

Titan T5000

VMware vSphere Best Practices



Dell EMC Unity: VMware vSphere Best Practices

All-flash arrays

Abstract

This document provides best practices for deploying VMware® vSphere® with Dell EMC™ Unity All-Flash arrays including settings and configuration recommendations for vSphere hosts to achieve an optimal combination of performance and resiliency.

August 2020

Revisions

Date	Description
July 2017	Initial release for Dell EMC Unity OE version 4.2
June 2019	Update for Dell EMC Unity XT release
November 2019	vVols branding update
April 2020	vSphere 7.0 updates
August 2020	Minor corrections

Acknowledgments

This paper was produced by the following members of the Dell EMC storage engineering team:

Authored by: David Glynn

Updated by: Darin Schmitz

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2017–2020 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. [8/24/2020] [Best Practices] [H16319.3]

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	4
1 Storage configuration	5
1.1 Storage pools.....	5
2 Dell EMC Unity features	7
2.1 FAST VP.....	7
2.2 FAST Cache	7
2.3 Compression.....	7
2.4 Data at Rest Encryption.....	8
2.5 Host I/O limits	8
3 VMware vSphere considerations.....	9
3.1 Fibre Channel switch zoning	9
3.2 Mapping storage to a vSphere host.....	10
3.3 Multipathing	10
3.4 iSCSI LUNs.....	10
3.5 LUN sizing and creation	11
3.6 Guest virtual SCSI adapter selection	11
3.7 Thin provisioning and virtual disks.....	12
3.8 VMware NFS datastores	13
3.9 Virtual Volumes.....	13
4 Data protection	14
4.1 AppSync	14
4.2 Snapshots.....	14
4.3 Asynchronous replication	14
4.4 Synchronous replication	14
A Technical support and resources	16
A.1 Related resources	16

Executive summary

This guide provides best practices for customers using Dell EMC™ Unity All-Flash storage systems in a VMware® vSphere® environment. This paper focuses on optimizing system performance and maximizing the usability of Dell EMC Unity automated storage features when used with VMware vSphere. For general best practices on using Dell EMC Unity systems, see the [Dell EMC Unity: Best Practices Guide](#).

These guidelines are intended to cover the majority of use cases. They are strongly recommended by Dell EMC but are not strictly required.

This paper was developed using the Dell EMC Unity XT 480F All-Flash array, but is also applicable when using the 380F, 680F, or 880F Dell EMC Unity XT All-Flash arrays.

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

1 Storage configuration

Dell EMC™ Unity is a virtually provisioned, flash-optimized storage system designed for ease of use. This paper covers the all-flash array models. This section describes the foundational array technologies that support the application-specific sections that follow. Additional information for this section can be found in the *Dell EMC Unity Best Practices Guide*.

1.1 Storage pools

As of Dell EMC Unity OE version 4.2, Dell EMC Unity supports two types of storage pools on all-flash storage systems: traditional pools and dynamic pools. Traditional pools apply RAID protection to discrete groups of drives within the storage pool. Dynamic pools apply RAID to groups of drive extents from drives within the pool and allow for greater flexibility in managing and expanding the pool. Dynamic pools must be configured from all-flash drives; dynamic pools cannot be built with HDDs.

In general, it is recommended to use a small number of storage pools within Dell EMC Unity storage to reduce complexity and increase flexibility. However, it may be appropriate to configure additional storage pools in the following cases:

- Separate workloads with different I/O profiles
- Dedicate resources to meet specific performance goals
- Separate resources for multi-tenancy
- Create smaller failure domains

Refer to the [Dell EMC Unity: Dynamic Pools](#) white paper for additional information.

1.1.1 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 that is needed for snapshots of storage objects in the pool
- To track changes to replicated storage objects in that pool

Storage pools must maintain free capacity to operate properly. By default, Dell EMC Unity systems 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 EMC recommends that a storage pool always has at least 10% free capacity.

1.1.2 All-flash pool

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

- Dell EMC FAST™ Cache and FAST VP are not applicable to all-flash pools.
- Compression is only supported on an all-flash pool.
- Snapshots and replication operate most efficiently in all-flash pools.
- Dell EMC recommends using only a single drive size and a single RAID width within an all-flash pool.

For example, for an all-flash pool, use only 1.6 TB SAS flash 3 drives and configure them all with RAID 5 8+1.

1.1.3 Hybrid pool

Hybrid pools are not applicable when using Dell EMC Unity All-Flash arrays.

2 Dell EMC Unity features

This section describes some of the native features available on the Dell EMC Unity platform. Additional information on each of these features can be found in the *Dell EMC Unity: Best Practices Guide*. Features not applicable to the all-flash array models covered by this paper are noted.

2.1 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 flash system.

2.2 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 flash system. FAST Cache can only be created with SAS flash 2 drives, and is only applicable to hybrid pools.

2.3 Compression

Dell EMC Unity compression is available for block LUNs and VMFS datastores in an all-flash pool starting with Dell EMC Unity OE version 4.1. Compression is available for file systems and NFS datastores in an all-flash pool starting with Dell EMC Unity OE version 4.2.

Be aware that compression increases the overall CPU load on the system when storage objects service reads or writes of compressible data, and may increase latency when accessing the data. Before enabling compression on a storage object, it is recommended to monitor the system and ensure the system has available resources to support compression. (Refer to the “Hardware Capability Guidelines” section and Table 2 in the *Dell EMC Unity: Best Practices Guide*.) Enable compression 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 compression on more storage objects. Additional information regarding compression can be found in the *Dell EMC Unity: Compression* white paper.

Compression will only provide space savings if the data on the storage object is at least 25% compressible. Before enabling compression on a storage object, determine if it contains data that will compress; do not enable compression on a storage object if there will be no space savings. Contact your Dell EMC representative for tools that can analyze the data compressibility. Additional information regarding compression can be found in the *Dell EMC Unity: Compression* white paper.

2.4 Data at Rest Encryption

Data at Rest Encryption (D@RE) is controller-based encryption that does not impact system performance. Dell EMC recommends ordering Dell EMC Unity systems as encryption-enabled, if appropriate for your environment.

Note: Encryption can only be enabled at the time of system installation with the appropriate license.

If encryption is enabled, Dell EMC recommends making external backups of the encryption keys after system installation as well as immediately following any change in the system's drive configuration (such as creating or expanding a storage pool, adding new drives, or replacing a faulted drive). Additional information regarding Data at Rest Encryption can be found in the [Dell EMC Unity: Data at Rest Encryption](#) white paper.

2.5 Host I/O limits

Dell EMC recommends setting host I/O limits on workloads that might monopolize pool resources and starve other applications of their required performance.

For example: Limit the bandwidth available to large-block applications that may be increasing the latency on other small-block workloads.

Additional information can be found in the [Dell EMC Unity: Unisphere Overview](#) white paper.

3 VMware vSphere considerations

This section provides configuration recommendations and storage access guidelines for integrating VMware vSphere 6.x hosts with the Dell EMC Unity All-Flash array models.

3.1 Fibre Channel switch zoning

Zoning Fibre Channel switches to connect vSphere hosts to a Dell EMC Unity array is similar to most zoning processes. This section explains the fundamentals of recommended zoning.

3.1.1 Single initiator zoning

Each Fibre Channel zone created should contain a single initiator (host HBA port) and all applicable storage target ports. In other words, each HBA port on the host requires its own zone; containing itself and the applicable storage ports. For resiliency and availability, two fabrics should be used, each containing half of the ports on each host and half of the ports on each storage processor (SP) on the Dell EMC Unity array.

3.1.2 WWN zoning

World Wide Name (WWN) zoning provides the most flexible zoning configuration since it is based on the WWN of each port, regardless of the physical port on the switch where it is connected. For single initiator zoning with WWNs, the WWN of each host port would have its own zone, and would also contain the applicable storage port WWNs.

Table 1 Example of a single-host zones in each of the two Fibre Channel fabrics

Zone Name	WWN	Description
HOST1-HBA1 (fabric 1 zone)	20:01:00:0E:1E:C2:6A:94	Host1 HBA port 1
	50:06:01:60:47:E4:1E:49	Dell EMC Unity SPA FC Port 0
	50:06:01:61:47:E4:1E:49	Dell EMC Unity SPA FC Port 1
	50:06:01:68:47:E4:1E:49	Dell EMC Unity SPB FC Port 0
	50:06:01:69:47:E4:1E:49	Dell EMC Unity SPB FC Port 1
HOST1-HBA2 (fabric 2 zone)	20:01:00:0E:1E:C2:6A:95	Host1 HBA port 2
	50:06:01:62:47:E4:1E:49	Dell EMC Unity SPA FC Port 2
	50:06:01:63:47:E4:1E:49	Dell EMC Unity SPA FC Port 3
	50:06:01:6A:47:E4:1E:49	Dell EMC Unity SPB FC Port 2
	50:06:01:6B:47:E4:1E:49	Dell EMC Unity SPB FC Port 3

3.1.3 Port zoning

Port zoning refers to including physical ports on the Fibre Channel switch in each zone rather than the WWN as described above. Unless your environment requires the use of port zoning, WWN zoning is the recommended method.

3.2 Mapping storage to a vSphere host

Dell EMC Unisphere™, the graphical management interface for Dell EMC Unity systems, can be used to present storage to all host servers. For vSphere hosts, Unisphere utilizes integration points to reduce required steps and perform several host-side functions. To take advantage of the integration point, the vSphere host, or vCenter environment must be imported into the Unisphere interface (this requires credentials for the vSphere host or the vCenter server). Once imported, the presentation of VMware VMFS, NFS, or vVol datastore storage from Unisphere results in mounted datastore(s) on the hosts. Additional information and detail can be found in the *Dell EMC Unity: Unisphere Overview* and *Dell EMC Unity: Virtualization Integration* white papers.

3.3 Multipathing

Block storage (both iSCSI and FC) presented to vSphere hosts from Dell EMC Unity has the native Path Selection Policy (PSP) of round robin (RR) applied by default. While RR is the recommended PSP to apply to Dell EMC Unity block storage, the default number of I/Os between switching paths is 1000. By reducing this value, all paths are more efficiently utilized. Additional information on this topic can be found in the [Adjusting Round Robin IOPS limit from default 1000 to 1](#) article in the VMware Knowledge Base.

Internal lab testing has indicated that adjusting the RR policy to `iops=1` has shown up to a 121% increase in IOPS, and up to a 32% reduction in latency for vVol and VMFS datastores. It is recommended to test this in your environment to determine if this change positively effects your workloads.

A claim rule can be created to automatically set this value on future LUNs mapped to the host. This must be performed on all hosts and requires a reboot to take effect.

```
esxcli storage nmp satp rule add --satp "VMW_SATP_ALUA_CX" --vendor "DGC" --psp "VMW_PSP_RR" --psp-option "iops=1" --claim-option="tpgs_on"
```

3.4 iSCSI LUNs

Dell EMC Unity LUNs presented to vSphere hosts through iSCSI should utilize the following best practices.

3.4.1 Enabling jumbo frames

On the Dell EMC Unity array (in Unisphere click **Settings > Access > Ethernet**), as well as on all vSphere hosts accessing its iSCSI storage, the MTU size should be set to 9000 to enable the transmission of large blocks of data, which is more efficient than the default MTU size with block-based storage. Additionally, all Ethernet switches in the iSCSI data path must support jumbo frames. Refer to the specific switch configuration guide for instructions on enabling jumbo frames.

3.4.2 Disable delayed ACK

In some environments, periods of high network congestion can cause iSCSI transfer latency to exceed acceptable levels. To avoid this, VMware recommends disabling delayed ACK following the steps provided in the VMware Knowledge Base article [ESX/ESXi hosts might experience read or write performance issues with certain storage arrays](#).

3.5 LUN sizing and creation

Determining the size and number of LUNs to create and present to your vSphere environment is a complex task. This section does not provide specific recommendations for the size and number of LUNs because every environment is different, but it provides guidance for identifying the most effective configuration for your environment.

3.5.1 LUN size

VMware currently supports a maximum datastore size of 64 TB. However, in most circumstances, a much smaller, more manageable size is recommended to accommodate a reasonable number of virtual machines (VMs) per datastore (learn more in section 3.5.2). Since LUNs and vSphere datastores can be expanded to address future growth, the recommendation is to create LUNs and datastores with sizes in the range of 500–750 GB for most environments. This size of datastore accommodates 10–15 virtual machines with a 40 GB virtual disk and provides the additional capacity required for the various overhead files for each virtual machine.

Note: This sizing recommendation supports limitations on the number of virtual machines on each datastore to keep performance optimal. This recommendation does not take into consideration high-capacity virtual machines. Virtual machines requiring a large virtual disk would require a larger LUN/datastore size, and would not be exempt from these best practices.

3.5.2 Virtual machines per datastore

VMware currently supports a maximum of 2,048 powered-on virtual machines per VMFS datastore. However, in most circumstances and environments, a target of 15–25 virtual machines per datastore is the conservative recommendation.

By maintaining a smaller number of virtual machines per datastore, potential for I/O contention is greatly reduced, resulting in more consistent performance across the environment. For this reason, the recommended LUN/datastore size is 500–750 GB (as mentioned in section 3.5.1).

Note: The virtual-machines-per-datastore recommendation should be adjusted to support your environment. The appropriate number should be based on the I/O and capacity requirements for your specific environment.

3.5.3 VMFS partition alignment

Dell EMC Unity block LUNs where VMFS datastores are created do not require alignment because VMFS automatically aligns the partition at creation time.

3.6 Guest virtual SCSI adapter selection

When creating a new virtual machine, vSphere automatically selects the recommended virtual SCSI controller based on the operating system being installed to the VM. The best practice is to maintain the automatically selected virtual SCSI controller under most circumstances. In some cases, changing to the VMware Paravirtual SCSI controller may result in better I/O performance as well as reduced CPU utilization on the VM. Additional information and limitations can be found in the *vSphere Virtual Machine Administration Guide* at [VMware vSphere 6.x documentation](#) in the “About VMware Paravirtual SCSI Controllers” section.

3.7 Thin provisioning and virtual disks

Thin provisioning of storage, from both Dell EMC Unity storage and virtual disks created as part of a virtual machine, allows for increased space efficiency in the storage environment. Thinly provisioned storage can be configured with substantial sizes but will only occupy the storage capacity required to accommodate the actual storage needs. This ability reduces upfront storage costs in many cases, and allows for a more manageable and predictable storage growth over time.

3.7.1 Virtual disk formats

When creating a virtual machine, one or more virtual disks are created as part of the virtual hardware of that virtual machine. There are three virtual disk formats to choose from:

Thick provision lazy zeroed: A small amount of space is used for the virtual disk at the time of creation. New blocks of data are only allocated during write operations. However, before data is written to new blocks, vSphere will first zero the block to ensure write integrity. This process introduces additional I/O and latency when writes occur and could potentially affect latency-sensitive applications.

Thick provision eager zeroed: The space required for the virtual disk is fully allocated at the time of creation. All data blocks for the virtual disk are zeroed during creation. This format takes longer to prepare than other formats, but because all data blocks have been zeroed, there are no I/O penalties as found with other formats. However, there is no realized space efficiency with this format either, since all space has been consumed in the zeroing process.

Thin provisioned: This format does not allocate all the logical space for the virtual disk at creation. Rather, logical space is allocated on demand during the first write issued to the block. Like the thick provision lazy zeroed disk, blocks are zeroed prior to first write, introducing additional I/O and latency.

The default virtual disk format is thick provision lazy zeroed. The best practice is to use the default virtual disk format unless specific needs dictate the use of thick provisioned eager zeroed for performance or availability needs including Microsoft Cluster Service (MSCS) and VMware Fault Tolerance (FT).

With the release of vSphere 7.0, VMware has improved the first-write process when using thin provisioned VMDKs. While VMware still recommends using Eager Zeroed Thick (EZT) virtual disks for maximum performance, this new method should decrease much of the on-first-write performance overhead previously experienced with thin VMDKs. For an overview of the new feature, see the [VMware vSphere 7 Core Storage](#) whitepaper.

3.7.2 Dell EMC Unity thin provisioning or VMware thin provisioning

Dell EMC Unity storage offers thin provisioning as a recommended option when creating block storage and requires thin provisioning to enable compression. Utilizing thin provisioning within VMware on virtual disks does not initially result in additional space efficiency when thin provisioning is enabled on the array. However, the ability to reclaim space from within a compatible guest OS does require thin provisioning to be used on both the storage and the virtual disks. Additional requirements for space reclamation within the guest include the following:

Virtual machines on VMFS 6 datastores

- Reclaiming space on a guest operating system is only processed when the space to reclaim is equal to 1 MB or a multiple of 1 MB.

Virtual machines on VMFS 5 datastores

- Virtual machine hardware must be version 11 (ESXi® 6.0) or greater.
- The advanced setting EnableBlockDelete must be set to 1.
- The guest operating system must be able to identify the virtual disk as thin.

3.8 VMware NFS datastores

When creating NFS-based storage for vSphere in a Dell EMC Unity environment, Dell EMC recommends creating VMware NFS datastores using the VMware-specific tools within Unisphere, rather than using general-purpose NFS file systems. VMware NFS datastores are optimized to provide better performance with vSphere. Additionally, creating two NAS servers (one on each SP) on the Dell EMC Unity platform where NFS datastores are served is recommended for both availability and increased performance. Dell EMC also recommends using the default 8K Host I/O Size for VMware NFS datastores. Only choose a different host I/O size if all applications that will be hosted in the NFS datastore primarily use the selected I/O size. Additional NFS tuning information can be found in the VMware document [Best Practices for Running VMware vSphere on Network-Attached Storage \(NAS\)](#).

3.9 Virtual Volumes

When configuring block Virtual Volume (vVol) datastores over Fibre Channel or iSCSI, protocol endpoint devices are presented to the vSphere hosts rather than disk devices. These protocol endpoint devices are used for tuning as described in section 3.3. Best practices regarding block vVol datastores mirror those of the standard block storage described previously in this paper.

When configuring file-based (NFS) vVol datastores in a Dell EMC Unity environment, it is recommended to create at least two vVol-enabled NAS servers: one on SPA and the other on SPB. Additionally, committing two or more interfaces per SP can result in better performance, but this follows the same guidelines and limitations as NFS datastores described in *Best Practices for Running VMware vSphere on Network-Attached Storage (NAS)*.

4 Data protection

This section covers features or products used for data protection. It also covers how storage administrators or application administrators use or integrate application data protection features.

4.1 AppSync

Dell EMC AppSync™ is software that enables integrated Copy Data Management (iCDM) with Dell EMC primary storage systems, including Dell EMC Unity arrays.

AppSync simplifies and automates the process of creating and using snapshots of production data. By abstracting the underlying storage and replication technologies, and through application integration, AppSync empowers application owners to manage data copy needs themselves. The storage administrator, in turn, need only be concerned with initial setup and policy management, resulting in a more agile environment. Additional information on AppSync can be found in the [AppSync User and Administration Guide](#) and the [AppSync Performance and Scalability Guidelines](#).

4.2 Snapshots

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 and Table 2 in the *Dell EMC Unity Best Practices Guide*. Enable snapshots on a few storage objects at a time, and then monitor the system to be sure it is still within the recommended operating ranges before enabling more snapshots. Additional information can be found in the [Dell EMC Unity: Snapshots and Thin Clones](#) white paper.

4.3 Asynchronous replication

Dell EMC 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.

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

Setting smaller recovery point objective (RPO) values on replication sessions will not make them transfer data more quickly, but smaller RPOs will result in more snapshot operations. Choosing larger RPOs, or manually synchronizing during non-production hours, may provide more predictable levels of performance. Additional information can be found in the [Dell EMC Unity: Replication Technologies](#) white paper.

4.4 Synchronous replication

Dell EMC recommends including a flash tier in a hybrid pool where synchronous replication will be active.

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. Additional information can be found in the *Dell EMC Unity: Replication Technologies* white paper.

A Technical support and resources

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

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

[Dell EMC Unity Info Hub](#) is a source providing helpful links to document and tools to customers.

A.1 Related resources

Dell EMC publications:

- Dell EMC Unity: Best Practices Guide

Dell EMC titles on [Dell EMC Support](#):

- Dell EMC Unity XT: Introduction to the Platform
- Dell EMC Unity: Compression – Overview
- Dell EMC Unity: Data at Rest Encryption
- Dell EMC Unity: Dynamic Pools – Overview
- Dell EMC Unity: Performance Metrics – A Detailed Review
- Dell EMC Unity: Replication Technologies – A Detailed Review
- Dell EMC Unity: Snapshots and Thin Clones – A Detailed Review
- Dell EMC Unity: Unisphere Overview – Simplified Storage Management
- Dell EMC Unity: Virtualization Integration - Whitepaper

VMware publications:

- [Adjusting Round Robin IOPS limit from default 1000 to 1](#)
- [Best Practices for Running VMware vSphere on Network-Attached Storage \(NAS\)](#)
- [ESX/ESXi hosts might experience read or write performance issues with certain storage arrays](#)
- [VMware vSphere documentation](#)