



Technical Report

V-Series Best Practices Guide

Daniel Isaacs and Steve Lusnia, NetApp

April 2013 | TR-3461

Abstract

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

TABLE OF CONTENTS

CONTENTS

V-Series Best Practices Guide	1
Daniel Isaacs and Steve Lusnia, NetApp.....	1
April 2013 TR-3461.....	1
1 Introduction.....	5
2 Solution Overview	5
2.1 Data ONTAP 7-Mode Overview	5
2.2 Clustered Storage Overview	6
3 Installing a V-Series System.....	6
4 Back-End Setup	7
4.1 Aggregates.....	7
4.2 Summary of Best Practices.....	7
5 Native Disk Shelf Support.....	8
6 RAID Types.....	8
7 Disk Space Calculations	8
7.1 Right Sizing.....	8
7.2 Array LUN Sizing.....	9
7.3 Checksums and Capacity Overheads	9
8 LUN Paths.....	9
8.1 Paths from V-Series to Storage Array	9
8.2 Asymmetric Logical Unit Access (ALUA) (Clustered Data ONTAP Only).....	10
9 Zoning.....	10
9.1 Four-Array-Port LUN Groups	10
9.2 Multiple-Array LUN Groups	11
9.3 Multiple Arrays	12
9.4 Single Initiator to Multiple Arrays.....	13
9.5 Eight-Array-Port Configuration (Clustered Data ONTAP Only)	13
10 LUN Masking.....	15
10.1 HOST LUN IDs.....	15
11 Root Volume.....	15

11.1 Root Volume Location.....	15
11.2 Root Volume Sizes.....	15
12 Spare LUNS.....	15
13 Ports to Use on V-Series Systems.....	15
14 Test Procedures.....	16
14.1 Things That Should Be Tested.....	16
15 Operational Guidelines and Upgrades	17
15.1 Hardware Upgrades	17
15.2 Software Upgrades	17
15.3 Fabric or Array Maintenance.....	18
15.4 Expanding or Changing LUN Groups.....	18
Adding Paths.....	18
Adding LUNs	18
Removing LUNs	18
15.5 Powering Systems On and Off.....	18
16 V-Series MetroCluster	19
16.1 Supported MetroCluster Configuration Using Only Array LUNs	19
16.2 Using NetApp SAS Disk Shelves with V-Series MetroCluster.....	20
17 NetApp E5400.....	22
18 Array-Specific Notes	22
18.1 HDS USP, USP-VM, VSP Arrays.....	22
18.2 IBM DS4xxx	22

LIST OF TABLES

Table 1) Four-array-port LUN group zoning.	10
Table 2) Eight-array-port LUN group zoning.....	14
Table 3) Zoning forMetroCluster topology in Figure 9.	20

LIST OF FIGURES

Figure 1) 7-Mode overview.	5
Figure 2) Data ONTAP cluster with FAS and V-Series systems.....	6
Figure 3) Multiple storage array example.	7
Figure 4) Four-array-port LUN group.....	11
Figure 5) Multiple-array LUN groups.	11
Figure 6) Multiple arrays.....	12
Figure 7) Single initiator to multiple targets, different arrays.....	13
Figure 8) Eight-array-port configuration.....	14
Figure 9) V-Series MetroCluster supported configuration.....	19
Figure 10) V-Series MetroCluster with only NetApp SAS shelves and ATTO fibre bridge.....	21
Figure 11) V-Series MetroCluster with both NetApp shelves and third-party storage array LUNs.....	21

1 Introduction

The NetApp V-Series family is a scalable solution that virtualizes third-party storage arrays, allowing you to leverage the dynamic virtualization capabilities delivered in Data ONTAP® software across your existing Fibre Channel SAN infrastructure.

Like all NetApp storage appliances, the V-Series product line is based on the industry-tested Data ONTAP operating system. The 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 and 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 touches upon all the areas that need to be addressed when deploying V-Series systems with third-party storage arrays.

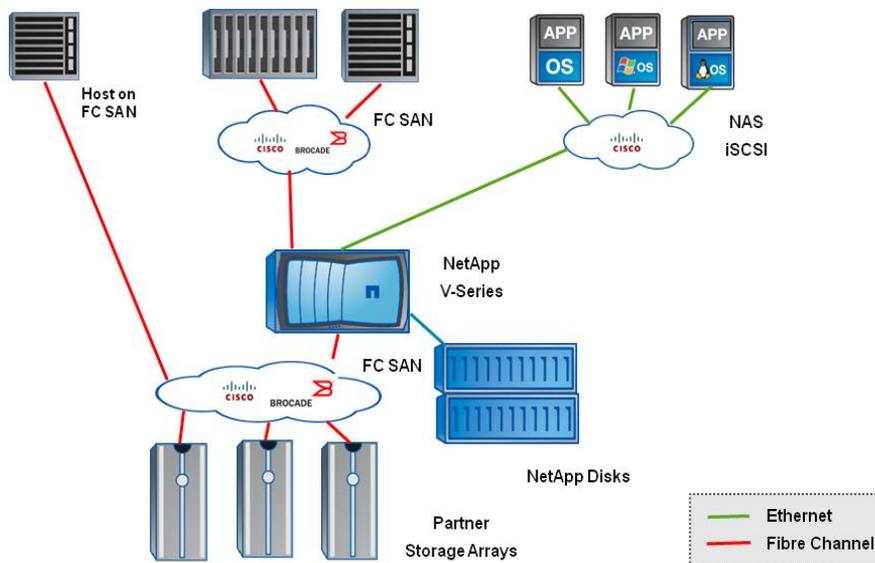
2 Solution Overview

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

2.1 Data ONTAP 7-Mode Overview

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, V-Series systems can use NetApp disk shelves. To see which storage arrays and switches are supported, consult the NetApp Interoperability Matrix Tool (IMT) on <http://support.netapp.com>.

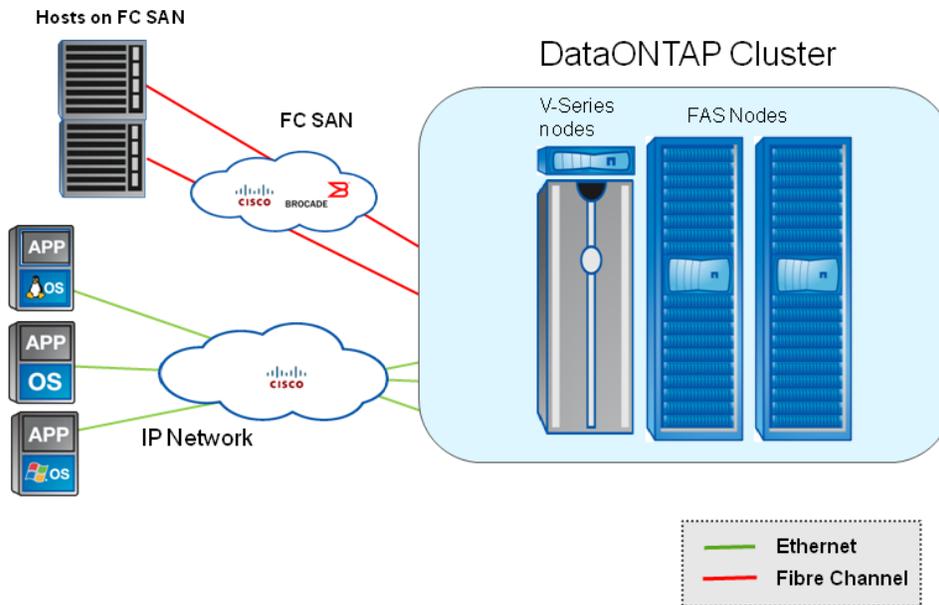
Figure 1) 7-Mode overview.



2.2 Clustered Storage Overview

V-Series supports the latest generation of clustered Data ONTAP, which provides scalable clustered storage through a single namespace. A Data ONTAP cluster may be built from FAS controllers, V-Series controllers, or a mix of FAS and V-Series systems.

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



3 Installing a V-Series System

A 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 V-Series system to the point at which Data ONTAP is installed. These tasks include cabling, array LUN formatting, LUN security, zoning, assigning array LUNs to the V-Series controller, creating aggregates, and loading Data ONTAP. Back-end implementation is largely a SAN deployment exercise. Once the V-Series system can access array LUNs, the rest of the setup is a quick process.

The front-end configuration for a V-Series system is exactly the same as for a FAS implementation. Once an aggregate is built and Data ONTAP is installed, creating FlexVol[®] volumes, NFS/CIFS shares; iSCSI, FCP, or FCoE LUNs; 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 design includes both front-end and back-end architectures. However, this guide is focused primarily on the back end, because the front-end setup of a V-Series system is identical to that of any FAS system.

If the V-Series system was ordered with NetApp disks, the factory installs the root volume, licenses, and Data ONTAP software, greatly simplifying the installation. Clustered Data ONTAP requires a NetApp disk shelf for the node root volumes, but its installation is otherwise identical to 7-Mode installation.

4 Back-End Setup

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

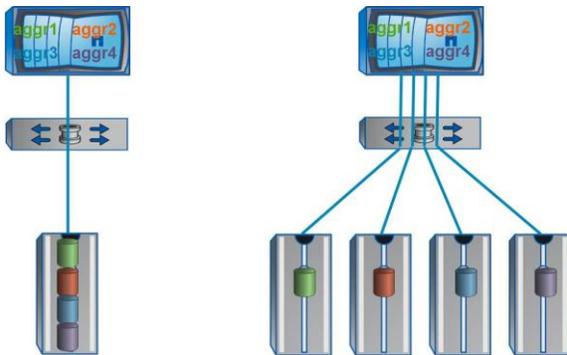
4.1 Aggregates

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 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 several arrays is possible provided the arrays are of the same type, the LUNs are built on array RAID groups of the same type, and guidelines for attaching to numerous arrays are followed (per the drawing in Figure 6).

Figure 3) Multiple storage array example.



4.2 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.
- Use dedicated array ports. Sharing array ports can adversely affect performance.
- NetApp recommends 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.
- Changes to a stable system introduce the risk of disruption. Changes to systems should be made during maintenance windows.
- Follow proper procedures when performing maintenance on the V-Series system, the back-end fabric, or the storage array.

5 Native Disk Shelf Support

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

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

6 RAID Types

Data ONTAP supports a variety of RAID types used by storage arrays such as RAID 1, 10 5, or 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. V-Series with a third-party array uses RAID 0 to stripe data across the array LUNs, which splits data evenly across two or more array LUNs. Performance is improved because 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 so that there is an adequate level of protection on the storage array and disk failure does not result in loss of data or loss of access to data.

When using NetApp disks, 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 8.1 and later or the most recent “Implementation Guide” for your vendor for releases prior to 8.1.

7 Disk Space Calculations

With any 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, the WAFL[®] file system, and Snapshot copies.

7.1 Right Sizing

The maximum array LUN size supported with V-Series (as of the time of this writing) is:

Data ONTAP 8.1.2: 3904GB (4TB)

Data ONTAP 8.0: 2916GB (3TB)

The latest updates can be found in the NetApp Hardware Universe (HWU) located on the NetApp Support site.

<https://hwu.netapp.com/Controller/Index>

7.2 Array LUN Sizing

On the storage array, you should build LUNs on numerous dedicated array RAID groups. Ideally, from a given RAID group, you should build two 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, and each LUN is counted toward the maximum device count limitations for a given V-Series platform, so a few dozen large LUNs are generally better than hundreds of small LUNs.

Each V-Series controller has a maximum capacity and a maximum number of devices (array LUNs or disks) it can support. Using small array LUNs may result in a V-Series system reaching its device count limit before reaching its capacity limit. Note that the maximum limits include both array LUNs and native disk drives; if you have 130 LUNs and 110 disks they are combined to count as 240 toward the max spindle count. For capacity, LUN counts, and other metrics, see the [Hardware Universe \(HWU\)](#).

7.3 Checksums and Capacity Overheads

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

Block Checksum (BCS)

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

Zoned Checksum (ZCS)

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

File System Overheads

When calculating the raw usable capacity of an array LUN as seen on a 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 the previous section for checksum overheads.

Additional user-configurable overheads include aggregate Snapshot reserves and FlexVol volume Snapshot reserves. These overheads are tunable based on customer SLAs around snapshot retention.

8 LUN Paths

8.1 Paths from V-Series to Storage Array

A 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 still spread the ports across two physically separate fabrics. For any given LUN, Data ONTAP assigns one path as the active and the other path as the passive, or failover, path for each V-Series controller. However, multiple LUNs will have different primary paths. NetApp recommends no more than a single "hop" between the 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. We also recommend 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 V-Series systems from accessing LUNs owned by non 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.

8.2 Asymmetric Logical Unit Access (ALUA) (Clustered Data ONTAP Only)

ALUA can improve the pathing and load balancing between the V-Series controller and some active-passive storage arrays. In Data ONTAP 8.2, ALUA is supported and enabled by default on the NetApp controllers. To use ALUA, it must also be enabled on the storage array. Contact your storage array provider to see if the array supports ALUA and how to enable this feature. We recommend enabling ALUA only on new configurations, not on existing configurations.

ALUA-supported arrays on V-Series (as of the time of this writing) are EMC: CX and VNX (8.2); and HP: EVA (8.1/8.2). See the IMT for the latest updates.

9 Zoning

Configuring zoning on a Fibre Channel switch enables you to restrict visibility and connectivity between devices connected to a Fibre Channel SAN.

- V-Series generally supports configurations in which a single initiator is connected to a single target port.
- Create separate zones for each initiator/target pair.
- We do not recommend shared array ports for performance reasons.
- Both software (WWPN) and hardware (port-based) zoning are supported.

Most of the configurations documented below are valid for both 7-Mode and clustered Data ONTAP. The configurations that are valid only for clustered Data ONTAP are explicitly stated.

9.1 Four-Array-Port LUN Groups

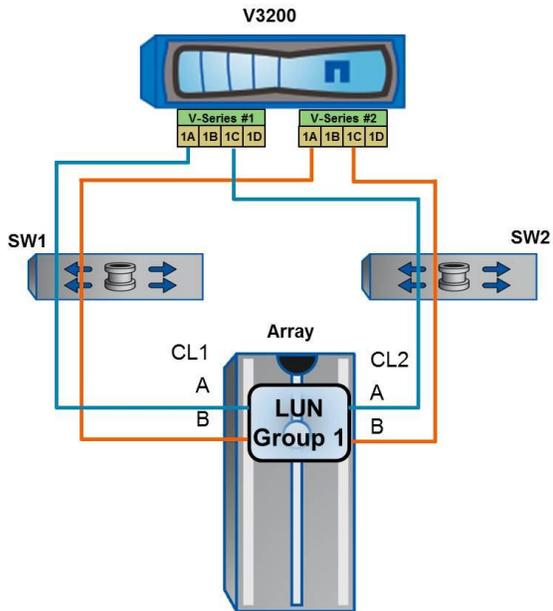
This is the recommended configuration. It maximizes both performance and resiliency.

Table 1 and Figure 4 show an example of how zoning should be set up with a 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. V-Series1 accesses the LUNs on array A ports and V-Series2 accesses the LUNs on array B ports.

Table 1) Four-array-port LUN group zoning.

SW1	Members	SW2	Members
Zone1	V-Series1:Port 1A, Controller CL1:Port A	Zone1	V-Series1:Port 1C, Controller CL2:Port A
Zone2	V-Series2:Port 1A, Controller CL1:Port B	Zone2	V-Series2:Port 1C, Controller CL2:Port B

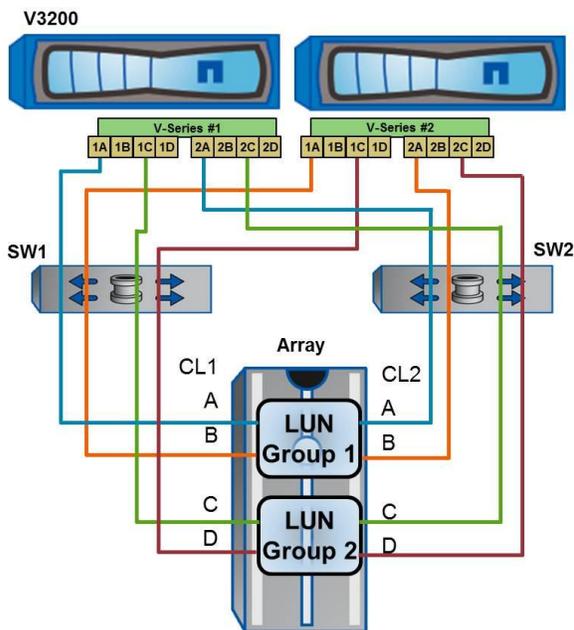
Figure 4) Four-array-port LUN group.



9.2 Multiple-Array LUN Groups

V-Series can attach to numerous LUN groups in the same storage array, allowing larger systems and greater system performance. 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.

Figure 5) Multiple-array LUN groups.



Each V-Series system has two paths to a given LUN group through separate fabrics and separate array controllers. V-Series1 port pair 1A/2A accesses LUN Group 1 and V-Series2 port pair 1A/2A also accesses LUN Group 1. V-Series ports 1C/2C are used for LUN Group 2.

The array is configured with numerous four-array-port LUN groups. All of the LUNs in this array LUN group are assigned to four-array ports, as shown. The array LUNs from Group 1 are assigned to array controller CL1 ports A and B and to array ports CL2 ports A and B.

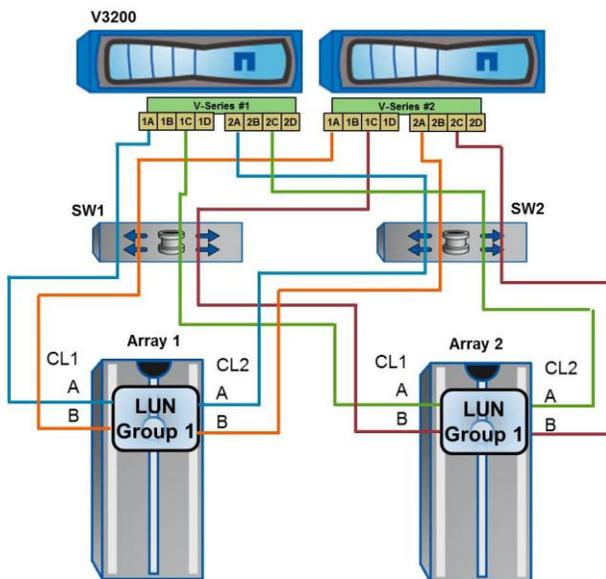
Zoning best practices are single-initiator-to-single-target zoning. V-Series1:1A is zoned to array CL1:A, and V-Series1:1C is zoned to array CL2:A. Similarly, V-Series2:1A is zoned to array CL1:B, and V-Series2:2A is zoned to array CL2:B. This pattern is repeated for additional LUN groups using additional initiators.

To maximize performance with multiple-array LUN groups, half the LUNs of each LUN group should be owned by each V-Series controller in the HA pair. Half the LUNs in LUN Group 1 above would belong to V-Series1, the other half to V-Series2. This pattern is repeated for LUN Group 2, and for any additional LUN groups.

9.3 Multiple Arrays

NetApp supports multiple arrays behind a single V-Series system or HA pair. This setup is very similar to the multiple-array LUN groups example in the previous section. The same dedicated-initiator-to-dedicated-target zoning rules apply, as well as the same two paths per V-Series system limit (four per HA pair) per LUN group. Figure 6 shows this model.

Figure 6) Multiple arrays.



Note that where 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.

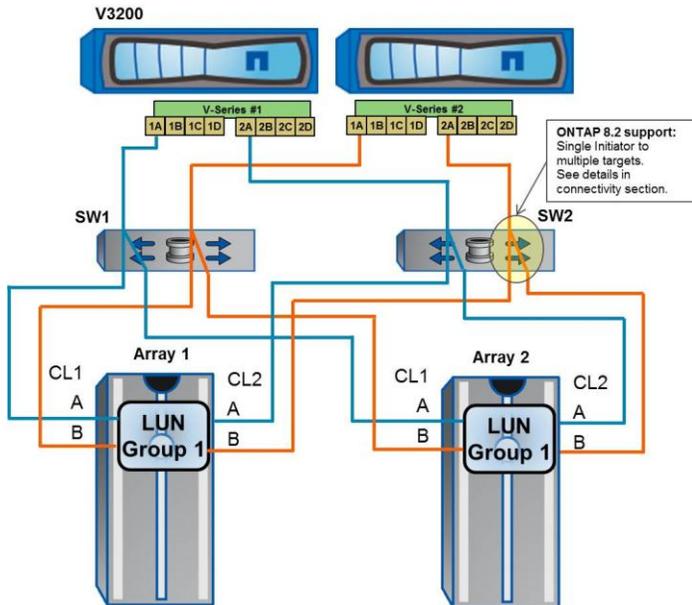
When connecting to multiple storage arrays you can spread an aggregate across the storage arrays only if the arrays are from the same family and use the same disk type. If this is not the case, performance is much less predictable.

9.4 Single Initiator to Multiple Arrays

In Data ONTAP 8.2, V-Series supports an expanded configuration in which an initiator port can attach to target ports from separate arrays (see Figure 7, below). Although our recommendation is to continue using single-initiator-to-single-target configurations described previously, we do also support this expanded configuration. Use it only for these conditions:

- The FC initiator port is attached to up to 4 target ports from different arrays of the same type.
- Use separate zones for each initiator-to-target port.
- The total LUN count for the initiator port does not exceed the maximum supported LUNs count per group. In Data ONTAP 8.2, the maximum supported LUNs count is increased from 100 to 254 per LUN group or per initiator.

Figure 7) Single initiator to multiple targets, different arrays.



9.5 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. Figure 8 shows a 2-node cluster. Since 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 clustered Data ONTAP as well. We simply have the ability to configure more paths. Note that the connections from V-Series ports 'B' and 'D' are intentionally crossed between the switch and the array ports 'C' and 'D.' This makes better use of both switch ports and array ports.

In Data ONTAP 8.1, the number of LUNs we support is 100 in a single LUN group. This configuration allows greater back-end bandwidth at the expense of initiator ports, but it does not enable you to see more storage.

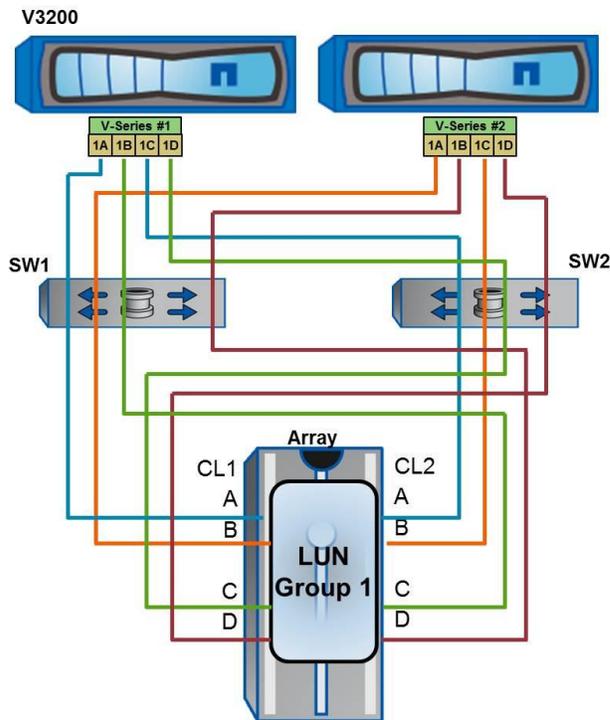
In Data ONTAP 8.2, the maximum number of LUNs supported in a LUN group increases from 100 to 254, making this a recommended configuration for V-Series. The eight-array port configuration with LUN counts up to 254 is the recommended configuration for large V-Series deployments in clustered Data ONTAP.

When attaching to an array using the eight-port LUN group as shown in Figure 8, pay close attention to how the array is attached to the switches and then to the V-Series. V-Series 1 uses switch 1 when attaching to the array CL1:A, B and switch 2 when attaching to array ports CL1:C, D. This configuration optimizes both switch and array port utilization.

Table 2) Eight-array-port LUN group zoning.

SW1	Members	SW2	Members
Zone1	V-Series1:Port 1A, Controller CL1:Port A	Zone1	V-Series1:Port 1C, Controller CL2:Port A
Zone2	V-Series2:Port 1A, Controller CL1:Port B	Zone2	V-Series2:Port 1C, Controller CL2:Port B
Zone3	V-Series1:Port 1B, Controller CL2:Port C	Zone3	V-Series1:Port 1D, Controller CL1:Port C
Zone4	V-Series2:Port 1B, Controller CL2:Port D	Zone4	V-Series2:Port 1D, Controller CL1:Port D

Figure 8) Eight-array-port configuration.



10 LUN Masking

LUN masking is similar to switch zoning but it is performed on the storage array. It is used to control which hosts can access which array LUNs. 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 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 V-Series system.

10.1 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. It is not supported, and it will not work.

11 Root Volume

11.1 Root Volume Location

When using NetApp storage, the preferred option for the root volume is to be on the NetApp disks. If you have a system that has only third-party storage arrays and no NetApp disks, you should install the root volume as a FlexVol volume on a single array LUN.

11.2 Root Volume Sizes

The minimum size of the root volume varies, depending on the 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 Hardware Universe.

12 Spare LUNS

If NetApp disks are not being used, NetApp recommends that you have a spare LUN for each V-Series system in an HA pair. For all 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 is not treated as a spare LUN in the sense that a spare disk with RAID-DP is treated. It is not used to replace a lost LUN since a spare disk in RAID-DP would be used to replace a failed disk. For spare LUN size details, see the Hardware Universe.

13 Ports to Use on V-Series Systems

V6000 systems have eight on-board Fibre Channel ports; V6200, V3000, and V3100 systems have four on-board Fibre Channel ports; and V3200 systems have two Fibre Channel ports. The on-board 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 in which initiator cards or target cards can be installed. V3200 systems may require an I/O expansion module (IOXM) 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 on-board V-Series FC ports 0a and 0b are on

one ASIC, ports 0c and 0d are on another ASIC, and so on. V3000 and V3100 systems use ports 0a and 0c as a port pair; similarly, V60x0 systems use ports 0a and 0e. 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, requires 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 on-board 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.

14 Test Procedures

Before putting a 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. This also provides an opportunity for documentation and customer education. This list covers V-Series-specific tests. Front-end testing should also be completed, but it is identical to FAS and therefore not covered here.

Use Data ONTAP commands to capture information about the environment (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 the status of each adapter
- `sysconfig -r` for aggregate and volume information

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. These new commands are to be run in the cluster instance, and they 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.

14.1 Things That Should Be Tested

You need to verify that you have the correct number of paths for each LUN. There are two commands that 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

- Disable an initiator
 - Assuming that 0a is used, `fcadmin offline 0a`, verify paths (`storage show disk -p all`). Manually execute CFO or SFO (if running clustered Data ONTAP) and giveback. Enable and repeat for each path.
- Switch failure
 - Disable all paths to a given switch for both 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.
- 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 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](#).

15 Operational Guidelines and Upgrades

15.1 Hardware Upgrades

Before starting any hardware upgrade, make sure that you have the AutoSupport™ (ASUP™) 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 an ASUP + `disk show -v` and `storage show disk -p`)
- For clustered Data ONTAP, the commands above 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 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 V-Series system. A new 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 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 this is not possible, you must redo zoning (in case of soft zoning) and host groups, because the new HBA cards will have different WWNs.

15.2 Software Upgrades

When updating Data ONTAP you should 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 would recommend sending just an AutoSupport message, 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.

15.3 Fabric or Array Maintenance

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

Fabric maintenance might include firmware upgrades or switch replacements. V-Series systems must have redundant fabrics, and each fabric must be upgraded separately. You should not begin maintenance on the second fabric until you've verified that the V-Series systems can see all of their LUNs via 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, determine that at least one of these paths is available at all times. With clustered Data ONTAP, you can simply run the `storage error show` command to confirm redundant paths.

As with any maintenance operation, 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 performing the maintenance.

Array maintenance is generally in the form of firmware upgrades. Most arrays support nondisruptive firmware upgrades, but check the V-Series Support Matrix to verify both that your new array firmware or microcode is supported and that array NDU is supported.

15.4 Expanding or Changing LUN Groups

Adding Paths

If you are upgrading to the four-array-port model, do not exceed the limit of two paths per V-Series system or four paths for a clustered Data ONTAP system.

It is also important that you have the LUNs shared via the target ports before the initiator is connected.

Adding LUNs

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

- In Data ONTAP 8.2, the maximum number of LUNs supported in a LUN group is 254. Previous versions of Data ONTAP supported 100 LUNs per group.
- When adding 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 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](#).

15.5 Powering Systems On and Off

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

Power on:

- Storage array and any SAN switches
- Boot the V-Series system

Power off:

- Halt the V-Series system
- Storage array and any SAN switches

16 V-Series MetroCluster

MetroCluster™ software 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 is able to transparently recover from failures, including complete loss of a data center, enabling business-critical applications to continue uninterrupted.

Installing a V-Series MetroCluster system is like installing two separate V-Series systems at different locations, then enabling HA and enabling SyncMirror® software between the two sites. FAS rules must be followed for SyncMirror. MetroCluster systems can be configured as fabric attached with distance support up to 200 km, or as stretch with support up to 500 m. For detailed guidelines for systems running Data ONTAP 8.1 and later, see the “Data ONTAP 7-Mode High-Availability Guide.” For detailed guidelines for systems running Data ONTAP releases prior to 8.1, see the “V-Series MetroCluster Configuration Guide.”

16.1 Supported MetroCluster Configuration Using Only Array LUNs

Figure 9 shows the supported V-Series MetroCluster topology with array LUNs. V-Series (V-Series1) 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. V-Series11 accesses these two sets of LUNs through 2A/3A for the local pool and 2B/3B 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 it is also synchronously mirrored between Site A and Site B.

Figure 9) V-Series MetroCluster supported configuration.

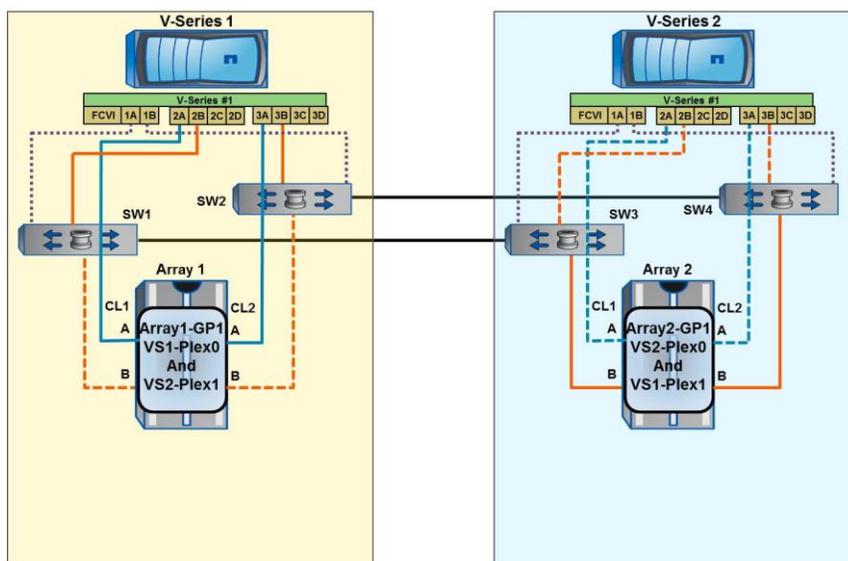


Figure 9 shows a V6200 MetroCluster, each controller with a local and a remote storage pool. Larger systems can be built by using additional storage pools, new sets of V-Series initiators ports, array target ports, and switch ports.

Table 3) Zoning for MetroCluster topology in Figure 9.

Name	SW1-SW3 Fabric	Name	SW2-SW4 Fabric
Zone 1:	VS1:2A, Array1:CL1:A	Zone 1:	VS1:3A, Array1:CL2:A
Zone 2:	VS1:2B, Array2:CL1:B	Zone 2:	VS1:3B, Array2:CL2:B
Zone 3:	VS2:2A, Array2:CL1:A	Zone 3:	VS2:3A, Array2:CL2:A
Zone 4:	VS2:2B, Array1:CL1:B	Zone 4:	VS2:3B, 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 determine 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.

V-SERIES METROCLUSTER SUPPORTS:

- Data ONTAP running in 7-Mode only
- Symmetric configurations between sites (array family, switch types, firmware versions)
- Brocade and Cisco® switches (see the NetApp Interoperability Matrix Tool for specific models and firmware)
- 254 LUNs in each array LUN group (in Data ONTAP 8.1 and earlier the number is 100)

SAN REQUIREMENTS INCLUDE:

- Dedicated fabric
- 200 km maximum supported distance for a V-Series MetroCluster (with Data ONTAP 8.2)
- Using single-initiator zoning, as shown in Table 3
- Having available long-distance ISL shares FCVI (cluster interconnect) and data traffic (sufficient bandwidth 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.” For detailed guidelines for systems running Data ONTAP releases prior to 8.1, see the “V-Series MetroCluster Configuration Guide.”

16.2 Using NetApp SAS Disk Shelves with V-Series MetroCluster

In Data ONTAP 8.2, V-Series MetroCluster can use NetApp SAS shelves with the ATTO fibre bridge by themselves or in addition to LUNs from supported storage arrays. Using V-Series MetroCluster with NetApp SAS shelves is identical to using a FAS controller with the same NetApp SAS shelves (see TR-3548 <http://www.netapp.com/us/media/tr-3548.pdf>).

MetroCluster uses SyncMirror for local and remote replication of data. When creating the NetApp aggregates, we use SyncMirror to mirror local LUNs to remote LUNs. Separately, we use SyncMirror to mirror local NetApp disks to remote NetApp disks.

The ISLs between the two sites are shared for mirroring the local and remote disks to disks and LUNs to LUNs. The cabling and zoning for the storage arrays are outlined in the table above and in (xx list the implementation document). Details on cabling and zoning for the NetApp SAS shelves are beyond the scope of this document and are detailed in TR-3548. FCVI connections between local and remote sites

are shown using the same ISL as data SyncMirror. TR-3548 shows alternative supported configurations in which separate ISLs are used for FCVI traffic.

Figure 10 is a diagram of a V-Series MetroCluster system with only NetApp SAS shelves with the ATTO fibre bridge. Figure 11 is a combined diagram with both NetApp disks and array LUNs.

Figure 10) V-Series MetroCluster with only NetApp SAS shelves and ATTO fibre bridge.

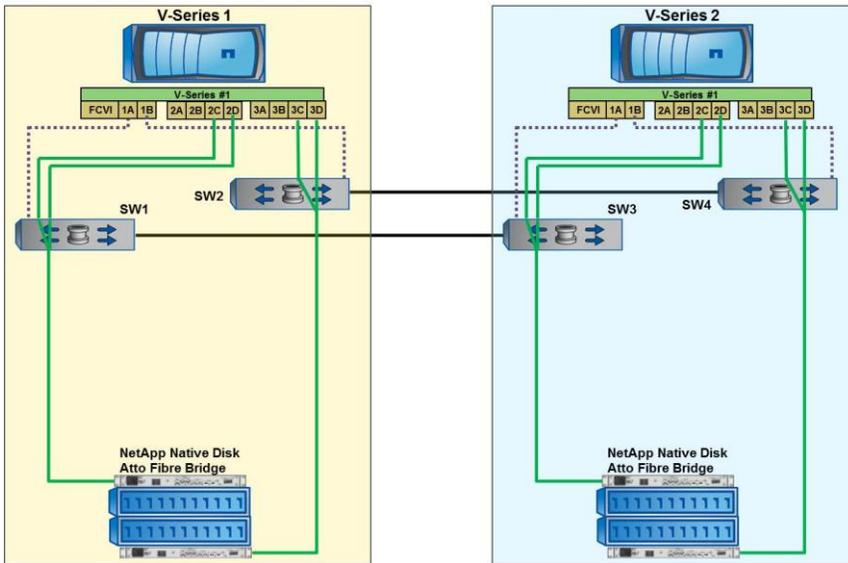
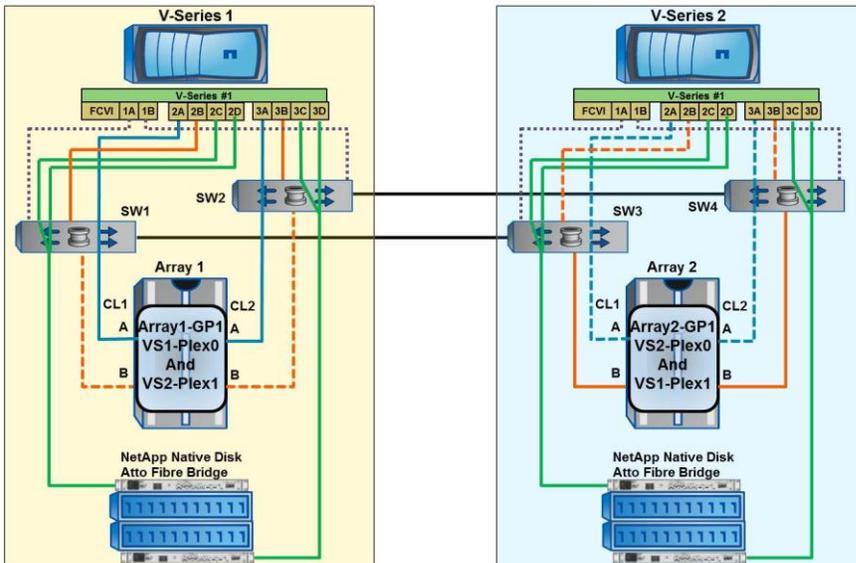


Figure 11) V-Series MetroCluster with both NetApp shelves and third-party storage array LUNs.



17 NetApp E5400

V-Series controllers support NetApp branded E-5400 arrays with the release of Data ONTAP 8.1.2P2 in 7-Mode only. Our standard configuration guidelines apply equally to the E5400 (choosing V-Series controllers, LUN sizing, zoning, and so on).

To find out 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 E5400, consult the [E5400 series product documentation](#) on the NetApp Support site.

18 Array-Specific Notes

V-Series systems do not support any advanced array 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. Use instead the NetApp software features in Data ONTAP (Snapshot, SnapMirror[®] software, SyncMirror, SnapMover[®] software, FlexVol, FlexClone[®], and so on). For information about how to properly configure array options, see the “V-Series Implementation Guide” for your vendor for releases prior to 8.1 and the “Implementation Guide for Third-Party Storage” for 8.1 and later. Always read and follow the notes for each array in the V-Series Support Matrix. The rest of this section summarizes some highlights.

18.1 HDS USP, USP-VM, VSP Arrays

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

- V-Series root volume 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 V-Series settings are the default settings.

18.2 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 V-Series system. Similarly, any access LUN that is created in a host group should not be mapped to the V-Series system.

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.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®



www.netapp.com

© 2013 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, ASUP, AutoSupport, Data ONTAP, FlexClone, FlexVol, MetroCluster, RAID-DP, SnapMirror, SnapMover, Snapshot, SyncMirror, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Cisco is a registered trademark of Cisco Systems, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-3461-0413