

Dell EMC Unity: Best Practices Guide

Best practices for performance and availability

Abstract

This white paper provides recommended best practice guidelines for installing and configuring Dell EMC Unity systems for best performance.

March 2020

Revisions

Date	Description
May 2016	Initial release – Unity OE 4.0
December 2016	Updated for Unity OE 4.1
July 2017	Updated for Unity OE 4.2
March 2018	Updated for Unity OE 4.3
August 2018	Updated for Unity OE 4.4
January 2019	Updated for Unity OE 4.5
June 2019	Updated for Unity OE 5.0
August 2019	Minor corrections
February 2020	Branding update for VMware® vSphere® Virtual Volumes™ (vVols)
March 2020	Minor corrections

Acknowledgements

This paper was produced by the following:

Author: Stephen Wright

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 © 2016-2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [3/20/2020] [Best Practices] [H15093.8]

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	5
1 System configuration	6
1.1 Essential guidelines	6
1.1.1 Maximize Flash capacity	6
1.1.2 Distribute workloads across all available resources.....	6
1.1.3 Simplify the configuration	6
1.1.4 Design for resilience	6
1.1.5 Maintain latest released Dell EMC Unity operating environment version	7
1.2 SAS cabling and drive placement.....	7
1.3 Hot spares	8
1.4 Front-end connectivity	8
1.4.1 Fibre Channel	9
1.4.2 iSCSI.....	9
1.4.3 NAS	9
1.5 Hardware capability guidelines.....	10
1.5.1 CPU utilization	10
1.5.2 Connectivity options	11
1.5.3 Drives.....	13
1.6 Sizing Dell EMC Unity for performance requirements.....	14
1.7 Dell EMC UnityVSA	14
2 Storage configuration	15
2.1 General recommendations for storage pools	15
2.1.1 Storage pool capacity	15
2.1.2 RAID protection	15
2.2 Traditional pools	15
2.2.1 RAID protection	16
2.2.2 All-Flash pool.....	16
2.2.3 Hybrid pool.....	16
2.2.4 System drives	17
2.3 Dynamic pools	17
2.3.1 RAID protection	17

2.3.2	Spare capacity	18
2.4	Storage object types	18
3	Features	19
3.1	FAST VP	19
3.2	FAST CACHE	19
3.3	Data reduction	19
3.3.1	Advanced deduplication	20
3.4	Snapshots	20
3.5	Thin clones	21
3.6	Asynchronous replication	21
3.7	Synchronous replication / Metrosync.....	21
3.8	SAN Copy	22
3.9	NDMP	22
3.10	Data At Rest Encryption (D@RE).....	22
3.11	Host I/O limits	22
4	Application considerations	23
4.1	Host alignment for block LUNs	23
4.2	VMware ESXi.....	23
4.2.1	Multipathing	23
4.2.2	iSCSI.....	23
4.2.3	VMware NFS datastores	23
4.2.4	vVols	23
4.3	Transactional applications	23
4.4	Sequential applications.....	23
5	Conclusion	25
A	Technical support and resources	26
A.1	Related resources.....	26

Executive summary

The best practices guide delivers straightforward guidance to customers using Dell EMC Unity™ storage systems in a mixed business environment. The focus is on system performance and maximizing the ease of use of the automated storage features, giving guidance to achieve optimal performance, while avoiding non-optimal configurations.

These guidelines are intended to cover the majority of use cases. They are strongly recommended by Dell Technologies; however, they are not strictly required. In some cases it is completely acceptable not to adhere to them.

Some exception cases are addressed in this guide; however, less commonly encountered edge cases are not covered by general guidelines, and are instead addressed in use-case-specific white papers.

If you have questions about the applicability of these guidelines in your environment, contact your Dell Technologies™ representative to discuss the appropriateness of the recommendations.

This guide occasionally presents example configuration choices in order to demonstrate how a certain recommendation might be applied. Examples will be presented as follows:

EXAMPLE: Here is one possible way to apply the recommendation...

Examples presented in this way are intended to show one possible configuration option, and do not necessarily represent a preferred configuration.

1 System configuration

1.1 Essential guidelines

Dell EMC Unity is a virtually provisioned, Flash optimized storage system designed for ease of use. This guide introduces specific configuration recommendations that enable good performance from Dell EMC Unity. At the highest level, good performance design follows a few simple rules. The main principles of designing a Dell EMC Unity system for performance are:

- Maximize Flash capacity
- Distribute workloads across available storage resources
- Simplify the configuration
- Design for resilience
- Maintain latest released Dell EMC Unity Operating Environment version

1.1.1 Maximize Flash capacity

Dell EMC Unity utilizes the latest technology in multicore processors, capable of high levels of storage performance. In order to realize the performance potential, Dell Technologies recommends the use of Flash drives in all Dell EMC Unity systems. The best way to harness the power of Flash is to ensure as much dataset as possible resides on Flash drives. Therefore, Dell Technologies recommends provisioning as much Flash capacity as possible. Use the largest Flash drives appropriate for your solution.

1.1.2 Distribute workloads across all available resources

In order to maximize system performance, eliminate individual bottlenecks by utilizing all hardware resources. Spread Flash drives across back-end buses. Build storage pools with many drives. Balance LUN assignment and NAS Servers across both Storage Processors (SPs). Spread host connections across as many front-end ports as possible.

1.1.3 Simplify the configuration

Designing for simplicity also increases the system's flexibility, and leads to higher, more consistent performance. Using fewer storage pools means that more drives can be devoted to any single workload in the pool, possibly improving performance. Standardizing on a single capacity drive per tier, rather than provisioning a system with multiple drive sizes for different purposes, can reduce the number of required hot spares (see the Hot Spares section for more details), make it possible to build larger pools, and enable greater flexibility for future reconfiguration, if needed.

When in doubt about the impact of certain configuration options, selecting the recommended default is typically the best choice.

1.1.4 Design for resilience

Hardware can fail, and Dell EMC Unity is architected to continue providing storage services under such conditions. Understand the hardware limitations of different components of the Dell EMC Unity system in order to design a solution that continues to provide good performance under failure conditions. (Hardware capabilities of different components are provided later in this document.)

1.1.5 Maintain latest released Dell EMC Unity operating environment version

Dell Technologies regularly updates the Dell EMC Unity Operating Environment in order to improve performance, enhance functionality, and provide new features. Dell Technologies recommends running the latest version on your Dell EMC Unity system.

1.2 SAS cabling and drive placement

Dell EMC Unity systems have 2 on-board 12Gb SAS ports in each of the SPs in the Disk Processor Enclosure (DPE), for connecting to Disk Access Enclosures (DAEs). A 4-port 12Gb SAS I/O Module can also be provisioned in some Dell EMC Unity systems in order to provide additional back-end buses. In general, the on-board SAS ports provide sufficient IOPS and bandwidth capabilities to support most workloads. The SAS I/O Module should be considered in systems which will require extremely high bandwidth performance (greater than 10GB/s), or will attach more DAEs or drives that are supported with the on-board buses, or will attach an 80-drive dense enclosure with high-bandwidth option. The SAS I/O Module consumes an expansion slot that could otherwise be used for front-end connectivity.

When cabling DAEs to the DPE, balance them as evenly as possible across all available buses. When only the 2 SAS ports on each DPE are available, it is recommended to connect DAEs in the following order:

- DAE 1 connects to SAS Bus 1 (on-board port 1)
- DAE 2 connects to SAS Bus 0 (on-board port 0)
- DAE 3 connects to SAS Bus 1 (on-board port 1)

Connect additional DAEs in the same rotation.

When provisioning the 4-port SAS I/O Module, it is recommended to connect DAEs in the following order:

- DAE 1 connects to SAS Bus 1 (on-board port 1)
- DAE 2 connects to SAS Bus 2 (I/O Module port 0)
- DAE 3 connects to SAS Bus 4 (I/O Module port 2)
- DAE 4 connects to SAS Bus 3 (I/O Module port 1)
- DAE 5 connects to SAS Bus 5 (I/O Module port 3)
- DAE 6 connects to SAS Bus 0 (on-board port 0)

Once all buses have at least 1 connected DAE, connect additional DAEs in the same rotation.

80-drive dense enclosures can connect and intermix with other DAEs. When connecting multiple 80-drive enclosures to a system, place them on as many separate buses as possible, before daisy chaining them on a single bus.

80-drive dense enclosures can also be connected with a high bandwidth option that utilizes 2 SAS buses per enclosure, ganging them together for higher performance. This option is only available on Dell EMC Unity systems that support the SAS I/O Module. Dell Technologies recommends connecting these enclosures with this high bandwidth option, when possible.

Flash drives have the highest performance potential of any drive type. Therefore, Dell Technologies recommends spreading Flash drives across all available buses, if possible, to ensure the best IOPS and service times.

1.3 Hot spares

Note that dynamic pools do not require dedicated hot spares. See the Dynamic Pools section for additional considerations.

When using traditional storage pools, Dell EMC Unity automatically reserves 1 unbound drive out of every 31 drives of a given type for use as a hot spare. Note that hot spare drives have no special configuration, they simply remain unbound. Dell EMC Unity will not allow a traditional storage pool to be created or expanded unless the available number of unbound drives is sufficient to create the requested pool and continue to satisfy the hot spare policy. Consider spare drive count requirements when designing traditional storage pool layouts.

Required hot spare count can be reduced by decreasing the number of individual drive types within the system. Table 1 illustrates this point—

Table 1 Hot spare reservation example

	Drives for traditional storage pools	Total requires spare drives
System 1 25TB usable	9x 600GB SAS 15K 9x 1200GB SAS 10K 9x 1800GB SAS 10K	3 total: 1x 600GB SAS 15K 1x 1200GB SAS 10K 1x 1800GB SAS 10K
System 2 25TB usable	27x 1200GB SAS 10K	1 total: 1x 1200GB SAS 10K

Note that Dell EMC Unity will throttle hot spare rebuild operations to reduce the impact to host I/O. Rebuilds will occur more quickly during periods of low system utilization.

1.4 Front-end connectivity

Dell EMC Unity provides multiple options for front-end connectivity, via on-board ports directly on the DPE or optional mezzanine card, and via optional I/O Modules. This section discusses recommendations for the different types of connectivity.

In general front-end ports need to be connected and configured symmetrically across the 2 storage processors (SPs), in order to facilitate high availability and continued connectivity in case of SP failure.

EXAMPLE: A NAS Server is configured so that NAS clients connect via port 0 of the first I/O Module on SPA; therefore port 0 of the first I/O Module on SPB must be cabled so that it is accessible to the same networks.

For best performance, it is recommended to use all front-end ports that are installed in the system, so that workload is spread across as many resources as possible.

EXAMPLE: If configuring the 4-port Fibre Channel I/O Module, zone different hosts to different ports so that all 8 ports across the 2 SPs are utilized; don't simply zone all hosts to the first port of each I/O Module.

1.4.1 Fibre Channel

When configured for Fibre Channel, Dell EMC Unity CNA ports and I/O Module ports can be configured with 8Gb or 16Gb SFPs. All Fibre Channel ports can negotiate to lower speeds. 16Gb FC is recommended for the best performance.

Dell Technologies recommends single-initiator zoning when creating zone sets. For high availability purposes, a single host initiator should be zoned to at least 1 port from SPA and 1 port from SPB. For load balancing on a single SP, the host initiator can be zoned to 2 ports from SPA and 2 ports from SPB. When zoning additional host initiators, zone them to different SP ports when possible, to spread the load across all available SP ports.

Utilize multi-pathing software on hosts connected via Fibre Channel, such as Dell EMC PowerPath, which coordinates with the Dell EMC Unity system to provide path redundancy and load balancing.

1.4.2 iSCSI

Dell EMC Unity supports iSCSI connections on multiple 1Gb/s, 10Gb/s, and 25Gb/s (on Dell EMC Unity XT) port options. 10GBase-T ports can auto-negotiate to 1Gb/s speeds. Use the highest speed possible for the best performance. If possible, configure Jumbo frames (MTU 9000) on all ports in the end-to-end network, in order to provide the best performance.

To achieve optimal iSCSI performance, use separate networks and VLANs to segregate iSCSI traffic from normal network traffic. Configure standard 802.3x Flow Control (Pause or Link Pause) on all iSCSI Initiator and Target ports that are connected to the dedicated iSCSI VLAN.

Dell EMC Unity supports 10GbE and 1GBase-T ports that provide iSCSI offload. Specifically, the CNA ports (when configured as 10GbE or 1GBase-T) and the 2-port 10GbE I/O Module ports provide iSCSI offload. Using these modules with iSCSI can reduce the protocol load on SP CPUs by 10-20%, so that those cycles can be used for other services. Note that iSCSI offload is not available for Dell EMC Unity XT models.

Utilize multi-pathing software on hosts connected via iSCSI, such as Dell EMC PowerPath, which coordinates with the Dell EMC Unity system to provide path redundancy and load balancing.

1.4.3 NAS

Dell EMC Unity supports NAS (NFS, FTP, and/or SMB) connections on multiple 1Gb/s, 10Gb/s, and 25Gb/s (on Dell EMC Unity XT) port options. 10GBase-T ports can auto-negotiate to 1Gb/s speeds. Use the highest speed possible for the best performance. If possible, configure Jumbo frames (MTU 9000) on all ports in the end-to-end network, in order to provide the best performance.

Dell Technologies recommends configuring standard 802.3x Flow Control (Pause or Link Pause) on all storage ports, switch ports, and client ports that are used for NAS connectivity.

Dell EMC Unity provides network redundancy for NAS via Link Aggregation Control Protocol (LACP) and Fail Safe Networking (FSN). Combine FSN and LACP with redundant switches to provide the highest network availability. In addition to redundancy, LACP can also improve performance with multiple 1GBase-T connections, by aggregating bandwidth. LACP can be configured across any Ethernet ports that have the same speed, duplex, and MTU. (Note that LACP cannot be enabled on ports that are also used for iSCSI connections.)

While LACP creates a link aggregation with multiple active links, FSN provides redundancy by configuring a primary link and a standby link. The standby link is inactive unless the entire primary link fails. If FSN is configured with links of different performance capability (such as a link aggregation of 10Gb/s ports, and a stand-alone 1Gb/s port), it is recommended to configure the highest performing link as the primary.

NAS Servers are assigned to a single SP, and all file systems serviced by that NAS Server will have I/O processed by the SP on which the NAS Server is resident. For load-balancing, it is recommended to create at least 2 NAS Servers per Dell EMC Unity system, 1 on SPA, and 1 on SPB. Assign file systems to each NAS Server such that front-end workload is roughly the same for each SP.

1.5 Hardware capability guidelines

The following sections provide information on expected maximum performance from different hardware components of Dell EMC Unity. Dell Technologies does not recommend continuously operating a hardware resource near its maximum potential. When utilization is below the maximum level, average response times will be better, the system will be able to handle activity bursts without becoming overloaded, and performance may still be maintained during hardware failure scenarios.

For some of the following hardware sections, both IOPS and bandwidth (MB/s) capabilities are shown. A different type of workload is assumed when determining IOPS capabilities versus when determining bandwidth capabilities:

- IOPS workloads assume random small-block access. The following numbers are specifically modeled on a workload with a mix of I/O sizes less than 64KB, containing both read and write operations, and accessing randomly selected blocks within the datasets.
- Bandwidth workloads assume sequential large-block access. The following numbers are specifically modeled on a workload with a mix of I/O sizes greater than 64KB, containing both read and write operations, and accessing contiguous blocks of data with minimal seek.

This data is offered as a guideline for what most customers can expect for maximum performance capabilities from certain hardware. Note that individual components in your specific system may be able to provide higher or lower numbers, based on differences in workload characteristics and environment.

Most hardware components in the Dell EMC Unity system can be monitored in Unisphere, via the System Performance page.

1.5.1 CPU utilization

Dell EMC Unity is designed to achieve the maximum performance possible from today's multi-core processors. The Dell EMC Unity architecture results in non-linear CPU utilization increases as the workload scales.

EXAMPLE: A Dell EMC Unity system reporting 50% CPU might still be capable of providing more than 3 times the current workload.

It is recommended to balance workload across the 2 SPs, such that CPU utilization is roughly equivalent on each SP.

Table 2 provides recommended operating ranges, in terms of sustained CPU utilization. Brief spikes of high utilization are normal and expected on any system. The following values refer to average sustained levels over long periods of time. This table also provides the expected system behavior in the case that a single SP

would need to service the entire workload (such as coordinated reboots during upgrades, or if a single SP fails, etc.).

Table 2 Recommended sustained CPU utilization

Reported CPU utilization	Analysis	Expected single-SP behavior	Enable data reduction?	Enable snapshots and/or replication
Below 50% (Low utilization)	System is capable of accepting additional features and workloads.	A single SP should be able to service the entire workload while maintaining existing IOPS and response time.	Data Reduction can be enabled on the storage objects in this system.	Snapshots and Replication can be enabled on the storage objects in this system.
50%-70% (Normal utilization)	System is operating normally, and may be capable of accepting additional features and workloads.	A single SP should be able to service the entire workload while maintaining existing IOPS and response time.	Choose the best candidate storage objects, and enable Data Reduction on only a few at a time (see the Data Reduction section)	Enable Snapshots or Replication on only a few storage objects at a time.
70%-90% (High utilization)	System is nearing saturation; carefully consider whether additional features and workload should be added to this system.	A single SP should be able to maintain the existing IOPS load; however, increases in response time may be experienced.	Data Reduction should not be enabled on any additional storage objects in this system.	Choose the best candidate storage objects, and enable Snapshots or Replication on only a few at a time (see the Snapshots section).
Above 90% (Extremely high utilization)	System is saturated; additional workload should not be applied to the system; consider moving some work to other systems.	A single SP will likely <u>not</u> be able to maintain the existing IOPS load.	Do not enable Data Reduction on any additional storage objects in this system.	Do not enable Snapshots or Replication on any additional storage objects in this system.

1.5.2 Connectivity options

The following tables provide maximum expected IOPS and bandwidth from the different ports which are available in the Dell EMC Unity system. (Note that the capability of a port does not guarantee that the system can reach that level, nor does it guarantee that performance will scale with additional ports. System capabilities are highly dependent on other configuration parameters.)

SAS ports are used by the SPs to move data to and from the back-end drives; all other ports can be used to provide access to hosts.

Table 3 provides maximum expected IOPS and bandwidth from a 12Gb SAS port. The base Dell EMC Unity configuration contains 4 ports.

Table 3 Port capabilities for back-end SAS connectivity

Port	Maximum IOPS per port	Maximum MB/s per port
12Gb SAS	250,000	2,500
On-board or 4-port I/O Module		

Note that the DPE drives in most Unity systems are serviced by the pair of SAS ports that control SAS Bus 0. In the Dell EMC Unity XT 480/F, 680/F, and 880/F, the DPE drives are directly accessed by the underlying SAS controllers via internal connections. Therefore, DPE-only systems of these models can see higher bandwidth than other DPE-only systems.

Table 4 provides maximum expected IOPS and bandwidth from the front-end ports which provide Block protocols (FC and iSCSI).

Table 4 Port capabilities for front-end block connectivity

Port	Maximum IOPS per port	Maximum MB/s per port
16Gb FC	45,000	1,500
CNA or 4-port I/O Module		
8Gb FC	45,000	750
CNA or 4-port I/O Module		
25GbE iSCSI	30,000	1,700
4-port mezzanine or I/O Module		
10GbE iSCSI	30,000	1,100
4-port I/O Module		
10GbE iSCSI	25,000	900
CNA or 2-port I/O Module		
10GBase-T iSCSI	30,000	1,100
On-board, mezzanine, or 4-port I/O Module		
1GBase-T iSCSI	3,000	110
On-board, CNA, or 4-port I/O Module		

Table 5 provides maximum expected IOPS and bandwidth from the front-end ports which provide File protocols (NFS and SMB).

Table 5 Port capabilities for front-end file connectivity

Port	Maximum IOPS per port	Maximum MB/s per port
25GbE NAS 4-port mezzanine of I/O Module	60,000	2,900
10GbE NAS 4-port I/O Module	60,000	1,100
10GbE NAS CNA or 2-port I/O Module	60,000	1,100
10GBase-T NAS On-board, mezzanine, or 4-port I/O Module	60,000	1,100
1GBase-T NAS On-board, CNA, or 4-port I/O Module	6,000	110

1.5.3 Drives

Individual drives provide varying levels of performance based on the workload which is applied. Workload attributes such as I/O size, access patterns, queue depth, etc., can result in a drive providing more or less IOPS at acceptable response times.

Table 6 provides maximum recommended IOPS per drive for the drive types supported in Dell EMC Unity. In general, Dell Technologies recommends the following values for determining when a drive may be approaching its maximum performance, and therefore may become a bottleneck to system performance. If the drives in a storage pool are observed to sustain an IOPS load near these values for long periods of time, consider adding additional drives to the pool in order to spread the load across more resources.

Table 6 Recommended maximum sustained IOPS per drive

Drive type	Maximum recommended IOPS per drive
SAS Flash (all types)	20,000
SAS 15K	350
SAS 10K	250
NL-SAS	150

1.6 Sizing Dell EMC Unity for performance requirements

Dell Technologies and its partners can design the appropriate Dell EMC Unity configuration to meet your applications' performance and capacity requirements. Contact your Dell Technologies representative for assistance with sizing an appropriate solution.

1.7 Dell EMC UnityVSA

Dell EMC UnityVSA™ is a software-defined storage solution which provides most of the same features as Dell EMC Unity, but can be deployed on a customer's VMware infrastructure. Because Dell EMC UnityVSA does not have dedicated hardware, and is installed on user-provided equipment, there are recommendations that are specific to this system.

It is recommended to use Dell EMC UnityVSA HA for better system availability. UnityVSA provides redundancy by clustering 2 virtual SPs.

When provisioning storage to Dell EMC UnityVSA, Dell Technologies recommends using "Thick Provision Eager Zeroed" to guarantee the capacity and provide the best performance.

When provisioning storage for Dell EMC UnityVSA, it is recommended to use storage for the Dell EMC UnityVSA Operating Environment that is physically separate from the storage that will be used for storage pools.

2 Storage configuration

2.1 General recommendations for storage pools

Dell EMC Unity supports 2 different types of storage pools, traditional pools and dynamic pools. The following recommendations are applicable to both types of pool.

In general, it is recommended to use fewer storage pools within Dell EMC Unity, as this reduces complexity and increases flexibility. However, it may be appropriate to configure multiple storage pools, in order to:

- Separate workloads with different I/O profiles
- Separate pools where FAST Cache is and is not active
- Dedicate resources to meet specific performance goals
- Separate resources for multi-tenancy
- Create smaller failure domains

2.1.1 Storage pool capacity

Storage pool capacity is used for multiple purposes:

- To store all data written into storage objects — LUNs, file systems, datastores, and VMware® vSphere® Virtual Volumes™ (vVols) — in that pool
- To store data which is needed for Snapshots of storage objects in that pool
- To track changes to replicated storage objects in that pool
- To perform efficient data relocations for FAST VP

Storage pools must maintain free capacity in order to operate properly. By default, Dell EMC Unity will raise an alert if a storage pool has less than 30% free capacity, and will begin to automatically invalidate Snapshots and Replication sessions if the storage pool has less than 5% free capacity. Dell Technologies recommends that a storage pool always have at least 10% free capacity.

2.1.2 RAID protection

Dell EMC Unity applies RAID protection to the storage pool in order to protect user data against drive failures. Choose the RAID type that best suits your needs for performance, protection, and cost:

- RAID-1/0 provides the highest level of performance from a given set of drive resources, with the lowest CPU requirements; however, only 50% of the total drive capacity is usable
- RAID-5 provides the best usable capacity from a set of drive resources, but at lower overall performance and availability than RAID-1/0
- RAID-6 provides better availability than RAID-5 and better usable capacity than RAID-1/0, but has the lowest performance potential of the three RAID types

2.2 Traditional pools

Traditional Storage Pools apply RAID protection to discreet groups of drives within the storage pool.

Traditional pools are the only type of pool available on Dell EMC Unity hybrid systems, and are also available on all-Flash systems.

2.2.1 RAID protection

For traditional pools, Dell Technologies generally recommends RAID-5 for drives in Extreme Performance and Performance tiers, and RAID-6 for drives in the Capacity tier.

Assuming that roughly the same number of drives will be configured in a traditional pool, Dell Technologies generally recommends smaller RAID widths as providing the best performance and availability, at the cost of slightly less usable capacity.

EXAMPLE: When configuring a traditional pool tier with RAID-6, use 4+2 or 6+2 as opposed to 10+2 or 14+2.

When choosing RAID-1/0, 1+1 can provide better performance with the same availability and usable capacity as larger RAID widths (assuming that the same total number of drives are used), and also provides more flexibility.

2.2.2 All-Flash pool

All-Flash pools provide the highest level of performance in Dell EMC Unity. Use an all-Flash pool when the application requires the highest storage performance at the lowest response time.

Snapshots and Replication operate most efficiently in all-Flash pools.

Data Reduction is only supported in an all-Flash pool.

FAST Cache and FAST VP are not applicable to all-Flash pools.

Dell Technologies recommends using only a single drive size and a single RAID width within an all-Flash pool.

EXAMPLE: For an all-Flash pool, use only 1.6TB SAS Flash 3 drives, and configure them all with RAID-5 8+1.

2.2.3 Hybrid pool

Hybrid pools can contain HDDs (SAS and NL-SAS drives) and Flash, and can contain more than one type of drive technology in different tiers. Hybrid pools typically provide greater capacity at a lower cost than all-Flash pools, but also typically have lower overall performance and higher response times. Use hybrid pools for applications that do not require consistently low response times, or that have large amounts of mostly inactive data.

Dell Technologies recommends provisioning a Flash tier in hybrid pools. The Flash tier helps enable pool performance efficiencies, and improves response times when using Snapshots and/or Replication. The minimum recommended Flash capacity is at least 5% of the pool capacity.

Performance of a hybrid pool can be improved by increasing the amount of capacity in the Flash tier, so that more of the active dataset resides on and is serviced by the Flash drives. See the FAST VP section.

Hybrid pools can have up to 3 tiers (Extreme Performance, Performance, and Capacity). Dell Technologies recommends using only a single drive speed, size, and RAID width within each tier of a hybrid pool.

EXAMPLE:

- For the Extreme Performance tier, use only 800GB SAS Flash 2 drives, and configure them all with RAID-5 8+1.

- For the Performance tier, use only 1.2TB SAS 10K RPM drives, and configure them with RAID-5 4+1.
- For the Capacity tier, use only 6TB NL-SAS drives, and configure them all with RAID-6 6+2.

2.2.4 System drives

The first 4 drives of every Dell EMC Unity system are called the system drives (DPE Disk 0 through DPE Disk 3). Dell EMC Unity uses capacity from these 4 drives to store copies of configuration information and other critical system data. Therefore, the available capacity from each of these drives is about 107GB less than from other drives. System drives can be added to storage pools like any other drive, but offer less usable capacity due to the system partitions. To reduce the capacity difference when adding the system drives to a pool, Dell Technologies recommends using a smaller RAID width for the tier which will contain the system drives.

EXAMPLE: Choose RAID-5 4+1 for a tier containing the system drives, instead of 12+1.

Due to the capacity difference of the system drives, Dell EMC Unity will not add them to storage pools until all other non-spare drives of the same type have been used.

When used in traditional storage pools, the system drives are counted by Dell EMC Unity in determining how many hot spare drives need to be reserved, but note that unbound system drives cannot be used as hot spare drives.

The system drives also service a modest IOPS workload related to system management functions. For systems with Flash system drives, this workload is negligible. Dell EMC Unity Hybrid Flash arrays can use HDD (SAS or NL-SAS) as system drives. When used in a storage pool, these HDD system drives may have a lower performance capability than other HDD drives of the same type.

Consider not using the system drives in storage pools for large configurations with high drive counts and many storage objects, if the system drives are HDD.

2.3 Dynamic pools

Dynamic Storage Pools apply RAID protection to groups of drive extents from drives within the pool, and allow for greater flexibility in managing and expanding the pool. Dynamic pools are only available on Dell EMC Unity all-Flash systems, and therefore must be all-Flash pools; dynamic pools cannot be built with HDDs.

2.3.1 RAID protection

At the time of creation, dynamic pools will use the largest RAID width possible with the number of drives specified, up to the following maximum widths:

- RAID-1/0: 4+4
- RAID-5: 12+1
- RAID-6: 14+2

With dynamic pools, there is no performance or availability advantage to smaller RAID widths. Therefore, to maximize usable capacity with parity RAID, it is recommended to initially create the pool with enough drives to guarantee the largest possible RAID width.

- For RAID-5, initially create the pool with at least 14 drives
- For RAID-6, initially create the pool with at least 17 drives.

2.3.2 Spare capacity

Hot spares are not needed with dynamic pools. A dynamic pool will automatically reserve spare space in the pool at a rate of 1 drive's worth of capacity per every 32 drives. In the event of a drive failure, the data that was on the failed drive is rebuilt into the spare capacity on the other drives in the pool. Additionally, un-bound drives of the appropriate type can be used to replenish a pool's spare capacity, after the pool rebuild has occurred.

2.4 Storage object types

Dell EMC Unity supports two types of storage objects, thin or thick.

By default, Dell EMC Unity creates thin storage objects. Thin storage objects are virtually provisioned and space efficient. In general, Dell Technologies recommends using thin storage objects, as they provide the best capacity utilization, and are required for most features. Thin storage objects are recommended when any of the following features will be used:

- Data Reduction
- Snapshots
- Thin Clones
- Asynchronous Replication

Thick storage objects will reserve capacity from the storage pool, and dedicate it to that particular storage object. Thick storage objects guarantee that all advertised capacity is available for that object. Thick storage objects are not space efficient, and therefore do not support the use of space-efficient features. If it is required to enable a space-efficient feature on a thick storage object, it is recommended to first migrate the thick storage object to a thin storage object, and enable the feature during the migration (for Data Reduction) or after migration has completed (for Snapshots, Thin Clones, and Asynchronous Replication).

In addition to capacity for storing data, storage objects also require pool capacity for metadata overhead. The overhead percentage is greater on smaller storage objects. For better capacity utilization, Dell Technologies recommends configuring storage objects that are at least 100GB in size, and preferably at least 1TB in size.

3 Features

3.1 FAST VP

FAST VP accelerates performance of a specific storage pool by automatically moving data within that pool to the appropriate drive technology, based on data access patterns. FAST VP is only applicable to hybrid pools within a Dell EMC Unity hybrid system.

The default and recommended FAST VP policy for all storage objects is “Start High then Auto-tier”. This policy places initial allocations for the storage object in the highest tier available, and monitors activity to this storage object to determine the correct placement of data as it ages.

FAST VP is most effective if data relocations occur during or immediately after normal daily processing. Dell Technologies recommends scheduling FAST VP relocations to occur before backups or nightly batch processing. For applications which are continuously active, consider configuring FAST VP relocations to run constantly.

Dell Technologies recommends maintaining at least 10% free capacity in storage pools, so that FAST VP relocations can occur efficiently. FAST VP relocations cannot occur if the storage pool has no free space.

3.2 FAST CACHE

FAST Cache is a single global resource that can improve performance of one or more hybrid pools within a Dell EMC Unity hybrid system. FAST Cache can only be created with SAS Flash 2 drives, and is only applicable to hybrid pools. It is recommended to place a Flash tier in the hybrid pool before configuring FAST Cache on the pool. FAST Cache can improve access to data that is resident in the HDD tiers of the pool.

Enable FAST Cache on the hybrid pool if the workload in that pool is highly transactional, and has a high degree of locality that changes rapidly.

For applications that use larger I/O sizes, have lower skew, or do not change locality as quickly, it may be more beneficial to increase the size of the Flash tier rather than enable FAST Cache.

FAST Cache can increase the IOPS achievable from the Dell EMC Unity system, and this will most likely result in higher CPU utilization (in order to service the additional I/O). Before enabling FAST Cache on additional pools or expanding the size of an existing FAST Cache, monitor the average system CPU utilization to determine if the system can accommodate the additional load. Refer to Table 2 for recommendations.

3.3 Data reduction

Dell EMC Unity Data Reduction through compression is available for Block LUNs and VMFS datastores in an all-Flash pool starting with Dell EMC Unity OE 4.1. Data reduction via compression is available for file systems and NFS datastores in an all-Flash pool starting with Dell EMC Unity OE 4.2. Beginning with Dell EMC Unity OE 4.3, data reduction includes both compression and deduplication.

Be aware that data reduction increases the overall CPU load on the system when storage objects service reads or writes of reduceable data, and may increase latency. Before enabling data reduction on a storage object, it is recommended to monitor the system and ensure that the system has available resources to support data reduction (refer to the Hardware Capability Guidelines section, Table 2). Enable data reduction

on a few storage objects at a time, and then monitor the system to be sure it is still within recommended operating ranges, before enabling data reduction on more storage objects.

For new storage objects, or storage objects that will be populated by migrating data from another source, it is recommended to create the storage object with data reduction enabled, before writing any data. This provides maximum space savings with minimal system impact.

3.3.1 Advanced deduplication

Dell EMC Unity Advanced Deduplication is an optional extension to Data Reduction, that can be enabled to increase the capacity efficiency of data reduction enabled storage objects. Advanced deduplication is not available on all models. For information on which systems can enable advanced deduplication, please refer to the Dell EMC Unity: Data Reduction white paper available on Dell Support.

As with data reduction, advanced deduplication is only applied to data when it is written to the storage object. LUN Move can be utilized to deduplicate existing data on Block storage objects.

For new storage objects, or storage objects that will be populated by migrating data from another source, it is recommended to create the storage object with advanced deduplication enabled, before writing any data. This provides maximum space savings with minimal system impact.

3.4 Snapshots

Dell Technologies recommends including a Flash tier in a hybrid pool where snapshots will be active.

Be aware that snapshots increase the overall CPU load on the system, and increase the overall drive IOPS in the storage pool. Snapshots also use pool capacity to store the older data being tracked by the snapshot, which increases the amount of capacity used in the pool, until the snapshot is deleted. Consider the overhead of snapshots when planning both performance and capacity requirements for the storage pool.

Before enabling snapshots on a storage object, it is recommended to monitor the system and ensure that existing resources can meet the additional workload requirements (refer to the Hardware Capability Guidelines section, Table 2). Enable snapshots on a few storage objects at a time, and then monitor the system to be sure it is still within recommended operating ranges, before enabling more snapshots.

It is recommended to stagger snapshot operations (creation, deletion, etc.). This can be accomplished by using different snapshot schedules for different sets of storage objects. It is also recommended to schedule snapshot operations after any FAST VP relocations have completed.

Snapshots are deleted by the system asynchronously; when a snapshot is in the process of being deleted, it will be marked "Destroying". If the system is accumulating "Destroying" snapshots over time, it may be an indication that existing snapshot schedules are too aggressive; taking snapshots less frequently may provide more predictable levels of performance. Note that Dell EMC Unity will throttle snapshot delete operations to reduce the impact to host I/O. Snapshot deletes will occur more quickly during periods of low system utilization.

3.5 Thin clones

Dell Technologies recommends including a Flash tier in a hybrid pool where Thin Clones will be active.

Thin Clones use snapshot technology to provide space-efficient clones of Block objects. Consider the overhead of snapshots when planning performance and capacity requirements for a storage pool which will have Thin Clones.

3.6 Asynchronous replication

Dell Technologies recommends including a Flash tier in a hybrid pool where asynchronous replication will be active. This is applicable to both the source and the destination pools.

Dell Technologies recommends configuring multiple replication interfaces per SP, and distributing replication sessions across them. LACP can also be used to aggregate bandwidth for a replication interface. Configure Jumbo frames (MTU 9000) when possible.

Asynchronous replication takes snapshots on the replicated storage objects in order to create the point-in-time copy, determine the changed data to transfer, and maintain consistency during the transfer. Consider the overhead of snapshots when planning performance and capacity requirements for a storage pool which will have replicated objects.

When possible, fill the source storage object with data before creating the replication session. The data will then be transmitted to the destination storage object during initial synchronization. This is typically the fastest way to populate the destination storage object with asynchronous replication.

Setting smaller RPO values on replication sessions will not make them transfer data more quickly; but smaller RPOs will result in more frequent snapshot operations. Choosing larger RPOs, or manually synchronizing during non-production hours, may provide more predictable levels of performance.

3.7 Synchronous replication / Metrosync

Dell Technologies recommends including a Flash tier in a hybrid pool where synchronous replication will be active. This is applicable to both the source and the destination pools.

Synchronous replication transfers data to the remote system over the first Fibre Channel port on each SP. When planning to use synchronous replication, it may be appropriate to reduce the number of host connections on this port. When the CNA ports are configured as FC, CNA port 4 is defined as the replication port. If the CNA ports are configured as Ethernet, then port 0 of the lowest numbered FC I/O Module is the replication port.

When possible, create the synchronous replication session before filling the source storage object with data, as this alleviates the need to perform initial synchronization of the replication session. This is typically the fastest way to populate the destination storage object with synchronous replication.

When sizing a disaster recovery solution using synchronous replication, consider provisioning a destination system that has similar performance capabilities as the source system. This can help maintain the same level of application performance after a failover event.

3.8 SAN Copy

SAN Copy provides one-time migration of Block resources from a 3rd-party array, using either iSCSI or FC connections. When using FC, note that SAN Copy must use different ports than the FC ports which are designated for Synchronous Replication. This is true even if Synchronous Replication is not actively being used.

To lessen the impact of SAN Copy migrations on other host activity, consider reducing the number of host connections on the FC ports used for SAN Copy.

3.9 NDMP

Dell EMC Unity supports 2-way NDMP for File data, which allows the system to send File data directly to a backup device using FC connections. Note that NDMP must use different ports than the FC ports which are designated for Synchronous Replication. This is true even if Synchronous Replication is not actively being used.

To lessen the impact of 2-way NDMP backups on other host activity, consider reducing the number of host connections on the FC ports used for NDMP.

3.10 Data At Rest Encryption (D@RE)

D@RE is Controller Based Encryption that does not impact performance; therefore Dell Technologies recommends ordering Dell EMC Unity systems as encryption-enabled, if appropriate for your environment. (Note that encryption can only be enabled at the time of system installation with the appropriate license.) If encryption is enabled, Dell Technologies recommends making external backups of the encryption keys after system installation, and immediately following any change in the system's drives (such as creating or expanding a storage pool, adding new drives, replacing a faulted drive, etc.).

3.11 Host I/O limits

Dell Technologies recommends setting Host I/O Limits on workloads which might monopolize pool resources and starve other applications of their required performance. Consider some of the following opportunities to utilize Host I/O Limits:

- Limit the bandwidth available to large-block applications, such as backup jobs, which may be increasing the latency on other small-block workloads.
- Limit the IOPS capability of Thin Clones which are used in Test/Dev environments, so that they do not impact the capability of the associated source objects.
- Limit the IOPS / bandwidth that is available to non-essential applications that are hosted on the same Dell EMC Unity system as your critical applications.

4 Application considerations

4.1 Host alignment for block LUNs

Alignment only needs to be done for host operating systems which still use a 63-block disk header. If alignment is required, perform the alignment using a host-based method, and align with a 1MB offset.

See the “Host Connectivity Guide” on <https://support.emc.com> to determine if alignment is required for your operating system, and how to perform the alignment.

4.2 VMware ESXi

4.2.1 Multipathing

Dell Technologies recommends configuring ESXi with Round Robin Path Selection Plug-in with IOPS limit of 1. See VMware Knowledge Base article 2069356 -- <https://kb.vmware.com/kb/2069356>

4.2.2 iSCSI

When configuring LUNs on ESXi that are accessed via iSCSI, disable “DelayedACK” on ESXi. See VMware Knowledge Base article 1002598 -- <http://kb.vmware.com/kb/1002598>

4.2.3 VMware NFS datastores

When creating NFS-based storage for ESXi, Dell Technologies recommends creating VMware NFS Datastores instead of general-purpose NFS file systems. VMware NFS Datastores are optimized to provide better performance with ESXi. Install the VAAI plug-in before creating NFS Datastores, to enable these optimizations. When creating VMware NFS Datastores, Dell Technologies recommends using the default 8K Host IO Size. Only choose a different Host IO Size if all applications that will be hosted in the NFS Datastore primarily use the selected I/O size.

4.2.4 vVols

When configuring vVol NFS Datastores, it is recommended to create at least 2 vVol-enabled NAS Servers, 1 on SPA and 1 on SPB.

4.3 Transactional applications

Dell EMC Unity systems require high concurrency to deliver the maximum performance (IOPS). This is naturally achieved when connecting many hosts with many LUNs. For systems that will be configured with only a few hosts and/or LUNs, host HBA settings may need to be adjusted to increase concurrency. Consult the documentation for your OS or HBA on how to adjust LUN queue depth settings.

4.4 Sequential applications

For workloads which require high bandwidth for sequential streaming data, it may be beneficial to use thick storage objects. Thick storage objects fully allocate the capacity before application use, and in a consistent manner, which can improve subsequent sequential access.

Note that thick storage objects are not compatible with most features (Data Reduction, Snapshots, Thin Clones, Asynchronous Replication), so only use thick storage objects if these features will not be utilized.

5 Conclusion

This best practices guide provides configuration and usage recommendations for Dell EMC Unity storage systems in general use cases.

For detailed discussion of the reasoning or methodology behind these recommendations, or for additional guidance around more specific use cases, see the documents in the References section, or contact your Dell Technologies representative.

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

A.1 Related resources

Refer to the following resources available on Dell Support:

- 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: Data Reduction
- 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: Replication Technologies
- Dell EMC Unity: Snapshots and Thin Clones
- Dell EMC Unity: Operating Environment (OE) Overview
- 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