

Dell EMC PowerMax: Data Reduction

Inline compression and deduplication

Abstract

Dell EMC™ PowerMax Data Reduction is a storage-efficiency feature that combines inline compression and inline deduplication. Using both storage efficiency features together enhances capacity savings while maintaining great performance and reliability.

February 2021

Revisions

Date	Description
May 2018	Initial release
September 2020	Updates for PowerMaxOS Q3 2020 release
February 2021	Updated for PowerMaxOS Q1 2021 release

Acknowledgments

Author: Robert Tasker

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.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

This document may contain language from third party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's own content. When such third party content is updated by the relevant third parties, this document will be revised accordingly.

Copyright © 2018–2021 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. [2/5/2021] [Technical White Paper] [H17072.3]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Executive summary.....	4
1 Data Reduction overview	5
1.1 Adaptive Compression Engine overview.....	5
1.2 Deduplication overview.....	5
1.3 Terminology	5
2 Configuration details.....	7
3 Adaptive Compression Engine	8
3.1 Hardware acceleration.....	8
3.2 Optimized data placement.....	8
3.3 Activity Based Compression.....	8
3.4 Fine grain data packing	8
3.5 Extended Data Compression.....	9
4 Deduplication.....	10
4.1 Hardware acceleration.....	10
4.2 Deduplication algorithm	10
4.3 Hash table.....	10
4.4 Dedupe Management Object.....	10
4.5 Data Reduction I/O flow.....	10
5 Capacity usage.....	12
6 Data Reduction Settings.....	13
6.1 System and Storage Resource Pool	13
6.2 Storage Group	13
7 System Efficiency	14
7.1.1 System Resource Usage.....	16
7.1.2 Storage group compression in Unisphere	19
8 Supported data services.....	21
8.1 Local replication (SnapVX).....	21
8.1.1 Nocopy sessions (SnapVX, VP Snap)	21
8.1.2 Copy sessions (SnapVX Full Copy Linked Targets, Clone, Mirror)	21
8.2 Remote replication (SRDF)	21
8.3 Data at Rest Encryption (D@RE).....	22
8.4 Virtual Volumes.....	22
9 Conclusion.....	23
A Technical support and resources	24

Executive summary

Data reduction with the [Dell EMC PowerMax](#) system delivers a boost to system efficiency by combining inline compression, inline deduplication and pattern detection. Paring these techniques of capacity savings creates a system where users can achieve great capacity savings on reducible data. Data Reduction not only compresses data it also eliminates redundant copies of compressed data and delivers great performance. The contents of this technical white paper are intended to inform the reader how Data Reduction functions within the Dell EMC™ PowerMax systems.

1 Data Reduction overview

In PowerMax data storage systems data reduction combines the Adaptive Compression Engine (ACE) and inline deduplication to provide a highly performing space efficient platform. Data reduction allows users to write more host data than the total amount of usable capacity available. Compression and deduplication are two different functions that work together. Compression reduces the size of datasets and deduplication identifies identical datasets and stores a single instance. Performing these functions in parallel allows the system to be capacity efficient and deliver optimal performance.

1.1 Adaptive Compression Engine overview

The Adaptive Compression Engine (ACE) is the combination of multiple components to deliver the performance expected from an all flash storage system and maintain data storage efficiency. Incoming data is compressed inline using hardware-based compression with software compression in place and used as needed. Intelligent algorithms learn from the incoming workload to dynamically create a customized backend catering to the incoming workload. The Adaptive Compression Engine changes the backend compression pool layout as needed to ensure the system operates at optimal levels for both performance and space efficiency. Using internal statistics, algorithms identify the busiest data in the system allowing it to skip the compression process. The result minimizes decompression overhead on the system for data that is accessed the most. Working together, these functions allow data reduction to deliver great performance and efficiently manage back-end capacity usage.

1.2 Deduplication overview

Deduplication (dedupe) is a capacity-savings method that identifies identical copies of data and stores a single instance. There are a few facets of deduplication that are needed for it to provide efficient capacity savings.

- **Hash ID:** The Hash ID is a unique identifier for incoming data that is used to determine if a dedupe relationship is needed. The system uses a SHA-256 algorithm to generate the Hash ID.
- **Hash ID Table:** Hash Tables are an allocation of system memory distributed between the system directors. These tables are a catalog of the Hash IDs used by the dedupe process to determine if a dedupe relationship is needed or if the data can be stored on disk.
- **Dedupe Management Object (DMO):** The DMO manages the pointers for deduped data between front-end devices and the data stored on disk. This also manages what Hash Table the Hash IDs are stored in when dedupe relationships exist.

1.3 Terminology

Data Reduction: The use of compression, deduplication and pattern detection to reduce capacity usage and the cost of physical storage. In systems prior to the release of Dell EMC PowerMax, data reduction is compression only.

Data Reduction Reservation: A system Cache reservation used to determine the maximum savings that can be achieved when data reduction is enabled. This reservation acts as a system guardrail for management of resources such as disk capacity and cache in order to maintain a balance between performance and efficient capacity usage.

Storage Group Compression Ratio: The compression ratio displayed for allocations related to a specific storage group. This value may be greater than or less than the system or storage resource pool data reduction ratio displayed in management application displays.

Compressibility: The maximum compression ratio that may be achieved for either a storage group or a device. This value may be presented as a higher value than current savings due to the design activity-based compression (ABC refer to section 3.3).

Data Reduction Ready: The state of the system when the default storage resource pool (SRP) can store compressed data. For a system to be able to compress data it must have one compression I/O module per director, have compression enabled and have a system data reduction reservation ratio set.

Data Reduction Capable: A system installed with at least the Q1 2018 PowerMaxOS where the data reduction reservation applied to the system IMPL is 1.0:1.

Compression Pool: The collection of data devices configured within the physical disks where the track size is the same. For example, the 64KB pool is made up of data devices where all the device's tracks are 64KB in size.

Terabytes Usable (TBu): The backend usable storage capacity in the absence of compression referring to the amount of physical storage in the system.

Example: 50 TBu is 50 Terabytes of usable physical storage.

Terabytes Effective (TBe): The front-end effective storage capacity in the presence of data reduction. This represents the potential maximum amount of host or application data that can be written to the array.

Example: 50 TBu of physical storage with a data reduction reservation of 3:1 translates to 150 TBe. The total TBe value can be achieved assuming the data consuming capacity on the array is reducible at a level equal to or greater than the data reduction reservation set in the system.

2 Configuration details

PowerMaxOS is supported on PowerMax and VMAX™ All Flash data storage arrays. There are a few different scenarios for the two storage arrays. See Table 1 for details.

Table 1 Configuration details by Storage platform

	PowerMax	VMAX All Flash
Adaptive Compression Engine (ACE)	Yes	Yes
Extended Data Compression (EDC)	Yes	No
Inline Deduplication	Yes	No
Data Reduction I/O Module	Yes	No
Compression I/O Module	No	Yes
DEFLATE Compression Algorithm	Yes	No
LZS Compression Algorithm	No	Yes
FBA Storage Resource Pool (SRP)	Yes	Yes
CKD Storage Resource Pool (SRP)	No	No

3 Adaptive Compression Engine

The Adaptive Compression Engine (ACE) is the combination of multiple core components that work together to achieve maximum system efficiency and deliver optimized performance. These core components are:

3.1 Hardware acceleration

Each system is equipped with data reduction hardware that handles the actual compressing and decompressing of data. For PowerMax systems where deduplication applies the data reduction hardware also generates a unique hash ID needed for the deduplication process. The arrays are configured with one module per director which equates to 2 for each engine. The use of the modules reduces data reduction processing overhead. As a secondary function, software compression is automatically applied in the event of a fault or failure with one or more of the data reduction modules.

3.2 Optimized data placement

To maximize data reduction efficiency the system needs to accommodate multiple sizes of compressed data. To support a variety of compression sizes, multiple compression pools are used to create an optimal backend. Optimized data placement is the function responsible for dynamically changing the compression pools as needed. This alters the backend by creating various compression pools that cater to the incoming data. The result is an evolving layout of compression pools that dynamically change to match the reducibility of data sent to the system.

Compression pools are identified by labels which represent the track size for the data devices within the pool. For example, the 128KB pool is made up of data devices where the tracks are all 128KB in size. The 8KB pool is made up of data devices where the tracks are all 8KB in size. In comparison the capacity of the data devices between the pools is the same, however, the 8KB pool has 16 times the number of tracks. Here is a complete list of possible compression pools; 8KB, 16KB, 24KB, 32KB, 40KB, 48KB, 56KB, 64KB, 72KB, 80KB, 88KB, 96KB, 104KB, 112KB, and 128KB. Due to the dynamic design each compression-enabled system may have a different combination of compression pools that reduced data populates.

3.3 Activity Based Compression

Activity Based Compression (ABC) aims to prevent constant compression and decompression of data that is accessed frequently. This function allows the busiest data to avoid being compressed. It differentiates busy data from less busy data and accounts for up to 20% of the allocations in the SRP. Allowing the busiest allocations to skip compression is a benefit to the system as well as to end users. This ensures optimal performance and reduced overhead that can result from constantly decompressing frequently accessed data. The mechanism used to determine the busiest data does not add additional load on the system. ABC leverages statistics collected from incoming I/O to the frontend devices to determine what data sets are the busiest and best candidates to skip compression. It allows the system to maintain balance the system resources providing an optimal environment for both data reduction savings and performance.

3.4 Fine grain data packing

The Adaptive Compression Engine uses data reduction hardware to process incoming data which is divided into four sections. Each section is compressed in parallel which maximizes the efficiency of the data reduction module. The sum of the four compressed sections is the final compressed size and determines where the data is to be stored. In PowerMax systems where deduplication applies a unique hash id is applied to the

compressed data set. This process includes pattern detection, a non-zero allocate function. Pattern detection prevents the allocation of any of the four sections that contain all zeros. This behavior results in an efficient data reduction process that has minimal cost to performance.

Another benefit of dividing the extents into four sections comes when there are partial read or write operations. In this case only the sections that contain the requested data are processed. This means each section can be handled independently.

The efficiency of data compression is measured in terms of the compression ratio. This is the ratio between the original size of the data and its size after being compressed. For example, a 128KB dataset is compressed to 64KB, resulting in a compression ratio of 2:1.

3.5 Extended Data Compression

PowerMax systems include an additional function that will compress already compressed data to gain further capacity savings. The goal of extended data compression (EDC) is to apply additional compression savings to already compressed data. This is accomplished by identifying data that has not been accessed for a set amount of time. The factors that make data a candidate for EDC are the following:

- The data belongs to a data reduction enabled storage group.
- The data has not been accessed for 30 days.
- The data is not already compressed by EDC.

Data that qualifies for EDC is compressed using the Def9_128_SW algorithm and moved to the appropriate compression pool. This is an automated background process within the system. Additional savings are included in the storage group level achieved compression ratio. EDC is only available with PowerMax storage arrays.

4 Deduplication

Deduplication (dedupe) is the process of reducing redundant copies of data that consume storage capacity. The redundant copies are replaced with pointers. The pointers provide the access for the subsequent requests of that shared data by multiple sources. In PowerMax systems, dedupe is accomplished through a series of functions and components including hardware acceleration, dedupe algorithm, hash table and dedupe management object (DMO).

4.1 Hardware acceleration

Dedupe is an inline process that uses the same data reduction hardware as compression. All data reduction enabled incoming data is passed through data reduction hardware. In a single pass, the data reduction hardware handles compression, pattern detection and generates a hash ID for deduplication. This produces compressed data with a unique hash ID. Leveraging data reduction hardware for this process allows system resources to be focused on host I/O and other system operations.

4.2 Deduplication algorithm

PowerMax systems use the SHA-256 hashing algorithm implemented in the data reduction hardware to find duplicate data. The data is then stored as a single instance for multiple sources to share. This provides enhanced data efficiency while maintaining a long history of data integrity.

The SHA-256 algorithm generates a 32-byte code for each 32KB block of data. Consider a system with 1 PB of written data with 5% updated per day. In 1 million years of operation, there is a 20% likelihood of a hash collision. As each 128KB track is handled as 4 blocks of 32KB there would need to be a hash collision on all four blocks in the same 128KB track to have an actual hash collision. The odds of having all 4 collide makes this only theoretical (less than a 1% chance in a trillion years of operation).

4.3 Hash table

During the data reduction process a hash ID is generated when data is passed through the data reduction hardware. The hash table stores the unique hash IDs that are used for comparison as part of the dedupe process. Hash IDs stored in the table are a unique representation of data in a dedupe relationship. Hash IDs generated by the data reduction hardware and the SHA-256 algorithm for new writes are compared against the IDs already populating the hash table. If a matching hash ID already exists in the hash table, then a dedupe relationship is generated for the newly written data. During the comparison if the hash ID does not exist the table is updated, and that hash ID is added.

4.4 Dedupe Management Object

The dedupe management object (DMO) is a 64-byte object within system memory. DMOs only exist when dedupe relationships exist. These objects store and manage the pointers between front-end devices and the deduplicated data that consumes backend capacity in the array.

4.5 Data Reduction I/O flow

All I/O is passed through cache and then processed by the system. This means data reduction actions are performed after the data is received by the system, but before it is placed on disk. Using an inline process requires additional checks within the I/O flow where data reduction applies. The system uses these checks to

determine whether incoming data needs to pass through the data reduction hardware or not. Incoming data for a storage group with data reduction enabled will follow the data reduction flow. However due to the activity-based compression (ABC) function, active data for a storage group with data reduction enabled will skip the data reduction flow for performance optimization. Data not compressed due to ABC may be compressed later and moved to a compression pool. Data for a storage group with data reduction disabled will ignore the data reduction flow and will be written to the system unreduced.

There are a few different I/O types to consider, Read, Write, Write-update.

- Read - A request to access data that is already populating the array.
- Write - Incoming I/O that will consume disk space.
- Write-update - Incoming I/O that can change data that is allocated to disk space on the array.

Figure 1 below describes the path the I/O will follow which is determined by characteristics of the dataset or the related storage group.

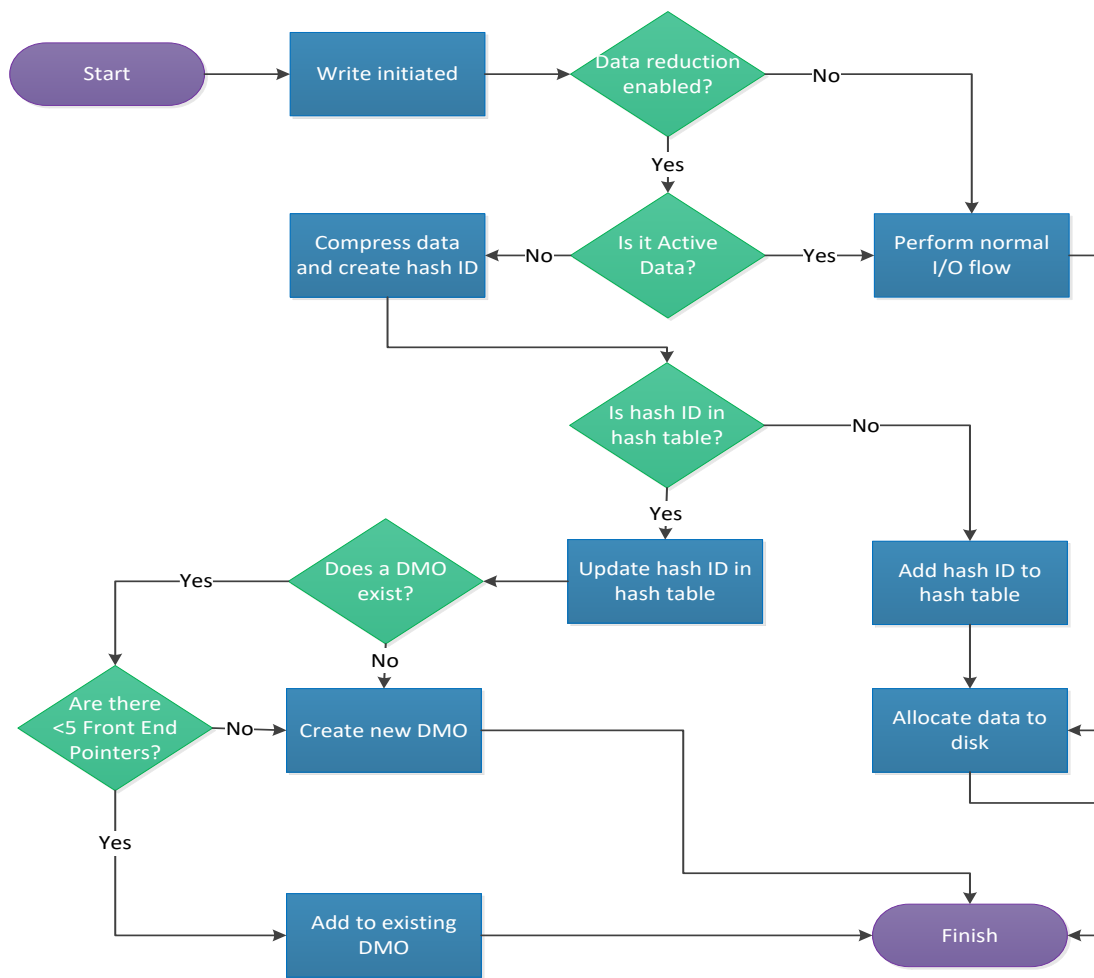


Figure 1 Data Reduction I/O flow for PowerMax enterprise storage systems

5 Capacity usage

Data Reduction is a feature intended to offer long-term space savings. Data Reduction uses machine learning resulting in efficient use of the available system resources. The use of statistics collected from the incoming data determines what is active and what is idle. Therefore, the activity-based compression function does not apply to net new writes. They are compressed and allocated to a compression pool. This also applies to dedupe as net new writes may not be consuming drive capacity yet. This would be the first entry of a hash ID into the hash table. Continued access to data generates statistics that are used to differentiate the activity level of data.

Capacity usage is represented in two ways; a physical capacity used percent and an effective capacity used percent. While both percentage values reflect how much host written data is consuming the system, the effected system resource is different. When the physical capacity usage percent is greater than effective used percent it indicates that there is potential for usable capacity to reach 100% full. This is also an indication the achieved data reduction ratio is lower than the systems data reduction reservation. When the effective used percent is greater than the physical used percent it is an indication there may be impact to the cache that supports the compression pools. The common variable related to either percent used is the current data reduction ratio.

For example, assume a data reduction reservation of 3:1. When the achieved data reduction ratio is less than 3:1 the physical used percent will be greater than the effective used percent. Likewise, when the data reduction ratio is greater than 3:1 the effective used percent will be greater than the physical used percent. Using the same example of 3:1 as the system reservation, applied to 100TB of usable disk capacity. This indicates the system will manage the resources to achieve 300TB of host data consuming 100TB of usable capacity. When the achieved ratio is less than 3:1 the system is less likely to accommodate 300TB of host data on 100TB of usable capacity. When the achieved ratio is greater than the system reservation the system will accommodate 300TB of host data on 100TB of disk capacity.

6 Data Reduction Settings

6.1 System and Storage Resource Pool

Data reduction is set at the system level per storage resource pool (SRP). The data reduction reservation set in the system is used to determine the potential savings from data reduction that can be supported by the available system resources. The data reduction reservation is used by the system to determine how much cache is needed to support the potential effective capacity. The cache used to support effective capacity is allocated as backend metadata in order to support the layout of the storage resource pool. This also determines how the capacity will be used to store data reduced by data reduction. As reduced data fills the compression pools, they expand automatically in order to accommodate more reduced data. Reaching the potential effective capacity relies on data being written to the system is reducible to a level equal to or greater than the data reduction reservation. For example; a system with 100TB of usable capacity where the reservation set is 3.5:1 has a potential effective capacity of 350TB. The written data needs to be reducible to 3.5:1 or better for the system to accommodate 350TB of host data and fit it into 100TB of usable capacity.

6.2 Storage Group

For application workloads to achieve capacity savings from data reduction it must be enabled at the storage group level. This is supported with Unisphere and Solutions Enabler. The feature is enabled by default when creating storage groups. There are two I/O flows for incoming data, data reduction enabled, where data is sent through the data reduction hardware and reduced, or data reduction disabled in which data bypasses the data reduction hardware and is written to disk unreduced. The storage group data reduction setting determines which I/O path the data will follow. In both cases, setting enabled or disabled is done using the data reduction option when provisioning storage. The option to enable or disable data reduction for individual storage groups can be changed at any time. However, changing the setting simply informs the system which I/O path the data will follow. Changing the setting does not immediately inflate already reduced data or attempt to reduce data already consuming capacity.

7 System Efficiency

Data reduction savings are presented as ratios and are available in both Unisphere for PowerMax and Solutions Enabler. The capacity report provides a single location to view system efficiency, capacity and system resource usage. The data is displayed in three sections, Array Usage, Efficiency and System Usage. There are two levels of detail available. The default view (see figure 2) offers a high-level view of efficiency in the form of ratios and capacity usage displayed as bar graphs. The detailed view expands the information provided under array usage revealing a more detail of the capacity usage (see figure 3). The detailed view also reveals system usage in the form of percentage used categorized as metadata.

As part of the Q1 2021 PowerMaxOS release there is a further breakdown under the efficiency section for the data reduction ratio. A flyover display reveals additional information on the data reduction ratio. The data presented relates specifically to data reduction enabled allocations. This is divided into two sections, unreduceable capacity and reduceable capacity. This information is available in both the default and detailed views of the capacity report.

- **Unreduceable Capacity:** Represents data reduction enabled allocations the system was not able to reduce. In most cases this is due to data already compressed and/or encrypted at the host or application level. In addition, there are data sets that simply do not reduce well, for example, image files (jpeg, pdf, etc.) or audio files.
- **Reduceable capacity:** This represents data reduction enabled allocations the system has determined are reduceable. The values are displayed as capacity representing reduceable data that will be reduced. Adding compression and dedupe savings with pattern detection savings may not equal the reduceable capacity value due to performance optimization (refer to section 3.3). The enabled for data reduction percentage will also be a factor when less than 100.

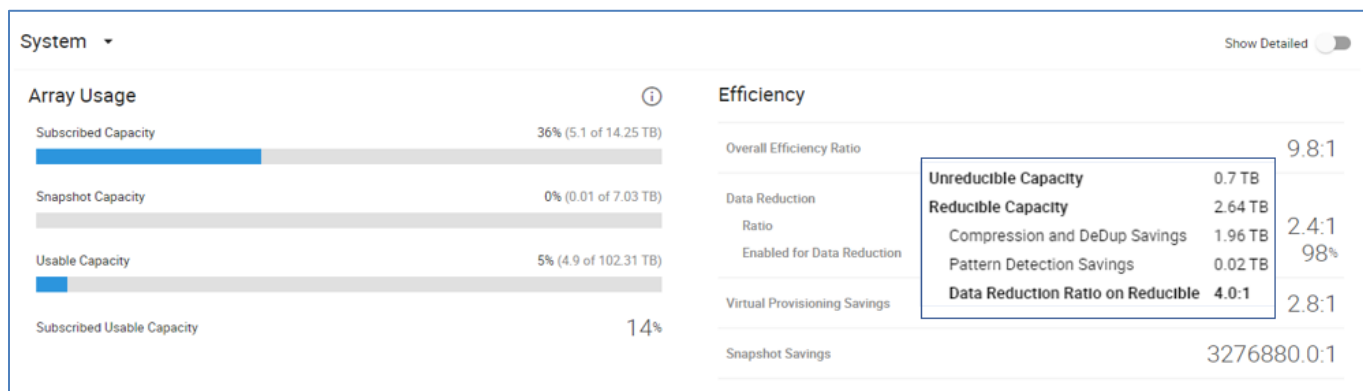


Figure 2 System efficiency report as seen in Unisphere for PowerMax (Default high-level view).

Calculating Efficiency Ratios: This data is revealed in the array usage section when switching the capacity report to the detailed view (see figure 3 below). It is needed to calculate the ratios displayed in the efficiency section. The data available from the detailed view can be used in the formulas below to calculate the ratios displayed in the efficiency section.

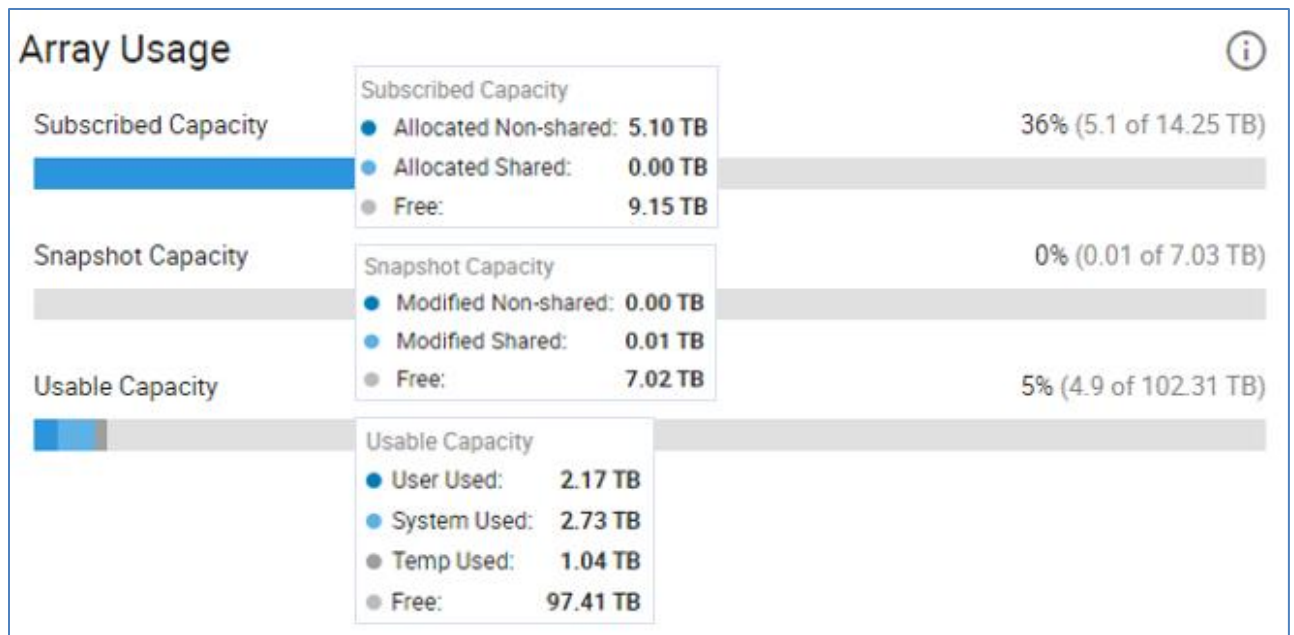


Figure 3 System efficiency report detailed view of Array Usage in Unisphere for PowerMax

- **Overall Efficiency Ratio:** The range of values that describe the capacity space savings that a user may experience regarding data reduction or other data services that offer capacity savings, such as Data Reduction, non-zero allocation, over provisioning and SnapVX.

$$\frac{\text{Subscribed Total} + \text{Snapshot total}}{\text{User Used}}$$

- **Data Reduction Ratio:** Savings that represents the combination of inline compression and inline deduplication presented as a ratio. To calculate the data reduction ratio the user needs to toggle to the detailed view shown in figure 3 above. Additional information needed is revealed by hovering over the Usage bars. When calculating the data reduction ratio using the values presented in the usage portion of the capacity report. The ratio may reflect a different value due to the performance optimization leaving compressible data uncompressed. The enabled percent being less than 100 may also be a factor.

$$\frac{\text{Subscribed Allocated Non Shared} + \text{Modified Non Shared}}{\text{User Used}}$$

- **Data Reduction Ratio on Reducible:** Represents the data reduction savings using only data reduction enabled allocations that have been reduced.

$$\frac{\text{ReducibleCapacity}}{\text{ReducibleCapacity} - (\text{CompressionAndDedupeSavings} + \text{PatternDetectionSavings})}$$

- **Enabled Percent:** The amount of subscribed host allocations that have Data Reduction enabled.
- **Virtual Provisioning Savings:** Savings achieved relative to provisioned capacity and total usable capacity displayed as a ratio. This may exceed the maximum usable capacity.

$$\frac{\text{Subscribed Total Capacity}}{\text{Allocated non – shared}}$$

- **Snapshot Savings:** A representation of savings resulting from the use of SnapVX to create local replication data.

$$\frac{\text{SnapShot Capacity Total}}{\text{Modified NonShared}}$$

7.1.1 System Resource Usage

In this section system resource usage refers to the two main components of the system relative to data reduction, capacity and cache.

Capacity is displayed as subscribed, and usable. Subscribed capacity represents the amount of capacity presented to hosts or applications. Subscribed capacity exceeding the total usable capacity is commonly referred to as subscription or virtual provisioning. This allows users to present more capacity to applications or hosts than the system can store. This is displayed in Unisphere as the Subscribed Usable Capacity Percent. Usable capacity is the amount of disk capacity available to store application data. The total amount of usable capacity is determined by the capacity of the physical disks that are configured in the system.

Like previous generations, PowerMax operates as a cache centric architecture. All data is passed through cache before being stored on disk. It is used to support multiple functions within the system not simply host I/O. Provisioning, local replication and data reduction also use cache. Cache is divided into two main sections, data cache and metadata cache. (see figure 4). Cache usage is displayed in the Unisphere capacity report under the System usage section, represented as metadata used (see figure 6).

- **Data cache:** Represents the amount of cache available for host IO, reads and writes from hosts or applications. The system configuration ensures there is always data cache available for host I/O.
- **Metadata cache:** This is comprised of three sections, Front-End, Replication and Back-End. Each section represents an amount of metadata cache it can consume.
 - **Front-End Metadata:** At initial install the percentage used will show zero as there is no subscribed capacity. There are two factors that will cause front-end metadata usage to increase, provisioning capacity to hosts or applications as well as host allocations. In PowerMax systems the increase is primarily due to host allocations.
 - **Replication Metadata:** At initial install the percentage used will show zero as there is no local replication activity. As local replication is used the percentage will increase up to 100 percent. When replication metadata has reached 100 percent, local replication has reached its limit. (for more information refer to [dell-emc-powermax-vmx-all-flash-timefinder-snapvx-local-replication.pdf](#))
 - **Back-End Metadata:** At initial install the percentage used represent the initial layout of compression pools. As the compression pools are expanded to support more effective capacity the usage can grow up to 100%. When Back-End Metadata shows 100% used it indicates the system has expanded usable capacity to the maximum effective capacity the system can support. This has no impact on Front end metadata growth or the ability to support Host I/O.

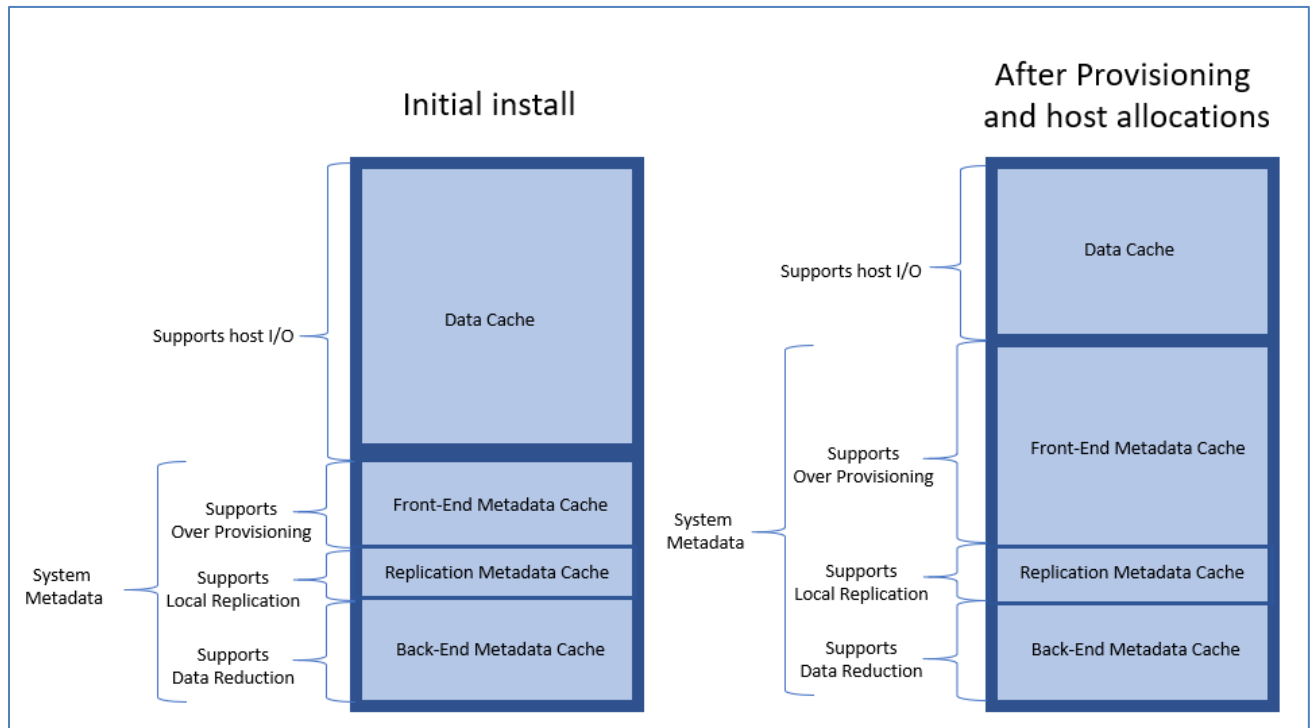


Figure 4 Cache breakdown

Array usage:

- **Subscribed Capacity:** The bar presented represents the total provisioned capacity. The dark shaded portion of the bar indicates host allocations of the presented capacity.
- **Snapshot Capacity:** The total represents the sum of all existing snapshots. The dark shaded portion of the bar represents the amount of existing snapshot capacity that has been modified. The modified capacity also represents additional usable capacity that is consumed by snapshot data.
- **Usable Capacity:** The total amount of usable disk space available. The dark shaded portion represents the amount of disk space that is consumed.
- **Subscribed Usable Capacity:** The percent displayed represents the amount of subscribed capacity in relation to the total amount of usable capacity.




Figure 5 Array usage from the Unisphere for PowerMax Capacity report

System Usage: The capacity report in Unisphere for PowerMax displays meta data usage in the form of percentage used. The values displayed represent the amount of metadata used for each function. These values are also available in Solution Enabler and REST API.

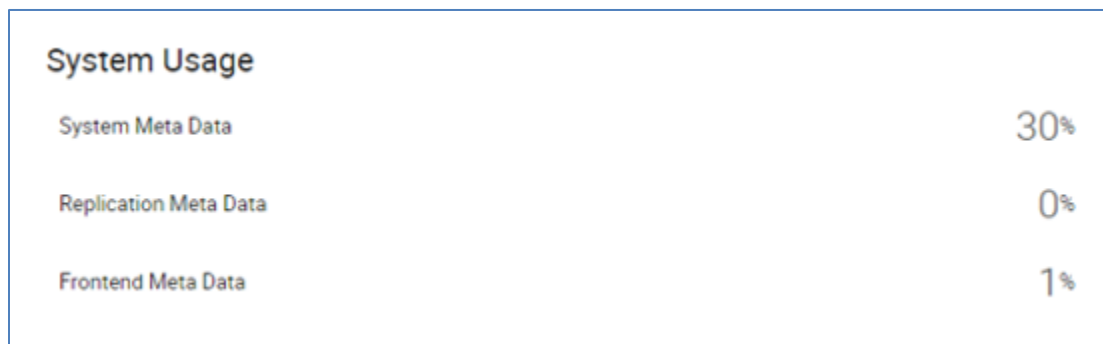


Figure 6 System Usage as seen in Unisphere for PowerMax

- System Metadata:** Represents the total metadata usage for the system. The amount of cache used by the system for all functions supported by metadata. The system used percentage represents the usage encompassing the total amount of metadata cache available.
- Replication Metadata:** A cache resource used in the form of metadata to support replication data pointers used with local replication. At initial install the percent used starts at zero as there is no local replication activity. The total amount of cache available for replication metadata is based on the configuration of the system and will not increase with the use of local replication.

When replication metadata has reached its maximum the use of local replication has reached its limit.

- Front-End Metadata:** A cache resource used in the form of metadata to support subscribed capacity and host allocations. As subscribed capacity is increased, the amount of front-end metadata increases. In VMAX All Flash systems provisioning will cause this to increase at the time of device creation. In PowerMax system the increase is primarily due to host allocations. In both cases the increase of Front-End metadata can consume data cache.

7.1.2 Storage group compression in Unisphere

Data reduction savings displayed at the Storage Group level represents only compression savings. This information can be viewed with the Storage Group list, the detailed view and the Storage Group Demand report. The ratio displayed shows the compression savings for data specific to the storage group being viewed. In addition to the compression ratio the amount of unreduceable data is shown. The amount of unreduceable data shown represents the amount of data the Storage group has allocated that the system has determined is not reducible. See examples below in figures 7, 8, 9 and 10.

Name	Compliance	SRP	Service Level	Capacity (GB)	Emulation	Compression ...	Unreduceable ...
WM_T50	✓	SRP_1	Diamond	800.00	FBA	3.4:1	41.90
WM_T66	✓	SRP_1	Diamond	4,000.02	FBA	2.7:1	0.00

Figure 7 Storage group list view in Unisphere for PowerMax.

Properties	
SRP	SRP_1
Compliance	✓
Service Level	Diamond
Workload Type	NONE
Tags	—
Emulation	FBA
Masking Views	0
Is Child	—
Child Storage Groups	0

Capacity	
Capacity (GB)	800
Volumes	1
Allocated Capacity	100% (800 of 800 GB)
VP Saved	0% (0.00 of 100%)
Unreduceable (GB)	41.9
Data Reduction	✓
Compression Ratio	3.4:1

Figure 8 Storage group details view in Unisphere for PowerMax.

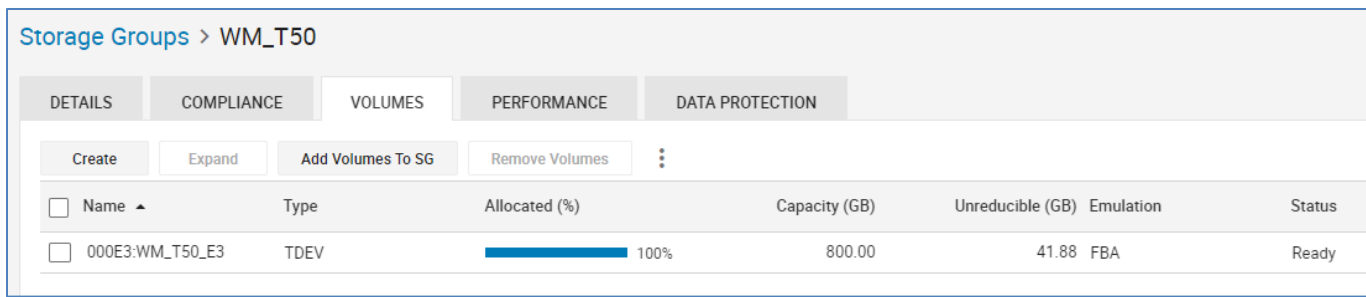


Figure 9 Storage group details view volume tab view in Unisphere for PowerMax.




Figure 10 Storage group demand report in Unisphere for PowerMax.

8 Supported data services

Data reduction is supported for FBA storage. Mixed FBA/CKD systems are supported however data reduction will only apply to the FBA storage resource pool/s. All other data services offered in both the PowerMax and VMAX All Flash systems are supported. This includes local replication (SnapVX), remote replication (SRDF), D@RE, and VMware® vSphere® Virtual Volumes™ (vVols).

8.1 Local replication (SnapVX)

Data reduction is supported with the use of local replication features; there are multiple variations and use cases for local replication. Below are the details regarding the different local replication sessions that can exist. For more detail regarding local replication and SnapVX see the TimeFinder and HYPERMAX OS Local Replication Technical Note available at DellEMC.com.

8.1.1 Nocopy sessions (SnapVX, VP Snap)

Uncompressed source data remains uncompressed when becoming snapshot data, and may be compressed later as it becomes less active. Activity to snapshot data through a linked target may prevent uncompressed data from being compressed. Compressed source data remains compressed when becoming snapshot data. Read activity to a snapshot through a linked target may cause the compressed data to be uncompressed.

The compression setting of a linked target only affects data written directly to the linked target and does not affect the snapshot data.

8.1.2 Copy sessions (SnapVX Full Copy Linked Targets, Clone, Mirror)

The compression settings for both the source and target are taken into account for copy sessions.

When compression is enabled on the source the data is decompressed before copying to the target. When compression is enabled on the target the data is compressed before being allocated to the target. Likewise, when compression is enabled on both the source and target the data is decompressed before the copy and then compressed to allocate for the target.

Copy times may vary due to decompression and compression of the data. It is not recommended to change the compression settings in between differential operations (that is, disabling compression before each differential operation and then again after the copy completes) as this causes data to go through needless compress/decompress cycles.

8.2 Remote replication (SRDF)

Compression for SRDF is already supported and known as SRDF compression. SRDF compression is a feature designed to reduce bandwidth consumption while sending data to and from systems connected using remote replication. SRDF compression and the Adaptive Compression Engine (ACE) both use the same compression module; however, they serve different purposes. Data that has been compressed using ACE is uncompressed before being sent across the SRDF link. If SRDF compression and inline compression apply, the data is uncompressed by the module and then compressed using the SRDF compression function and then sent to the remote site.

8.3 Data at Rest Encryption (D@RE)

D@RE provides hardware-based, on-array, back-end encryption, Data Reduction provides inline compression and Deduplication. Data is passed through the Data Reduction hardware before being sent through the encryption hardware. Therefore, data is compressed, deduped, or both before being encrypted by the D@RE process. On a D@RE enabled system data encrypted on disk has already been compressed, deduped, or both.

8.4 Virtual Volumes

Data reduction is supported for the allocation of data to vVols and follows the same I/O path as all other data. The IO path can be seen in Figure 71. Data Reduction as a feature is not included as a vVols resource to be configured at the host.

9 Conclusion

The use of physical storage capacity is a common concern of storage administrators across the storage industry. The constant and ever-growing amounts of data have created the need for more efficiency in the use of physical capacity. [Dell EMC PowerMax](#) and VMAX All Flash data storage systems take this to the next level. Combining inline compression with inline dedupe provides exceptional capacity savings with negligible cost to performance. This delivers on capacity savings, which leads to a smaller data center footprint and an overall reduction in TCO. In addition to the savings, using data reduction is as simple as a single click to enable or disable. The system handles all the work.

A Technical support and resources

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

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