



Technical Report

FlexArray and V-Series Best Practices Guide

Michael Peppers, NetApp
May 2019 | TR-3461

Abstract

This technical report covers best practices for the configuration of NetApp® V-Series and FlexArray® licensed FAS controllers, Fibre Channel SANs, and the storage arrays to which they connect.

TABLE OF CONTENTS

1	Introduction	5
2	Solution Overview	5
2.1	Data ONTAP 7-Mode Overview	5
3	Whats New - FlexArray Enhancements	6
3.1	Non-version specific enhancements	6
3.2	New ONTAP 9 FlexArray Enhancements	7
4	Installing a FlexArray and V-Series System	7
5	Back-End Setup	7
5.1	Aggregates	8
5.2	Summary of Best Practices	8
6	Native Disk Shelf Support	8
7	RAID Types	9
8	Disk Space Calculations	9
8.1	Right Sizing	9
8.2	Array LUN Sizing	10
8.3	Checksums and Capacity Overhead	10
9	LUN Paths	11
9.1	Shared Initiator Configuration	11
9.2	Shared Target Configuration	11
9.3	Crossed Back-End Configuration	11
9.4	Recommendation for Direct-Attached Configuration	12
10	Zoning	12
10.1	Two-Array-Port LUN Groups	12
10.2	Four-Array-Port LUN Groups	13
10.3	Eight-Array-Port Configuration (Clustered Data ONTAP Only)	14
10.4	Multiple-Array LUN Groups	14
10.5	Multiple Arrays	14
11	LUN Masking	15
11.1	HOST LUN IDs	15
12	Root Volume	16
12.1	Root Volume Location	16

12.2 Root Volume Sizes	16
13 Spare LUNs	16
14 Ports to Use on V-Series and FlexArray Systems.....	16
15 Test Procedures.....	17
15.1 7-Mode and Local Node in Clustered Data ONTAP	17
15.2 Clustered Data ONTAP	17
15.3 Things That Should Be Tested.....	17
16 Operational Guidelines and Upgrades	18
16.1 Hardware Upgrades	18
16.2 Software Upgrades	18
16.3 Fabric or Array Maintenance	19
16.4 Expanding or Changing LUN Groups.....	19
16.5 Powering Systems On and Off.....	19
17 FlexArray and V-Series MetroCluster	20
17.1 Installing a FlexArray and V-Series MetroCluster system	20
17.2 Supported MetroCluster Configuration.....	20
18 Array-Specific Notes	22
18.1 NetApp E-Series	22
18.2 HDS	22
18.3 IBM DS4xxx	22
18.4 EMC DMX/VMAX.....	23
19 FlexArray Virtualization Resiliency Enhancements.....	23
19.1 Delaying All-Path-Fail Condition Outage for Intermittent Failures	23
19.2 Consistent failover times on path failures	24
19.3 Automatic path quiesce if the number of intermittent failures on the path exceeds a threshold.	25
Version History	27

LIST OF TABLES

Table 1) Vendor math.....	9
Table 2) Checksum and WAFL overhead.....	10
Table 3) Two-array-port LUN group zoning.....	13
Table 4) Four-array-port LUN group zoning	14
Table 5) Zoning for topology in Figure 8.....	21

LIST OF FIGURES

Figure 1) FlexArray and V-Series clustered storage overview.....	6
Figure 2) Data ONTAP cluster with FAS and V-Series systems.....	6
Figure 3) Multiple storage array example.	8
Figure 4) Crossed back-end configuration.....	12
Figure 5) Two-port LUN group (CL1 active/CL2 passive).....	13
Figure 6) Four-array-port LUN groups.....	13
Figure 7) Multiple arrays.....	15
Figure 8) FlexArray and V-Series MetroCluster supported configuration.	21

1 Introduction

A NetApp FlexArray license can be added to any NetApp FAS8xxx controller. FlexArray replaces the NetApp V-Series family as NetApp's scalable solution to virtualize third-party arrays. Although FlexArray can now be ordered, V-Series is no longer available for sale; however, it is still a supported solution. Both FlexArray and V-Series utilize existing Fibre Channel infrastructure to stand as an initiator in front of a third-party array or several arrays. In that way, FlexArray and V-Series can leverage storage provisioned through the third-party arrays as though it is native NetApp disk. FlexArray and V-Series can create raid groups, aggregates, volumes, and LUNs. FlexArray and V-Series can therefore present all the protocols that the NetApp Data ONTAP® operating system supports: CIFS, NFS, FCoE, iSCSI, and FCP.

FlexArray is different from V-Series in that FlexArray is a feature that can be licensed on any FAS8xxx system running Data ONTAP 8.2.1 or later. FlexArray is the functional equivalent to the previously available array virtualization solution that NetApp offered for almost a dozen years: V-Series. V-Series is a hardware-specific solution that includes a V-Series license that allows the solution to virtualize third-party arrays and be functionally equivalent to FAS array controllers. As of Data ONTAP 7.3, V-Series can support native NetApp disks.

Like all NetApp storage appliances, the FlexArray and V-Series product line is based on the industry-tested Data ONTAP operating system. The FlexArray and V-Series product line is the first and only storage virtualization solution on the market that unifies block and file storage networking protocols (NAS, SAN, and IP SAN) under a common architecture. The line provides the complete suite of NetApp advanced data management capabilities for consolidating, protecting, and recovering mission-critical data for enterprise applications and users.

This guide focuses on all the areas that need to be addressed when deploying V-Series systems with third-party storage arrays.

2 Solution Overview


FlexArray and V-Series systems integrate into mission-critical SAN environments to provide a simple and elegant data management solution. This solution increases storage efficiency while decreasing management complexity, improving asset utilization, and streamlining operations to increase business agility and reduce total cost of ownership.

FlexArray and V-Series systems can be used as a NAS gateway to deliver CIFS or NFS access to a third-party array or to extend the full NetApp suite of solutions for FCP, iSCSI, and FCoE clients to your existing array. In addition to connecting to third-party storage arrays by using SAN or direct-attached storage, V-Series systems running Data ONTAP 7.3 and later can use NetApp disk shelves. To see which storage arrays and switches are supported, consult the NetApp Interoperability Matrix Tool (IMT). FlexArray licenses can be added to any FAS8xxx controller running Data ONTAP 8.2.1 and can use either or both native NetApp and third-party array-provided LUNs.

2.1 Data ONTAP 7-Mode Overview

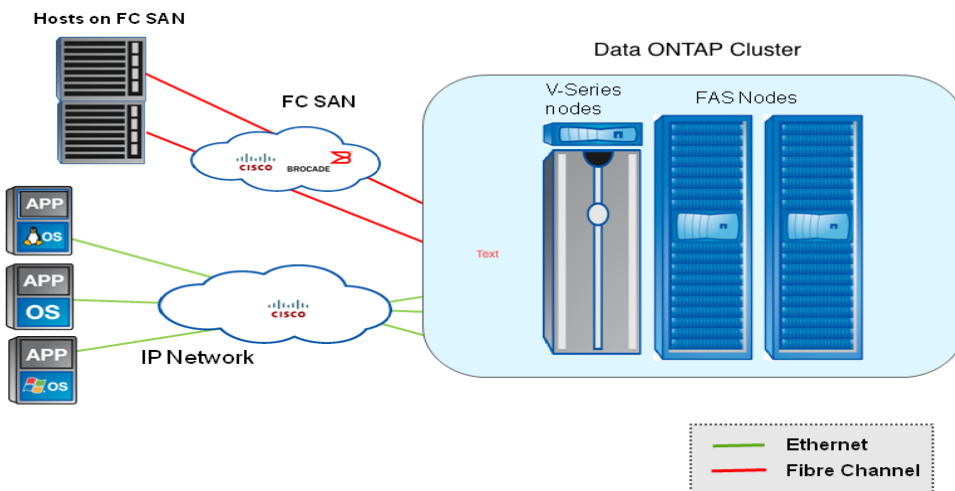
Both V-Series and FlexArray support Data ONTAP 7-Mode. The graphic below shows an example of where FlexArray and V-Series controllers would fit in your Fibre Channel environment.

Figure 1) FlexArray and V-Series clustered storage overview.



Both FlexArray and V-Series also support clustered Data ONTAP®, which provides scalable clustered storage through a single namespace. A Data ONTAP cluster can be built from FAS controllers, FlexArray, V-Series controllers, or a mix of FAS, FlexArray, and V-Series systems.

Figure 2) Data ONTAP cluster with FAS and V-Series systems.



3 Whats New - FlexArray Enhancements

3.1 Non-version specific enhancements

Direct attach support for FlexArray + E-Series Stretch Metrocluster

Direct attach support for FlexArray + E-Series Stretch Metrocluster was added to clustered ONTAP 8.3.1 and later releases. Review the [Stretch MetroCluster Installation and Configuration Guide](#)

3.2 New ONTAP 9 FlexArray Enhancements

FlexArray Virtualization Resiliency Enhancements

The enhancements include resiliency enhancements to reliably handle disruptions caused by back-end path failures, for example, link and switch failures. The resiliency enhancements include the following:

- Consistent I/O failover timing in case of path failures
- Temporary suspension of I/O for a tunable period of time in case of all-path fail conditions
- Automatic path quiescence in case of intermittent path failures

For more detailed information about these enhancements please go to [section 20](#) of this document.

Array LUN maximum size has increased to 16 TB

The maximum size of a back-end LUN (Array LUN) has gone from 6 TB to 16 TB.

4 Installing a FlexArray and V-Series System

A FlexArray and V-Series implementation using third-party storage can be thought of as having two parts: a front-end implementation and a back-end implementation.

The back-end implementation includes all the tasks required to set up the FlexArray and V-Series system to the point where Data ONTAP is installed, assuming that there are no native NetApp disks. If the FlexArray and V-Series system also contains native NetApp disks, then the best practice is to install Data ONTAP on those disks. Then you should provision the third-party array LUNs to provide additional storage space to the FlexArray and V-Series system. Initial FlexArray and V-Series implementation tasks include cabling, array LUN formatting, LUN security, zoning, assigning array LUNs to the FlexArray and V-Series controller, creating aggregates, and loading Data ONTAP. Back-end implementation is largely a SAN deployment exercise. Once the FlexArray and V-Series system can access array LUNs, the rest of the setup is a quick process.

The front-end configuration for a FlexArray and V-Series system is exactly the same as for a FAS implementation. Once an aggregate is built and Data ONTAP is installed, creating NetApp FlexVol® volumes; NFS/CIFS shares; iSCSI, FCP, or FCoE LUNs; NetApp Snapshot® schedules; and so on is exactly the same as for a FAS system. For clustered Data ONTAP deployments, the process for joining a node to an existing cluster or creating a new cluster is the same as for FAS systems. For help with these front-end tasks, consult the Data ONTAP documentation or search the [NetApp Technical Library](#).

The overall V-Series system or FlexArray licensed system design includes both front-end and back-end architecture. However, this guide is focused primarily on the back end, because the front-end setup of a FlexArray and V-Series system is identical to that of any FAS system.

If the FlexArray and V-Series system was ordered with NetApp disks, the factory installs the root volume, licenses, and Data ONTAP software, greatly simplifying the installation.

FlexArray and V-Series systems released with Data ONTAP 8.2.1 and later can use array LUNs if the proper license is installed. Refer to the implementation guide for further details.

5 Back-End Setup

The FlexArray and V-Series controller can attach to one or more storage arrays in addition to NetApp disk shelves. When you attach a controller to multiple storage arrays, these arrays can be from the same vendor or from different vendors. The NetApp [Interoperability Matrix Tool \(IMT\)](#) lists the supported arrays and array microcode that are certified for product use.


5.1 Aggregates

Data ONTAP creates pools of storage called aggregates. An aggregate can be built from NetApp disks or from array LUNs. You cannot mix disks and LUNs into the same aggregate. These aggregates, or pools of storage, are used to create dynamically resizable storage containers called FlexVol volumes. A FlexVol volume can serve as a mount point for NFS, a share for CIFS, or as a container for an array LUN.

NetApp recommends that any single aggregate be made from the same type and size of LUNs from the same storage array. For example, if the array has both SATA and FC disks, you should create two separate aggregates, one using LUNs built on the array SATA disks and a second aggregate built on LUNs from the array FC disks. Data ONTAP does not prevent you from using both types of LUNs in any aggregate, so care must be taken when provisioning storage.

The best practice is to use storage from a single array to create an aggregate. Mixing array types or LUN types is not supported. However, creating an aggregate from multiple arrays is possible provided the arrays are of the same type or family, the LUNs are built on array RAID groups of the same type, and the guidelines for attaching to multiple arrays are followed (per Figure 3).

Figure 3) Multiple storage array example.



5.2 Summary of Best Practices

The rest of this document focuses on this simple set of requirements and best practices:

- Use a NetApp disk shelf for the root volumes.
- Use the four-array-port LUN group described in section 9.1 if it is supported.
- Use dedicated array ports. Sharing array ports can adversely affect performance.
- Use single-initiator-to-single-target configuration and zoning.
- Use dedicated array RAID groups or storage pools.
- Assign a spare LUN. If using a NetApp disk shelf, the spare disk meets this requirement.
- Limit changes to back-end configurations once they are built. Changes to a stable system introduce the risk of disruption.
- Follow proper procedures when performing maintenance on the FlexArray and V-Series system, the back-end fabric, or the storage array.

6 Native Disk Shelf Support

V-Series systems also support NetApp disk shelves starting in Data ONTAP 7.3. FlexArray licensed controllers all can support native NetApp disks also. These disk shelves can be managed in accordance

with the “Data ONTAP Physical Storage Management Guide.” Consult this guide for installation and configuration of native disk shelves with V-Series systems. This guide can be found on the [NetApp Support site](#).

Disk shelves can be installed on new and existing FlexArray and V-Series systems. Review the [NetApp Interoperability Matrix Tool \(IMT\)](#) for the latest support requirements. You can order disk shelves with your new FlexArray system and add a new disk shelf to an existing FlexArray or V-Series system. You can also move a disk shelf from a FAS system to a FlexArray and V-Series system.

NetApp recommends that the root volume, when using native disk shelves, reside on the native disks.

7 RAID Types

Data ONTAP supports a variety of RAID types used by storage arrays—RAID 1, 10, 5, and 6—but requires some level of RAID protection on the array. Storage arrays cannot use RAID 0 for the LUNs that they make available to Data ONTAP. FlexArray and V-Series systems with a third-party array use RAID 0 to stripe data across the array LUNs, which splits data evenly across two or more array LUNs. Performance is improved the more disk spindles are used. But RAID 0 provides no data protection. Therefore, when creating RAID groups on storage arrays, follow the best practices of the storage array vendor. Then there will be an adequate level of protection on the storage array and disk failure will not result in loss of data or loss of access to data.

When using NetApp disks, NetApp RAID DP® technology is the default and recommended RAID level.

For the most recent information specific to your storage array, see the “Implementation Guide for Third-Party Storage” for Data ONTAP releases 8.3.2 and later. Or see the most recent implementation guide for your vendor for releases of Data ONTAP before 8.1.

8 Disk Space Calculations

With any FlexArray and V-Series deployment, there are a few things to take into consideration when planning storage provisioning. Among these are the differences in how vendors calculate a gigabyte; the maximum sizes that Data ONTAP supports; and overheads such as checksums, NetApp WAFL® (Write Anywhere File Layout), and Snapshot copies.

8.1 Right Sizing

Each vendor has its own calculation for capacity. These sizes can vary because of the differences in gigabyte definitions between the various storage vendors. Table 1 shows how various vendors define a gigabyte and the maximum size that Data ONTAP supports.

Table 1) Vendor math.

Vendor	How They Calculate a Gigabyte	500GB, 750GB, 1TB, 2TB, 3TB, and 4TB LUN Sizes
NetApp	1,000 x 1,024 x 1,024 bytes	n/a
IBM DS4000, IBM XIV, HDS, Fujitsu, EMC, 3PAR, HP EVA, IBM DS8000* (most arrays use 1,024 math)	1,024 x 1,024 x 1,024 bytes	488GB, 732GB, 976GB, 1952GB, 2916GB, 3904GB
IBM ESS800 (Shark), IBM DS8000*	1,000 x 1,000 x 1,000 bytes	512GB, 768GB, 1024GB, 2048GB, 3072GB, 4096

*Arrays that support both binary and decimal calculations are listed twice.

8.2 Array LUN Sizing

On the storage array, you should build LUNs on multiple dedicated array RAID groups. Ideally, from a given RAID group, you should build one or more large LUNs. When assigning the array LUNs to array ports, you should have at least 16 LUNs per port to make use of all the available disk queues. Data ONTAP treats each LUN as a separate device. Also, each LUN incurs some WAFL overhead in addition to counting toward the maximum device count limitations for a given FlexArray and V-Series platform. Therefore, a few dozen large LUNs are generally better than hundreds of small LUNs.

Each FlexArray and V-Series controller has a maximum capacity and a maximum number of devices (array LUNs or disks) that it can support. Using small array LUNs might result in a FlexArray and V-Series system reaching its device count limit before reaching its capacity limit. The maximum limits include both array LUNs and native disk drives; if you have 100 LUNs and 140 disks, they are combined to count as 240 toward the maximum spindle count. For capacity, LUN counts, and other metrics, see the [V-Series technical specifications](#) or the [Hardware Universe](#).

8.3 Checksums and Capacity Overhead

A checksum is a form of redundancy check, a simple measure for protecting the integrity of data through error detection. Checksums are used mainly in data storage and networking protocols. They add up the basic components of a message, typically the bytes, and store the resulting value. Later, checksums can be used to verify that the message was not corrupted by doing the same operation on the data and checking the sum.

Block Checksums (BCSs)

Block checksums are the default checksum type, optimized for performance drives. Most NetApp disks smaller than 3TB support only this type of checksum, and it is the preferred checksum for virtualized array LUNs. BCSs work well with all types of workloads. BCSs uses 12.5% of the raw device capacity.

Advanced Zoned Checksums (AZCSs)

Advanced Zoned Checksums are optimized for large capacity drives, starting with Data ONTAP 8.1.1. On FlexArray and V-Series systems with third-party storage, NetApp does not recommend AZCSs for high-performance random read workloads. This caution does not apply to NetApp disks. AZCSs use ~1.56% of the raw device capacity.

File System Overhead

When calculating the raw usable capacity of an array LUN as seen on a FlexArray and V-Series system, checksums and WAFL overheads are taken into account. When an aggregate is created, an additional 11% is used for WAFL (Data ONTAP's file system). See section 7.3 for checksum overheads.

Additional user-configurable overheads include the aggregate Snapshot reserve and FlexVol volume Snapshot reserve. These overheads can be tuned based on customer SLAs for snapshot retention.

Table 2) Checksum and WAFL overhead.

Checksum Type	Overhead	WAFL Overhead	Total Overhead
BCS	12.5%	11%	23.5%
AZCS	1.56%	11%	12.56%

9 LUN Paths

A FlexArray and V-Series system generally supports two paths from each controller in an HA pair to a given array LUN. These two paths should be on separate and redundant fabrics. Clustered Data ONTAP systems can support up to four paths per array LUN, but you would still spread the ports across two physically separate fabrics. For any given LUN, Data ONTAP assigns one path as the active and the other paths as the passive, or failover, path for each FlexArray and V-Series controller. Multiple LUNs might have different primary paths.

NetApp recommends no more than a single “hop” between the FlexArray and V-Series system and the array, though up to three hops are supported. A hop equals one ISL connection between edge-to-core, core-to-core, or core-to-edge switches. NetApp also recommends using dedicated array ports with one-to-one initiator-to-target port mapping. Use fabric zones and array host groups or LUN access control on the storage array to prevent FlexArray and V-Series systems from accessing LUNs owned by non-FlexArray and V-Series hosts and vice versa. Both nodes in an HA pair need access to the other’s array LUNs for storage failover to work.

Starting with Data ONTAP 8.3, I/O requests to an array LUN can be distributed through all the available optimized paths to the LUN. The I/O requests are no longer sent over a single active optimized path. See the [“Support for Load Distribution Over Multiple Paths to an Array LUN” section of the Clustered Data ONTAP 8.3.2 Release Notes](#) for more information.

9.1 Shared Initiator Configuration

Starting with Data ONTAP 8.2, connecting an FC initiator port on a Data ONTAP system to up to four target ports on separate storage arrays is supported. Certain rules must be followed while sharing an initiator; for example, all storage arrays must be in the same vendor model family,

Connecting a single FC initiator port to multiple target ports is supported for NetApp MetroCluster™ configurations with Data ONTAP 8.3 and later. For details, refer to the [FlexArray Virtualization Installation Requirements and Reference Guide](#) and [Fabric-attached MetroCluster Installation and Configuration Guide](#)

9.2 Shared Target Configuration

Starting with Data ONTAP 8.3, connecting a maximum of two Data ONTAP FC initiator ports to a single target port on the storage array is supported. Each target port is zoned into two FC initiator ports, one from each cluster node.

The rules for this configuration are as follows:

- When the Data ONTAP systems are in a HA pair, each node can share at the most one FC initiator port with the same target port.
- All storage arrays must be from the same vendor and model family.
- Connecting a single target port to multiple FC initiator ports is not supported for MetroCluster configurations.
- The best practice for zoning is to have each FC initiator target port pair in a separate zone (1:1).

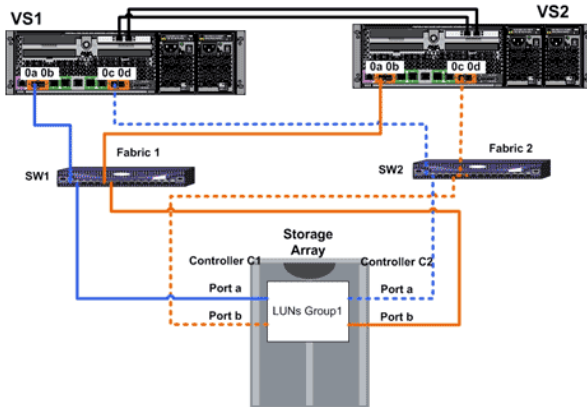
9.3 Crossed Back-End Configuration

A crossed back end, in which one storage array controller is attached to both fabrics, is supported on arrays that support ALUA and active-active arrays. For two-path active-passive arrays, a crossed back end is not supported. However, a crossed back end is supported on four-path active-passive arrays.

Some storage arrays are set up to access their fabrics using a crossed back end. A crossed back end design can better utilize the hardware during a hardware failure or during maintenance. During normal

operation in which all hardware is working, there is no benefit to the crossed back end. There are utilization advantages only during maintenance.

Figure 4) Crossed back-end configuration.



9.4 Recommendation for Direct-Attached Configuration

Starting with Data ONTAP 8.2.1, you can connect Data ONTAP systems in direct-attached configurations with E-Series storage arrays. You must ensure redundancy in connections between the Data ONTAP systems and the storage arrays to avoid the possibility of a single point of failure. The Interoperability Matrix contains information about the specific E-Series array models that are supported in direct-attached configurations with Data ONTAP systems. Data ONTAP supports the following direct-attached configurations with E-Series storage arrays:

- Stand-alone basic configuration
- Four-port-array LUN group
- Eight-port-array LUN group

Refer to [TR-4461i: NetApp Clustered Data ONTAP 8.3 FlexArray with E-Series and SANtricity 11.20](#) for more information.

10 Zoning

Configuring zoning on a Fibre Channel switch enables you to restrict visibility and connectivity between devices connected to a Fibre Channel SAN. FlexArray and V-Series systems support configurations in which a single initiator is connected to a single target port. Create separate zones for each initiator and target pair. NetApp does not recommend using shared array ports for performance reasons. Both software (WWPN) and hardware (port-based) zoning are supported.

End-to-end link speed between the NetApp initiator and storage array target ports should be consistent in a fabric.

10.1 Two-Array-Port LUN Groups

A typical FlexArray and V-Series implementation utilizes two separate back-end fabrics between the FlexArray and V-Series system and their back-end storage array(s). Figure 5 shows a two-port LUN group in which there are two paths from each FlexArray and V-Series system to each of the two back-end storage array controllers.

Figure 5) Two-port LUN group (CL1 active/CL2 passive).




Table 3) Two-array-port LUN group zoning.

SW1	Members	SW2	Members
Zone1	VS1:Port 0a, Controller CL1:Port A	Zone2	VS1:Port 0c, Controller CL2:Port A
Zone1	VS2:Port 0a, Controller CL1:Port B	Zone2	VS2:Port 1c, Controller CL2:Port B

10.2 Four-Array-Port LUN Groups

This is the recommended configuration. This configuration maximizes both performance and resiliency.

Table 3 shows an example of how zoning should be set up with a FlexArray and V-Series HA pair connected to a storage array using a four-array-port configuration. The array LUNs are assigned to all four array ports. The FlexArray and V-Series VS1 accesses the LUNs on array A ports; VS2 accesses the LUNs on array B ports.

Figure 6) Four-array-port LUN groups.

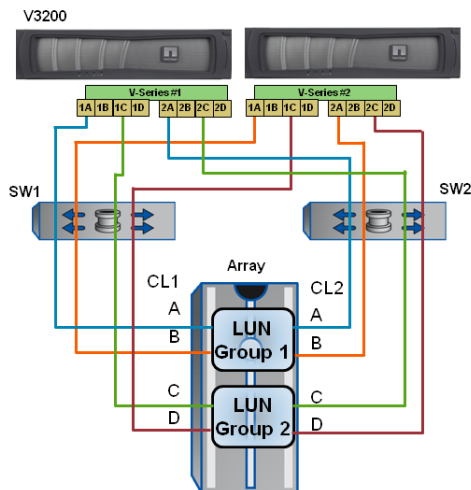


Table 4) Four-array-port LUN group zoning.

SW1	Members	SW2	Members
Zone1	VS1:Port 1a, Controller CL1:Port A	Zone1	VS1:Port 1c, Controller CL2:Port A
Zone2	VS2:Port 1a, Controller CL1:Port B	Zone2	VS2:Port 1c, Controller CL2:Port B

10.3 Eight-Array-Port Configuration (Clustered Data ONTAP Only)

Beginning with clustered Data ONTAP 8.1, V-Series can support four paths to a single LUN. All FlexArray controllers support four paths. Figure 8 shows a two-node cluster. Because a cluster is built from one or more HA pairs, a larger cluster would simply repeat this pattern (to a different LUN group or other arrays). The same considerations offered above for multiple array or multiple LUN group configurations follow for Cluster-Mode as well. There is simply the ability to configure more paths.

The number of LUNs supported in a single LUN group was up to 100 until Data ONTAP 8.2. Since then the number of LUNs supported in a single LUN group is up to 254.

When using a version earlier than Data ONTAP 8.2, our recommendation is to use the eight-array-port configuration sparingly, primarily for high-performance storage in which port bandwidth is likely to be a bottleneck. LUNs backed by SATA disks, for example, are not good candidates for this configuration. Because of the increase in the number of supported LUNs, we no longer recommend avoiding using eight-array-port configurations sparingly because they no longer reduce the number of supported LUNs per array.

10.4 Multiple-Array LUN Groups

Data ONTAP 7.3 introduced broader support for attaching to multiple LUN groups in the same storage array, enabling greater system performance. Before 7.3, many arrays were limited to a single LUN group. Figure 5 shows the supported topology for attaching a V-Series system to an array with two array LUN groups. Expanding this example to enable attaching to more than two array LUN groups is supported, provided that best practices are followed.

Each FlexArray and V-Series system has two paths to a given LUN group through separate fabrics and separate array controllers. VS1 port pair 1a and 2a accesses LUN group1 and VS2 port pair 1a and 2a also accesses LUN group1. V-Series ports 1c and 2c are used for LUN group2.

The array is configured with a four-array-port LUN group; all of the LUNs in this array LUN group are assigned to four-array ports, as shown. The array LUNs from group1 are assigned to array controller CL1 ports A and B and to CL2 ports A and B.

Fabric zoning best practices mandate single-initiator-to-single-target zoning. VS1:1a is zoned to array CL1:A, and VS1:2a is zoned to array CL2:A. Similarly, VS2:1a is zoned to array CL1:B, and VS2:2a is zoned to array CL2:B. This pattern is simply repeated for additional LUN groups using additional initiators.


To maximize performance with multiple-array LUN groups, half the LUNs of each LUN group should be owned by each FlexArray and V-Series controller in the HA pair. Half the LUNs in LUN group 1, above, would belong to VS1; the other half would belong to VS2. This pattern is repeated for LUN group 2, and for any additional LUN groups. Doing so allows all the ports to be utilized for data traffic and not simply as HA paths.

10.5 Multiple Arrays

NetApp supports multiple arrays behind a single FlexArray and V-Series system or HA pair. This setup is very similar to the multiple-array LUN groups example in section 9.3. The same dedicated-initiator-to-

dedicated-target zoning rules apply, as well as the same two paths per FlexArray and V-Series system limit (four per HA pair) per LUN group. Figure 7 shows this model.

Figure 7) Multiple arrays.



Although the ports were connected to LUN group 2 in the previous example, they are connected to Array 2 here. Arrays from different vendors, or different families from a single vendor, are supported.

When connecting to multiple storage arrays, you can spread an aggregate across the storage arrays only if the arrays are from the same family and use the same disk type. If this is not the case, performance is much less predictable.

11 LUN Masking

LUN masking is similar to switch zoning, but it is performed on the storage array. LUN masking controls which hosts can access which array LUNs. FlexArray and V-Series systems are hosts as seen by the storage array. Storage vendor terminology for LUN masking differs among vendors. Terms commonly used include host groups, LUN security, LUN mapping, and affinity groups.

On the storage array, you can create LUNs and assign the proper host type and any other array-specific settings as defined in the implementation guide for your array. You can create host groups on the storage array, assign the port on the array to the host group, assign the FlexArray and V-Series ports to the host group, and then assign LUNs to that host group.

NetApp recommends using both zoning and host groups for added protection for the FlexArray and V-Series system.

ALUA support is required for EMC CX and E-Series arrays from clustered Data ONTAP 8.3 onward.

11.1 HOST LUN IDs

Each unique array LUN must present the same LUN ID to the array ports. For example, the unique LUN A must be LUN number 15 on each array port. You cannot present the unique LUN A as LUN 15 on one port and LUN 37 for the same LUN on another port. Doing so is not supported and will not work.

12 Root Volume

12.1 Root Volume Location

When you use NetApp storage, NetApp recommends that the root volume be placed on NetApp native disks.

If you have a system that has only third-party storage arrays and no NetApp disks, install the root volume as a FlexVol volume on a single array LUN. A new menu will be listed that allows the user to select LUNs to install root volume. Refer to the [“Overview of Setting Up Data ONTAP to Use Array LUNs” section of the Physical Storage Management Guide](#) for further details.

12.2 Root Volume Sizes

The minimum size of the root volume varies depending on the FlexArray and V-Series model and the version of Data ONTAP. The underlying array LUN used for the root volume should be sized properly to provide a FlexVol volume of the required size. For root volume size details, see the legacy [NetApp Interoperability Matrix Tool](#) for V-Series controllers and the [NetApp Hardware Universe](#) for minimum root array LUN sizes.

13 Spare LUNs

If NetApp disks are not being used, NetApp recommends that you have a spare LUN for each FlexArray and V-Series system in an HA pair. For all FlexArray and V-Series systems that have FCP or iSCSI licensed, a spare LUN is required. The spare LUN is used to save a core file when needed. This LUN is not treated as a spare LUN in the same sense that a spare disk with RAID DP is treated. The LUN is not used to replace a lost LUN the way a spare disk in RAID DP would be used to replace a failed disk.

The size of the spare LUN depends on the memory and NVRAM size of a given FlexArray and V-Series system and a small buffer. For example, if the system has 20G of RAM and 4G of NVRAM, then a 30G array LUN would suffice. Each FlexArray and V-Series controller would need a spare LUN. Review the [NetApp Hardware Universe](#) for the minimum spare array LUN size for your FlexArray and V-Series controller(s).

14 Ports to Use on V-Series and FlexArray Systems

V6000 systems have eight onboard Fibre Channel ports; V6200, V3000, and V3100 systems have four onboard Fibre Channel ports; and V3200 systems have two Fibre Channel ports. FlexArray on the FAS8020 can use the two onboard UTA2 ports; the other three FAS8xxx controllers all have four onboard UTA2 ports that can be used for FlexArray. In addition to the onboard ports, an additional expansion card can be installed if additional initiator ports are needed. The onboard ports or UTA2 ports can be configured either as initiators or as targets and should be set to initiator mode for back-end LUN access. All models also have expansion slots for additional cards where initiator cards or target cards can be installed. V3200 and FAS8020 systems might require an I/O expansion module to support larger configurations.

When choosing FC initiator port pairs to communicate with an array LUN group, it is best to choose ports from different HBAs or separate ASICs within an HBA. The onboard V-Series FC ports 0a and 0b are on one ASIC, ports 0c and 0d are on another ASIC, and so on. In systems using PCI initiator cards, you should take advantage of separate PCI cards. For example, if cards are in slots 2 and 3, 2a and 3a would be a port pair. The V3200 series, except the V3210, require the use of PCIe FC initiator cards to attach to array LUNs. NetApp recommends, though does not require, using a PCIe card for array attaching on the V3210.

V-Series systems configured for Fibre Channel SANs can use the onboard ports or expansion slot initiators. Similarly, the target ports can be internal or expansion slot ports, depending on the number of required ports. See the “SAN Configuration Guide” for your version of Data ONTAP for information about recommended front-end SAN topologies. The guide is part of the Data ONTAP documentation.

15 Test Procedures

Before putting a FlexArray and V-Series system into production, it is important to make sure that everything is configured properly. The idea is to identify problems now rather than in production. Identifying problems early also provides an opportunity for documentation and customer education. This list covers FlexArray and V-Series specific tests. Front-end testing should also be completed, but because this testing is identical to FAS it is not covered here.

Use Data ONTAP commands to capture information about the environment.

15.1 7-Mode and Local Node in Clustered Data ONTAP

- **sysconfig -a** to get system information
- **sysconfig** to show disk counts and capacity for each port
- **disk show -v** to see which LUNs are assigned
- **storage show disk -p all** for primary and secondary path information
- **storage show adapter** for adapter details, including FC node name and so on
- **fcadmin config** for status of each adapter
- **sysconfig -r** for aggregate and volume information

15.2 Clustered Data ONTAP

Clustered Data ONTAP introduces a new set of commands to better organize configuration data. The commands above are still available, but only at the node level. The new commands that follow should be run in the cluster instance and return data for every node in a cluster.

- **storage path show-by-initiator** shows you which initiators are connected to which switch ports.
- **storage disk show** shows you the storage devices attached, who owns them, and to which aggregate they belong.
- **storage array config show** shows information about the connected arrays for each node.
- **storage errors show** shows back-end configuration errors.
- **storage load show** shows information about how IOPS are spread across devices.

15.3 Things That Should Be Tested

You need to ensure that you have the correct number of paths for each LUN. Two commands provide the information needed to verify the paths. In clustered Data ONTAP, these commands are run in the local node scope for the nodes attached to the storage.

- Check **storage show disk -p** to verify that you have a failover path for each LUN. If you disable a path, only the failover path should be listed.
- Use **sysconfig** to verify that each port pair has the same number of disks and the same capacity.
- Path failover
 - To disable an initiator:

- Assuming that 0a is used, **fcadmin offline 0a**; to verify paths, **storage show disk -p a11**. Manually execute CFO or SFO (if running clustered Data ONTAP) and giveback. Enable and repeat for each path.
- For switch failure:
 - Disable all paths to a given switch for both FlexArray and V-Series controllers to simulate a switch failure. Manually execute CFO or SFO (if running clustered Data ONTAP) and giveback. Enable and repeat for each switch.
- For array controller failure:
 - Disable all the switch ports from the array controller to the switch or reboot the array controller (if possible). Manually execute CFO or SFO (if running clustered Data ONTAP) and giveback. Enable and repeat for each array controller.

For Data ONTAP operating in 7-Mode, see the information in the “Data ONTAP 7-Mode Setup Guide” about testing FlexArray and V-Series setup. For clustered Data ONTAP, see the “Data ONTAP Cluster-Mode Physical Storage Management Guide.” These documents are available on the [NetApp Support site](#).

16 Operational Guidelines and Upgrades

16.1 Hardware Upgrades

Before starting any hardware upgrade, make sure that you have the NetApp AutoSupport® tool enabled or record the following information:

- **printenv** to show any variables that were set at the **CFE>** or **LOADER>** prompt (**fc-no-mgmt?** and **fc-non-array-adapter-list**).
- Output of **disk show -v**, **storage show disk -p**, **sysconfig -a**, **sysconfig -r**, storage show adapter, fcadmin config, license commands (or a NetApp ASUP® notification + **disk show -v** and **storage show disk -p**).
- For clustered Data ONTAP, the above commands should be run at the local node level, in addition to these commands run in the cluster instance: **storage path show-by-initiator**, **storage disk show**, **storage array config show**.

If you are upgrading both the FlexArray and V-Series controller model and the Data ONTAP version, perform the Data ONTAP upgrade before any hardware upgrade. Make sure that the new version of Data ONTAP works with the existing setup before moving on to upgrading the FlexArray and V-Series system. A new FlexArray and V-Series system requires changes to fabric zones, array LUN security, array host groups, and disk ownership because the system ID and WWNs will change.

There are a number of things to consider when upgrading FlexArray and V-Series controllers. If you are able to move the HBA cards from the old head to the new one, there is no need to make any zoning or host group changes because the WWNs of the HBAs stay the same. If you cannot move the cards, you must redo zoning (in case of soft zoning) and host groups because the new HBA cards will have different WWNs.

16.2 Software Upgrades

When updating Data ONTAP, first verify, by checking the [NetApp Interoperability Matrix Tool \(IMT\)](#), that the new version of Data ONTAP is supported with the firmware and microcode levels on the array and fabric. The upgrade procedure is identical to that of FAS, but because FAS recommends sending just an AutoSupport message, FlexArray and V-Series should have the ASUP and **disk show -v** output. If you do not have AutoSupport enabled, collect the output from the commands listed in section 15.1, above. Environment variables are not required for software upgrades.

16.3 Fabric or Array Maintenance

From time to time, the back-end components might need firmware or microcode upgrades. FlexArray and V-Series systems are sensitive to changes on the back-end fabrics and arrays. To prevent service disruption, it is important for any new firmware version for any switches or arrays the FlexArray and V-Series systems use to be listed on the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Fabric maintenance might include firmware upgrades or switch replacements. FlexArray and V-Series systems must have redundant fabrics, and each fabric must be upgraded separately. Do not begin maintenance on the second fabric until you verify that the FlexArray and V-Series systems can see all of their LUNs through the first fabric. Failure to do so could result in an outage.

You can verify that redundant paths exist by checking the output from the `storage show disk -p` command, run on the local node. Because each LUN name typically includes the switch name, make sure that at least one of these paths is available at all times. With clustered Data ONTAP, simply run the `storage error show` command to confirm redundant paths.

As with any type of maintenance, care should be taken to minimize exposure to service disruptions, preferably by using scheduled maintenance windows. It is also important to send an AutoSupport message (options `autosupport.doit [eventname]`) or to collect the data listed in [section 16](#). This information enables NetApp Technical Support to quickly diagnose and resolve any problems that result from the maintenance.

Array maintenance is generally in the form of firmware upgrades. Most arrays support nondisruptive firmware upgrades (NDUs), but check the [FlexArray/V-Series Support Matrix](#) to verify that your new array firmware or microcode is supported and that array NDU is supported.

16.4 Expanding or Changing LUN Groups

Adding Paths

If you are upgrading to the four-array-port model, make sure that you do not exceed the limit of two paths per FlexArray and V-Series system or four paths for a clustered Data ONTAP system. It is also important that you have the LUNs shared through the target ports before the initiator is connected.

Adding LUNs

When adding storage to a FlexArray and V-Series system, follow these guidelines:

- The maximum number of LUNs supported in a LUN group is up to 254.
- When you add to an existing aggregate, the new LUNs should be the same size as existing LUNs.
- A given LUN must have the same host ID on every port on which it is shared.

Removing LUNs

When removing an array LUN from a FlexArray and V-Series system, it is important to use the correct command so that the SCSI reservation is released for that LUN. Some types of arrays do not allow administrators to remove this reservation, complicating the removal process. This procedure involves advanced Data ONTAP commands. For details, see the [V-Series Maintenance Guide](#) or the [FlexArray Virtualization Installation Requirements and Reference Guide](#).

16.5 Powering Systems On and Off

Follow a sequence to power on and power off the various components in a FlexArray and V-Series deployment. Follow this order when powering on or off the whole deployment:

- Power on:

- Storage array and any SAN switches
- Boot the FlexArray and V-Series system
- Power off:
 - Halt the FlexArray and V-Series system
 - Storage array and any SAN switches

17 FlexArray and V-Series MetroCluster

MetroCluster is a unique solution that combines NetApp HA pairs with synchronous mirroring to deliver continuous availability and zero data loss at the lowest cost. As a self-contained solution at the NetApp storage controller level, MetroCluster can transparently recover from failures, including complete loss of a data center, enabling business-critical applications to continue uninterrupted.

17.1 Installing a FlexArray and V-Series MetroCluster system

Installing FlexArray and V-Series MetroCluster is like installing two separate FlexArray and V-Series systems at different locations, then enabling HA and enabling NetApp SyncMirror® technology between the two sites. FAS rules must be followed for SyncMirror. MetroCluster systems can be configured as fabric attached with distance support up to 300km or as stretch with support up to 500m. For detailed guidelines for systems running Data ONTAP releases 8.1 and later, see the [Clustered Data ONTAP 8.3 MetroCluster Installation Express Guide](#).

Please review the [MetroCluster FAQ](#) for additional information.

17.2 Supported MetroCluster Configuration

Figure 8 shows the supported FlexArray and V-Series MetroCluster topology. V-Series (VS1) at Site A has local storage from Array1-GP1 LUNs. These LUNs are synchronously mirrored to the remote array at Site B in the Array2-GP1 LUNs. VS1 accesses these two sets of LUNs through 2a/2c for the local pool and 2b/2d for the remote pool. Each NetApp initiator port is zoned to a single array target port. Likewise, the HA partner at Site B has a local and a remote pool attached through different initiator port pairs and is also synchronously mirrored between Site A and Site B.

Figure 8) FlexArray and V-Series MetroCluster supported configuration.

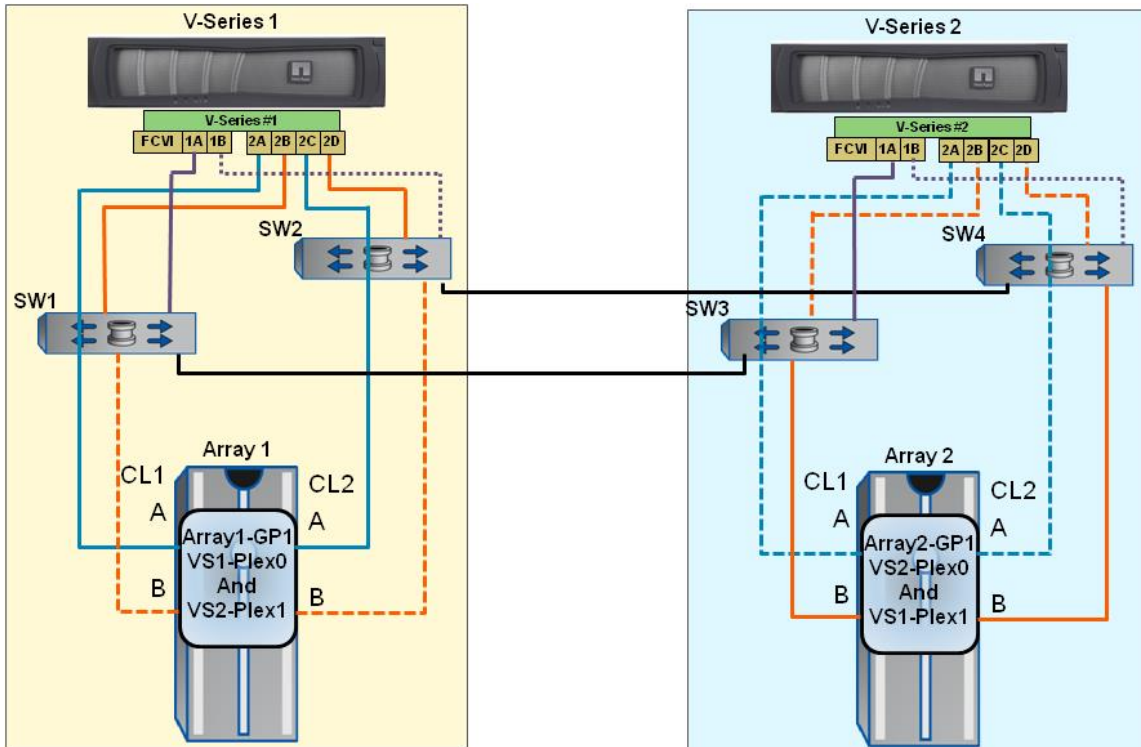


Figure 8 shows a V3200 MetroCluster system, each controller with a local and a remote storage pool. Larger systems can be built by using additional storage pools, new sets of V-Series initiator ports, array target ports, and switch ports.

Table 5) Zoning for topology in Figure 8.

Name	SW1-SW3 Fabric	Name	SW2-SW4 Fabric
Zone 1:	VS1:2a, Array1:CL1:A	Zone 1:	VS1:2c, Array1:CL2:A
Zone 2:	VS1:2b, Array2:CL1:B	Zone 2:	VS1:2d, Array2:CL2:B
Zone 3:	VS2:2a, Array2:CL1:A	Zone 3:	VS2:2c, Array2:CL2:A
Zone 4:	VS2:2b, Array1:CL1:B	Zone 4:	VS2:2d, Array1:CL2:B
Zone 5:	VS1:FCVI:a, VS2:FCVI:a	Zone 5:	VS1:FCVI:b, VS2:FCVI:b

When using an existing SAN that is shared with other fabric hosts, it's important to make sure that a supported configuration is possible. NVRAM mirroring uses FCVI and is much less tolerant of variations in SAN performance and design than FCP. These limitations are explained next.

FlexArray and V-Series MetroCluster Supports:

- Data ONTAP running in 7-Mode or clustered Data ONTAP.
- Symmetric configurations between sites (array family, switch types, firmware versions).
- Brocade and Cisco switches (check the support matrix for specific models and firmware).
- Number of LUNs supported in each array LUN group is 254.
- SAN requirements include:

- Dedicated fabric.
- Maximum supported distance for a FlexArray and V-Series MetroCluster system is up to 300km (with clustered Data ONTAP 8.3.2).
- Use single-initiator zoning, as shown in Table 5.
- Long-distance ISL shares FCVI (cluster interconnect) and data traffic (sufficient bandwidth must be available to support NVRAM mirroring).

For detailed guidelines for systems running Data ONTAP 8.1 and later, see the “Data ONTAP 7-Mode High-Availability Guide” or the [MetroCluster Installation and Configuration Guide](#). For detailed guidelines for systems running Data ONTAP releases before 8.1, see the “V-Series MetroCluster Configuration Guide.”

18 Array-Specific Notes

FlexArray and V-Series systems do not support any array advanced services or features, such as array-based Snapshot copies, thin provisioning, or mirroring schemes such as TrueCopy, FlashCopy, PPRC, Metro Mirror, Global Copy, and so on. Instead, use the NetApp software features in Data ONTAP (Snapshot, NetApp SnapMirror® software, SyncMirror, NetApp SnapMover® software, FlexVol, NetApp FlexClone® volumes, and so on). For information about how to properly configure array options, see the “V-Series Implementation Guide” or the [FlexArray Virtualization Implementation Guide for Third-Party Storage](#). Read and follow the notes for each array in the V-Series Support Matrix. The rest of this section summarizes some highlights.

18.1 NetApp E-Series

FlexArray and V-Series controllers support NetApp branded E-Series arrays. The specific models supported varies with the versions on ONTAP being run. Our standard configuration guidelines apply equally to both array families (choosing FlexArray, V-Series controllers, LUN sizing, zoning, and so on).

For details about supported configurations, consult the [NetApp Interoperability Matrix Tool \(IMT\) and the Implementation Guide for NetApp E-Series Storage](#). For additional information on configuring and managing the E-Series, consult the relevant E-Series family documentation which can be found in the Documentation menu on the [Documentation menu on the NetApp Support site](#).

18.2 HDS

FlexArray and V-Series systems can attach to HDS (USP, USP-VM, VSP) arrays that can virtualize LUNs presented from other storage arrays. When using USP tiered storage features:

- FlexArray and V-Series root volumes and spare core volumes should be on USP internal disks.
- Do not mix LUN types when building the aggregates. For example, do not mix LUNs based on FC and SATA drives, and don't mix USP internal and virtualized LUNs into the same aggregate.
- Follow HDS recommendations for external storage array attach and for queue depth settings on the arrays. The FlexArray and V-Series settings are the default settings.

18.3 IBM DS4xxx

On IBM DS4000 arrays, Storage Manager is used to manage the array in band (over Fibre Channel) or out of band (over Ethernet). If the array is being managed in band, Storage Manager creates a LUN 31 for its internal use. This LUN 31 should not be mapped or seen by the FlexArray and V-Series system. Similarly, any access LUN that is created in a host group should not be mapped to the FlexArray and V-Series system.

18.4 EMC DMX/VMAX

On DMX, use the LUN masking capability to associate logical devices to all host (channel) director ports, but allow only the FlexArray and V-Series FC initiator ports to see the LUNs for FlexArray and V-Series systems. Do not present the volume configuration management database (VCMDB) to all hosts by default. Configure the global setting to restrict visibility to the VCMDB unless it has been specifically made visible to a particular host. The benefit is that masking is at the logical level (initiator port WWN) rather than at the physical or switch zone level.

19 FlexArray Virtualization Resiliency Enhancements

Beginning with ONTAP 9.0, FlexArray Virtualization (V-Series) includes resiliency enhancements to reliably handle disruptions caused by back-end path failures, for example, link and switch failures.

The resiliency enhancements include the following:

- Consistent I/O failover timing in case of path failures
- Temporary suspension of I/O for a tunable period time in case of all-path fail conditions
- Automatic path quiescence in case of intermittent path failures

Use-cases where the tunable parameters can be used to handle link and path failures like those described below:

19.1 Delaying All-Path-Fail Condition Outage for Intermittent Failures

All Path Fail condition occurs when last path to an array LUN goes down. Events that can make all path or last path to go down are:

- Array side disruptions
- Upgrades/replacements of storage components like switches, array controllers etc...
- Intermittent component failures

ONTAP 9.0 provides an optional capability to enable FlexArray to identify all-path fail scenario and delay the outage for a short duration of time. Delaying all path fail condition will make ONTAP suspend IO for an extended duration before resulting in Multi-Disk panic (MDP), this gives a chance for IO to resume if any path comes back online within the specified duration.

Option 'all-path-fail-delay' is provided for user to specify the duration for which all-path fail condition needs to be delayed. This can take a value in the range of 30-90 seconds, or 0 to disable the option. This option is disabled by default. Once this is enabled it is applicable to all the nodes that see the array within the cluster and across clusters in MetroCluster configuration. Value set for option 'all-path-fail-delay' will be persistent across reboots.

Option	all-path-fail-delay
Default value	0
Range	0,30-90

Steps:

1. List the arrays connected to the cluster using 'storage array show' command.

```
Cluster01::> storage array show
Prefix          Name  Vendor      Model Options
-----
NET-1           NETAPP_INF_1  NETAPP      INF
```


- Set the delay for all path fail condition to <delay> sec on the array using 'storage array modify' command:

```
storage array modify -name <Array-name> -all-path-fail-delay <in secs>
```

```
Cluster01::> storage array modify -name NETAPP_INF_1 -all-path-fail-delay 75
```

- Ensure the value is set by executing 'storage array show' command.

```
Cluster01 ::> storage array show -name NETAPP_INF_1
```

```
Name: NETAPP_INF_1
```

```
Prefix: NET-1
```

```
Vendor: NETAPP
```

```
Model: INF
```

```
.....
```

```
Path Failover Time (sec): 0
```

```
Extend All Path Failure Event (secs): 75
```

19.2 Consistent failover times on path failures

Prior to ONTAP 9.0, the path failure times were different based on where the failure occurred.

- Failure between the switch and the backend array.
- Failure between the switch and the FlexArray initiator port.

Starting ONTAP 9.0, if there is a path failure, I/O failover duration is consistent irrespective of where the failure has happened.

User tunable parameter '**path-failover-time**' will allow the user to tune the failover time for an array and can be used with applications which require predictable failover times. This can take a value in the range of 0-30 seconds. Value of "0" means failover immediately after receiving a non retrievable error on the path or a max value of "30" seconds indicates that the IO waits for 30 seconds before it is retried on another path. Default value for path failover time will be 0. Once this is enabled it is applicable to all the nodes that see the array within the cluster and across clusters in MetroCluster configuration. Value set for option 'path-failure-time' will be persistent across reboots.

Option **path-failover-time**

Default value **0**

Range **0-30**

Steps:

- List the arrays connected to the cluster using 'storage array show' command.

```
Cluster01::> storage array show
```

```
Prefix          Name  Vendor      Model Options
```

```
-----
```

```
NET-1          NETAPP_INF_1  NETAPP      INF
```

- Set the 'path-failover-time' to <delay> sec on the array using 'storage array modify' command:

```
storage array modify -name <Array-name> -path-failover-time <in secs>
```



```
Cluster01::> storage array modify -name NETAPP_INF_1 -path-failover-time 25
```

8. Ensure the value is set by executing 'storage array show' command.

```
Cluster01 ::> storage array show -name NETAPP_INF_1
```

```
Name: NETAPP_INF_1
Prefix: NET-1
Vendor: NETAPP
Model: INF
....
Path Failover Time (sec): 25
Extend All Path Failure Event (secs): 0
```

19.3 Automatic path quiesce if the number of intermittent failures on the path exceeds a threshold.

Intermittent failures (due to faulty components) in the fabric might result in disruptions in the path resulting in frequent IO errors. User tunable parameters 'path-failure-threshold' and 'wait-duration' can be used to monitor such paths and suspend I/Os to all the LUNs on the path by quiescing the path.

User can monitor a path by specifying the initiator port(HBA) and the Target port WWPN(on the array) with tunable options path-failure-threshold and wait-duration. Option 'path-failure-threshold' can take a value ranging 1-100 and is to be used in combination with option 'wait-duration' where the duration is specified in minutes which ranges from 1 to 1440. User can specify the maximum errors which can occur (path-failure-threshold) during a period (wait-duration) after which I/O on the path should be suspended. On exceeding the threshold provided, all the LUNs through the path specified will be quiesced. Values set for options 'path-failure-threshold' and 'wait-duration' will not be persistent across reboots.

For example, if a path is to be quiesced when it encounters more than 5 errors within duration of 60 minutes, follow below steps:

Option	path-failure-threshold
Default value	0
Range	1-100

Option	wait-duration
Default value	0
Range	1-1440

Steps:

9. List all the paths to the array using 'storage array config show' command.

```
Cluster01: ::> storage array config show -array-name NETAPP_INF_2
```

```
      LUN  LUN
Node   Group Count      Array Name   Array Target Port Initiator
-----
```

```

vgv8080f147a  9  50      NETAPP_INF_2  XYXYXYXYXYXYXYXY  11c
                XYXYXYXYXYXYXYXY      11d
vgv8080f147b  9  50      NETAPP_INF_2  XYXYXYXYXYXYXYXY  11c
                XYXYXYXYXYXYXYXY      11d

```

4 entries were displayed.

10. Monitor a path for intermittent failures using command 'storage path quiesce' and specify the limits using '-path-failure-threshold' and '-wait-duration'

storage path quiesce -node <node_name> -initiator <initiator_port> -target-wwpn <target_WWPN> -path-failure-threshold <count> -wait-duration <duration>

```

Cluster01::> storage path quiesce -node vgv8080f147a -initiator 11c -target-wwpn
XYXYXYXYXYXYXYXY -path-failure-threshold 5 -wait-duration 60

```

Resuming from the Path Quiesce state will be provided to the end user.

The quiesced path can be resumed with 'storage path resume' command:

storage path resume -node <node_name> <initiator_port> -target-wwpn <target_WWPN>

```

Cluster01::> storage path resume -node vgv8080f147a -initiator 11c -target-wwpn
XYXYXYXYXYXYXYXY

```

If the path does not encounter the said number of errors within the specified duration, counter resets and monitoring continues for the next cycle. To explicitly disable the monitoring, 'path-failure-threshold' and 'wait-duration' values should be set to '0'.

Steps:

1. Disable path monitoring using command 'storage path quiesce' and specify the limits using '-path-failure-threshold' and '-wait-duration'

storage path quiesce -node <node_name> -initiator <initiator_port> -target-wwpn <target_WWPN> -path-failure-threshold 0 -wait-duration 0

```

Cluster01::> storage path quiesce -node vgv8080f147a -initiator 11c -target-wwpn
XYXYXYXYXYXYXYXY -path-failure-threshold 0 -wait-duration 0

```

A path can be manually quiesced for maintenance purposes and resumed when the activity is complete.

Steps:

2. To quiesce a path manually, command 'storage path quiesce' can be used with no limits specified.

storage path quiesce -node <node_name> -initiator <initiator_port> -target-wwpn <target_WWPN>

```

Cluster01::> storage path quiesce -node vgv8080f147a -initiator 11c -target-wwpn
XYXYXYXYXYXYXYXY

```

3. The quiesced path can be resumed with 'storage path resume' command:

storage path resume -node <node_name> <initiator_port> -target-wwpn <target_WWPN>

```

Cluster01::> storage path resume -node vgv8080f147a -initiator 11c -target-wwpn
XYXYXYXYXYXYXYXY

```

In an event of all paths exhaustion (If there are no paths available to the LUN except the quiesced path), quiesced path will be automatically resumed and I/O will resume through this path.

- When all-path-fail-delay and path-failover-time are both set for a given array, path-failover-time takes precedence. However, in case of the last path going down resulting in an all-path fail condition, all-path-fail-delay will be considered.
- Path quiesce event won't happen for LUNs with only one path.

Version History

Version	Date	Document Version History
Version 1	October 2011	Last version of V-Series only TR
Version 2	April 2016	Updated TR-3461, added FlexArray
Version 3	June 2016	Updated for ONTAP 9
Version 4	October 2017	Minor updates for ONTAP 9.3
Version 5	April 2019	Minor updates for ONTAP 9.6

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2019 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.