESH	REFRESH FROM	IS ALLOWED
BASE LUN	THIN CLONE	No
BASE LUN	SNAPSHOT (Any snapshot of a Base Lun or Thin Clone snapshot)	Yes
BASE LUN SNAPSHOT	THIN CLONE	No
BASE LUN SNAPSHOT	BASE LUN	Yes
THIN CLONE X	THIN CLONE Y	No
THIN CLONE	BASE LUN	No
THIN CLONE	SNAPSHOT (Any snapshot of a Base Lun or Thin Clone snapshot)	Yes
THIN CLONE SNAPSHOT	BASE LUN	No
THIN CLONE SNAPSHOT	THIN CLONE	Yes

EDIT PROPERTIES

Certain settings of a Thin Clone cannot be changed directly. Changing the FAST VP policy, enabling or disabling Data Reduction, and changing the SP ownership need to be done through the Base LUN. After a change is made, it will be reflected on all Thin Clones that are derived from that specific Base LUN.

Other settings of a Thin Clone can be changed directly. For instance, the user can change the size of the Thin Clone so it is larger than the original Base LUN. This change will not affect the Base LUN. Note that the Thin Clone's size cannot be decreased to be less than the Base LUN.

The Thin Clone's **Properties** page is similar to the LUN/Consistency Group and VMFS Datastore's page. However, the Thin Clone's Allocated space will not be displayed. Users can check the allocated space from the Base LUN. Figure 29 shows an example of the Thin Clone's properties page. A **Host I/O Limit** applies to Thin Clones just as with snapshots.

Clone LUN_Test-4

✓ Populate
• **Configure**
Access
Snapshot
Replication
Summary
Results

Configure Thin Clone

Thin Clone Base: LUN_Test-4
Original Parent Resource: LUN_Test-4
Name: *
Thin_Clone_01
Description:
Pool: Pool-01
Size(GB): 25.0
Thin: Yes
Data Reduction: Yes
Advanced Deduplication: No
Host I/O Limit:
No Limit
Create I/O Limit

Thin clone will inherit the configuration from the base LUN including pool, size, data reduction, tiering policy, SP ownership. Size can be modified later, after thin clone creation.
What are Host I/O Limits?

Cancel Back Next

Figure 29. Thin Clone properties

DELETE

Deleting a Thin Clone is like deleting a LUN or other storage resource. To delete a Thin Clone, the user needs to make sure that it is not associated with an attached snapshot or other Thin Clones. The source snapshots can be deleted. However, the Base LUN/Consistency Group and VMFS Datastore cannot be deleted if it is associated with other Thin Clones because this base is the parent of other Thin Clones.

THIN CLONE INTEROPERABILITY

CONSISTENCY GROUP

When creating a Thin Clone from a Consistency Group, all the LUNs within the consistency group are cloned. Similar to creating a Thin Clone from a stand-alone LUN, users can configure, schedule snapshots and replications, and give a consistency group access to a host. LUNs within the Consistency Group cannot be removed, added or deleted. Figure 30 shows the Consistency Group **Clone** wizard.

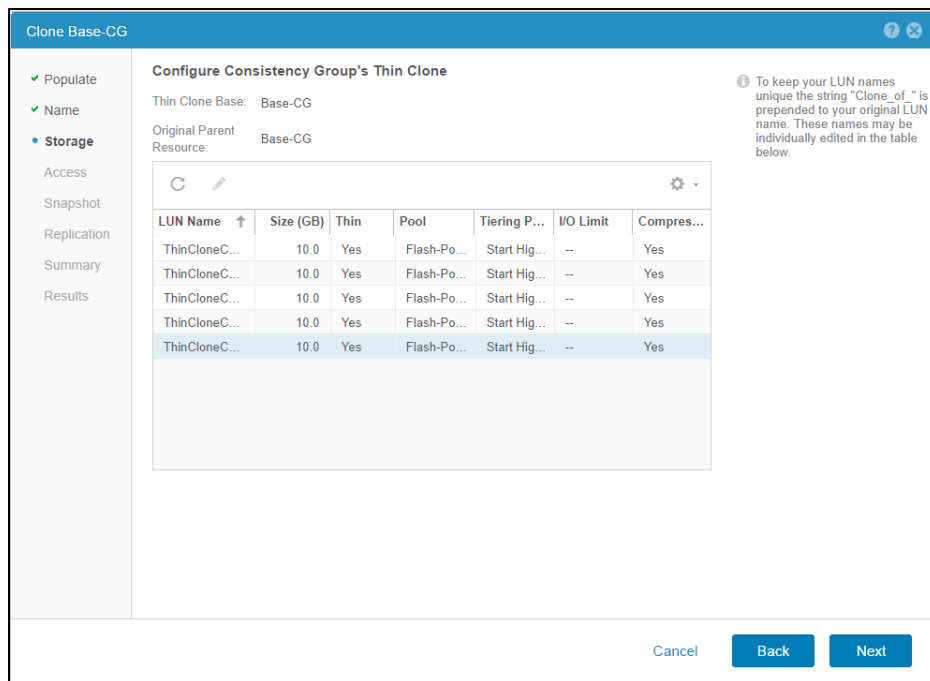


Figure 30. Thin Clone with Consistency Group

REPLICATION

Thin Clone supports synchronous and asynchronous replication. When a Thin Clone is replicated to the destination, the Thin Clone becomes a full copy of a LUN, Consistency Group, or VMFS Datastore. After being replicated, the Thin Clone is a full independent LUN, and it can have its own independent settings. This feature offers advantage for users, for example, apply the move operation to the replicated LUN. Figure 31 shows the replication process of Thin Clone.

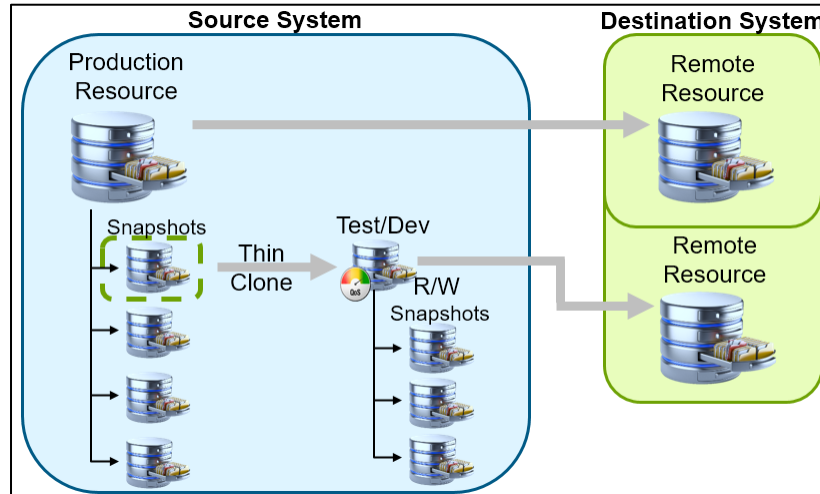


Figure 31. Thin Clone replication

CONCLUSION

Dell EMC Unity snapshots introduces unified snapshots, a fully unified method that protects all types of storage resources locally, whether they are block or file storage. This method allows administrators to protect data by taking point-in-time copies of data that can be restored at a later time when necessary. Using redirect-on-write technology, new writes to a snapped storage resource are written to a new location in the same pool and pointers are updated, which removes the need to copy data to a predefined save volume before overwriting. These snapshots can be created manually or through user-defined schedules which can differ in terms of frequency and retention, allowing for custom schedules to meet any service level. Because snapshots are compatible with other Dell EMC Unity features and data services, they can be used to easily protect users' data in a wide array of environments and use cases.

The Thin Clone in Dell EMC Unity gives power to testing and development environment to be easily deploy and ready to be tested. Thin Clone takes a read-write copy of the block resource and act individually after it creates. Any changes to the Base LUN do not affect the Thin Clone and vice versa. To take advantage of Thin Clone the user needs to create a Read-Only or unattached snapshot and then take a Thin Clone of the desired block resource. To rollback to previous copies, Thin Clone provides Refresh functionality.

REFERENCES

For additional information regarding any of the topics covered in this white paper, refer to the following resources, which are available on Dell EMC Online Support:

- *Dell EMC Unity: NAS Capabilities*
- *Dell EMC Unity: Replication Technologies*
- *Dell EMC Unity: Introduction to the Platform*
- *Dell EMC Unity: Unisphere Overview*
- *Dell EMC Unity: Compression*
- *Dell EMC Unity: Data Reduction*
- *Dell EMC Unity: Cloud Tiering Appliance (CTA)*
- *Dell EMC Unity Family : Unisphere Command Line Interface User Guide*
- *Dell EMC Unity Simple Support Matrix*