

# Dell EMC Unity: Data Reduction

Overview

## Abstract

This white paper is an introduction to the Dell EMC™ Unity Data Reduction feature. It provides an overview of the feature, methods for managing data reduction, and interoperability with other Dell EMC Unity features. Data Reduction exists in Dell EMC Unity OE version 4.3 and later.

June 2019

## Revisions

Date	Description
March 2018	Initial Release – Unity OE 4.3
January 2019	Updated for Unity OE 4.5
June 2019	Updated for Unity OE 5.0

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 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. [6/27/19] [White Paper] [H16870.2]

# Table of Contents

1	Dell EMC Unity Data Reduction Licensing .....	7
2	Overview .....	8
2.1	Supported Configurations .....	8
3	Theory of Operation .....	10
3.1	Writes .....	10
3.2	Reads .....	14
3.3	Overwrites .....	14
3.4	Creating Data Reduction Enabled Storage Resources .....	14
3.5	Enabling Data Reduction on an Existing Storage Resource .....	15
3.6	Local LUN Move .....	15
3.7	Disabling Data Reduction on a Resource .....	15
4	Management .....	17
4.1	Creating a Data Reduction Enabled Storage Resource .....	17
4.2	Enabling and Disabling Data Reduction on an Existing Storage Resource .....	22
4.3	How to Determine Which Storage Resources Have Data Reduction Enabled .....	29
4.4	Local LUN Move .....	33
4.5	Savings Reporting .....	35
5	Interoperability .....	40
5.1	Data at Rest Encryption .....	40
5.2	Replication .....	40
5.3	Snapshots .....	40
5.4	Thin Clones .....	41
5.5	Dell EMC Unity Native File and Block Import .....	41
5.6	Pool Expansion .....	42
6	Conclusion .....	43
A	Technical support and resources .....	44
A.1	Related resources .....	44

## Executive summary

Data reduction technologies play a critical role in environments in which storage administrators are attempting to do more with less. Dell EMC Unity Data Reduction aids in this effort by attempting to reduce the amount of physical storage needed to save a dataset, which helps reduce the Total Cost of Ownership of a Dell EMC Unity storage system. Dell EMC Unity Data Reduction provides space savings through the use of data deduplication and compression. Data reduction is easy to manage, and once enabled, is intelligently controlled by the storage system. Configuring data reduction and reporting savings is simple, and can be done through Unisphere, Unisphere CLI, or REST API.

In Dell EMC Unity OE version 4.3 and later, the Dell EMC Unity Data Reduction feature replaces Dell EMC Unity Compression. Data reduction includes deduplication, compression, and zero-block detection which potentially increases the amount of space savings that can be achieved. Once the Dell EMC Unity OE has been upgraded, data reduction enabled storage resources can be created, or data reduction can be enabled on existing storage resources which support data reduction. All resources with compression enabled previously will automatically begin using the data reduction algorithm. All compression references, including ways to manage compression and view space savings information, have been updated to utilize the data reduction terminology. For more information on Dell EMC Unity Compression and features prior to the OE version 4.3 release, refer to the Dell EMC Unity: Compression white paper found on Dell EMC Online Support.

This white paper discusses the Dell EMC Unity Data Reduction feature, including technical information on the underlying technology of the feature, how to manage data reduction on supported storage resources, how to view data reduction savings, and the interoperability of data reduction with other features of the storage system. Best Practices information for using Dell EMC Unity Data Reduction, along with information on when to enable data reduction, can be found in the Dell EMC Unity: Best Practices Guide white paper found on Dell EMC Online Support.

## Audience

This white paper is intended for customers, partners, and employees who are planning to utilize Dell EMC Unity Data Reduction. It assumes familiarity with Dell EMC Unity and Dell EMC Unity's management software.

# Terminology

**Advanced Deduplication** – A dynamic deduplication algorithm which reduces storage consumption by eliminating duplicate 8KB blocks of data within a storage resource.

**All Flash Pool** – A Pool which contains only Flash Drives. An All Flash Pool can be a Traditional Pool or a Dynamic Pool.

**Asynchronous Replication** – A replication method which allows you to replicate data over long distances, and maintain a replica at a destination site. Updates to the destination image can be issued manually, or automatically based on a customizable Recovery Point Objective (RPO).

**Block Storage Resources** – LUNs, LUNs within a Consistency Group, and VMware VMFS Datastores.

**Compression** – A data reduction method which reduces the physical amount of storage required to save a dataset.

**Consistency Group** – A storage instance which contains one or more LUNs within a storage system. Consistency Groups help organize the storage allocated for a particular host or hosts. Data protection configurations, such as replication and snapshot settings, 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.

**Data at Rest Encryption (D@RE)** – The process of encrypting data and protecting it against unauthorized access unless valid keys are provided. This prevents data from being accessed and provides a mechanism to quickly crypto-erase data.

**Deduplication** – A data reduction method which reduces the physical amount of storage required to save a dataset.

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

**Flash drive (SSD)** – A Flash based storage device used to store data.

**Hard Disk Drive (HDD)** – A storage device based on spinning platters used to store data.

**Hybrid Pool** – A Pool which does not contain only Flash Drives. A Hybrid Pool typically contains more than one type of drive technology, such as Flash, SAS, and NL-SAS.

**LUN** – A block based storage resource which a user provisions. It represents a SCSI logical unit.

**Pool** – A set of drives that provide specific storage characteristics for the resources that use them, such as LUNs, VMware Datastores, and File Systems.

**REST API** – An application programming interface that utilizes familiar HTTP operations like GET, POST, and DELETE. REST architecture includes certain constraints that ensure that different implementations of REST conform to the same guiding principles, thereby allowing developers the ease of application development when working with different REST API deployments.

**Snapshot** – A snapshot, also called a Dell EMC Unity Snapshot, is a point-in-time view of a storage resource. When a Snapshot is taken, the snapshot is an exact copy of the source storage resource, and shares all blocks of data with it. As data changes on the source, new blocks are allocated and written to. Dell EMC Unity Snapshot technology can be used to take a snapshot of a Block or File storage resource.

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

**System Cache (DRAM Cache)** – Dell EMC Unity software component which leverages DRAM memory to improve host read and write performance.

**Thin Clone** – A read-write copy of a Thin Block storage resource (LUN, Consistency Group, or VMware VMFS Datastore) that shares blocks with the parent resource.

**Unisphere** – A web-based management environment 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.

# 1 Dell EMC Unity Data Reduction Licensing

Dell EMC Unity Data Reduction is licensed with all physical Dell EMC Unity systems at no additional cost. Data reduction is not available on the Dell EMC UnityVSA version of the Dell EMC Unity platform as data reduction requires write caching within the system. To use data reduction with Block and File Storage Resources such as Thin LUNs, Thin LUNs within a Consistency Group, Thin File Systems, and Thin VMware VMFS and NFS Datastores, the system must be running Dell EMC Unity OE version 4.3 or later. For systems running a version of Dell EMC Unity OE older than this revision, Dell EMC Unity Compression is available. After upgrading the system to Dell EMC Unity OE version 4.3 or later, the data reduction feature will automatically be licensed, with no additional steps needed. Storage resources with compression enabled prior to the upgrade will automatically have data reduction enabled after upgrading to Dell EMC Unity OE version 4.3 or later.

To verify which version of Dell EMC Unity OE your system is running, simply select the **View System Status** icon found on the top blue menu bar of Unisphere. Alternatively, you can view the license status for Dell EMC Unity Data Reduction by clicking the **Update System Status** icon, denoted by a gear icon on the top blue menu bar, and finding **Data Reduction** in the **License Management** list. An entry of **Data Reduction** and a green checkmark besides it confirms the feature is licensed on the system.

## 2 Overview

The Dell EMC Unity family of storage systems are feature-rich, easy-to-use, and deliver full Block and File unified environments starting in a single 2U enclosure. To help reduce the Total Cost of Ownership and increase the efficiency of a Dell EMC Unity storage system, Dell EMC Unity Compression was added in Dell EMC Unity OE version 4.1 for Thin Block storage resources. Thin File storage resource support was added in Dell EMC Unity OE version 4.2. In Dell EMC Unity OE version 4.3, the Dell EMC Unity Data Reduction feature replaces compression, and provides more space savings logic to the system with the addition of zero block detection and deduplication. 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. With data reduction, the amount of space required to store a dataset for data reduction enabled storage resources is reduced when savings are achieved. This space savings reduces the amount of physical storage required to store a dataset, which can lead to cost savings. Data reduction savings are not only achieved on the storage resource it is enabled on, but space savings are also realized on Snapshots and Thin Clones of those resources as well. Snapshots and Thin Clones inherit the data reduction setting of the source storage resource, which helps to increase the space savings that they can provide.

### 2.1 Supported Configurations

In Dell EMC Unity OE version 4.3 and later, data reduction is supported on Thin LUNs, Thin LUNs within a Consistency Group, Thin File Systems, and Thin VMware VMFS and NFS Datastores. To support data reduction, Thin File storage resources must be created on Dell EMC Unity OE version 4.2 or later and be created within an All Flash Pool. File storage resources created on Dell EMC Unity OE versions prior to version 4.2 do not support data reduction. Within a Consistency Group, data reduction enabled LUNs can be mixed with LUNs which have data reduction disabled. In Dell EMC Unity OE 4.5 and later, Advanced Deduplication is included as an optional feature to the Data Reduction algorithm for configurations that support it. Table 1 below outlines the supported configurations for data reduction and Advanced Deduplication.



Table 1. Supported Configurations for Data Reduction and Advanced Deduplication

Dell EMC Unity OE	Technology	Supported Pool Type	Supported Models
4.3 / 4.4	Data Reduction	All Flash Pool*	300   400   500   600 300F   400F   500F   600F 350F   450F   550F   650F
4.5	Data Reduction	All Flash Pool*	300   400   500   600 300F   400F   500F   600F 350F   450F   550F   650F
	Data Reduction + Advanced Deduplication	All Flash Pool**	450F   550F   650F
5.0	Data Reduction	All Flash Pool*	300   400   500   600 300F   400F   500F   600F 350F   450F   550F   650F 380   480   680   880 380F   480F   680F   880F
	Data Reduction + Advanced Deduplication	All Flash Pool*	450F   550F   650F 380   480   680   880 380F   480F   680F   880F

\* Resource can be created on either a Traditional or a Dynamic Pool (For systems that support Dynamic Pools)

\*\* Resource can be created on a Dynamic Pool Only

Dell EMC Unity Data Reduction and the Advanced Deduplication option, for configurations supporting the Advanced Deduplication feature, can be enabled on supported storage resources at the time of the resource's creation, or enabled or disabled at a later time. Advanced Deduplication requires Data Reduction to be enabled on the resource, but can be enabled or disabled independently to the Data Reduction setting. Local LUN Move can be leveraged to move a resource's data into a data reduction, and optionally an Advanced Deduplication enabled resource within an All Flash Pool. For File storage resources, there is no direct method available to convert from a Thick File resource or pre-Dell EMC Unity OE 4.2 Thin File resource to a 4.2 or later Thin File resource. Methods to move File data include host-based migration to a Dell EMC Unity OE 4.2 or later Thin File System, VMware vMotion for VMs created on NFS Datastores, or Dell EMC Unity Asynchronous Replication. More information on migration options can be found in the Dell EMC Unity: Migration Technologies white paper found on Dell EMC Online Support.

Dell EMC Unity Data Reduction can also be enabled on Block and File storage resources participating in replication sessions. The source and destination storage resources in a replication session are completely independent, and data reduction with or without the Advanced Deduplication option can be enabled or disabled separately on the source and destination resource. The availability of enabling data reduction, Advanced Deduplication, or compression, in the case of codes prior to Dell EMC Unity OE version 4.3, on a source and/or a destination resource depends on the Dell EMC Unity OE version, the system type, and the Pool configuration.

Pools containing data reduction enabled storage resources cannot be expanded with SAS or NL-SAS. For more information about expanding Pools and how to convert to a Hybrid Flash Pool, please review the *Pool Expansion* section found under *Interoperability*.

## 3 Theory of Operation

### 3.1 Writes

Dell EMC Unity Data Reduction works the same for both Block and File storage resources. Data reduction utilizes a software algorithm to analyze and achieve space savings within a storage resource. 100 below is a high-level diagram of a storage resource with data reduction enabled residing within an All Flash Pool. As shown in 0, data reduction occurs inline between System Cache and the storage resource on an All Flash Pool.

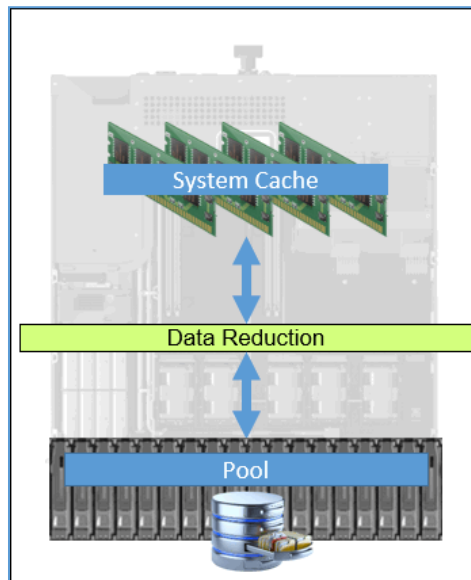


Figure 1. Dell EMC Unity Data Reduction Overview

When data is written to the system, the data is saved in System Cache, and the write is acknowledged with the host. The data reduction algorithm is not invoked for write I/Os at this point in time in order to provide the fastest response to the host. 0 below outlines an example of a write to a storage resource with data reduction enabled. No data has been written to the drives within the Pool at this time.

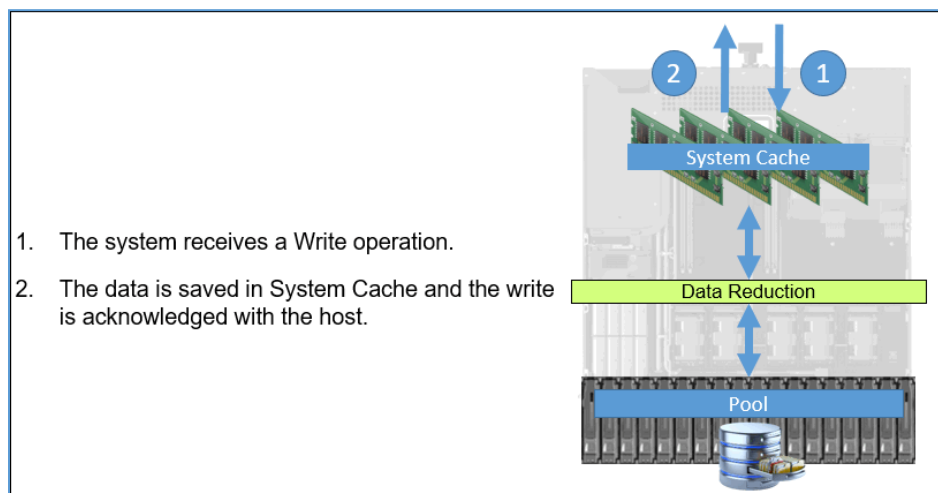


Figure 2. Host Write Example

In Dell EMC Unity, before a write is saved in System Cache, the system ensures space is available and allocated for the I/O within the target storage resource. As all back-end allocations and lookups within the target resource are deferred until after writes are accepted into System Cache and the host is acknowledged, a portion of the private space within the storage resource's overhead is tracked and utilized as a possible location to store the I/O when accepting data into cache. A storage resource's private space is fixed in size, and allocated at time of the storage resource's creation. After the I/O is acknowledged, the normal cache cleaning process occurs. Space within the storage resource is utilized or allocated, if needed, and the data is saved to disk. This caching behavior not only applies to data reduction enabled resources, but it is also applicable to Block and File storage resources (excluding VVols) created on All Flash Pools.

For data reduction enabled storage resources, the data reduction process occurs during the System Cache's proactive cleaning operations or when System Cache is flushing cache pages to the drives within the Pool. The data in this scenario may be new to the storage resource, or the data may be an update to existing blocks of data currently residing on disk. In either case, the data reduction algorithm occurs before the data is written to the drives within the Pool. During the data reduction process, multiple blocks are aggregated together and sent through the algorithm. After determining if savings can be achieved or data needs to be written to disk, space within the Pool is allocated if needed and the data is written to the drives. A high-level diagram of this operation is displayed in 0 below.

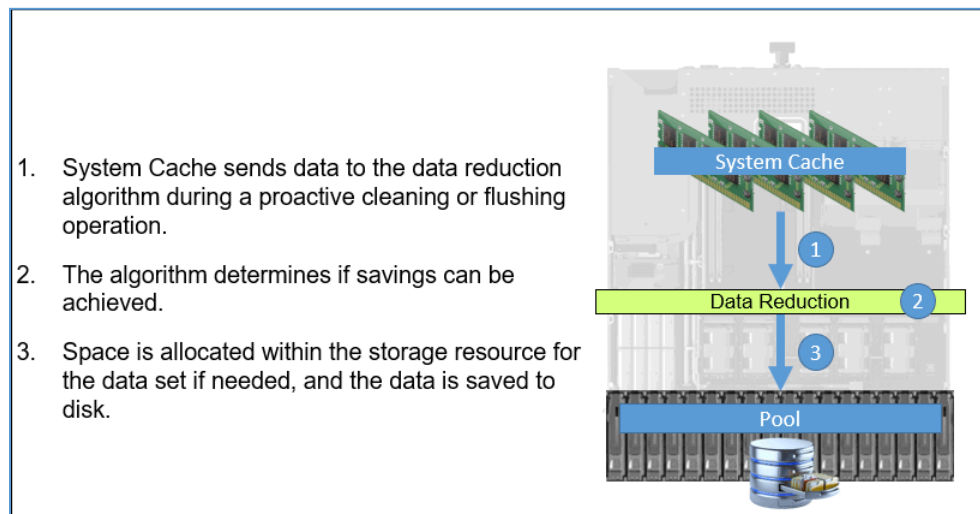


Figure 3. Dell EMC Unity Data Reduction Example

Dell EMC Unity's Data Reduction feature includes multiple space efficiency algorithms to help reduce the total space occupied by a dataset. Included in the Data Reduction feature is deduplication, compression, and optionally Advanced Deduplication algorithms. 0 below is an overview of the data reduction feature with Advanced Deduplication enabled. Before data is sent to the Data Reduction algorithm, it is first segmented into 8KB blocks. As an 8KB block of data passes through the algorithm, it may or may not touch all efficiency algorithms within data reduction. If a block can be deduplicated, the remainder of the data reduction algorithms are skipped, saving time and processing overhead. Each of the algorithms within data reduction feature will be discussed in detail later in this section.

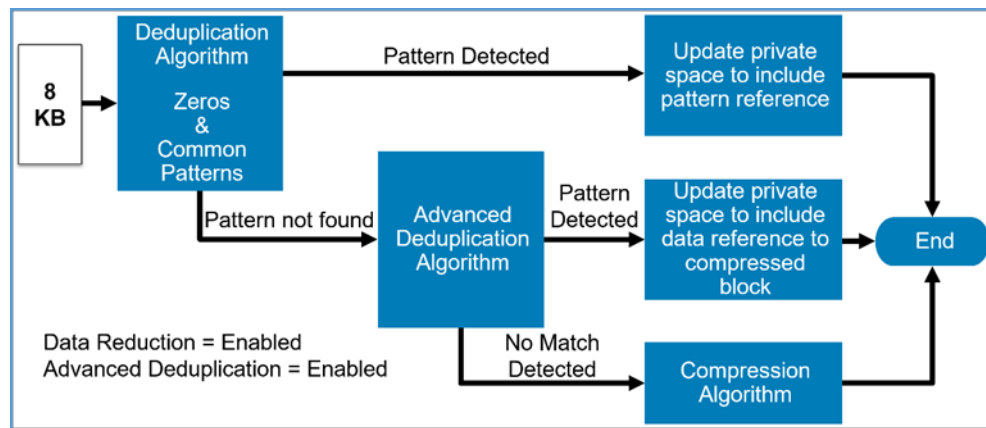


Figure 4. Dell EMC Unity Data Reduction – Advanced Deduplication Enabled

With Advanced Deduplication disabled, a block of data entering the data reduction feature is only passed through the deduplication and compression algorithms. The compression algorithm is only reached when zeros or common patterns are not detected on the block of data. An example of data reduction with Advanced Deduplication disabled is shown in 0.

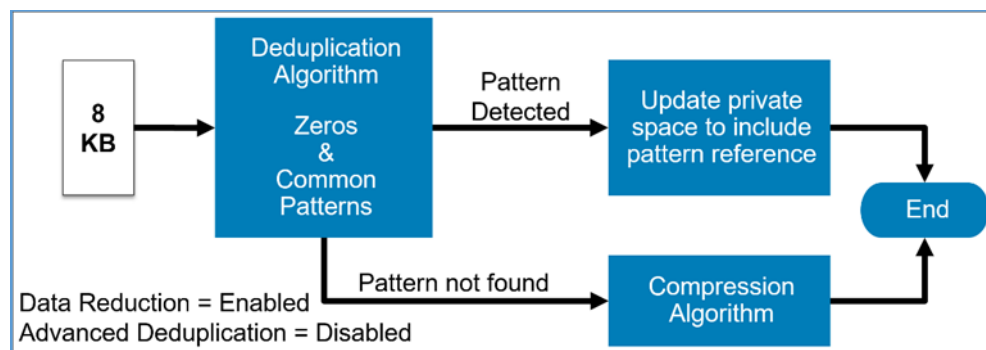


Figure 5. Dell EMC Unity Data Reduction – Advanced Deduplication Disabled

### 3.1.1 Deduplication – Zeros and Common Patterns

When new data first enters the data reduction logic, it is first passed through the deduplication algorithm. The deduplication algorithm is a lightweight software algorithm which analyzes the blocks of data for known patterns. The patterns may be a block of zeros written by the host, or common patterns found in Dell EMC Unity's many use cases, such as virtual environments. If a pattern is detected, the private space within the storage resource is updated to include that the particular block is a pattern and information on how to recreate the data block if it is accessed in the future. No data is written to disk in this scenario, which helps reduce storage consumption and drive wear. Also, when deduplication finds a pattern match, the remainder of the data reduction feature is skipped for those blocks which saves system resources. In the instance where no pattern is detected, the data is passed to Advanced Deduplication if it is enabled. If the Advanced Deduplication option is disabled, the data is passed through the compression logic within the data reduction algorithm.

### 3.1.2 Advanced Deduplication

If Advanced Deduplication is enabled, and deduplication did not detect a pattern, the data is passed through Advanced Deduplication. Advanced Deduplication is a dynamic deduplication algorithm which reduces storage consumption by eliminating duplicate 8KB blocks within a storage resource. Advanced Deduplication will only compare and detect duplicate data found within a single storage resource, such as a LUN or File

System. The Advanced Deduplication algorithm utilizes fingerprints created for each block of data to quickly identify duplicate data within the dataset. 0 below shows the Advanced Deduplication algorithm in detail.

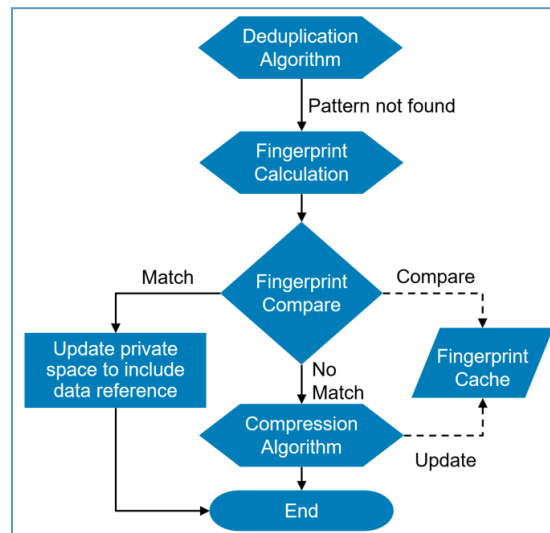


Figure 6. Dell EMC Unity Data Reduction – Advanced Deduplication Algorithm

The fingerprint cache is a component of the Advanced Deduplication algorithm. The fingerprint cache is a region in system memory reserved for storing fingerprints for each storage resource with Advanced Deduplication enabled. There is one fingerprint cache per storage processor, and it contains the fingerprints for storage resources residing on that SP. Through machine learning and statistics, the fingerprint cache determines which fingerprints to keep, and which ones to replace with new fingerprints. The fingerprint cache algorithm learns which resources have high deduplication rates and allows those resources to consume more fingerprint locations.

If an 8KB block is not deduplicated by the zero and common pattern deduplication algorithm, the data is passed into the fingerprint calculation portion of the Advanced Deduplication algorithm. Each 8KB block receives a fingerprint, which is compared to the fingerprints for the storage resource. If a matching fingerprint is found, deduplication occurs and the private space within the resource is updated to include a reference to the block of data residing on disk. No data is written to disk at this time. Storage resource savings are compounded as deduplication can reference compressed blocks on disk. If a match is not found, the data is passed to the compression algorithm.

### 3.1.3 Compression

As blocks enter the compression algorithm, they are passed through the compression software. If savings can be achieved, space is allocated within the Pool which matches the compressed size of the data, the data is compressed, and the data is written to the Pool. When Advanced Deduplication is enabled, the fingerprint for the block of data is also stored with the compressed data on disk. The fingerprint cache is then updated to include the fingerprint for the new data. Compression will not compress data if no savings can be achieved. In this instance, the original block of data will be written to the Pool. Waiting to allocate space within the resource until after the compression algorithm is complete helps to not over-allocate space within the storage resource.

## 3.2 Reads

When a read operation is sent to a storage resource which has data reduction enabled, the system first needs to determine where the data is currently located. A lookup is performed to determine if the block is currently stored within System Cache or on the Pool in its original size, or was previously deduplicated or compressed. If the data resides in System Cache in its original form, the data is sent to the host requesting the data. If the data does not reside in System Cache, and is stored in its original form on the Pool, a normal read operation occurs as if data reduction is disabled on the storage resource. The data is copied into System Cache and then sent to the host requesting the data.

If data reduction achieved space savings within the block, the data must be recreated within System Cache in its original form so it can be sent to the host. If the block was previously deduplicated, the block is either recreated if the block contained a common pattern or copied into System Cache if the block was deduplicated by Advanced Deduplication using the information within the private space of the resource, and the host is sent the data. If the data is compressed, it must first be uncompressed before the data is sent to the host. If the compressed data already resides in System Cache, the data is uncompressed to a temporary location, the data is sent to the host, and the temporary location is released. If the compressed data being requested resides on a drive, the data is first read into System Cache, uncompressed to a temporary location, and the host is sent the data. Data is never uncompressed on disk due to a read operation, as this would reduce the amount of savings on the storage resource.

## 3.3 Overwrites

When an update is received for a previously written block of data, the system determines if the overwrite is for a block which has space savings or not. The data is also passed through the data reduction logic to determine if any space savings can be achieved. If the new block of data deduplicates to a known pattern, the private space within the resource is updated with the new pattern information. If the now outdated block of data had compression savings or was written to the Pool in its original form, these blocks are freed within the resource for reuse. If deduplication savings cannot be achieved the data is sent though the compression logic.

If compression can reduce the size of the data, the system needs to determine where to store the block of data within the Pool. If the amount of compression savings is now less than the last time the data was written, then new space must be allocated within the storage resource to store the new data size. If the dataset size hasn't changed or is smaller than it was previously, then a write to an already allocated block may occur. This logic prevents causing fragmentation in the resource, which helps with performance and space savings. If a new block is allocated, the previously used block is freed for reuse. If no deduplication or compression savings can be achieved, the data is written in the Pool at its original size, and may overwrite the original data for the resource.

In the background, the old locations that are no longer needed are freed by a cleanup process and can be reused. This process also frees blocks no longer in use by the storage resource and its Snapshots or Thin Clones. If enough space is freed within a 256 MB slice, the slice can be freed back to the Pool.

## 3.4 Creating Data Reduction Enabled Storage Resources

Dell EMC Unity Data Reduction enabled Block and File storage resources with or without Advanced Deduplication enabled can be created using Unisphere, Unisphere CLI (UEMCLI), or REST API. When creating data reduction enabled LUNs in Unisphere, one or more LUNs can be created at the same time using the Create LUNs wizard, which is found on the Block page. Once the Data Reduction checkbox is enabled, the Advanced Deduplication option become available. When creating storage resources using

Unisphere CLI or REST API, extra options for enabling data reduction and Advanced Deduplication are available. Data reduction enabled LUNs are also supported within Consistency Groups. A mix of LUNs with and without data reduction enabled can reside within the same Consistency Group. Data reduction enabled File Systems can also reside within a NAS Server with File Systems with data reduction disabled.

## 3.5 Enabling Data Reduction on an Existing Storage Resource

Dell EMC Unity Data Reduction and/or Advanced Deduplication can be enabled on a supported storage resource at any point in time. This can be completed in Unisphere from the properties window of the storage resource, or by utilizing Unisphere CLI or REST API. Enabling data reduction on File storage resources created prior to Dell EMC Unity OE 4.2 is not supported. When data reduction and/or Advanced Deduplication is enabled on an existing resource, no data on the storage resource is passed through the data reduction algorithm at that time. All previously written data is left in its original state, and only new writes or overwrites are subject to the data reduction algorithm. To achieve savings on existing data within a Block storage resource, the user can utilize the Local LUN Move option, which is outlined below. For File resources, the methods to move File data include host-based migration to a Dell EMC Unity OE 4.2 or later Thin File System, VMware vMotion for VMs created on NFS Datastores, or Dell EMC Unity Asynchronous Replication.

## 3.6 Local LUN Move

Local LUN Move, also referred to as Move, is used to move data from a supported source Block storage resource to a target Block storage resource within the same system. The Move operation is completely transparent to any associated hosts, and no interruption to access is seen. Move can be utilized to move a storage resource from one Pool to another, or within the same Pool. When utilizing the Move feature, the Session Priority and destination Pool may be customized. You can also move to a data reduction enabled resource if the destination Pool is an All Flash Pool and the storage resource supports data reduction. Move can also migrate to a resource with Advanced Deduplication enabled if it is supported. A Move session can be cancelled at any time.

Move can also be used on a storage resource with data reduction enabled to migrate the data to a resource with data reduction enabled. This is also true for Advanced Deduplication enabled resources. During the migration all data passes through the data reduction algorithm, and additional savings may be achieved. This process is most often utilized when data has been written to a resource before data reduction was enabled, and the user is looking to ensure all data is subject to the data reduction algorithms. If Dell EMC Unity Compression was previously used, Move can be used to pass all data through the new space savings logic instead of waiting for the data to be overwritten. For more information on Local LUN Move and any restrictions of its usage, refer to the white paper titled *Dell EMC Unity: Migration Technologies* on Dell EMC Online Support.

As with other features, Local LUN Move can be managed in Unisphere, Unisphere CLI, and REST API.

## 3.7 Disabling Data Reduction on a Resource

As with enabling data reduction, data reduction can be disabled at any point in time on a supported storage resource. Advanced Deduplication can also be disabled at any time without disabling data reduction if required. These operations can be completed in Unisphere from the properties window of the storage resource, or by utilizing Unisphere CLI or REST API. If Advanced Deduplication is disabled, all deduplication savings achieved by Advanced Deduplication are unchanged and the data is left in its current state. As data is overwritten, the data is subject to the data reduction algorithm if still enabled, or written in its original form if data reduction is disabled. If data reduction is disabled, all data for the storage resource is left in its current

state within its Pool, whether it is deduplicated, compressed, or not. Data written after disabling data reduction will be stored in its original form. As data that has been deduplicated or compressed is overwritten, data reduction savings are reduced on the storage resource. To fully remove data reduction savings from a Block storage resource, Local LUN Move can be utilized by specifying a non-data reduction enabled destination.



## 4 Management

Creating and managing Dell EMC Unity Data Reduction from Unisphere, Unisphere CLI, and REST API is easy and intuitive. The following sections outline how to create a data reduction enabled storage resource, how to enable and disable data reduction, how to determine the current state of data reduction on a resource, and where data reduction savings are reported. Unisphere examples for each of these areas will be shown. For more information on using the Unisphere CLI, refer to the *Unisphere Command Line Interface User Guide* on Dell EMC Online Support. For information on managing Dell EMC Unity Data Reduction from REST API, consult the REST API documentation which can be accessed directly from any Dell EMC Unity system:

**REST API Programmer's Guide** – [https://<Management\\_IP>/apidocs/programmers-guide/index.html](https://<Management_IP>/apidocs/programmers-guide/index.html)

**REST API Reference Guide** – [https://<Management\\_IP>/apidocs/index.html](https://<Management_IP>/apidocs/index.html)

Where <Management\_IP> is the management IP of your system.

### 4.1 Creating a Data Reduction Enabled Storage Resource

Dell EMC Unity Data Reduction is supported on Thin LUNs, whether standalone or within a Consistency Group, Thin File Systems, and Thin VMware VMFS and NFS Datastores created on All Flash Pools. Advanced Deduplication can be enabled on resources which support Data Reduction, and reside on a supported configuration. Each of these storage resources can be created using Unisphere, Unisphere CLI, or REST API. The following sections outline how you can easily create data reduction enabled storage resources using Unisphere and Unisphere CLI. Creating resources with Advanced Deduplication enabled will also be covered.

#### 4.1.1 LUNs

In Unisphere, standalone Thin and Thick LUNs are created using the **Create LUNs** wizard, which can be found on the **Block** page. The **Block** page can be accessed by selecting **Block** under **Storage** in the left Unisphere pane. Figure 7 below shows the **Create LUNs** wizard. To create a data reduction enabled LUN or multiple data reduction enabled LUNs, ensure the **Thin** and **Data Reduction** checkboxes are checked in the **Configure** step. Once **Data Reduction** is enabled the **Advanced Deduplication** checkbox is shown and can be enabled for resources which support it. After customizing the other settings of the LUN, click **Next**.

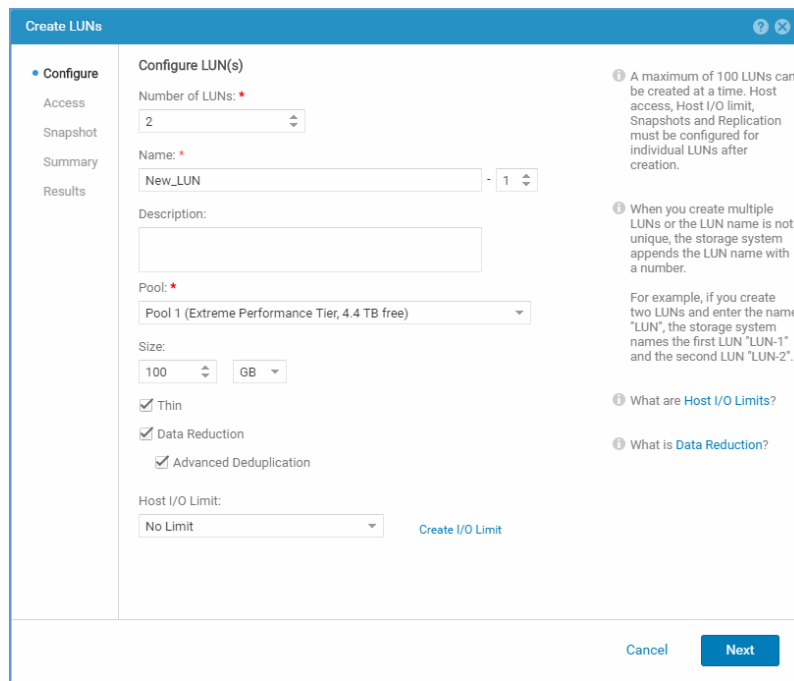


Figure 7. Create LUN Wizard – Configure Step

After configuring the other settings of the LUN(s), the **Summary** step is displayed. An example of the **Summary** step when creating multiple LUNs with data reduction and Advanced Deduplication enabled is shown in Figure 8. The Summary screen includes the **Name(s)** of the LUNs being created, and if **Data Reduction** and **Advanced Deduplication** will be enabled on the new LUNs.

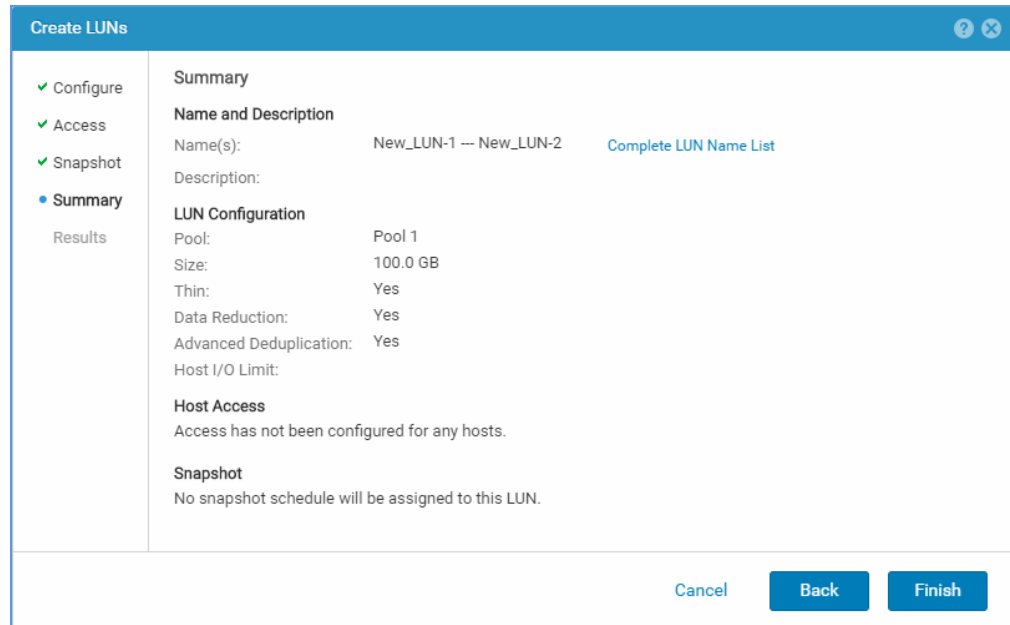


Figure 8. Create LUN Wizard – Summary Step

### 4.1.2 Consistency Groups

When creating a new Consistency Group, Thin LUNs can also be created with data reduction and Advanced Deduplication enabled. Consistency Groups can be created and configured by navigating to the **Consistency**

**Group** tab within the **Block** page. The **Configure LUNs** window, shown below in Figure 9, has an option to create LUNs with data reduction and Advanced Deduplication enabled. To create LUNs with data reduction and Advanced Deduplication enabled within a Consistency Group, ensure the **Thin**, **Data Reduction**, and **Advanced Deduplication** boxes are checked before clicking **OK**.

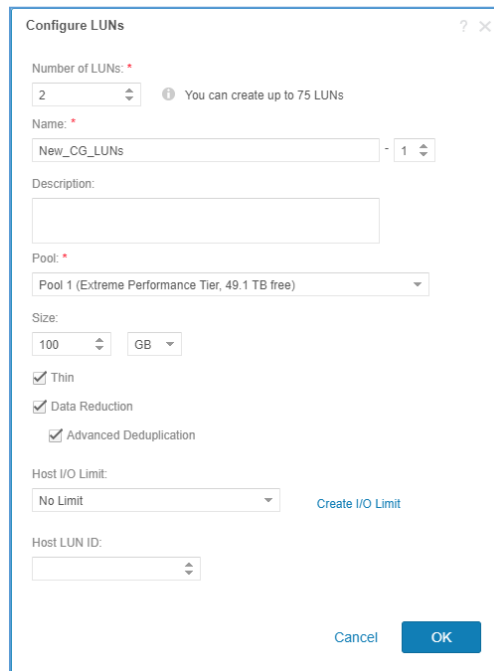


Figure 9. Configure LUNs step within the Create a Consistency Group Wizard

After configuring LUNs within the Consistency Group, the **Storage** screen within the **Create a Consistency Group** wizard is populated. An example of this window is shown in Figure 10. From this screen you can see if the LUNs being created have data reduction and Advanced Deduplication enabled or not by reviewing the **Data Reduction** and **Advanced Deduplication** columns. The **Data Reduction** and **Advanced Deduplication** columns are not enabled by default, but can be added to the view by clicking the **Gear Icon** and clicking the checkbox next to **Data Reduction** and **Advanced Deduplication** under the **Columns** option.

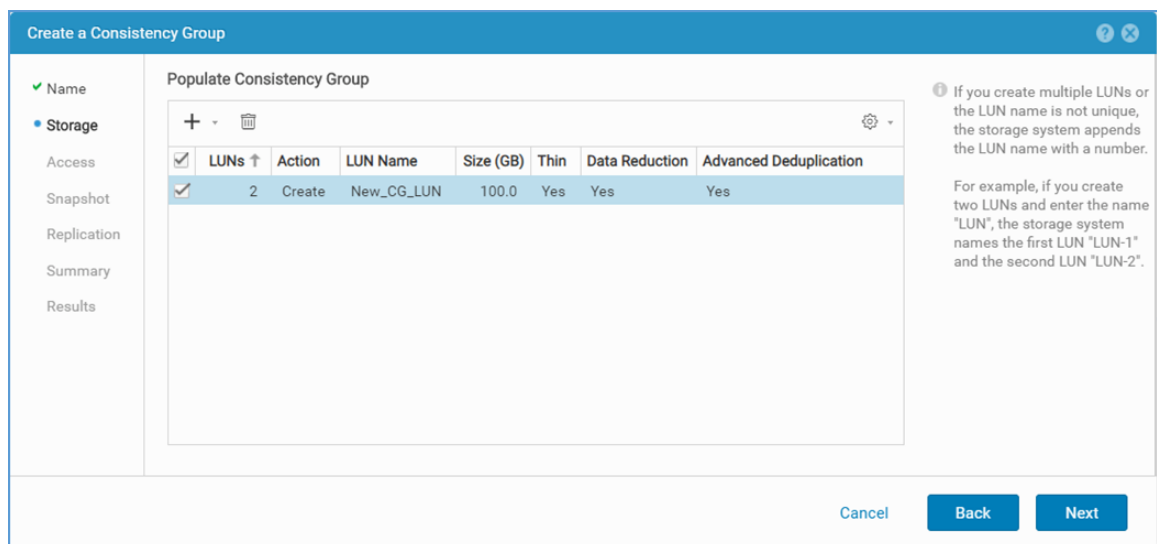


Figure 10. Create a Consistency Group - Populate a Consistency Group screen

### 4.1.3 File Systems

Data reduction is supported on Thin File Systems created on Dell EMC Unity OE version 4.2 and later, and Advanced Deduplication can also be enabled if the resource resides on a supported configuration. In Unisphere, Thin File Systems are created using the **Create a File System** wizard, found on the **File** page. The **File** page can be accessed by selecting **File** under **Storage** in the left pane of Unisphere. Figure 11 below shows the **Storage** step within the **Create a File System** wizard. To enable data reduction, first ensure the **Thin** checkbox is checked, as data reduction is only available on Thin File Systems. Second, select the **Data Reduction** checkbox, which is disabled by default. If supported, once Data Reduction is enabled **Advanced Deduplication** can also be enabled. After customizing the other settings of the File System, click **Next**.

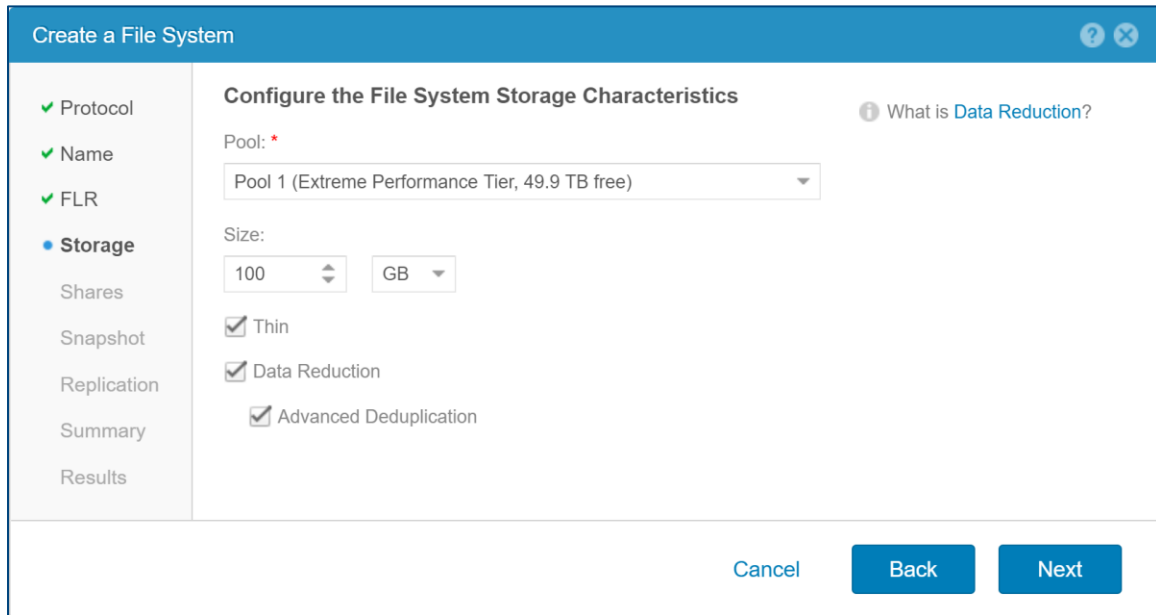


Figure 11. Create a File System Wizard – Storage Step

After configuring the other File System settings, the **Summary** step is displayed. This is the last step before the creation of the File System. An example of the **Summary** step can be seen in Figure 12. The **Summary** screen includes entries for **Data Reduction** and **Advanced Deduplication**, which signifies if the File System will be configured with those features enabled (**Yes**) or disabled (**No**).

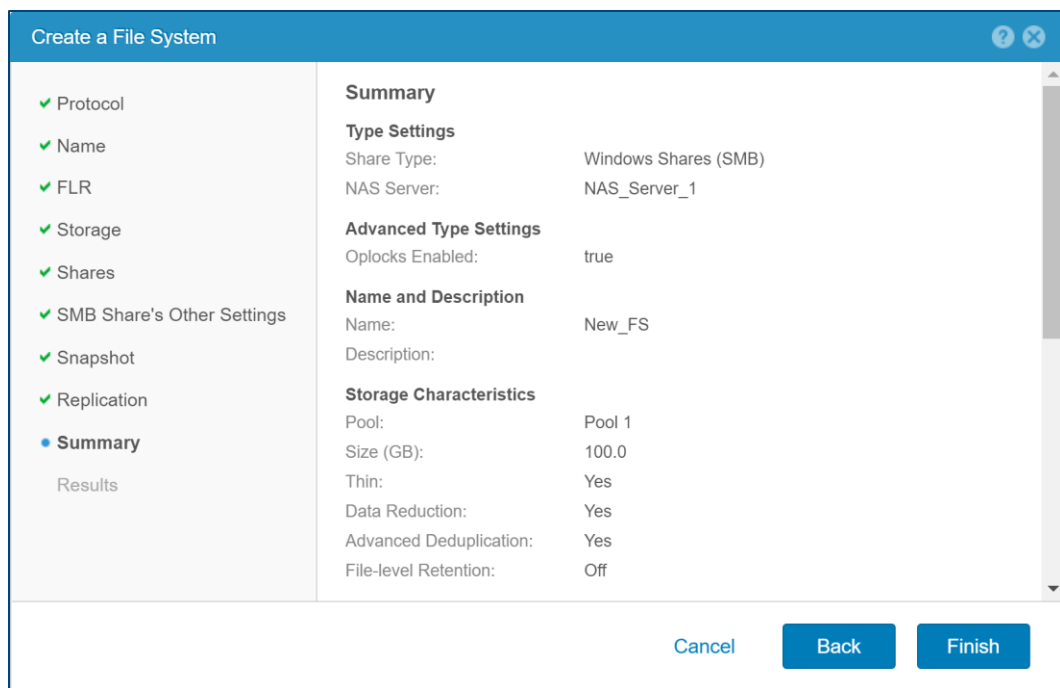


Figure 12. Create a File System Wizard – Summary Step

#### 4.1.4 VMware VMFS And NFS Datastores

To create a VMware Datastore, first navigate to the **VMware** page, found by selecting **VMware** under **Storage** in the left Unisphere pane. While on the **VMware** page, select the + symbol on the **Datastores** tab to launch the **Create VMware Datastore** wizard. On the **Type** step, select **File** to create a VMware NFS Datastore, or **Block** to create a VMware VMFS Datastore. On the **Storage** step you can select which Pool to create the VMFS or NFS Datastore on, the size of the Datastore, if it will be Thin, and if data reduction and Advanced Deduplication are enabled. The Pool selected must be an All Flash Pool and the Thin checkbox must be selected for the Data Reduction option to be available. Data Reduction must be enabled and the configuration must be supported for the Advanced Deduplication option to be available. An example of the **Storage** step when creating a VMware VMFS Datastore is found in Figure 13. The **Storage** step contains similar information when creating a VMware NFS Datastore.

Figure 13. Create VMware VMFS Datastore Wizard – Storage Step

## 4.2 Enabling and Disabling Data Reduction on an Existing Storage Resource

Dell EMC Unity Data Reduction, along with Advanced Deduplication, can be enabled or disabled on supported storage resources at any point in time. This can easily be accomplished using either Unisphere, Unisphere CLI, or REST API. The following sections outline how you can easily enable and disable the options on existing storage resources using Unisphere and Unisphere CLI.

### 4.2.1 LUNs

To enable and disable data reduction and Advanced Deduplication on an existing LUN, review the properties of the LUN from the **Block** page. The **Block** page can be accessed by selecting **Block** under **Storage** in the left Unisphere pane. After double clicking the Name of the LUN, or after selecting a LUN and clicking the **Pencil** (View/Edit) icon, the LUN Properties window is displayed. On the **General** tab, a checkbox for **Data Reduction** exists. Depending if data reduction is currently disabled or enabled on the storage resource, the box will either be unchecked or checked. To change the state of data reduction, simply check or uncheck the **Data Reduction** box and click **Apply**. Advanced Deduplication is only available once Data Reduction is enabled and the configuration supports it. Advanced Deduplication can be enabled or disabled independently to the Data Reduction setting. To change the state of Advanced Deduplication, simply check or uncheck the **Advanced Deduplication** box and click **Apply**.

Figure 14 below shows the Properties window of a LUN which currently has Data Reduction disabled. The **Data Reduction** box has been selected to enable data reduction on the resource, but Apply has not yet been selected. When enabling data reduction on a resource, an informational message is displayed after checking the **Data Reduction** box. A similar message appears when Data Reduction is enabled, and the Advanced Deduplication box has been checked but Apply has not been selected. The message when enabling Data Reduction states:

*Only the newly written data will have data reduction applied. Existing data will remain unchanged.*

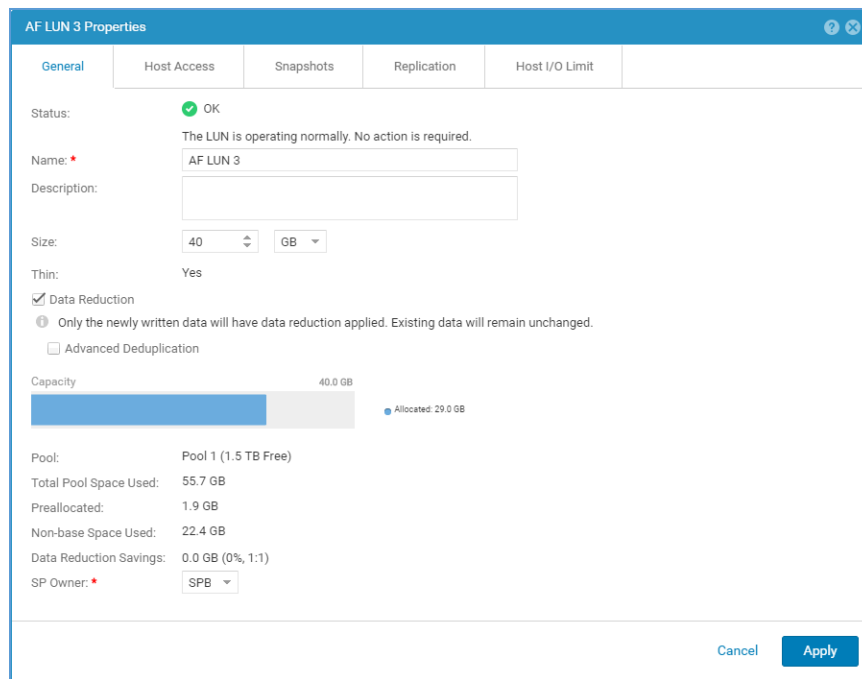


Figure 14. LUN Properties Window. Enabling Data Reduction Example

Figure 15 below shows the Properties window of a LUN which currently has Data Reduction and Advanced Deduplication enabled. The **Advanced Deduplication** box has been deselected to disable Advanced Deduplication on the resource, but Apply has not yet been selected. When disabling Advanced Deduplication on a resource, an informational message is displayed after unchecking the **Advanced Deduplication** box. A similar message appears when Data Reduction is disabled, but Apply has not been selected. The message when disabling Advanced Deduplication states:

*Newly written data will not have advanced deduplication applied. Existing data will remain unchanged.*

This message implies that the Local LUN Move option must be used to remove all Advanced Deduplication savings from the existing data within the storage resource if desired. If the Local LUN Move option is not utilized, then only overwrites to the LUN will cause Advanced Deduplication savings to decrease.

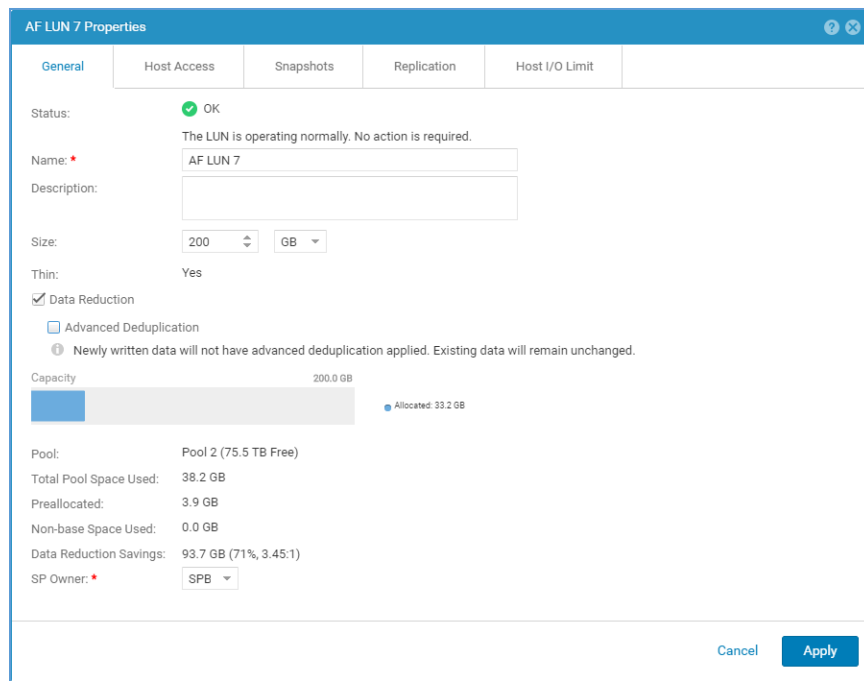


Figure 15. LUN Properties Window. Disabling Advanced Deduplication Example

## 4.2.2 Consistency Group LUNs

Enabling and disabling data reduction, along with Advanced Deduplication, on a LUN within a Consistency Group can be accomplished by first navigating to the **Consistency Group** tab within the **Block** page. From this page either double click the Name of a Consistency Group, or after selecting a Consistency Group click the Pencil (View/Edit) icon. This will open the Consistency Group Properties window. An example of this window can be seen in Figure 16.

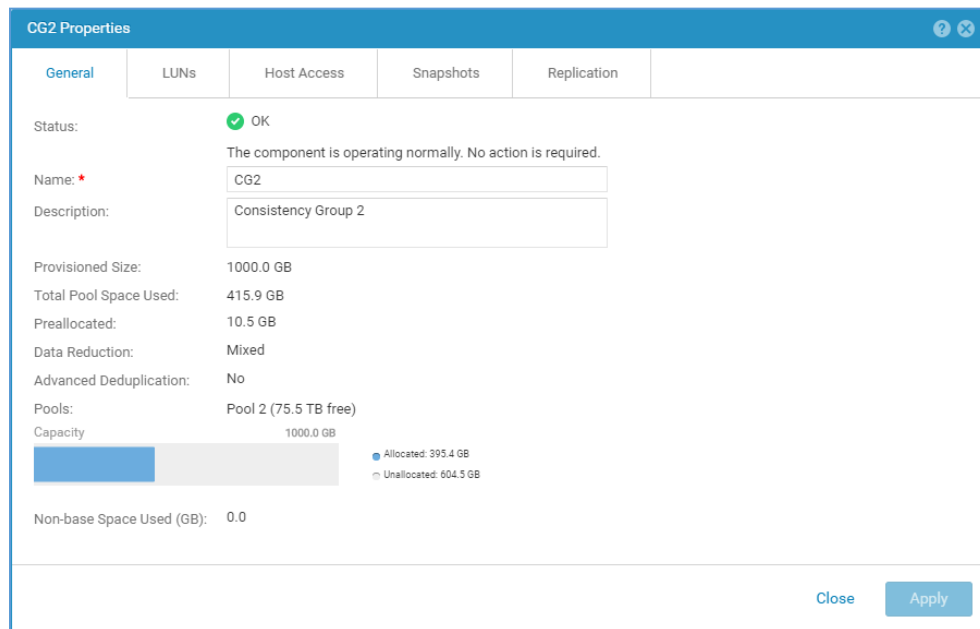


Figure 16. Consistency Group Properties Window



After opening the Consistency Group Properties window, navigate to the LUNs tab. From this screen you can view the current LUNs within the Consistency Group, and the current state of data reduction and Advanced Deduplication if the **Data Reduction** and **Advanced Deduplication** columns are displayed. A sample of the LUNs tab is shown in Figure 17. To edit the settings of one of the LUNs, double click the LUN or select the LUN and click the **Pencil** icon.

	!	LUN	↑	Size (GB)	Allocated (%)	Non-base Space Used (GB)	Thin	Pool	Data Reduction	Advanced Deduplication
<input type="checkbox"/>	✓	CG2 - LUN 1		250.0	<div style="width: 100%;"></div>	0.0	Yes	Pool 2	No	No
<input type="checkbox"/>	✓	CG2 - LUN 2		250.0	<div style="width: 100%;"></div>	0.0	Yes	Pool 2	No	No
<input type="checkbox"/>	✓	CG2 - LUN 3		250.0	<div style="width: 100%;"></div>	0.0	Yes	Pool 2	Yes	Yes
<input type="checkbox"/>	✓	CG2 - LUN 4		250.0	<div style="width: 100%;"></div>	0.0	Yes	Pool 2	Yes	Yes

Figure 17. Consistency Group – LUNs Tab

The LUN Properties window is now shown. This is the exact same Properties window as a LUN not currently in a Consistency Group. An example of this window is shown in Figure 18. As mentioned previously, to enable or disable data reduction, along with Advanced Deduplication on a LUN, simply check or uncheck the **Data Reduction** or **Advanced Deduplication** box and click **Apply**. The new state for data reduction and Advanced Deduplication are reflected in the LUNs tab within the Consistency Group Properties window.

**CG2 - LUN 3 Properties**

**General** | Host Access | Host I/O Limit

Status: ✓ OK  
The LUN is operating normally. No action is required.

Name: \* CG2 - LUN 3

Description:

Size: 250 GB

Thin: Yes

Data Reduction

Advanced Deduplication

Capacity: 250.0 GB  
Allocated: 57.9 GB

Pool: Pool 2 (75.5 TB Free)  
Total Pool Space Used: 63.4 GB  
Preallocated: 3.5 GB  
Non-base Space Used: 0.0 GB  
Data Reduction Savings: 124.2 GB (66%, 2.96:1)  
SP Owner: \* SPB

Close Apply

Figure 18. LUN Properties Window

## 4.2.3 File Systems

Data reduction can be enabled or disabled at any time on File Systems created on Dell EMC Unity OE version 4.2 or later if they reside within an All Flash Pool. Advanced Deduplication can also be enabled or disabled if the resource resides on a supported configuration. To modify these settings, first navigate to the **File** page by selecting **File** under **Storage** in the left Unisphere pane. In this window, all existing File Systems on the system are displayed. To modify the data reduction or Advanced Deduplication setting on a File System, double click on a supported File System to modify it, or select a File System and select the **Pencil** icon. The Properties window is now shown. To enable or disable data reduction, simply check or uncheck the Data Reduction box and click Apply. Likewise, to enable or disable the Advanced Deduplication setting on supported configurations, check or uncheck the Advanced Deduplication box and hit apply. A sample of the Properties window for a data reduction enabled File System is shown in Figure 19. The File System currently has Data Reduction disabled, and the Data Reduction box has been selected to enable data reduction on the resource, but Apply has not yet been selected. When enabling data reduction on a resource, an informational message is displayed after checking the **Data Reduction** box. A similar message is displayed when enabling Advanced Deduplication. The message when enabling Data Reduction on a File System states:

*Only the newly written data will have data reduction applied. Existing data will remain unchanged.*

This message outlines that only new writes to the File System or overwrites to existing data within the File System are considered for data reduction.

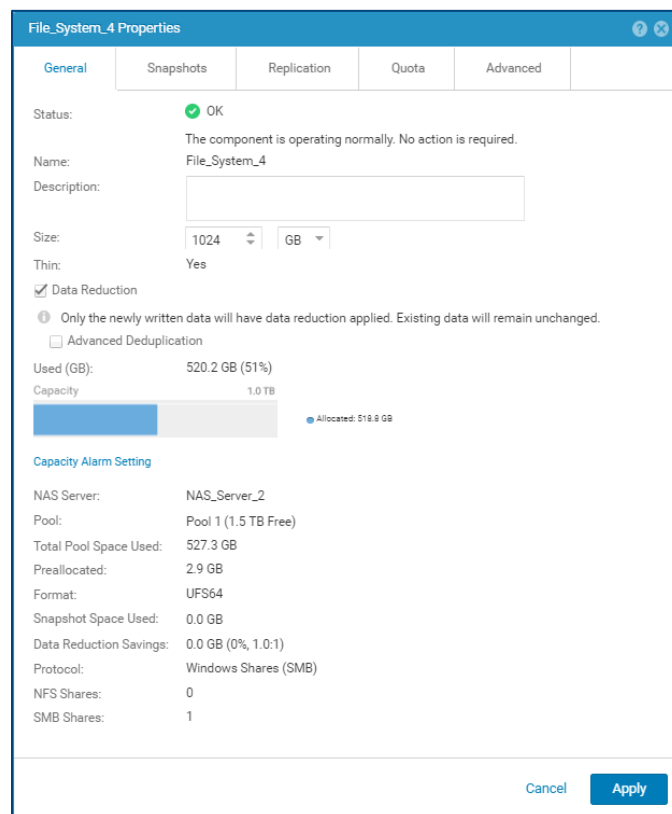


Figure 19. File System Properties Window – Enabling Data Reduction

When reviewing the Properties window for a File System which does not support data reduction, no option to enable data reduction will be shown. This is true for Thick File Systems, and any File Systems created on a code previous to Dell EMC Unity OE version 4.2. Figure 20 below shows an example of the File System Properties window for a Thick File System. As you can see no option for data reduction is shown.

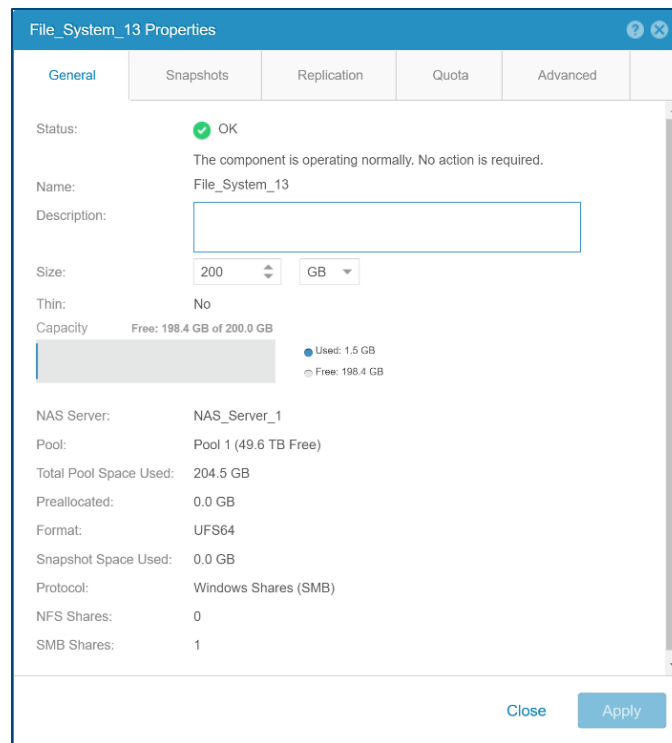


Figure 20. File System Properties Window – Data Reduction Not Supported

Figure 21 below shows the Properties window of a File System which currently has data reduction enabled. The Advanced Deduplication box has been deselected to disable Advanced Deduplication on the resource, but Apply has not yet been selected. When disabling Advanced Deduplication on a resource, an informational message is displayed after unchecking the Advanced Deduplication box. The message states:

*Newly written data will not have advanced deduplication applied. Existing data will remain unchanged.*

This message implies that only overwrites to the File System will cause data reduction savings to decrease.

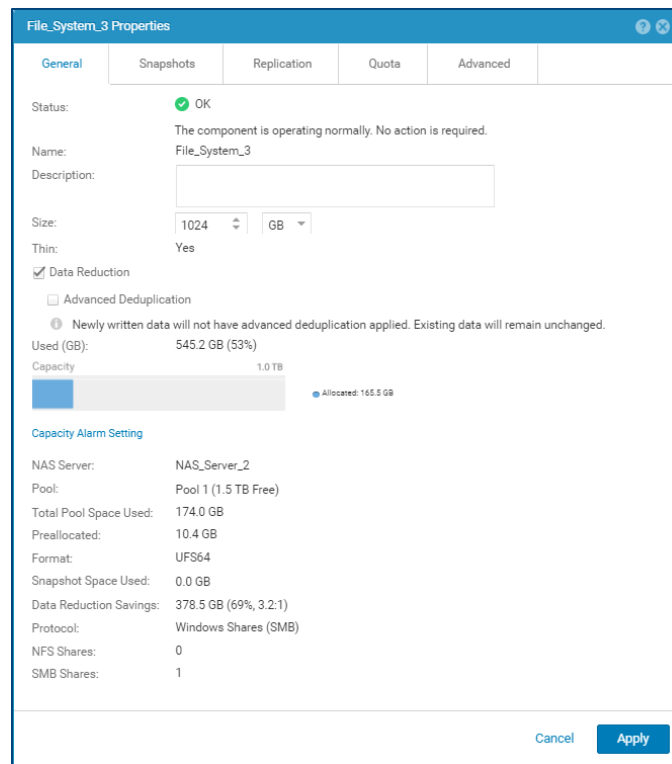


Figure 21. File System Properties Window – Disabling Advanced Deduplication

#### 4.2.4 VMware VMFS and NFS Datastores

Enabling or disabling data reduction and Advanced Deduplication on existing VMware VMFS or NFS Datastores can be accomplished at any time through Unisphere, Unisphere CLI, or REST API. To enable or disable these features in Unisphere, first navigate to the **VMware** page within Unisphere. To do so, select **VMware** under **Storage** in the left Unisphere pane. All existing VMware Datastores are displayed on this page. To modify the VMware Datastore, double click on a Datastore, or select the Datastore and select the **Pencil** icon. The Properties window is now shown. To enable or disable data reduction, simply check or uncheck the **Data Reduction** box and click **Apply**. To enable or disable Advanced Deduplication, check or uncheck the Advanced Deduplication box and hit **Apply**. A sample of the Properties window for a VMware VMFS Datastore is shown in Figure 22.

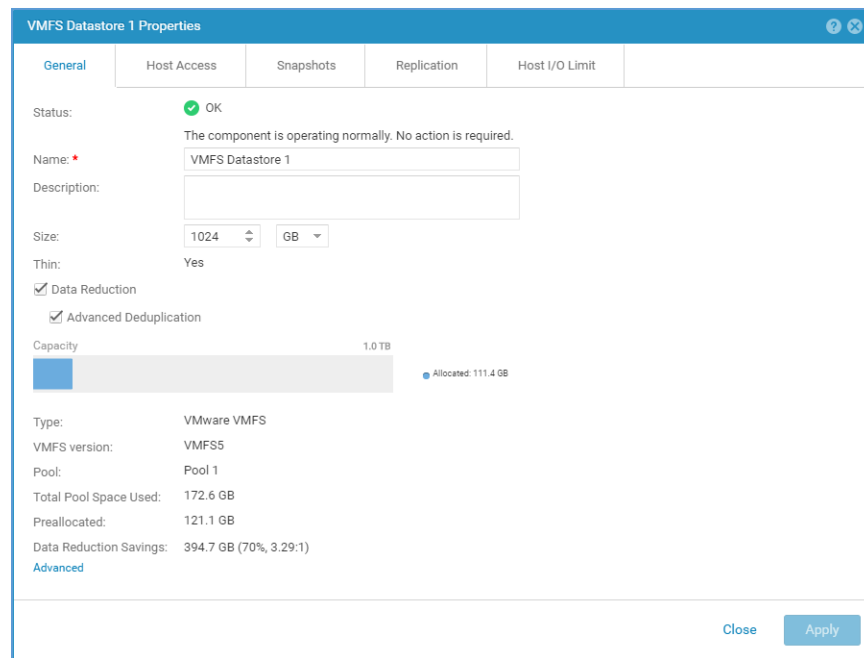


Figure 22. VMware VMFS Datastore – Properties Window

For VMware NFS Datastores, data reduction is only supported if the resource is Thin, resides on an All Flash Pool, and is created while the system is running Dell EMC Unity OE version 4.2 or later. If the resource supports data reduction, data reduction can be enabled or disabled at any time. If the resource is Thick, or was created on an earlier code, Unisphere will not display an option to enable data reduction. When attempting to enable or disable data reduction on a resource which does not support data reduction via Unisphere CLI or REST API, an error will be returned. Advanced Deduplication is only supported on resources which support Data Reduction, and currently reside on a configuration which supports Advanced Deduplication.

## 4.3 How to Determine Which Storage Resources Have Data Reduction Enabled

Dell EMC Unity Data Reduction is supported on Thin LUNs, whether standalone or within a Consistency Group, Thin File Systems, and Thin VMware VMFS and NFS Datastores. For configurations that support it, Advanced Deduplication can be enabled on resources with Data Reduction enabled. The following sections outline how you can easily determine what the current state of data reduction and Advanced Deduplication are on existing resources using Unisphere. The current state of data reduction and Advanced Deduplication can also be determined via Unisphere CLI and REST API, but won't be covered in the following sections.

### 4.3.1 LUNs

To review the status of data reduction and Advanced Deduplication on each of the LUNs created on the system, navigate to the **Block** page, which can be accessed by selecting **Block** under **Storage** in the left Unisphere pane. This page contains three columns specific for data reduction. The columns are **Data Reduction** column, which shows if Data Reduction is enabled or not on the resource, the **Advanced Deduplication** column, which shows if Advanced Deduplication is enabled, and the **Data Reduction Savings (GB)** column, which shows the amount of savings in GBs for the resource. To add these and other columns to the view, simply click the **Gear** Icon in the top right portion of the LUNs tab and select the columns to add under the **Columns** option. An example of this screen is shown in Figure 23.

Data reduction information has also been added to the quick properties view of the **LUN** tab on the **Block** page. After selecting a LUN, the right portion of the screen is populated with more information about the storage resource. In Figure 23, a data reduction enabled storage resource is selected. In the information provided in the right portion of the screen, you can determine if **Data Reduction** and Advanced Deduplication are enabled and the current **Data Reduction Savings** on the selected resource.

	Name	Size (GB)	Allocated (%)	Pool	Data Reduction	Data Reduction Savings (GB)	Advanced Deduplication
✓	AF LUN 1	20.0		Pool 1	Yes	19.7	Yes
✓	AF LUN 2	20.0		Pool 1	Yes	10.2	Yes
✓	AF LUN 3	40.0		Pool 1	No	0.0	No
✓	AF LUN 4	40.0		Pool 1	No	0.0	No
✓	AF LUN 5	100.0		Pool 2	Yes	362.2	Yes
✓	AF LUN 6	100.0		Pool 2	No	0.0	No
✓	AF LUN 7	200.0		Pool 2	Yes	91.5	Yes
✓	AF LUN 8	200.0		Pool 2	No	0.0	No
✓	AF LUN 9	1,024.0		Pool 1	Yes	242.7	Yes
✓	AF LUN 10	1,024.0		Pool 1	Yes	367.0	Yes
✓	AF LUN 11	1,024.0		Pool 2	No	0.0	No
✓	AF LUN 12	1,024.0		Pool 2	No	0.0	No

Figure 23. Block Page – LUNs Tab

### 4.3.2 Consistency Group LUNs

To review which Consistency Groups contain data reduction enabled LUNs, view the **Consistency Group** tab, which is found on the **Block** page. On this page, columns named **Data Reduction** and **Advanced Deduplication** can be added to the current view by clicking the **Gear Icon** and selecting **Data Reduction** or **Advanced Deduplication** under **Column**. An example of the Consistency Group tab with the **Data Reduction** and **Advanced Deduplication** columns displayed is shown in Figure 24. The **Data Reduction** and **Advanced Deduplication** columns have three potential entries, **No**, **Yes**, and **Mixed**. **No** is displayed if none of the LUNs within the Consistency Group have the option enabled. **Yes** is displayed if all LUNs within the Consistency Group have the option enabled. **Mixed** is displayed if there are a mix of LUNs with the option enabled and disabled within the Consistency Group. After selecting a Consistency Group, the quick properties view also displays the **Data Reduction** and **Advanced Deduplication** options, and **No**, **Yes**, or **Mixed**.

Name	Size (GB)	Allocated (%)	Data Reduction	Advanced Deduplication
CG1	1,000.0		No	No
CG2	1,000.0		Mixed	Mixed
CG3	1,000.0		Yes	No
CG3 - Dest - 3922	1,000.0		Yes	No

Figure 24. Block Page – Consistency Group Tab

### 4.3.3 File Systems

To review the status of data reduction on each of the File Systems, navigate to **File Systems** tab within the **File** page, which can be accessed by selecting **File** under **Storage** in the left Unisphere pane. This page contains three columns specific for data reduction. The columns are the **Data Reduction** column, which shows if data reduction is enabled or not on the resource, the **Advanced Deduplication** column, which shows if Advanced Deduplication is enabled, and the **Data Reduction Savings (GB)** column, which shows the amount of savings in GBs for the resource. To add these and other columns to the view, simply click the **Gear** Icon in the top right portion of the File Systems tab and select the new columns to add under the **Columns** option. An example of this screen is shown in Figure 25.

When reviewing the **Data Reduction**, **Advanced Deduplication**, and **Data Reduction Savings (GB)** columns, the information provided depends on if the storage resource supports data reduction or not. For Thick File Systems or File Systems created on a code prior to Dell EMC Unity OE version 4.2, -- is displayed in the columns to denote that the storage resource does not support data reduction. For Thin File Systems created on Dell EMC Unity OE version 4.2 or later, the **Data Reduction** column will display **Yes** or **No** depending on if data reduction is enabled or not. **No** is also displayed for File Systems created on Dell EMC Unity OE version 4.2 or later, and currently reside within a non-All Flash Pool. The **Advanced Deduplication** column is available for configurations which support Advanced Deduplication, and displays either **Yes** or **No** depending on the current state of Advanced Deduplication. The **Data Reduction Savings (GB)** column displays the amount of savings currently achieved within the File System. As data reduction savings are not removed when disabling data reduction on a storage resource, data reduction may be disabled, but savings still exist.

Data reduction information has also been added to the quick properties view of the **File Systems** tab on the **File** page. After selecting a File System, the right portion of the screen is populated with more information about the storage resource. In Figure 25, a data reduction enabled storage resource is selected. In the information provided in the right portion of the screen, you can determine if **Data Reduction** and **Advanced Deduplication** are enabled and the current **Data Reduction Savings** on the selected resource. This is an easy way to review the current state of data reduction on a specific resource if the **Data Reduction**, **Advanced Deduplication**, and **Data Reduction Savings (GB)** columns are not shown. The -- designation, which is explained above, is also used when the storage resource selected does not support data reduction.

Name	Size (GB)	Allocated (%)	Used (%)	Pool	Data Reduction	Data Reduction Savings (GB)	Advanced Deduplication
File_System_1	200.0			Pool 1	Yes	98.2	Yes
File_System_2	200.0			Pool 1	No	0.0	No
File_System_3	1,024.0			Pool 1	Yes	370.0	Yes
File_System_4	1,024.0			Pool 1	No	0.0	No
File_System_5	1,024.0			Pool 1	Yes	234.0	Yes
File_System_6	1,024.0			Pool 3	Yes	239.2	Yes
File_System_7	1,024.0			Pool 2	Yes	504.7	Yes
File_System_8	1,024.0			Pool 2	No	0.0	No
File_System_9	1,024.0			Pool 2	Yes	452.2	Yes
File_System_10	1,024.0			Pool 2	No	0.0	No
File_System_11	1,024.0			Pool 2	Yes	172.7	Yes
File_System_12	1,024.0			Pool 2	No	0.0	No
File_System_13	250.0			Pool 1	--	--	--

Figure 25. File Page – File Systems Tab

#### 4.3.4 VMware VMFS and NFS Datastores

To review the status of data reduction on each of the VMware VMFS and NFS Datastores created on the system, navigate to the **VMware** page, which can be accessed by selecting **VMware** under **Storage** in the left Unisphere pane. This page contains three columns specific for data reduction. The columns are the **Data Reduction** column, which shows if data reduction is enabled or not on the resource, the **Advanced Deduplication** column, which shows if Advanced Deduplication is enabled, and the **Data Reduction Savings (GB)** column, which shows the amount of savings in GBs for the resource. To add these and other columns to the view, simply click the **Gear** Icon in the top right portion of the Datastores tab and select the columns to add under the Columns option. An example of this screen is shown in Figure 26.

When reviewing the **Data Reduction**, **Advanced Deduplication**, and **Data Reduction Savings (GB)** columns, the information provided depends on if the storage resource supports data reduction or not. For Thick NFS Datastores or NFS Datastores created on a code prior to Dell EMC Unity OE version 4.2, -- is displayed in the columns to denote that the storage resource does not support data reduction. For Datastores which support data reduction, the **Data Reduction** column will display **Yes** or **No** depending on if data reduction is enabled or not. **No** is also displayed for Datastores which support data reduction, but currently reside within a non-All Flash Pool. The **Advanced Deduplication** column is available for configurations which support Advanced Deduplication, and displays either **Yes** or **No** depending on the current state of Advanced Deduplication. The **Data Reduction Savings (GB)** column displays the amount of savings currently achieved within the File System. As space savings are not removed when disabling data reduction on a storage resource, data reduction may be disabled, but savings still exist.

Data reduction information is also added to the quick properties view of the **Datastores** tab on the **VMware** page. After selecting a VMware Datastore, the right portion of the screen is populated with more information about the storage resource. In Figure 26, a data reduction enabled storage resource has been selected. In the information provided in the right portion of the screen, you can determine if **Data Reduction** and **Advanced Deduplication** are enabled, and the current **Data Reduction Savings** on the selected resource. This is an easy way to review the current state of data reduction on a specific resource if the **Data Reduction**, **Advanced Deduplication**, and **Data Reduction Savings (GB)** columns are not shown. The -- designation, which is explained above, is also used when the storage resource selected does not support data reduction.



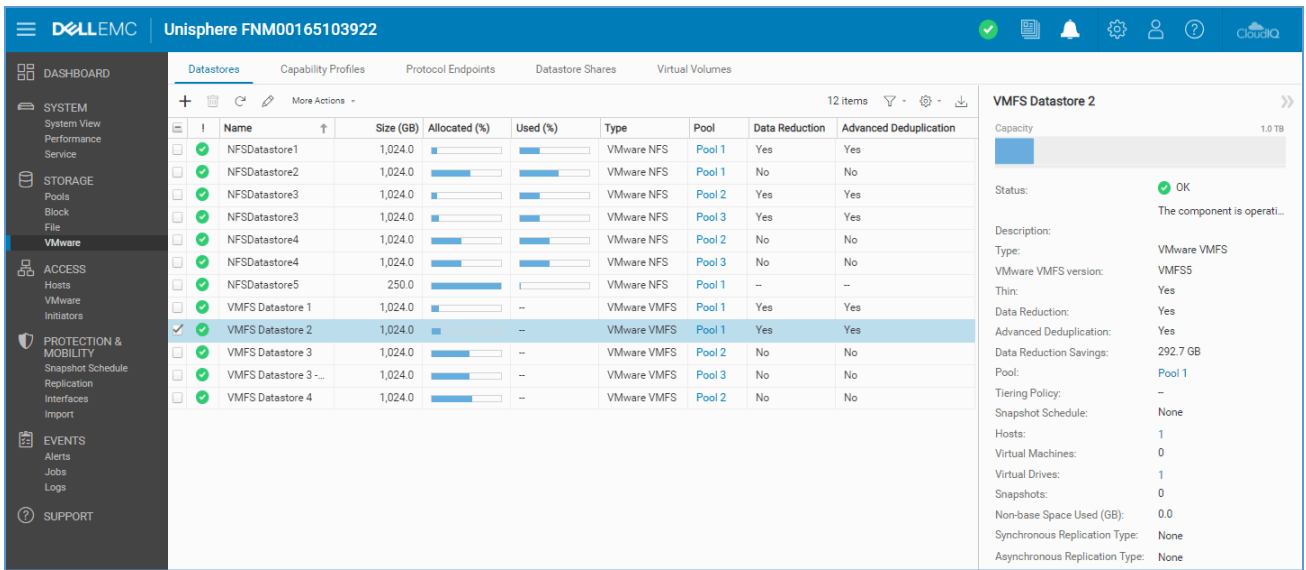


Figure 26. VMware Page – Datasets Tab

## 4.4 Local LUN Move

The Local LUN Move feature, also known as Move, provides native support for moving LUNs and VMFS Datastores online between pools or within the same pool. This ability allows for manual control over load balancing and rebalancing of data between pools. Local LUN Move leverages Transparent Data Transfer (TDX) technology, a multi-threaded, data copy engine. TDX is designed to help ensure a low impact on host access during a move session. In Dell EMC Unity OE version 4.3, the Compress Now option has been deprecated.

Local LUN Move can also be leveraged to migrate a Block resource’s data to or from a resource with data reduction and/or Advanced Deduplication enabled. When Move is utilized to migrate data to a resource with data reduction enabled, as the data is migrated via the TDX data copy engine, all data will pass through the data reduction logic. If Advanced Deduplication is supported and enabled, the data will also pass through the Advanced Deduplication algorithm. This allows space savings to be achieved during the migration. When migrating to a resource with data reduction disabled, all space savings achieved on the source will be removed during the migration.

The Move option can be found on the **Block** page for LUNs. After selecting a storage resource, select the **More Actions** drop down list, then **Move**. This launches the **Move** dialog box. An example of the **More Actions** dialog box and the **Move** option can be found in Figure 27.




Figure 27. Block Page – Move Under More Actions

The Move LUN dialog box is shown in Figure 28 below. In this window the user will select the destination Pool for the new storage resource. The user will also select if the resource will be Thin, and if it will have data

reduction and Advanced Deduplication enabled. The resource must be Thin for the data reduction option to be available. The Move operation is completely transparent to the host. To start the Move operation, simply select **Yes** in the Move LUN dialog box.

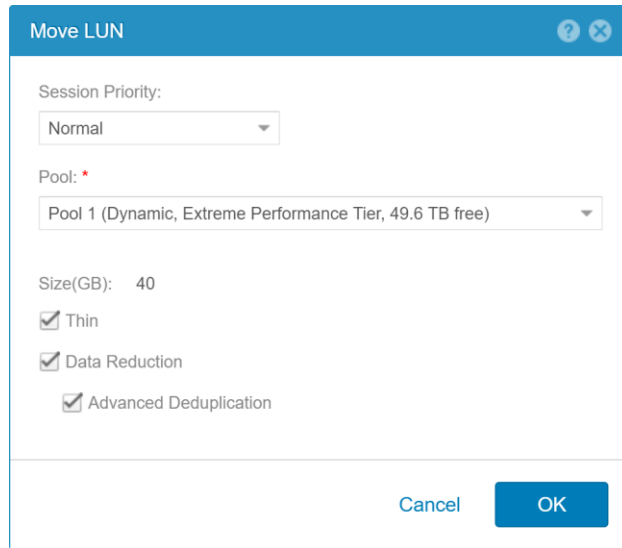


Figure 28. Move dialog box

Move can also be used for LUNs contained within a Consistency Group. The Move option for Consistency Groups is found on the LUNs tab within the Properties window of a Consistency Group. An example of this is shown in Figure 29. After selecting a LUN within the Consistency Group, select **More Actions**, then **Move**. A similar dialog box as previously shown in Figure 28 is displayed. After selecting **Yes**, the data is moved.

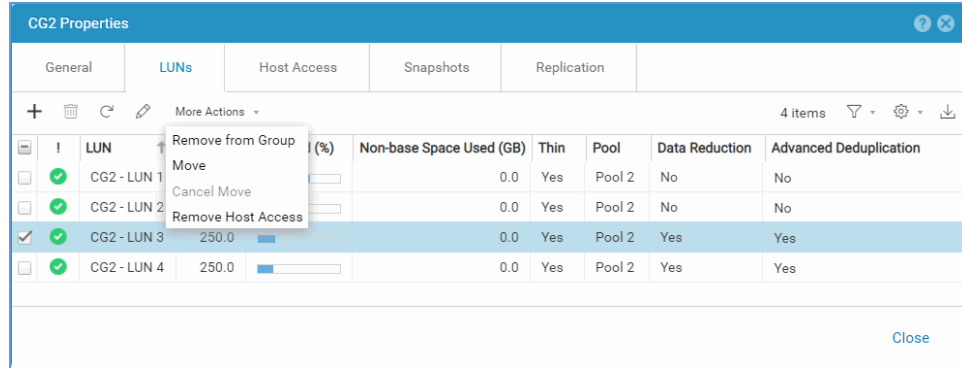


Figure 29. Move – Consistency Group Properties Window – LUN Tab

Move can also be used with VMware VMFS Datastores. For VMware VMFS Datastores, the Move option is found on the **Datastores** tab of the **VMware** page. As shown in Figure 30, after selecting a VMware VMFS Datastore, the **Move** option can be found under **More Actions**. After selecting Move, a dialog box similar to Figure 28 is displayed.

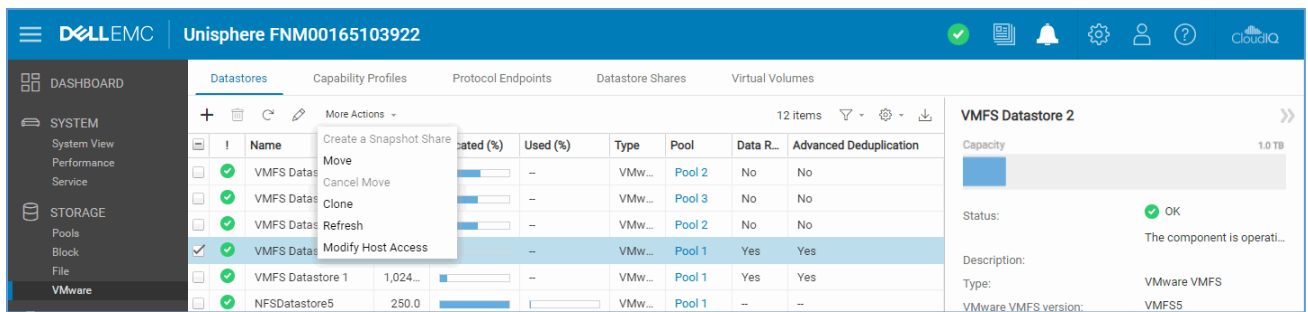


Figure 30. Move – VMware VMFS Datastores

## 4.5 Savings Reporting

Dell EMC Unity Data Reduction provides savings information at many different levels within the system, and in many different formats. Savings information is provided at the individual storage resource, the Pool, and the System level. For each of the levels, savings information is reported in GBs, percent savings, and as a ratio. The total savings not only includes savings due to data reduction on the storage resource, but also Advanced Deduplication savings and savings which are realized on any Snapshots and Thin Clones taken of the resource. All savings information is aggregated and then displayed at the Pool level and System level. The following sections outline where savings information can be viewed in Unisphere. Savings information is also accessible via Unisphere CLI and REST API, but won't be covered in the following sections.

### 4.5.1 Storage Resource Level

Space savings information in the three formats are available within the Properties window of the storage resource. For LUNs, you will either need to access the Properties page from the **Block** page, or on the **LUN** tab from within the **Consistency Group** Properties window. As shown in Figure 31, data reduction savings are shown on the **General** tab within the LUN Properties Window.

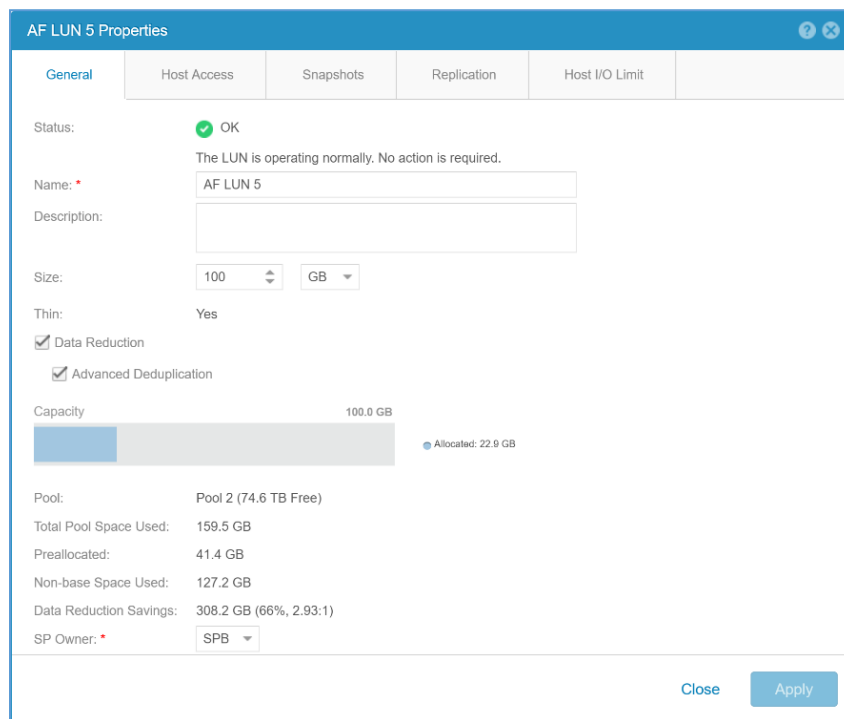


Figure 31. LUN Properties Window – Data Reduction Savings

File Systems also provide the same savings information as LUNs. Figure 32 below shows an example of the Properties window of a File System. As with LUNs, the data reduction savings are reported in GBs, % savings, and ratio.

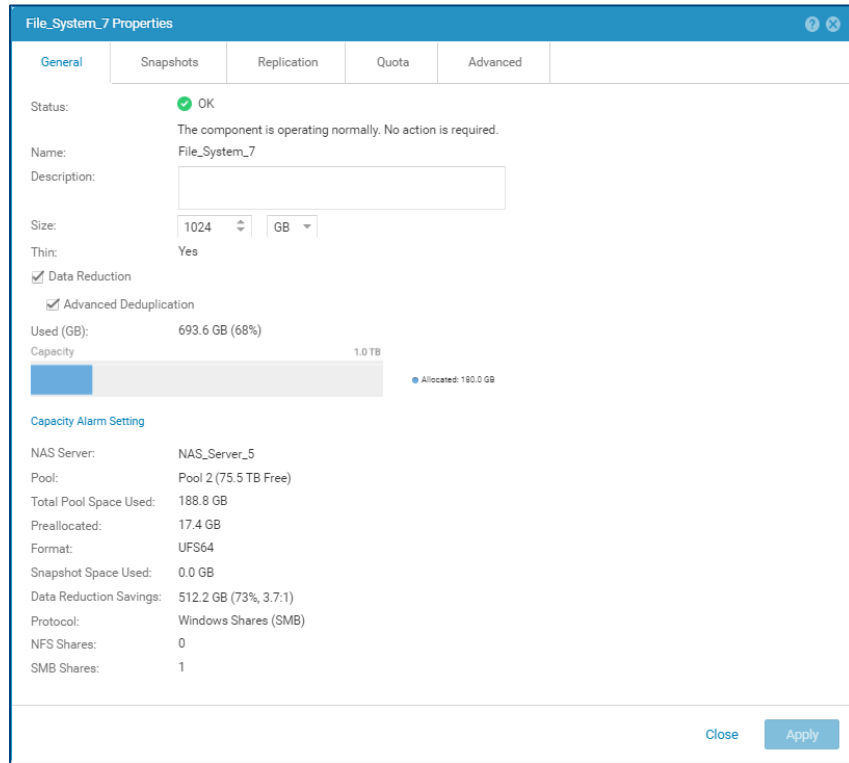


Figure 32. File System Properties Window – Data Reduction Savings

VMware VMFS Datastores display data reduction savings in the same manner as LUNs and File Systems. In the VMware VMFS Properties window on the General tab, the GBs saved, % savings, and ratio can be viewed. An example of this screen can be found in Figure 33.

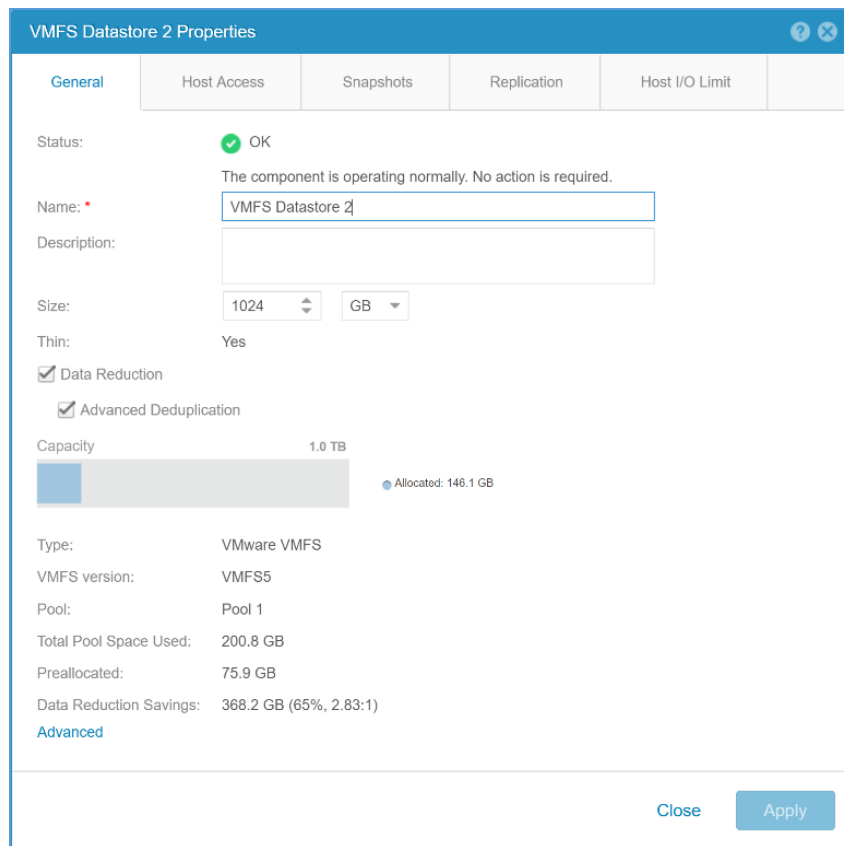


Figure 33. VMware VMFS Datastore Properties Window – Data Reduction Savings

VMware NFS Datastores display data reduction savings in the same manner as VMware VMFS Datastores. In the VMware NFS Properties window on the General tab, the GBs saved, % savings, and ratio can be viewed. An example of this screen can be found in Figure 34.

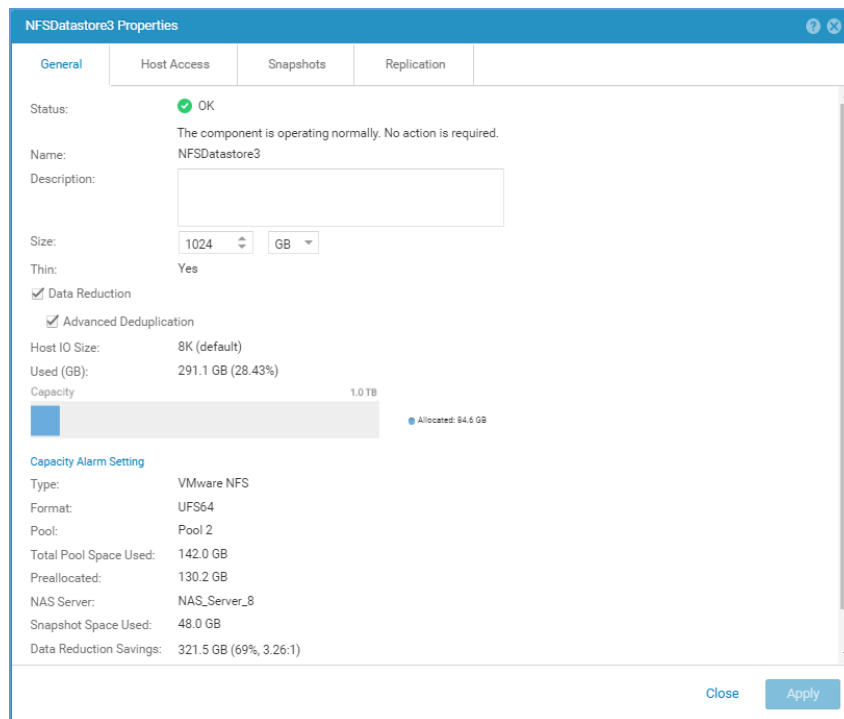


Figure 34. VMware NFS Datastore Properties Window – Data Reduction Savings

### 4.5.2 Pool Level

Data reduction information is also aggregated to the Pool level. Savings are reported in the three formats, which include the GBs saved, % savings, and ratio, on the **Usage** tab within the **Pool** Properties Window. The GBs savings reflects the total amount of space saved due to data reduction on storage resources and their Snapshots and Thin Clones. An example of the Usage tab can be seen in Figure 35.

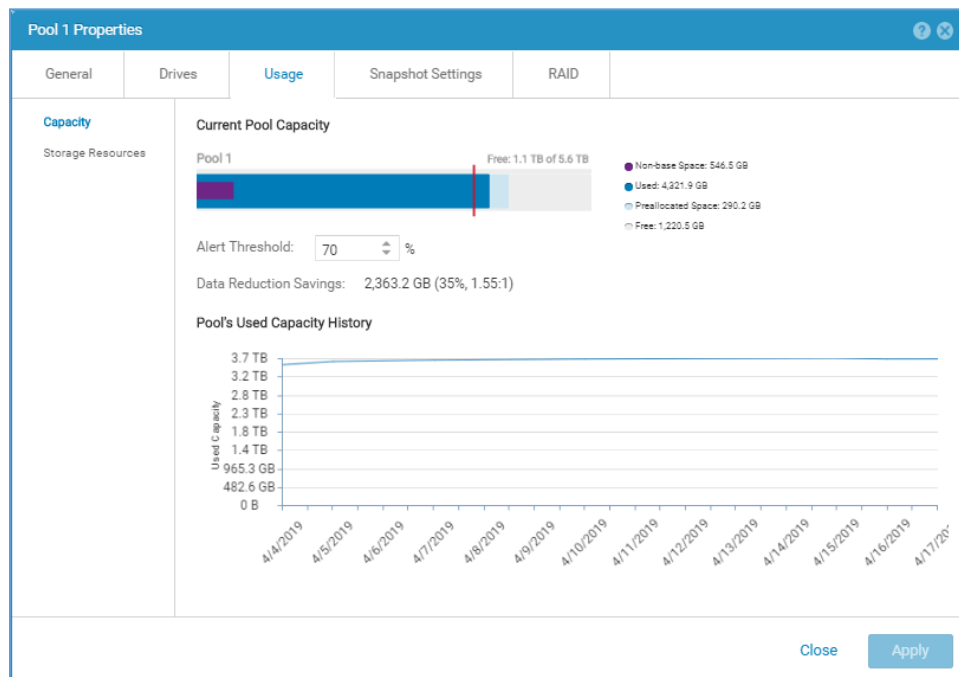


Figure 35. Pool Properties Window – Usage Tab – Data Reduction Savings

### 4.5.3 System Level

Data reduction savings information is also available at the System Level. The system level aggregates all savings across the entire system and displays them in the three formats available, GBs saved, % saved, and ratio. For the GBs saved, this value is the total amount of space saved due to data reduction, along with savings achieved by Snapshots and Thin Clones of data reduction enabled storage resources. In the Dell EMC Unity OE 4.5 release, the system level data reduction savings information is displayed within the System Efficiency view block found on the system Dashboard page. To see the data reduction savings in the three values mentioned previously, hover over the Data Reduction value. In codes prior to this release, data reduction savings information is displayed in the System Capacity view block. If the view block is not shown on your system, you can add it by selecting the **Main** tab, clicking **Customize**, and adding the view block. An example of the System Efficiency view block, which is available in Dell EMC Unity OE version 4.3 and later, is shown in Figure 36.

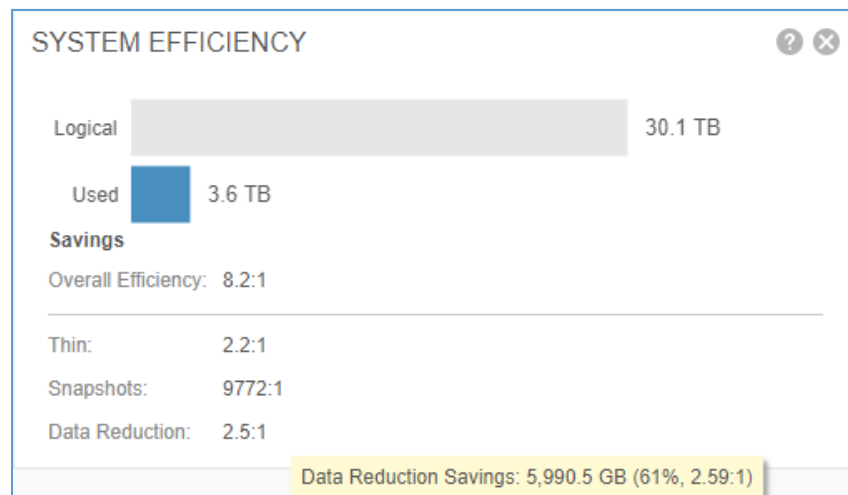


Figure 36. System Efficiency View Block – Unisphere Dashboard Page

### 4.5.4 Calculating Savings

In Dell EMC Unity OE version 4.5, multiple changes were made to space reporting within the system. The space reporting updates affect the System, Pool, and Storage Resource values. With these changes, users are able to utilize the formulas, displayed in Figure 37, to calculate and verify the data reduction savings percentage and ratio for the System, Pools, and Storage Resources.

$$\text{Savings Ratio} = \frac{(\text{Total Pool Space Used} + \text{Data Reduction Savings})}{\text{Total Pool Space Used}} : 1$$

$$\text{Savings Percentage} = \frac{\text{Data Reduction Savings}}{(\text{Data Reduction Savings} + \text{Total Pool Space Used})} * 100$$

Figure 37. Data Reduction Savings Formulas

## 5 Interoperability

Dell EMC Unity Data Reduction is supported on standalone LUNs, LUNs contained within a Consistency Group, File Systems, or VMware VMFS and NFS Datastores. All software features on a storage system are supported with data reduction. Data reduction algorithms also support Local LUN Move which leverages the Transparent Data Transfer (TDX) technology and offload data transfer operations sent to the system. The following sections talk specifically about certain features of the Dell EMC Unity storage system, and how they relate to data reduction.

### 5.1 Data at Rest Encryption

Data at Rest Encryption, also known as D@RE, is fully supported on systems utilizing Dell EMC Unity Data Reduction. Data reduction is not impacted by Data at Rest Encryption, as all data reduction operations occur on data residing in System Cache. For data being written to disk, the data is first passed through the data reduction algorithms within System Cache, then written through hardware-based encryption modules to the backend drives. For reads from disk, the data is first decrypted and saved into System Cache before the data reduction algorithm recreates the original data and sends it to the host.

More information on Data at Rest Encryption can be found in the *Dell EMC Unity: Data at Rest Encryption* white paper found on Dell EMC Online Support.

### 5.2 Replication

Storage Resources utilizing data reduction can be replicated using any supported replication software, such as Native Synchronous or Asynchronous Replication to any supported destination system. All data replicated, regardless if it is local replication or to a remote system, is first restored to its original size then replicated to the destination. This method of replicating data reduction enabled storage resources ensures that all replication topologies are supported as if data reduction is not enabled on the resource. Replicating to systems which do not support data reduction and/or Advanced Deduplication is also supported, such as replicating to Dell EMC UnityVSA or a physical Dell EMC Unity system with a configuration which does not support data reduction and/or Advanced Deduplication.

Dell EMC Unity Data Reduction and Advanced Deduplication can also be enabled on only the source, only the destination, or both the source and destination storage resources, depending on if the system and Pool configuration support Dell EMC Unity Data Reduction and/or Advanced Deduplication. This allows the user to fully control where to implement data reduction. One example of a supported replication configuration is when utilizing Asynchronous Local Replication. The source storage resource may reside on an All Flash Pool and have data reduction enabled, but the destination may be on a large capacity Hybrid Pool which does not support data reduction. Another example of a supported configuration is when replicating a storage resource from a Dell EMC UnityVSA system or a production system not utilizing data reduction, to a storage resource with data reduction and Advanced Deduplication enabled on a remote system. Replication can also occur between a resource created before the system was running Dell EMC Unity OE version 4.4, and one after.

More information on Replication can be found in the *Dell EMC Unity: Replication Technologies* white paper found on Dell EMC Online Support.

### 5.3 Snapshots

The Dell EMC Unity Snapshots feature is fully supported with data reduction and Advanced Deduplication. Snapshots also benefit 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 or deduplicated. The data is left in its current state, and the Snapshot inherits the savings achieved on the source storage resource.

When a snapshot is mounted and the source storage resource has data reduction enabled, data reduction is also utilized on any snapshot I/O. If a read is received for a compressed or deduplicated block of data, the data is restored to its original size and sent to the requestor. 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 savings is tracked and reported as part of the GBs saved for the source storage resource.

More information on Snapshots can be found in the *Dell EMC Unity: Snapshots and Thin Clones* white paper found on Dell EMC Online Support.

## 5.4 Thin Clones

A Thin Clone is a read-write copy of a Thin Block storage resource, such as a LUN, Consistency Group, or VMware VMFS Datastore, that shares blocks with the parent resource. Along with sharing blocks with the source, Thin Clones also share the data reduction and Advanced Deduplication settings, which can only be set on the base storage resource. When data reduction is enabled on the source storage resource, any writes to a Thin Clone are also passed through the data reduction algorithm. All savings achieved on Thin Clones are also reported as a cumulative total with the savings on the source storage resource.

Data reduction and Advanced Deduplication can be enabled or disabled on the source storage resource at any time. Changing the data reduction setting will also control if data reduction is enabled or disabled on all Thin Clones residing on the storage resource. This is also true for the Advanced Deduplication setting. When data reduction or Advanced Deduplication is enabled on the source, no existing data is changed unless overwritten. When data reduction or Advanced Deduplication is disabled, all data is left in its current state. While a Thin Clone exists for a storage resource, Local LUN Move is not available on the source. Also, Local LUN Move is not available for use on a Thin Clone.

More information on Thin Clones can be found in the *Dell EMC Unity: Snapshots and Thin Clones* white paper found on Dell EMC Online Support.

## 5.5 Dell EMC Unity Native File and Block Import

Dell EMC Unity Native File and Block Import was first introduced in Dell EMC Unity OE version 4.1. In this release, VNX1 or VNX2 Block and File (NFS) resources could be migrated to Dell EMC Unity using the native import feature. In Dell EMC Unity OE version 4.2, support for Import was extended to VNX1 or VNX2 File (CIFS) configurations. When configuring an Import Session, data reduction and Advanced Deduplication are supported on the destination as long as the destination system and Pool configuration supports it. When creating an Import Session, if the destination resource supports data reduction, a checkbox is available to enable it on the destination resource. An option for Advanced Deduplication is also shown for configurations which support it. As data is migrated from the source VNX system to the Dell EMC Unity system, it passes through the data reduction algorithm as it is written to the Pool.

More information on Dell EMC Unity Native File and Block Import can be found in the *Dell EMC Unity: Migration Technologies* white paper found on Dell EMC Online Support.

## 5.6 Pool Expansion

Expanding a Pool residing on a Hybrid Flash System which contains data reduction savings is only supported if it is being expanded by supported Flash Drives. For instance, a Pool containing SAS Flash 2 or SAS Flash 3 drives can be expanded by adding more SAS Flash 2 or SAS Flash 3 drives to the Pool. While storage resources exist within the Pool which have data reduction enabled or have had data reduction enabled in the past, adding SAS or NL-SAS drives to convert the Pool to a Hybrid Pool is not supported.

To expand and convert an All Flash Pool to a Hybrid Pool, all storage resources which have data reduction enabled or have used data reduction must be removed from the Pool. For Block resources, Local LUN Move can be used to move the resource's data to another resource on the same Pool with data reduction disabled, or to another Pool. When utilizing Local LUN Move to relocate a Block device within the same Pool, you must ensure data reduction is disabled on the destination device. For File resources, the data must be migrated to a new resource, either by leveraging Dell EMC Unity Asynchronous Replication or a host based migration tool. Once all resources which have utilized data reduction have been removed from the Pool, the expansion will be allowed.

More information on migration options can be found in the *Dell EMC Unity: Migration Technologies* white paper found on Dell EMC Online Support.

## 6 Conclusion

Dell EMC Unity storage systems offer a powerful capacity efficiency feature which can improve the effective capacity utilization of a Dell EMC Unity system. Data reduction is included with all physical Dell EMC Unity systems at no additional cost. Advanced Deduplication is also supported at no additional cost on configurations which support it. When data reduction is utilized, not only is space saved due to the storage resources being Thin, but savings are achieved by utilizing intelligent zero detect, deduplication, and compression algorithms as well. Dell EMC Unity Snapshots and Thin Clones also save space within the system, which can greatly reduce the amount of storage needed for a dataset. By reducing the amount of storage needed to store a dataset, Dell EMC Unity Data Reduction helps to further reduce the Total Cost of Ownership of a Dell EMC Unity system.

## A Technical support and resources

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

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

### A.1 Related resources

The following documents can be found on Dell EMC Online Support:

- Dell EMC Unity: Best Practices Guide
- Dell EMC Unity: Cloud Tiering Appliance (CTA)
- Dell EMC Unity: Compression
- Dell EMC Unity: Compression for File
- Dell EMC Unity: Data at Rest Encryption
- Dell EMC Unity: Data Integrity
- Dell EMC Unity: DR Access and Testing
- Dell EMC Unity: Dynamic Pools
- Dell EMC Unity: FAST Technology Overview
- Dell EMC Unity: File-Level Retention (FLR)
- Dell EMC Unity: High Availability
- Dell EMC Unity: Introduction to the Platform
- Dell EMC Unity XT: Introduction to the Platform
- Dell EMC Unity: NAS Capabilities
- Dell EMC Unity: MetroSync
- Dell EMC Unity: MetroSync and Home Directories
- Dell EMC Unity: MetroSync and VMware vSphere NFS Datastores
- Dell EMC Unity: Migration Technologies
- Dell EMC Unity: OpenStack Best Practices for Ocata Release
- Dell EMC Unity: Performance Metrics
- Dell EMC Unity: Snapshots and Thin Clones
- Dell EMC Unity: Operating Environment (OE) Overview
- Dell EMC Unity: Replication Technologies
- Dell EMC Unity: Unisphere Overview
- Dell EMC Unity: Virtualization Integration
- Dell EMC UnityVSA
- Dell EMC Unity Cloud Edition with VMware Cloud on AWS
- Dell EMC Unity Data Reduction Analysis
- Dell EMC Unity: Migrating to Dell EMC Unity with SAN Copy
- Dell EMC Unity Storage with Microsoft Hyper-V
- Dell EMC Unity Storage with Microsoft SQL Server
- Dell EMC Unity Storage with Microsoft Exchange Server
- Dell EMC Unity Storage with VMware vSphere
- Dell EMC Unity Storage with Oracle Databases
- Dell EMC Unity 350F Storage with VMware Horizon View VDI
- Dell EMC Unity: 3,000 VMware Horizon Linked Clone VDI Users
- Dell EMC Storage with VMware Cloud Foundation

Other resources:

- Dell EMC Online Support
- Dell EMC Unity InfoHub (<http://bit.ly/unityinfohub>)
- Unisphere Command Line Interface User Guide