



V-Series Best Practices Guide

Steve Lusnia, NetApp

Daniel Isaacs, NetApp

November 2010 | TR-3461

TABLE OF CONTENTS

1	INTRODUCTION	4
2	SOLUTION OVERVIEW	4
3	INSTALLING A V-SERIES SYSTEM	5
4	BACK-END SETUP	5
5	NATIVE DISK SHELF SUPPORT	6
6	RAID TYPES	6
7	DISK SPACE CALCULATIONS	6
7.1	RIGHT SIZING	7
7.2	LUN SIZING	7
7.3	CHECKSUMS	7
7.4	CAPACITY OVERHEADS	8
8	LUN PATHING	8
9	ZONING 8	
9.1	FOUR-ARRAY-PORT LUN GROUPS	8
9.2	TWO-ARRAY-PORT CONFIGURATION	9
9.3	MULTIPLE-ARRAY LUN GROUPS	10
9.4	MULTIPLE ARRAYS	11
10	LUN MASKING	12
10.1	HOST LUN IDS	13
11	ROOT VOLUME	13
11.1	ROOT VOLUME LOCATION	13
11.2	ROOT VOLUME SIZES	13
12	SPARE LUNS	13
13	PORTS TO USE ON V-SERIES SYSTEMS	13
14	DATA ONTAP 8.0 CLUSTER-MODE	14
15	TEST PROCEDURES	14
15.1	THINGS THAT SHOULD BE TESTED	14
16	OPERATIONAL GUIDELINES AND UPGRADES	15
16.1	HARDWARE UPGRADES	15
16.2	SOFTWARE UPGRADES	15
16.3	FABRIC OR ARRAY MAINTENANCE	15
16.4	EXPANDING OR CHANGING LUN GROUPS	16
16.5	POWERING SYSTEMS ON AND OFF	16
17	V-SERIES METROCLUSTER	16
17.1	INSTALLING A V-SERIES METROCLUSTER SYSTEM	17
17.2	SUPPORTED METROCLUSTER CONFIGURATION	17

18 ARRAY-SPECIFIC NOTES 18

- 18.1 HDS TAGMASTORE18**
- 18.2 IBM DS4000.....18**
- 18.3 EMC DMX19**

1 INTRODUCTION

The NetApp® V-Series family is a network-based, 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 decreasing management complexity, improving asset utilization, and streamlining operations to increase business agility and reduce total cost of ownership.

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. A V-Series system can attach to all supported arrays through a fabric. For many arrays, a V-Series system can also attach directly to the array. To see which storage arrays and switches are supported, consult the V-Series support matrix: http://now.netapp.com/NOW/knowledge/docs/V-Series/supportmatrix/V-Series_SupportMatrix.pdf.

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.

Figure 1) Overview of a V-Series system using third-party storage.



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 where 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. 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.

The overall V-Series system design includes both front-end and back-end architecture. However, this guide focuses mainly on the back end, because the front-end setup of a V-Series system is similar 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.

4 BACK-END SETUP

On the back end, the V-Series controller can attach to a single storage array or to multiple storage arrays. Multiple storage arrays from the same vendor as well as multiple storage arrays from different vendors are supported as long as those arrays are on the V-Series Support Matrix. For details, see section 9.4.

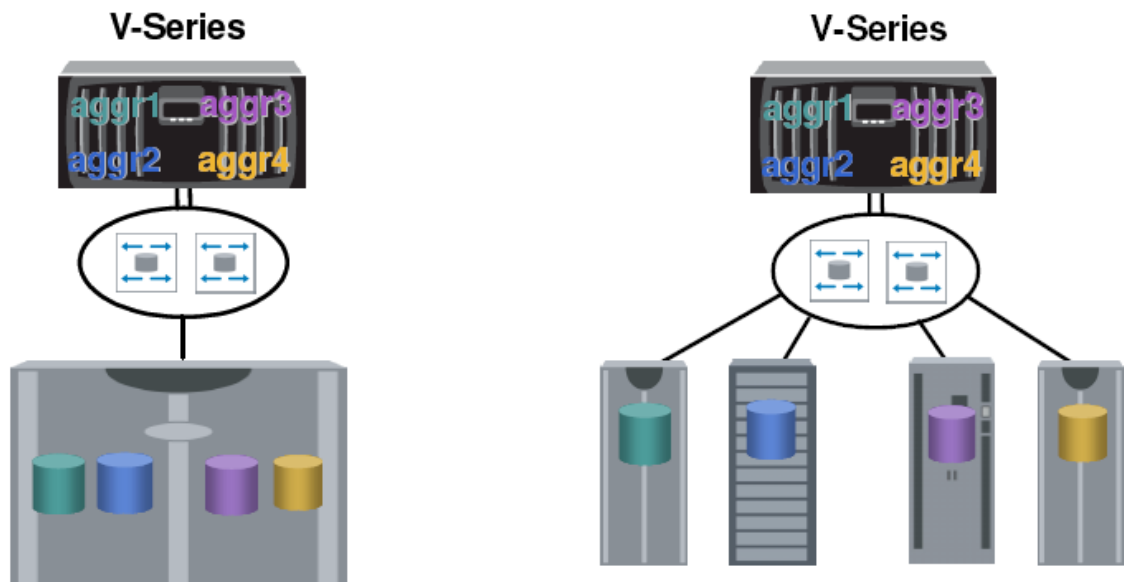


Figure 2) Multiple storage array example.

NetApp recommends that any single aggregate is made from one array, from LUNs supported by the same types of disks. For instance, if the array has both SATA and FC disks, a single Data ONTAP aggregate should be composed of a single type of disk: SATA only or FC only. Data ONTAP does not prevent you from spanning any aggregate across arrays or different types of LUNs or drives, so care must be taken when provisioning storage.

V-Series aggregates can span multiple arrays provided that the underlying storage (array family, disk technologies, and so on) is the same. NetApp recommends that aggregates are isolated to a single array. The best practice is not to span the aggregate across multiple arrays. If the arrays are different types, use different disk types, and so on; spanning aggregates across them is not supported.

The following sections expand on this simple set of best practices:

- Use a NetApp 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.
- Assign a spare LUN. If using a NetApp shelf, the spare disk meets this requirement.
- Don't make changes to configurations once they are built. Changes to a stable system introduce the risk of disruption.
- Follow proper procedures when performing maintenance on the V-Series system, the back-end fabric, or the storage array.

5 NATIVE DISK SHELF SUPPORT

The V-Series product line has always virtualized storage from third-party storage array vendors. Starting in Data ONTAP 7.3, V-Series systems support native disk shelves. Data ONTAP provides a unified storage software platform that simplifies managing both native disk shelves and LUNs on storage arrays. You can add storage when and where you need it, without disruption. These disk shelves can be managed in accordance with the NetApp "Data ONTAP Storage Management Guide." Consult this guide, along with the "V-Series Implementation Guide for Native Disk Shelves," for installation and configuration of native disk shelves with V-Series systems. These guides can be found on the NOW™ (NetApp on the Web) support site.

Disk shelves can be installed on new and existing V-Series systems. Review the [V-Series Support Matrix](#) for the latest support requirements. You can add disk shelves to 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.

Beginning with Data ONTAP 7.3.1.1, SnapLock® Compliance is also supported on V-Series, provided that the compliant storage is provided by native disk shelves.

NetApp recommends that the root volume, when using native disk shelves, reside on the native disks. This is a requirement for Data ONTAP 8.0 Cluster-Mode.

6 RAID TYPES

Data ONTAP supports a variety of RAID types used by storage arrays, but restricts storage arrays from using RAID 0 for the LUNs that they make available to Data ONTAP. The reason is that Data ONTAP uses RAID 0 to stripe across the array LUNs, which splits data evenly across two or more array LUNs. Performance is maximized because more disk spindles are used.

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.

For NetApp disk shelves, RAID-DP® is the default RAID level, but RAID 4 is also supported.

Always consult the latest array "Implementation Guide" for the most recent information specific to your array.

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, WAFL®, and Snapshot copies.

7.1 RIGHT SIZING

Each vendor has its own calculation for capacity. These sizes can vary, due to the differences in gigabyte definitions between the various storage vendors. Always consult the latest array “Implementation Guide” for the most recent information specific to your array.

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, and 2TB LUN Sizes
NetApp	1,000 x 1,024 x 1,024 bytes	
IBM DS4000, HDS, Fujitsu, EMC, 3PAR, HP EVA (most arrays use 1,024 math)	1,024 x 1,024 x 1,024 bytes	488GB, 732GB, 976GB, 1952GB
IBM ESS800 (Shark)	1,000 x 1,000 x 1,000 bytes	512GB, 768GB, 1024GB, 2048GB
IBM DS8000 (depending on setup)	1,000 x 1,000 x 1,000 bytes or 1,024 x 1,024 x 1,024 bytes	512GB, 768GB, 1024GB, 2048GB; or 488GB, 732GB, 976GB, 1952GB

7.2 LUN SIZING

On the storage array, you should build LUNs that take advantage of as many spindles as possible.

Ideally, from a given RAID group, you should build one or more large LUNs. However, this might not always be possible. It is better to have fewer larger LUNs than many smaller LUNs, but you should have at least 8 LUNs per port to make use of all the available disk queues. Data ONTAP treats each LUN as a separate disk, and each LUN incurs some WAFL overhead in addition to counting toward the maximum spindle count limitations for a given V-Series platform.

Each V-Series controller has a maximum capacity and a maximum number of LUNs it can support. Using small LUNs may result in a V-Series system reaching its LUN count limit before reaching its capacity limit. Note that the maximum limits include both LUNs and native disk drives; if you have 10 LUNs and 14 disks they are combined to count as 24 toward the max spindle count. For capacity, LUN counts, and other metrics, see the V-Series technical specifications at www.netapp.com/us/products/storage-systems/.

7.3 CHECKSUMS

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 authentic checksum can verify that the message was not corrupted by doing the same operation on the data and checking the sum.

BLOCK CHECKSUM (BCS)

Block checksum is the default checksum type. With block checksums, a checksum entry is appended to each WAFL block. In the case of V-Series systems, the entry uses up a sector—the sector following the WAFL block—for every WAFL block. Data ONTAP reserves 12.5% of the space in the LUN for the checksum. NetApp requires that the array LUNs supporting the root volume use BCS.

7.4 CAPACITY OVERHEADS

When calculating the capacity of an array LUN as seen on a V-Series system, the factors shown in Table 2 decrease the usable capacity. Except for checksum, these factors are identical to space calculations for FAS systems with FC drives. And they are identical in every respect to the overhead seen in FAS with SATA drives.

Table 2) Capacity overheads.

Fixed Overheads	Configurable Overheads
11% for WAFL reserve and core	5% for aggregate-level Snapshot reserve*
12.5% for block checksum	20% for volume-level Snapshot reserve*

* These are the default values; they are optional and are customer configurable.

8 LUN PATHING

A V-Series system supports only two paths to a given LUN. In a V-Series HA pair, the partner and local V-Series systems each have two paths to the array LUN. These two paths should be on separate and redundant fabrics. For any given LUN, Data ONTAP assigns one path as the active and the other path as the passive, or failover, path. NetApp recommends no more than a single “hop” between the V-Series system and the array. A hop equals one ISL connection between edge-to-core, core-to-core, or core-to-edge switches.

NetApp recommends using dedicated array ports with one-to-one initiator-to-target port mapping. Use 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 of a V-Series cluster should see all of the array LUNs. Use zoning and host groups (both are recommended) to control LUN access to the V-Series system.

9 ZONING

Configuring zoning on a Fibre Channel switch enables you to restrict visibility and connectivity between devices connected to a Fibre Channel SAN. Implementations can use either hard or soft zoning.

A V-Series best practice is to use single-initiator-to-single-target zoning. Create separate zones for each initiator and target. Shared array ports are supported but not recommended. Both software (WWPN) and hardware (port-based) zoning are supported.

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

9.1 FOUR-ARRAY-PORT LUN GROUPS

This is the recommended configuration. It offers greater performance and reliability than the two-array-port model.

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

SW1		SW2	
Zone1	VS1:Port 0a, Controller CL0:Port A	Zone1	VS1:Port 0c, Controller CL1:Port A
Zone2	VS2:Port 0a, Controller CL0:Port B	Zone2	VS2:Port 0c, Controller CL1:Port B




Figure 3) Four-array-port LUN groups.

9.2 TWO-ARRAY-PORT CONFIGURATION

Table 4 shows an example of how zoning should be set up with a V-Series cluster connected to a storage array using a basic two-array-port configuration. You would typically use this model only on smaller arrays or for demonstrations. Figure 4 shows how the cabling looks.

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

SW1		SW2	
Zone1	VS1:Port 1a, Controller C1:Port A	Zone1	VS1:Port 1c, Controller C2:Port A
Zone2	VS2:Port 1a, Controller C1:Port A	Zone2	VS2:Port 1c, Controller C2:Port A

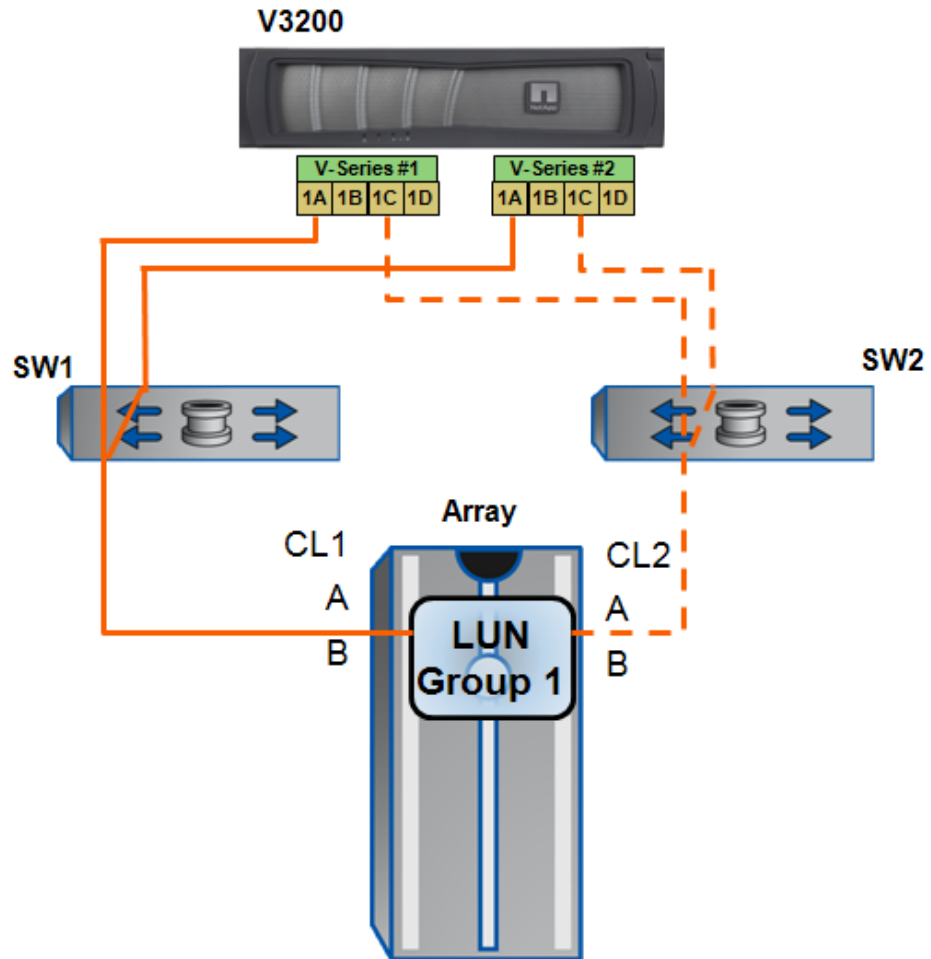


Figure 4) Two-array-port LUN group.

9.3 MULTIPLE-ARRAY LUN GROUPS

Data ONTAP 7.3 introduced broader support for attaching to multiple LUN groups in the same storage array, allowing greater system performance. Prior to 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.

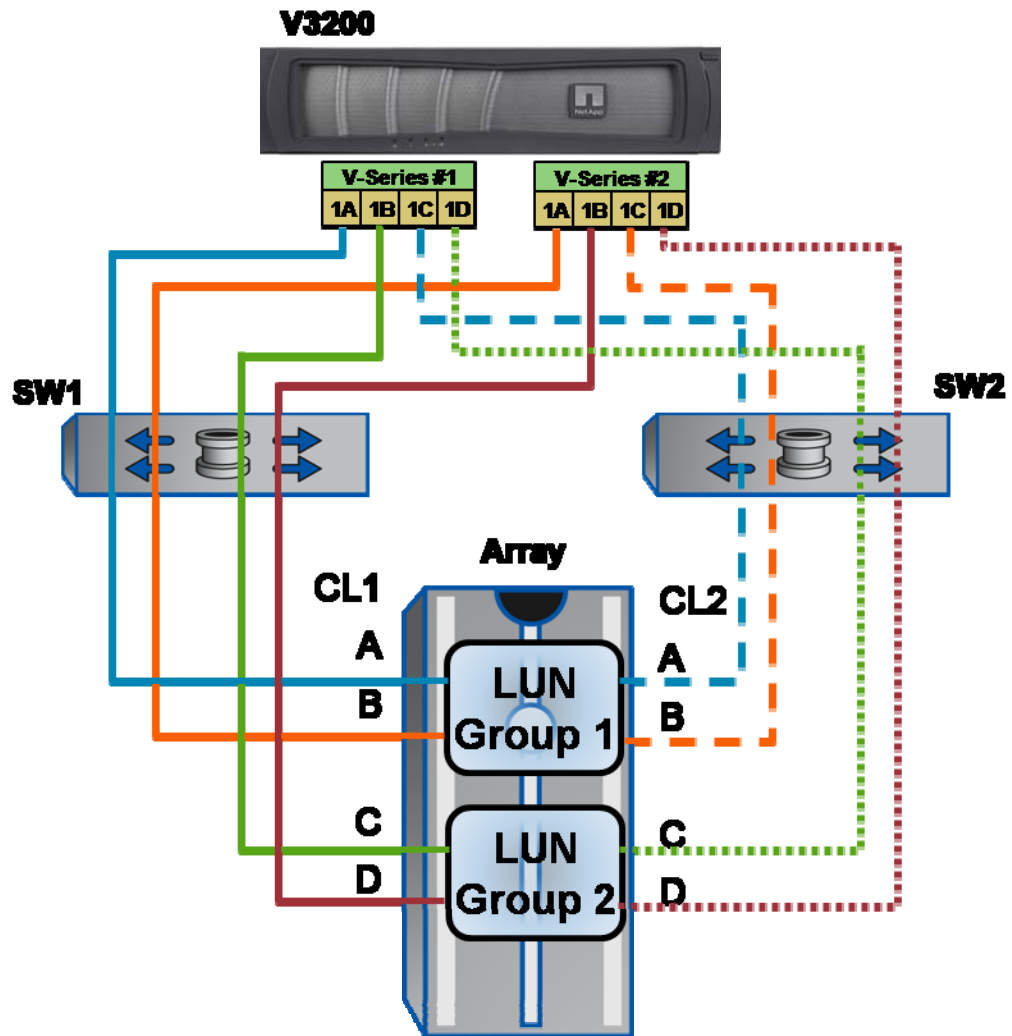


Figure 5) Multiple-array LUN groups.

Each V-Series system has two paths to a given LUN group through separate fabrics and separate array controllers. VS-1 port pair 1a/1c access LUN Group1 and VS-2 port pair 1a/1c also access LUN group1. V-Series ports 1b/1d 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 all four-array ports 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 VS1:0c is zoned to array CL1:A. Similarly, VS2:1a is zoned to array CL1:B, and VS2:1c is zoned to array CL2:B.

This layout is supported for some arrays with V-Series with Data ONTAP 7.3 or later. To determine whether your arrays support the four-array-port model, see the V-Series Support Matrix.

9.4 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 section 9.3. The same dedicated-initiator-to-dedicated-target zoning rules apply, as well as the same two paths per V-Series system (four per HA pair) limit per LUN group. Figure 6 shows this model.

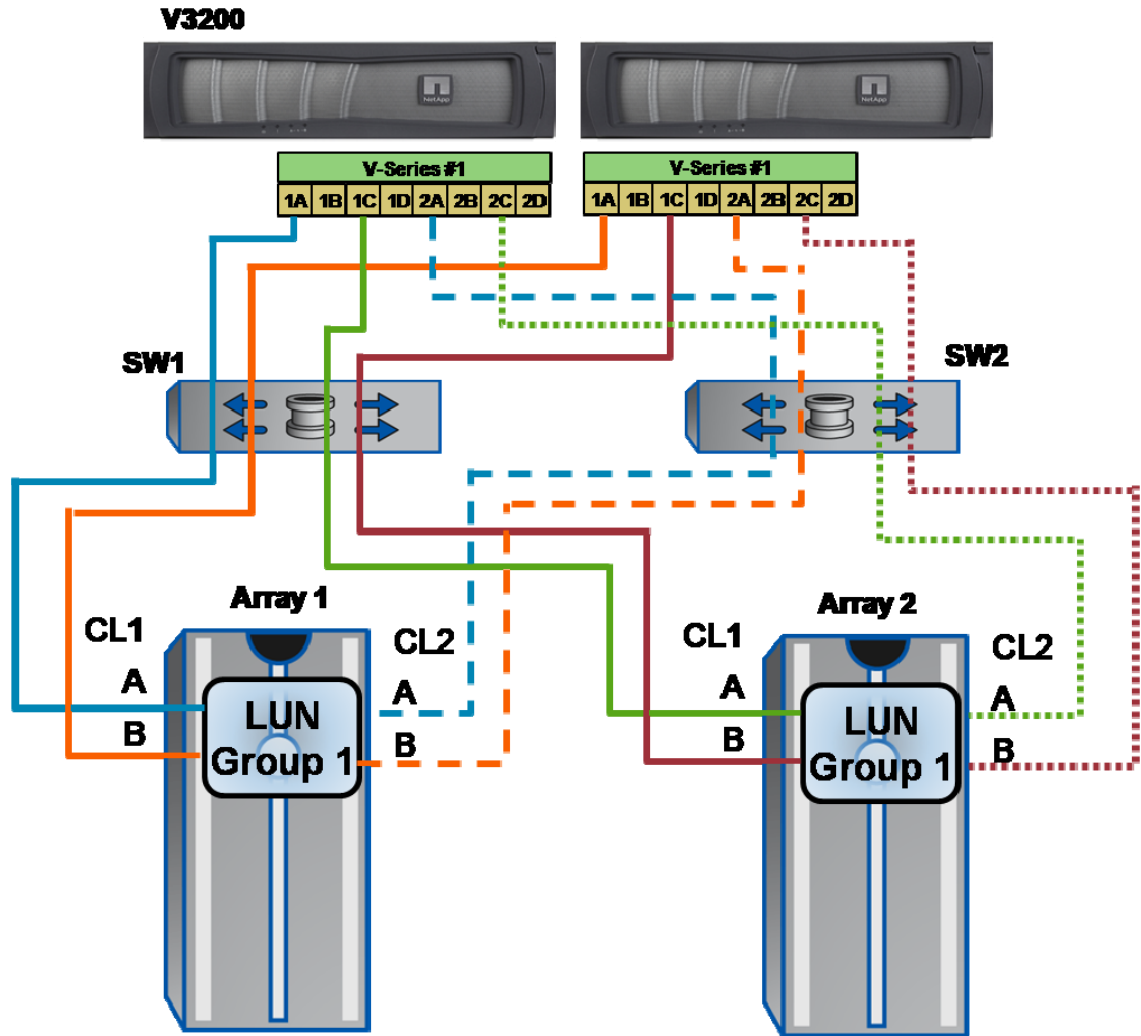


Figure 6) Multiple arrays.

Note that where ports were connected to LUN Group 2 in the previous example, they are connected to Array 2 here.

When connecting to multiple arrays, you should not spread an aggregate across the arrays unless they are of the same family and use the same disk types. Failure to observe this guideline can lead to uneven performance.

10 LUN MASKING

LUN masking is similar to switch zoning but 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 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 and redundancy 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. With Data ONTAP 8.0 Cluster-Mode, NetApp disks must be used for the root volume.

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 V-Series Support Matrix.

12 SPARE LUNS

If NetApp disks are not being used, NetApp recommends that you have a spare LUN for each V-Series system in a cluster. 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 rebuild a lost LUN as a spare disk in RAID-DP would be used to rebuild a failed disk.

The size of the spare LUN depends on the memory and NVRAM size of a given V-Series system and a small buffer. For details, see the V-Series Support Matrix.

13 PORTS TO USE ON V-SERIES SYSTEMS

V60x0 systems have eight on-board Fibre Channel ports, V3000 and V3100 systems have four on-board Fibre Channel ports, and V3200 systems have two Fibre Channel ports. The V-Series internal 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. The V3200 systems may require an I/O expansion module to support larger configurations.

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; 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 V3270/V3240 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 attach 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 “Fibre Channel and iSCSI Configuration Guide” on NOW for information about recommended front-end SAN topologies.

14 DATA ONTAP 8.0 CLUSTER-MODE

For the most part, building a V-Series system with Data ONTAP 8 Cluster-Mode is identical to building a system running previous versions of Data ONTAP. The same best practices (LUN sizes, LUN counts, zoning, cabling, and so on) are used. However, there are a few limitations that are specific to Cluster-Mode systems.

Initial array support is limited. For a list of supported arrays, platforms, and switches, see the V-Series Support Matrix. More arrays will be added, but until then you should file a PVR requesting support for other types of arrays.

A NetApp disk shelf is required for Cluster-Mode. The OS must reside on NetApp disks.

Data ONTAP 8.0 introduces a new volume type called *striped volumes*. These volumes span multiple nodes in the cluster. When V-Series nodes are involved, the member volumes of the striped volume must be built from array LUNs from the same family of arrays. For instance, a striped volume on a V-Series system using array LUNs must use only LUNs from a DMX or only from a CX (pending support of that platform).

15 TEST PROCEDURES

Before putting a V-Series system into production, it is important to make sure that everything works. The idea is to identify problems now rather than in production. This also provides an opportunity for documentation and customer education.

Use Data ONTAP commands to capture information about the environment:

- `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

15.1 THINGS THAT SHOULD BE TESTED

You need to make sure that you have the correct number of paths for each LUN. There are two commands that provide the information needed to verify:

- 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)`. Repeat for each initiator.
 - Disable all paths to a given switch for both V-Series controllers to simulate a switch failure.
- Cluster failover
 - For each of the path failure tests above, verify CFO and giveback.

See the chapters “Completing and Testing V-Series Specific Setup” and “Verifying an Active-Active Configuration Setup” in the “V-Series Software Setup, Installation, and Management Guide” for the complete steps to run for these tests.

16 OPERATIONAL GUIDELINES AND UPGRADES

16.1 HARDWARE UPGRADES

Before starting any hardware upgrade, make sure that you have AutoSupport™ 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`)

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 WWNNs.

16.2 SOFTWARE UPGRADES

When updating Data ONTAP, you should first verify, by checking the V-Series Support Matrix, that the new version of Data ONTAP is supported with the firmware and microcode levels on the array and fabric. The 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 16. Environment variables are not required for software upgrades.

16.3 FABRIC OR ARRAY MAINTENANCE

From time to time, the back-end components may need to have firmware or microcode upgrades. V-Series systems are very 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 V-Series Support Matrix.

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. Because each LUN name typically includes the switch name, you should make sure that at least one of these paths is available at all times.

As with any 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, but you should check the V-Series Support Matrix to verify both that your new array firmware or microcode is supported and that array NDU is supported.

16.4 EXPANDING OR CHANGING LUN GROUPS

16.4.1 Adding Paths

V-Series systems support two models for attaching to an array LUN group: the two-array-port model and the four-array-port model. These are covered above in [section 9](#). If you are upgrading from the two-array-port model to the four-array-port model, make sure that you do not exceed the limit of two paths per V-Series system.

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

16.4.2 Adding LUNs

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

- The maximum number of LUNs supported in a LUN group is 100.
- LUNs should not use 0 as a host ID.
- 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 it is shared on.

16.4.3 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](#).

16.5 POWERING SYSTEMS ON AND OFF

You should 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

17 V-SERIES METROCLUSTER

NetApp MetroCluster™ is a solution that combines NetApp clustering with synchronous mirroring to deliver continuous availability. MetroCluster works seamlessly with your host and storage environment to provide continuous data availability between two sites while eliminating the need to create and maintain complicated failover scripts. You will be able to serve data even if there is a complete site failure.

MetroCluster is a unique solution that combines NetApp clustering 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, so business-critical applications continue uninterrupted.

17.1 INSTALLING A V-SERIES METROCLUSTER SYSTEM

Installing V-Series MetroCluster is like installing two separate V-Series systems at different locations, then clustering and enabling SyncMirror[®] between the two sites. Normal FAS rules must be followed for SyncMirror. MetroCluster systems can be configured as fabric attached with distance support up to 100 km, or as stretch with support up to 500 m. For detailed guidelines, see the “V-Series MetroCluster Guide.”

17.2 SUPPORTED METROCLUSTER CONFIGURATION

Figure 7 shows the supported 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 0a/0c for the local pool and through 0b/0d for the remote pool. Each NetApp initiator port is zoned to a single array target port. Likewise, the cluster 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.

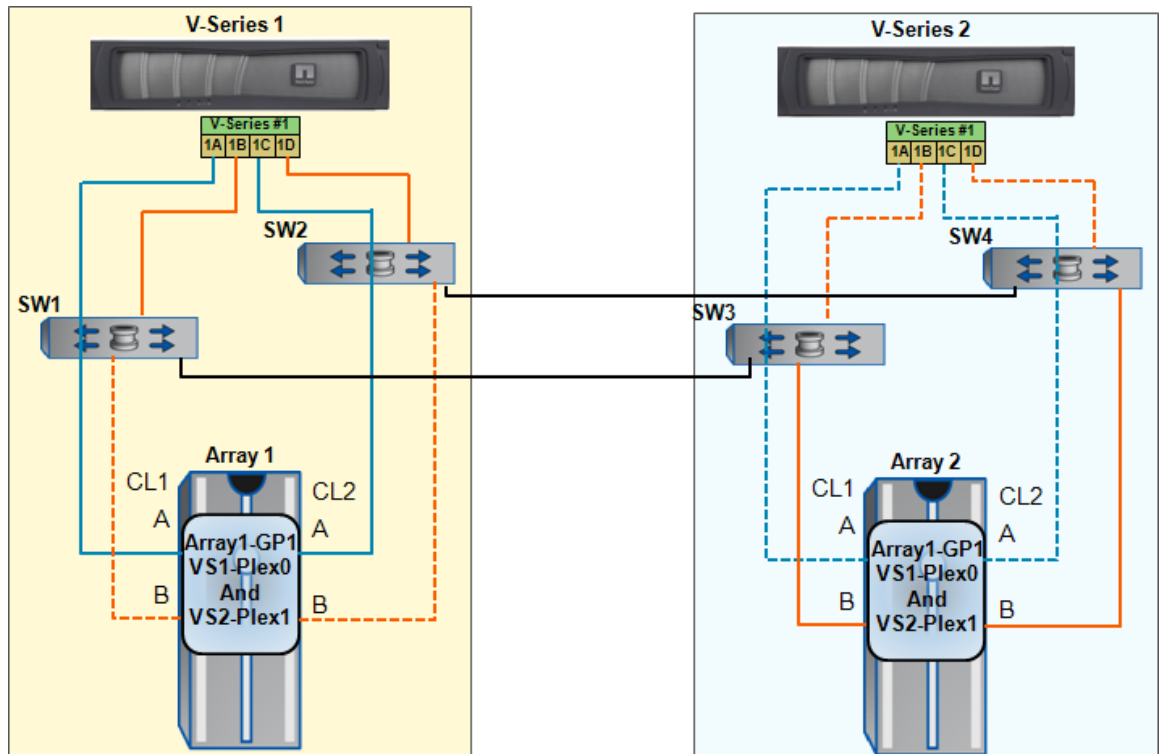


Figure 7) V-Series MetroCluster supported configuration.

Figure 7 shows a V3000 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 5) Zoning for topology in Figure 7.

Name	SW1-SW3 Fabric	Name	SW2-SW4 Fabric
Zone 1:	VS1:0a, Array1:CL1:A	Zone 1:	VS1:0c, Array1:CL2:A
Zone 2:	VS1:0b, Array2:CL1:A	Zone 2:	VS1:0d, Array2:CL2:A

Zone 3:	VS2:0a, Array2:CL1:B	Zone 3:	VS2:0c, Array2:CL2:B
Zone 4:	VS2:0b, Array1:CL1:B	Zone 4:	VS2:0d, 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.

V-Series MetroCluster does *not* support NetApp shelves. Only array LUNs are supported. If your solution requires NetApp shelves, you must use FAS MetroCluster.

V-SERIES METROCLUSTER SUPPORTS

- Symmetric configurations between sites (array family, switch types, firmware versions)
- Brocade switches (check the support matrix for specific models and firmware)
- Number of LUNs supported in each array LUN group is 100
- SAN requirements include:
 - Dedicated fabric
 - Maximum supported distance for a V-Series MetroCluster is 100 km
 - Use single initiator zoning, as shown in Table 5
- Cluster interconnect
 - Long-distance ISL shares FCVI (cluster interconnect) and data traffic (sufficient bandwidth must be available to support NVRAM mirroring)
 - The cluster interconnect requires in-order delivery (see the “V-Series MetroCluster Guide” for deployment examples)

For more detailed information about V-Series MetroCluster, see the “V-Series MetroCluster Guide” and the V-Series Support Matrix on NOW.

18 ARRAY-SPECIFIC NOTES

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. Use the NetApp software features in Data ONTAP (Snapshot, SnapMirror®, SyncMirror, SnapMover®, FlexVol, FlexClone®, and so on) instead. For information about how to properly configure array options, see the “V-Series Implementation Guide” for your array. 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 TAGMASTORE

V-Series systems can attach to HDS Tagmastore (USP and USP-V) arrays 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.
- Only BCS is supported (not ZCS) when using external drives behind the Tagmastore. Internal drives can be assigned as either BCS or ZCS.
- 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 DS4000

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.

18.3 EMC DMX

On DMX, use the LUN masking capability to associate logical devices to all host (channel) director ports, but allow only the V-Series FC initiator ports to see the LUNs for V-Series. 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.

© Copyright 2010 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, NOW, RAID-DP, SnapLock, SnapMirror, SnapMover, Snapshot, SyncMirror, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.