

Dell EMC PowerMax and VMAX All Flash: eNAS Best Practices

Applied best practices for PowerMax and VMAX PowerMaxOS version 5978.668.668 and Embedded NAS version 8.1.15-23

January 2022

H15188.7

White Paper

Abstract

This white paper outlines best practices for planning, implementing, and configuring Embedded NAS (eNAS) in Dell EMC PowerMax and VMAX All Flash storage arrays to obtain maximum operating efficiencies from the eNAS component of the array.

Dell Technologies

Copyright

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 © 2016-2022 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. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA January 2022 H15188.7.

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

Contents

Executive summary4

eNAS architecture5

Model comparison.....7

Best practices.....8

Conclusion.....15

Appendix: Technical support and resources16

Executive summary

Overview

Organizations need IT infrastructure that provides instant access to the huge volumes of data associated with traditional transaction processing or data warehousing as well as a new generation of applications built around the world of social, mobile, and big data. Dell Technologies is redefining data center cloud platforms to build the bridge between these two worlds to form the next generation of hybrid cloud.

Dell EMC PowerMax and VMAX All Flash unified storage extends the value of VMAX to file storage, enabling customers to deploy one infrastructure to easily manage mission-critical block and file resources. PowerMax and VMAX All Flash unified storage enables customers to significantly increase data availability, simplify operations, and improve productivity.

Audience

This document is intended for anyone who needs to understand Embedded NAS (eNAS) and the components and technology in the following members of the PowerMax and VMAX families:

- Dell EMC PowerMax family (2000 and 8000)
- Dell EMC VMAX All Flash family (250F/FX, 450F/FX, 850F/FX, and 950F/FX)

Revisions

| Date | Description |
|----------------|--------------------------------|
| May 2016 | Initial release |
| September 2016 | Updated edition |
| May 2017 | Updated edition |
| March 2018 | Updated edition |
| July 2018 | Updated edition |
| September 2019 | Content and template update |
| September 2020 | Updates for PowerMaxOS Q3 2020 |
| January 2022 | Updated template |

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Author: Rob Mortell

Contributor: Ramrao Patil

Note: For links to other documentation for this topic, see the [PowerMax and VMAX Info Hub](#).

eNAS architecture

Overview

The diagram below shows the eNAS architecture.

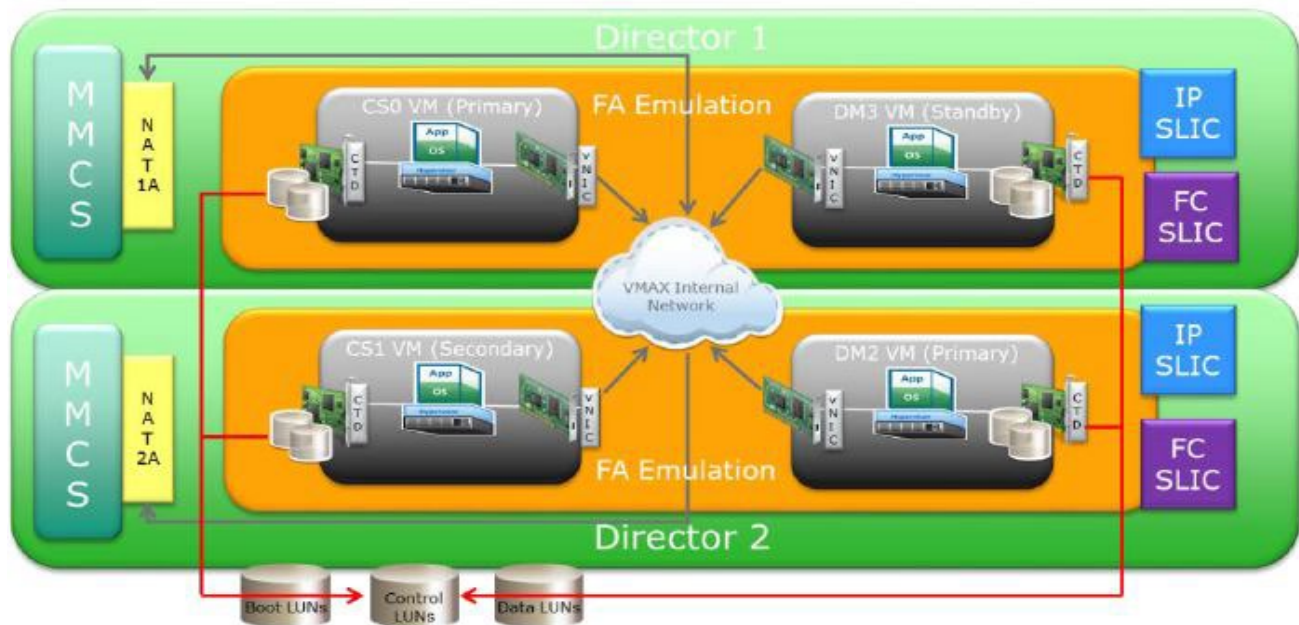


Figure 1. eNAS architecture

Management Module Control Station

The Management Module Control Station (MMCS) provides environmental monitoring capabilities for power, cooling, and connectivity. Each MMCS has two network connections that connect to the customer's network. One is to allow monitoring of the system, as well as remote connectivity for Dell Technologies Customer Support. The second is for use by the Network Address Translation (NAT) Gateway.

Network Address Translation gateway

The Network Address Translation (NAT) gateway provides translation services between external and internal IP addresses and utilizes a separate network connection on each of the two MMCS.

On PowerMax, there are two Network Address Translation (NAT_ IP addresses for eNAS, one per each NAT gateway. Both the IP addresses are mapped to different interfaces on the CS and source routed accordingly. So, for any incoming connections, the CS can be reached using either of the IP addresses unless there is a failure with any of the NAT gateways on the PowerMax.

There is also a default route (using NAT1 initially) that is used for any communication originating on the CS. In case of a failure of default gateway, the route is changed to use the other NAT gateway. There is a monitoring daemon that checks the health of the NAT gateways and updates the switch with the route as necessary.

The above works as is for IPv4. Hence, it is active/active for IPv4. For IPv6, there is no capability to use both the gateways at the same time. So, it is active/passive. For IPv6, CS can be reached only through one NAT IP address (NAT1 initially), and when there is a failure there, it can be reached using the other NAT.

Control station

The control station provides management functions to the file-side components referred to as Data Movers. The control station is responsible for Data Mover power cycle, as well as health and state monitoring, and failover. There is a primary control station (CS) which is configured with a matching secondary CS to ensure redundancy.

Data mover and I/O options

The data mover accesses data from the storage groups created on the PowerMax and provides host access using the I/O modules that support the NAS protocols, for example NFS and CIFS.

Data mover I/O options

Using PCI passthrough, the following I/O modules are available for the NAS services on the data mover:

- 4-port 1 Gb BaseT Ethernet copper module
- 2-port 10 Gb BaseT Ethernet copper module
- 2-port 10 Gb Ethernet optical module
- 4-port 8 Gb Fibre Channel module for NDMP back-up to tape use; there is a maximum of 1 per DM

Notes:

- The 4-port 8 Gb Fibre Channel module can be ordered separately and added at a later time.
- There must be at least one Ethernet I/O module for each Data Mover.
- At present, PowerMax and VMAX All Flash ships by default, with a single 2-port 10 Gb optical I/O module.
- To get others, an RPQ, or hardware upgrades or replacements need to be done.

Model comparison

The below table outlines the different components and capacity of the VMAX models offering eNAS.

Table 1. VMAX with eNAS model comparison

| VMAX Family | | VMAX3 | VMAX3 | VMAX All Flash | VMAX All Flash | VMAX All Flash | VMAX All Flash |
|-----------------------------|----------------------|--------------------|--|----------------|--------------------|--|--|
| Components | | VMAX 200K | VMAX 400K | VMAX 250F/FX | VMAX 450F/FX | VMAX 850F/FX | VMAX 950F/FX |
| Control Stations | Shared Logical Cores | 2 | 2 | 2 | 2 | 2 | 4 ^d |
| | Memory (GB) | 8 | 8 | 8 | 8 | 8 | 8 |
| Data Movers | Logical Cores | 10/20 | 16/32/48 ^a /64 ^a | 10/20 | 10/20 | 16/32/48 ^b /64 ^b | 16/32/48 ^b /64 ^b |
| | Memory (GB) | 48/96 | 48/96/144 ^a /192 ^a | 48/96 | 48/96 | 48/96/144 ^b /192 ^b | 48/96/144 ^b /192 ^b |
| | Max I/O Modules | 6/12 | 6/12/18 ^a /24 ^a | 6/12 | 6/12 | 6/12/18 ^b /24 ^b | 6/12/18 ^b /24 ^b |
| Max Data Movers | | 4 | 8 ^a | 4 | 4 | 8 ^b | 8 ^b |
| Max Capacity per Data Mover | | 512Tb ^c | 512Tb ^c | 512Tb | 512Tb ^c | 512Tb ^c | 512Tb ^c |
| Max eNAS capacity per VMAX | | 1.5Pb ^c | 3.5Pb ^c | 1.5Pb | 1.5Pb ^c | 3.5Pb ^c | 3.5Pb ^c |

- Support for more than four Software Data Movers in the VMAX 400K requires HYPERMAX OS 5977.813.785 or later.
- The 850F/FX and 950F/FX can be configured with a maximum of 4 Data Movers. However, by RPQ, that can be increased to either 6 or 8 Data Movers.
- Each VMAX 100K Software Data Mover can support up to 256 TB of usable capacity. Starting with HYPERMAX OS 5977.813.785 or later VMAX 200K, 400K, 450F, 850F, 950F each Software Data Mover can support up to 512 TB of usable capacity.

The below table outlines the different components and capacity of the PowerMax models offering eNAS.

Table 2. PowerMax with eNAS model comparison

| PowerMax Family | | | |
|--------------------------------|----------------------|---------------|--|
| Model | | PowerMax 2000 | PowerMax 8000 |
| Control Stations (2) | Shared Logical Cores | 2 | 4 ^b |
| | Memory (Gb) | 8 | 8 |
| Data Movers 2/4/6/8 | Logical Cores | 12/24 | 16/32/48 ^a /64 ^a |
| | Memory (Gb) | 48/96 | 48/96/144 ^a /192 ^a |
| | IO Modules (Max) | 6/12 | 6/12/18 ^a /24 ^a |
| Max Data Movers | | 4 | 8 ^a |
| Max Capacity per Data Mover | | 512Tb | 512Tb |
| Max eNAS Capacity per PowerMax | | 1.5Pb | 3.5Pb |

Best practices

Size for eNAS in advance

It is highly recommended to consider the potential to use eNAS when sizing and include this in upfront configuration.

Network connectivity

Port I/O module connectivity


eNAS I/O Modules are dedicated to eNAS and cannot be shared for other connectivity. When planning host to eNAS connectivity for performance and availability, connect at least two physical ports from each Data Mover to the network. Similarly, connect at least two ports from each host to the network. In this way, eNAS can continue to service the host I/O, in case of a component failure.

For best performance and availability, use multiple file systems and spread them across all the Data Movers serviced by different Ethernet interfaces. Each NAS share created on the Data Mover is accessible from all ports on the Data Mover, so it is essential that the host has connectivity to all the ports of the Data Mover. With SMB 3.0, the host can take advantage of load balancing and fault tolerance if multiple Ethernet ports are available on the host. For non-SMB 3.0 environments, create virtual network devices that are available to selected Data Movers, selecting a type of Ethernet channel, link aggregation, or Fail Safe Network (FSN). Do this before you create an interface.

[VNX documentation](#) provides information about configuring virtual Ethernet devices.

The cables that need to be connected are:

- Two network cables into the black customer connectors at the rear of the box for that engine (1 per director)



Use 10 Gb I/O module

- The NAS data is presented to clients by the Data Movers.
- For high-performance workloads, it is recommended to use 10 Gb I/O modules and jumbo frames. (MTU of 9000 bytes)

Note: The entire network infrastructure must also support jumbo frames.

- If link aggregation of Data Mover ports is required for bandwidth aggregation or load balancing, use the LACP protocol.

Storage group considerations

Create multiple storage groups/nas_pools in order to:

- Separate workloads
- Dedicate resources when you have specific performance goals

Service levels and eNAS

Service levels for PowerMaxOS address the challenge of ensuring applications have consistent and predictable performance by allowing users to separate storage groups based on performance requirements and business importance. PowerMaxOS gives the ability to set specified service levels to ensure the highest priority application response times are not impacted by lower priority applications. The available service levels are defined in PowerMaxOS and can be applied at the creation of a storage group or can be modified to an existing storage group at any time. eNAS is fully integrated with Service Levels introduced for PowerMax and VMAX All Flash. Service Levels can be associated with the storage group for the nas_pool which will determine the response time for the applications accessing data.

Service levels are offered with various ranges of performance expectations. The expectations of service levels are defined by their own characteristic of a target response time. The target response time is the average response time expected for the storage group based on the selected service level. Along with a target response time, service levels also have either an upper response time limit or both an upper and lower response time limit.

The service levels offered are:

| Service Level | DIAMOND  | PLATINUM  | GOLD  | SILVER  | BRONZE  | OPTIMIZED  |
|---------------------------|--|---|---|--|---|--|
| Target Response Time | X | X | X | X | X | |
| Upper Response Time Limit | X | X | X | X | X | |
| Lower Response Time Limit | | | X | X | X | |


HIGHEST
Priority and Performance
LOWEST


For more information about service levels, see the document [Dell EMC PowerMax Service Levels for PowerMaxOS](#).

Host I/O limits and eNAS

The Host I/O limits quality of service (QoS) feature was introduced in the previous generation of VMAX arrays. It enables customers place specific IOPS or bandwidth limits on any storage group, regardless of the service level assigned to that group. The I/O limit set on a storage group provisioned for eNAS applies to all file systems created from that storage group cumulatively. If the host I/O limits set at the storage-group level need to be transparent to the corresponding eNAS file system, there must be a one-to-one correlation between them. Assigning a specific host I/O limit for IOPS, for example, to a storage group (file system) with low-performance requirements can ensure that a spike in I/O demand does not saturate its storage, due to FAST inadvertently migrating extents to higher tiers, or overload the storage, affecting performance of more critical applications. Placing a specific IOPS limit on a storage group limits the total IOPS for the storage group, but it does not prevent FAST from moving data based on the service level for that group. For example, a storage group with Gold service level may have data in both EFD and hard-drive tiers to satisfy service level compliance yet be limited to the IOPS set by host I/O Limits.

File system considerations

Auto extending

You must specify the maximum size to which the file system should automatically extend. The upper limit for the maximum size is 16 TB. The maximum size you set is the file system size that is presented to administrators, if the maximum size is larger than the actual file system size.

Note: Enabling automatic file system extension does not automatically reserve the space from the storage pool for that file system. Administrators must ensure that adequate storage space exists, so that the automatic extension operation can succeed. If the available storage is less than the maximum size setting, automatic extension fails. Administrators receive an error message when the file system becomes full, even though it appears that there is free space in the file system. The file system must be manually extended.

For more information about auto extending a file system, refer to the [Managing Volumes and File Systems with VNX AVM document](#).

Use Automatic Volume Management for file system creation

Automatic Volume Management (AVM) is a Dell EMC storage feature that automates volume creation and volume management. By using the VNX command options and interfaces that support AVM, system administrators can create and extend file systems without creating and managing the underlying volumes.

The automatic file system extension feature automatically extends file systems that are created with AVM when the file systems reach their specified high water mark (HWM). Thin provisioning works with automatic file system extension and allows the file system to grow on demand. With thin provisioning, the space presented to the user or application is the maximum size setting, while only a portion of that space is actually allocated to the file system.

The AVM feature automatically creates and manages file system storage. AVM is storage-system independent and supports existing requirements for automatic storage allocation (SnapSure, SRDF, and IP replication).

It is recommended not to have any file system over 90% utilized. If a file system does reach this threshold, extend it. The reason for this is to allow at least 10% of the file system capacity for:

- File system extension
- Checkpoints or snapshots

Slice volume option when creating file systems

If you want to create multiple file systems from the same Storage Group/nas_pool, use the slice volumes option. Otherwise, the file system you are creating consumes all of the specified Storage Group/nas_pool, making it unavailable for any other file systems.

Use a high watermark and maximum file system size for auto extending file systems

If a user is employing auto extend of your file systems, set an HWM. The HWM specifies at what percent the file system should be auto extended. The default HWM is 90%. Once a file system reaches its HWM, it is auto extended, until the size of the free capacity is brought below the HWM again. The file system is extended, striping across all devices in the pool. This means that the used capacity might be brought down to just 1% below the HWM or more, depending on the amount and size of the volumes in the underlying storage pool.

Setting a maximum file system limit specifies the maximum size that a file system can grow to. To specify this maximum, type an integer and specify T for terabyte, G for gigabyte (the default), or M for megabyte.

If you do not set either of these variables, but DO setup auto extend of file systems:

- A file system auto extends once it reaches 90% of usable capacity.
- A file system auto extends to 16 TB, if the space is available.

For more information about auto extending a file system, see the document *Managing Volumes and File Systems with VNX AVN*.

Bandwidth-intensive applications and I/O Limits

If you are hosting an application that has a heavy I/O profile on eNAS, you can set a host I/O limit against the storage group that the application uses. This host I/O limit limits the amount of traffic that the array gives to this SG, regardless of the SLO that is set against the SG.

The I/O limit set on a storage group provisioned for eNAS applies to all file systems created on the volumes in that storage group. If the Host I/O limits set at the storage group level needs to be transparent to the corresponding eNAS file system, there must be a one-to-one correlation between them.

Assigning a specific host I/O limit for IOPS, for example, to a storage group (file system) with low-performance requirements can ensure that a spike in I/O demand does not saturate its storage, cause FAST inadvertently to migrate extents to higher tiers, or overload the storage, affecting performance of more critical applications.

Placing a specific IOPS limit on a storage group limits the total IOPS for the storage group. However, it does not prevent FAST from moving data based on the SLO for that group. For example, a storage group with Gold SLO may have data in both EFD and hard-drive tiers to satisfy SLO compliance yet be limited to the IOPS set by Host I/O Limits.

Data protection **SnapSure for checkpoints**

Filesystem checkpoints provide a point in time image of a given file system. This allows system administrators and eNAS users to return a file system to a given point in time, in order to recover lost or corrupt data, and to perform a variety of tasks on the data from this point in time. If checkpoints are required for eNAS, use SnapSure.

SnapSure enables you to create point-in-time logical images of a production file system. SnapSure uses a "copy on first modify" principle. A production file system consists of blocks. When a block within the production file system is modified, a copy containing the block's original contents is saved to a separate volume called the SavVol. Subsequent changes made to the same block in the production file system are not copied to the SavVol. The original blocks from the production file system in the SavVol and the unchanged production file system blocks remaining in the production file system are read by SnapSure according to a bitmap and blockmap data-tracking structure. These blocks combine to provide a complete point-in-time image called a checkpoint.

Checkpoints can be read-only or read/write. With SnapSure, you can restore a production file system to a point in time either type of checkpoint. Create checkpoints using the data protection tab in Unisphere for VNX. In Unisphere, you select the file system, the checkpoint name, and the pool to use for storing the checkpoints. Use Unisphere for VNX to schedule automated snapshots, allowing at least 15 minutes between snapshots.

If using NDMP backups, configure these two schedules so that they do not overlap.

Every checkpoint you create, is another file system, so it is important to take into account the number of file systems you are going to create when creating or scheduling checkpoints.

Virtual Data Movers and File Auto Recovery

eNAS supports Virtual Data Movers (VDMs). VDMs are used for isolating Data Mover instances within a secure logical partition. They are file system containers that isolate a Virtual Data Mover from other VDMs in the same Data Mover container. VDM is a security mechanism as well as an enabling technology that simplifies the disaster recovery failover process. It maintains file system context information (metadata) to avoid rebuilding these structures on failover. File systems can be mounted beneath VDMs that are logically isolated from each other. VDMs can be used to support multiple LDAP domains within a customer environment. They can also be used to rebalance file loads across physical Data Movers by moving VDMs and their underlying file systems between Data Movers. VDMs are important when deploying replication solutions.

Introduced in the HYPERMAX OS 5977.691.684, File Auto Recovery (FAR) allows you to manually failover or move a Virtual Data Mover (VDM) from a source eNAS system to a destination eNAS system. The failover or move leverages block-level Symmetrix Remote Data Facility (SRDF) synchronous replication, thus there is zero data loss in the event of

an unplanned operation. This feature consolidates VDMs, file systems, file system checkpoint schedules, CIFS servers, networking, and VDM configurations into their own `nas_pools`. This feature works for a recovery where the source is unavailable. For recovery support in the event of an unplanned failover, there is an option to recover and clean up the source system and make it ready as a future destination.

VDM failover can be run as a planned activity, or as a disaster recovery exercise. There are two different commands that can be run, depending on which scenario the administrator may be in:

- **Planned failover:** When running a planned failover, use the `--reverse` argument. This gracefully fails over the VDM in the specified session, and all its mounted file systems. The `--reverse` argument also includes a “Clean” of the original source site upon completion of the failover. The `--reverse` argument is to be run on the target site.
- **Disaster recovery:** In the event of an outage in the source site, and the host being down, the `--failover` argument is to be used. This is a reactive measure, and as the source host is down, a Clean operation is not run at this time. When the source site has been recovered, the session that has been failed over, requires a clean operation.

The manually initiated failover and reverse options can also be performed using the Dell EMC File Auto Recovery Manager (FARM) product. This product allows the ability to automatically failover Syncreplicable VDMs from a source to a target destination. FARM can be configured to monitor sync-replicated VDMs and trigger automated failover in the event of Data Mover, File System, Control Station or IP network unavailability, that would cause data unavailability to the eNAS client.

Pre-failover planning

If users are configuring your CIFS shares with Continuous Access (CA) you need to create an additional interface with connectivity to your domain controller (DC) to ensure that you do not lose access to your CIFS share when performing the first failover.

The Data Mover must always have access and be able to communicate with the DC to allow seamless failover of CA enabled CIFS shares. This can be achieved by configuring an additional interface on the DM of the remote site, prior to performing the first failover.

For all additional failovers, you need to ensure that the remote site, has an interface configured and active on it, prior to performing the failover.

FAR can be used for disaster recovery, maintenance and load-balancing operations. FAR has the following requirements:

- PowerMaxOS version 5978.144.144 or above
- SRDF is configured and enabled

See the support documentation for FAR: <https://www.dellemc.com/en-us/collaterals/unauth/technical-guides-support-information/products/storage/docu88921.pdf>

Backup solutions

eNAS supports traditional NAS backup applications, for example Dell EMC NetWorker or Dell EMC Avamar, which can both be found in the Dell EMC Data Protection Suite.

Further information about these solutions can be found at:

- https://support.emc.com/media52569_Data_Protection_Advisor_and_the EMC_Data_Protection_Suite.mp4?language=en_US
- https://support.emc.com/products/1095_NetWorker
- https://support.emc.com/products/27368_Backup-&-Recovery-Manager-Avamar

Note: A dedicated 8 Gb I/O module is required for NDMP backups.

eNAS upgrades

eNAS upgrades are performed during overall PowerMaxOS upgrades. During eNAS upgrades, the new code is deployed to the DMs, and a reboot of them is required to install the new code versions. The user has two options here:

- Allow the HYPERMAX/PowerMax Upgrade to automatically reboot the DM when required during the upgrade
- Manually reboot the DMs at a later date

There are many factors to consider here, including:

- Amount of file systems configured (each checkpoint is considered a file system)
- Size of each file system
- Amount, size, and type of I/O running to the file systems that are configured

Testing has shown that some I/O generators will not experience interruptions during these reboots. However, it is up to you to determine the type of workload that is running, and to determine if it can withstand the duration of outage that may be caused by the number of file systems to be failed over.

Conclusion

PowerMax and VMAX All Flash with eNAS provide consolidated platforms for both block and file storage. eNAS provides an easy way to provision, manage, and operate file and block environments while keeping in mind application performance needs.

This best practices guide provides configuration and usage recommendations for eNAS in general use cases. As always, every customer environment is different and may need different configuration from those mentioned in this document.

For more information about the reasoning for any of the recommendations made in this document, see the supporting documentation listed in the appendix.

Appendix: Technical support and resources

[Dell.com/support](https://support.emc.com) 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 products.

For additional information, see the following related resources:

- https://support.emc.com/products/38576_eNAS - eNAS support homepage
- https://support.emc.com/products/36656_VNX2-Series/Documentation/ - Dell EMC VNX support homepage
- https://support.emc.com/products/1095_NetWorker - Dell EMC NetWorker support homepage
- https://support.emc.com/products/27368_Backup-&-Recovery-Manager-Avamar - Dell EMC Avamar support homepage