

V-Series Systems

Implementation Guide for Native Disk Shelves

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089 U.S.A.
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>

Part number: 210-04521_A0
July 2009

Copyright and trademark information

Copyright information

Copyright © 1994-2009 NetApp, Inc. All rights reserved. Printed in the U.S.A.

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 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.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, Cryptainer, Cryptoshred, DataFabric, DataFort, Data ONTAP, Decru, FAServer, FilerView, FlexClone, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, NOW NetApp on the Web, SANscreen, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, StoreVault, SyncMirror, Topio, VFM, VFM (Virtual File Manager), and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The NetApp arch logo; the StoreVault logo; ApplianceWatch; BareMetal; Camera-to-Viewer; ComplianceClock; ComplianceJournal; ContentDirector; ContentFabric; EdgeFiler; FlexShare; FPolicy; Go Further, Faster; HyperSAN; InfoFabric; Lifetime Key Management, LockVault; NOW; ONTAPI; OpenKey, RAID-DP; ReplicatorX; RoboCache; RoboFiler; SecureAdmin; Serving Data by Design; Shadow Tape; SharedStorage; Simplicore; Simulate ONTAP; Smart SAN; SnapCache; SnapDirector; SnapFilter; SnapMigrator; SnapSuite; SohoFiler; SpinMirror; SpinRestore; SpinShot; SpinStor; vFiler; Virtual File Manager; VPolicy; and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of NetApp, Inc. in the U.S.A.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the U.S.A. and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks. NetApp, Inc. NetCache is certified RealSystem compatible.

Table of Contents

	Preface	vii
Chapter 1	Implementation Overview	1
	Expanding your virtual storage pool with native disk shelves	4
	Usage scenarios for native disk shelves with V-Series systems	6
	General requirements to use native disk shelves with V-Series systems.	8
	Supported disk types and disk shelves	10
	Understanding storage provided by native disk shelves	11
Chapter 2	Planning for Using Disk Shelves with V-Series	15
	Comparing how Data ONTAP functions work with disks and array LUNs	16
	Considerations when planning storage requirements.	20
	Guidelines for creating aggregates	21
	Considerations when planning for the root volume	24
	Guidelines for V-Series port usage	26
	Limits on a V-Series system using both disks and array LUNs	28
Chapter 3	Installing a New V-Series System with Disks	29
	Factory default setup of a V-Series system with disk shelves	30
	Installing a V-Series system with disk shelves	32
Chapter 4	Adding a New Disk Shelf to a V-Series System	35
	General requirements when adding a disk shelf	36
	Overview of adding a new disk shelf.	38
Chapter 5	Moving a Disk Shelf to a V-Series System	43
	Planning for moving a disk shelf	44
	Tasks to perform on the FAS system before disconnecting the shelf	48

	Process overview for physically moving the disk shelf	55
	Optional configuration after moving the disk shelf	56
	Moving a disk shelf that is installed on a V-Series system.	58
Chapter 6	Managing a System with Disks and Array LUNs	59
	Identifying disk and array LUN information in command output	60
	Assigning disks and array LUNs to the V-Series system.	64
Appendix A	Configuring V-Series Ports as Initiators.	69
Appendix B	Setting up and Managing Multipath Storage	71
	Description of Multipath Storage.	72
	V-Series specific requirements for Multipath Storage	73
	Setup overview—Multipath Storage on an HA pair	75
	Configuring the V-Series environment variable for Multipath Storage	77
Appendix C	Troubleshooting Configurations with Disks.	81
	Common problems	82
	Error messages	84
	Reinstalling the system	86
	Index	91

Preface

About this guide This guide provides information about how to set up your storage array to work with a V-Series system running Data ONTAP® software, including configuration guidelines and sample configurations. The information in this guide pertains to all supported V-Series platforms.

Note
Data ONTAP software runs on multiple hardware platforms. This documentation might describe features that are not supported on your platform.

Audience This guide is for system administrators who are familiar with operating systems such as UNIX® and Windows® and who will be installing V-Series systems. This guide does not discuss basic system or network administration topics, such as IP addressing, routing, and network topology; it emphasizes the characteristics of the V-Series system.

Relationship of this guide to other guides This guide is intended to be used in conjunction with other information in the V-Series and Data ONTAP libraries. The following table describes the relationships between this guide and other documentation.

Guide name	Information includes...
<i>Installation Requirements and Reference Guide</i>	<ul style="list-style-type: none">◆ General guidelines for creating and making array LUNs available to V-Series systems.◆ Quick start installation instructions for connecting devices together and for installing Data ONTAP on a V-Series system that uses only third-party storage◆ Reference information◆ Detailed background information including layout in aggregates and checksums
<i>Implementation Guides</i>	<ul style="list-style-type: none">◆ Vendor-specific details about how to set up a storage array to work with V-Series systems.◆ More detailed configuration examples than are provided in the <i>Installation Requirements and Reference Guide</i>.

Guide name	Information includes...
<i>Implementation Guide for Native Disk Shelves</i>	Information about setting up the storage on the native disk shelves connected to the V-Series system.
<i>Data ONTAP Software Setup Guide</i>	Detailed steps for setting up the V-Series system, including information about installing Data ONTAP software for installations using only third-party storage. This guide is most helpful to installers who are new to Data ONTAP setup and installation.
Data ONTAP guides	Detailed information about all Data ONTAP features used by all systems running Data ONTAP, for example, storage features and data protection features.

See the V-Series *Support Matrix* for details about Data ONTAP releases that support V-Series, supported switches, supported firmware, capacity, and maximum array LUN count.

Special messages

This guide contains special messages that are described as follows:

Note

A note contains important information that helps you install or operate the system efficiently.

Attention

Attention contains instructions that you must follow to avoid damage to the equipment, a system crash, or loss of data.

About this chapter

This chapter provides an overview of V-Series support for native disk shelves with V-Series systems. Later chapters discuss the requirements and procedures for setting up native disk shelves with a V-Series system.

Information provided in this guide

The majority of the processes for setting up and managing disk shelves are the same on all systems running Data ONTAP®. This book describes the differences for V-Series systems and provides information to help you plan for, set up, and manage a system that uses both native disk shelves and array LUNs from third-party storage arrays.

Note

For information about planning for and setting up array LUNs from third-party storage array vendors with V-Series systems, see other guides in the V-Series library, particularly the implementation guides.

Topics in this chapter

This chapter contains the following topics:

- ◆ [“Expanding your virtual storage pool with native disk shelves”](#) on page 4
- ◆ [“Usage scenarios for native disk shelves with V-Series systems”](#) on page 6
- ◆ [“General requirements to use native disk shelves with V-Series systems”](#) on page 8
- ◆ [“Supported disk types and disk shelves”](#) on page 10
- ◆ [“Understanding storage provided by native disk shelves”](#) on page 11

Terminology

This document refers to storage from third-party storage arrays and native disk shelves that are directly connected to the V-Series system. The following terminology is used to clarify the type of storage being discussed.

array LUN: The term *array LUN* is used in the V-Series library to refer to the storage that third-party storage arrays provide to the V-Series system. One array LUN is the equivalent of one disk on a disk shelf connected to a V-Series system.

front-end LUN or LUN: The V-Series system can virtualize the storage provided by third-party storage arrays and serve it up as *LUNs* to applications and customers outside the V-Series system (for example, through iSCSI). Clients are unaware of where such a LUN is stored. Data ONTAP software product guides

refer to this type of storage as a LUN. You can think about this type of storage as a *front-end LUN* or *Data ONTAP-served LUN*. It is not the same as an array LUN.

disks, disk shelves: The terms *disk* and *disk shelf* in this document refer to native storage connected to the V-Series system. These terms do not refer to disks or disk shelves on a third-party storage array.

HA pair: An HA pair is a pair of V-Series systems configured to serve data for each other if one of the two systems becomes impaired. In V-Series documentation, Data ONTAP documentation, and other information resources, HA pairs are sometimes also referred to as clusters or active/active configurations.

native disks, native disk shelves: Disks and disk shelves that are sold as local storage for systems that run Data ONTAP software.

Other documents with information about disks

Use this guide in conjunction with other guides in the Data ONTAP and hardware libraries that discuss setup and management of disks and disk shelves. In particular, refer to the documents in the following table.


For information about...	See...
Disk support, including supported disk speeds, and disk capacity	Search on the NOW™ site for this information for FAS systems.
Installing the V-Series system in a rack or system cabinet	On new systems, this task is typically performed by the factory. If you need instructions, see the guide for your cabinet.
Connecting a disk shelf to a stand-alone V-Series system	Hardware <i>Installation and Setup Overview</i> for your platform
Connecting an HA pair to a disk shelf	Data ONTAP <i>High Availability Configuration Guide</i>
Step-by-step procedures for adding a disk shelf	The appropriate guide for your disk shelf type.

For information about...	See...
Disk management	<ul style="list-style-type: none"> ◆ Chapter 6, “Managing a System with Disks and Array LUNs,” on page 59 ◆ Data ONTAP <i>Storage Management Guide</i>

Expanding your virtual storage pool with native disk shelves

Combining disk and array LUN storage with V-Series systems

Starting with Data ONTAP 7.3, V-Series expands its support matrix to include storage for native disk shelves and native disks. V-Series can virtualize storage from both third-party storage array vendors and native disk shelves. The Data ONTAP software 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.



Note

The preceding illustration shows single-path connections between the V-Series controllers and disk shelves. See Appendix B, “[Setting up and Managing Multipath Storage](#),” on page 71 for information about how to set up multipath connections between the V-Series controllers and disk shelves.

Disk shelves that meet V-Series requirements can be installed on new and existing V-Series systems. You can add disk shelves to your V-Series systems in any of the following ways:

- ◆ Order native disk shelves with a new V-Series system
- ◆ Add a new native disk shelf to an existing V-Series system
- ◆ Move a disk shelf from a FAS system to a V-Series system or between V-Series systems

Using only native disk shelves with V-Series systems

Although typical V-Series systems that use native disk shelves also use third-party storage, using only native disk shelves with a V-Series system is supported.

Expanded Data ONTAP feature set

A few Data ONTAP features, such as compliance features, require that Data ONTAP have complete control over the storage. Such features are now supported on native disk shelves used with V-Series systems, but not with the third-party storage. Not all features are supported with all versions of Data ONTAP. See the V-Series *Support Matrix* to determine which features are supported with your release of Data ONTAP.

Usage scenarios for native disk shelves with V-Series systems

This section describes a few ways that customers use native disk shelves with V-Series systems.

Expanding storage after reaching storage array limits

As an alternative to expanding your storage array, you can add disks shelves to the V-Series system when you need additional storage. For example, assume you have added disks to your Hitachi Tier 1 storage array. As a result, you have reached the limit for the maximum number of disks that the current configuration can support. Options for expanding storage are as follows:

- ◆ Add native disk shelves to the V-Series system
- ◆ Purchase more controller cards on the Hitachi storage array
- ◆ Add another storage array—for example, a Tier 2 storage array

Simplifying installation

When a V-Series system is ordered with disk shelves, initial setup of the system is easier because the factory performs some installation tasks that must be performed at the customer site for systems without native disk shelves. Tasks performed by the factory include installing the root volume, the Data ONTAP software, and the V-Series license.

Improving core file availability

When a native disk shelf is installed on a V-Series system and the root volume is installed on a disk, core dumps are likely to be more available than if the root volume is installed on an array LUN. See [“Considerations when planning for the root volume”](#) on page 24 for more details.

Using SnapLock Compliance functionality

If you need SnapLock Compliance functionality, you can add native disk shelves to your V-Series system and create SnapLock Compliance aggregates on those disks. SnapLock Compliance is not supported with all versions of Data ONTAP. See the V-Series *Support Matrix* to determine if it is supported with your release of Data ONTAP.

Note

Data ONTAP SnapLock Compliance aggregates with array LUNs is not supported.

See “[SnapLock Compliance aggregates](#)” on page 22 for more information.

General requirements to use native disk shelves with V-Series systems

Minimum system requirements

The following system requirements must be met to support native disk shelves on a V-Series system:

- ◆ Data ONTAP 7.3 or higher is required to use native disk shelves
- ◆ The V-Series *Support Matrix* must specify that the model supports native disk shelves
- ◆ Information about supported disk shelves is available in the V-Series *Support Matrix*

Note

Although V-Series systems support native disk shelves starting in Data ONTAP 7.3, V-Series systems continue to be the only systems running Data ONTAP that can support third-party storage arrays. You cannot convert a FAS system to a V-Series system or convert a V-Series system to a FAS system.

Storage arrays supported in configurations with native disk shelves

A V-Series system that supports native disk shelves can be used with any storage array that is supported by that model.

Licensing

No license is required to use native disk shelves with a V-Series system.

MetroCluster with native disk shelves on a V-Series system

V-Series systems on which native disk shelves are installed cannot be deployed in a MetroCluster configuration.

SyncMirror

Setting up the SyncMirror® feature between disks and array LUNs is not supported.

Where to find the most recent information

The V-Series *Support Matrix* at <http://now.netapp.com> is updated more frequently than this guide. The V-Series *Support Matrix* is the final authority about which V-Series models and Data ONTAP releases support disk shelves, which switches and storage array models are supported in a particular Data ONTAP release, and so on.

Supported disk types and disk shelves

Disk types that V-Series supports

V-Series supports the following disk types:

- ◆ FC (Fibre Channel)
- ◆ SATA (Serial ATA)

Note

Only AT-FCX modules shipped in December 2005 or later are supported. If you are unsure whether your module is the correct version, see the *Data ONTAP High Availability Configuration Guide* or the *Data ONTAP Active/Active Configuration Guide* for the procedure for checking the version of the AT-FCX module.

Native disk shelves that V-Series supports

V-Series models that support disk shelves support the same disk shelves and the same disk types per I/O module on the shelves as the equivalent FAS models. See the *System Configuration Guide* for more information.

Note

V-Series requires the software disk ownership scheme, which uses persistent reservations. Only disk shelves that support software disk ownership can be used with V-Series systems. You might also see the disk ownership scheme referred to as SANOWN or LUN ownership.

Understanding storage provided by native disk shelves

Connections and loops


Disks are inserted in disk shelves connected to a V-Series or FAS storage controller. The connection from a storage controller to the disk shelves, which might be daisy-chained, is sometimes called a *loop*. The loops are described as follows:

- ◆ The A loop or A channel is the connection from the storage controller to the A port on the disk shelf module (not the A port on the storage controller or host bus adapter).
- ◆ The B loop or B channel is the connection from the storage controller to the B port on the disk shelf module (not the B port on the storage controller or host bus adapter).

On a stand-alone system, you can use one or two connections between the storage controller and the disk shelf. Using two connections provides redundancy. A connection can be from any free FC initiator port on the storage controller; if you are using just one connection, that connection can be to either disk shelf module.

In an HA pair, you have the option to provide redundant connections through the Data ONTAP Multipath Storage feature. (See Appendix B, “[Setting up and Managing Multipath Storage](#),” on page 71 for more details.)

The following illustration shows a stand-alone system with two redundant connections to the disk shelf, one to each disk module.



See “[Supported disk types and disk shelves](#)” on page 10 for details about which disks and disk shelves that V-Series supports. For information about the number of disk shelves and disk types that you can connect to a controller, see the information provided on the NOW site for FAS systems.

RAID types supported

V-Series RAID support is different for native disk shelves and array LUNs, as follows.

RAID types supported by native disk shelves: For native disk shelves, Data ONTAP organizes disks into RAID groups, which are collections of data and parity disks, to provide parity protection. Data ONTAP supports the following RAID types on native disks:

- ◆ RAID4 technology: Within its RAID groups, Data ONTAP allots a single disk for holding parity data, which ensures against data loss due to a single disk failure within a group.
- ◆ RAID-DP™ technology (DP for double-parity): RAID-DP provides a higher level of RAID protection for Data ONTAP aggregates. Within its RAID groups, Data ONTAP allots one disk for holding parity data and one disk for holding double-parity data. Double-parity protection ensures against data loss due to a double disk failure within a group.

RAID types supported for array LUNs: For third-party storage arrays, the V-Series system stripes across array LUNs using RAID0, thereby maximizing performance because it is using more spindles. A V-Series system does not provide disk protection for storage from third-party storage arrays. The reason is that the array LUNs are already protected with the redundancy provided by the various RAID levels on the storage arrays.

Note

Do not use RAID0 on storage arrays for array LUNs for V-Series systems: you do not have redundancy for your data if you use RAID0.

Checksum type of a native disk

All new systems running Data ONTAP use block checksum (BCS) disks. FC BCS disks are formatted with 520 bytes per sector. Of those 520 bytes, 512 are available for data, while 8 bytes are used by RAID for the checksum. ATA BCS disks are formatted with 512 bytes per sector and use one sector for the checksum for every eight data sectors.

You can use only block checksums for native disk shelves. In contrast, you can use either block checksums or zoned checksums with array LUNs.

Note

If you are moving a disk shelf from an older FAS system to a V-Series system, it might have ZCDs, which have a disk format of 512 bytes per sector and use one sector for the checksum for every 64 data sectors. There are rules about combining BCS and ZCS in the same aggregate. See “[Restrictions on mixing storage types in an aggregate](#)” on page 21 for more information.

About this chapter

This chapter provides information to help you plan for using native disk shelves with V-Series systems.

Topics in this chapter

This chapter contains the following topics:

- ◆ [“Comparing how Data ONTAP functions work with disks and array LUNs”](#) on page 16
- ◆ [“Considerations when planning storage requirements”](#) on page 20
- ◆ [“Guidelines for creating aggregates”](#) on page 21
- ◆ [“Considerations when planning for the root volume”](#) on page 24
- ◆ [“Guidelines for V-Series port usage”](#) on page 26
- ◆ [“Limits on a V-Series system using both disks and array LUNs”](#) on page 28

Comparing how Data ONTAP functions work with disks and array LUNs

Data ONTAP provides a robust variety of features for managing storage, most of which work the same no matter what type of storage your V-Series system uses. However, when you have a V-Series system that uses both native disks and array LUNs, you need to think about the differences so that you can set up and manage your “mixed” system appropriately.

The following table compares how Data ONTAP storage features and functions work on a V-Series system that uses both native disks and array LUNs.

Data ONTAP feature or function	Native disk shelves	Array LUNs
checksum type	All currently shipping disks are set as block checksum. You cannot change the checksum type. If you are planning to move older zoned-checksum disks to your V-Series system, see the <i>Data ONTAP Storage Management Guide</i> for the rules for assigning zoned-checksum disks to aggregates.	Both block checksum and zone checksum are supported for array LUNs. You specify the checksum type of the LUN when you assign an array LUN to a V-Series system (using the Data ONTAP <code>disk assign</code> command).
disk sanitization	Supported for native disks	Not supported on array LUNs
disk ownership	V-Series supports only software-based disk ownership for all disks and array LUNs. You can only use disks and disk shelves that use software-based disk ownership with V-Series systems. Note — You might see disk ownership referred to as SANOWN or, in the V-Series library in particular, LUN ownership.	
disk sizing versus array LUN sizing	Fixed size	Customer-defined size, which must fall within the V-Series minimum and maximum LUN size requirements

Data ONTAP feature or function	Native disk shelves	Array LUNs
licensing	You need only a V-Series license. No special license is required to support a specific storage type on a V-Series system.	
limits for disks and array LUNs	When FAS disk limits and V-Series array LUN limits differ, the total limit of all the disks plus all array LUNs must be less than or the same as the V-Series platform LUN limit. See “Limits on a V-Series system using both disks and array LUNs” on page 28.	
MetroCluster	Not supported on the V-Series system if native disk shelves are installed on it.	
paths to storage	<ul style="list-style-type: none"> ◆ For stand-alone and HA pairs, you can establish one or two paths between the V-Series controller and the disk shelf. Two paths are required for redundancy. ◆ For an HA pair, you must use the Multipath Storage feature if you want redundant connections. See Appendix B, “Setting up and Managing Multipath Storage,” on page 71 for details about this feature, including restrictions. 	<p>The requirements are as follows:</p> <ul style="list-style-type: none"> ◆ Two redundant paths are required between the V-Series system and each array LUN. ◆ More than two paths to an array LUN is not supported.
RAID	RAID-DP and RAID4 are supported on native disks.	<p>V-Series uses RAID0 and relies on the RAID protection that the storage arrays provide through the various RAID levels on the storage arrays.</p> <p>RAID4 and RAID-DP are not supported for array LUNs</p>

Data ONTAP feature or function	Native disk shelves	Array LUNs
RAID group size	Maximum and default RAID group sizes vary according to the V-Series system model and the level of RAID group protection provided by either RAID-DP or RAID4. See the Data ONTAP <i>Storage Management Guide</i> for more information.	Maximum and default RAID group sizes vary according to the version of Data ONTAP software. See the V-Series <i>Support Matrix</i> for more information.
redundant connections to storage	You can use either one or two loops to a native disk. You must use two loops for redundancy. If you want redundant loops for an HA pair, you must use the Multipath Storage feature.	You must configure redundant paths to an array LUN.
root volume	The root volume can be installed either on a native disk or an array LUN. See “ Considerations when planning for the root volume ” on page 24 for more details.	
SnapLock Compliance	Only aggregates created from storage on native disk shelves support SnapLock Compliance. SnapLock Compliance is not supported in all releases. See the V-Series <i>Support Matrix</i> for more information.	Not supported for array LUNs.
SnapMover	You can move aggregates and volumes composed of native disks only between the nodes of the HA pair.	You can move aggregates and volumes composed of array LUNs between any of the V-Series systems in the neighborhood. See the V-Series <i>Support Matrix</i> to determine which versions of Data ONTAP support SnapMover.

Data ONTAP feature or function	Native disk shelves	Array LUNs
topology	Disk shelves are directly connected to the V-Series system through a loop-based technology.	The V-Series system is connected a storage array either directly or through a Fibre Channel switched fabric.
V-Series neighborhood	The concept of V-Series neighborhood does not apply to native disk shelves attached to a V-Series system. However, when disks are attached to a system in the neighborhood that uses both disks and array LUNs, you must count the disks toward the visible limit. See the <i>V-Series Installation Requirements and Reference Guide</i> for more information.	All V-Series systems in the neighborhood see each other's array LUNs.
V-Series FC ports	V-Series FC ports must be set as initiators to access storage on native disk shelves and array LUNs.	

Considerations when planning storage requirements

Before you start setting up your V-Series systems, you need to determine how much storage you need for your organization's requirements and where you want to locate data—what should go on disks and what should go on array LUNs. You need to know this information to configure your V-Series system.

Considerations include the following:

- ◆ Which LUNs on the storage arrays do you want this V-Series system to own and what should the checksum type of each be?
- ◆ Which native disks do you want this V-Series system to own?
- ◆ Where do you want to locate the root volume?
See “[Considerations when planning for the root volume](#)” on page 24 for information about your choices.
- ◆ Do you want your native disk shelf to be used only for the root volume, only for data, or for both?
- ◆ Do you want your array LUNs to be used only for the root volume, only for data, or for both?
- ◆ If you are moving a disk shelf from a FAS system to a V-Series system, do you want to preserve the data on that disk shelf?
- ◆ How do you want to use your aggregates and volumes?

Guidelines for creating aggregates

Restrictions on mixing storage types in an aggregate

Data ONTAP does not support mixing different types of storage in the same aggregate because it causes performance degradation. For example, you cannot mix the following storage types in the same aggregate:

- ◆ Native disks and array LUNs
- ◆ Different types of disk drives

Note

Avoid mixing drives with different disk speeds in the same aggregate; mixing drives of different speeds in the same aggregate directly affects disk input/output operations per second and response time.

- ◆ Native disks with block checksum type and native disks with zoned checksum type
- ◆ Array LUNs from Fibre Channel disks and array LUNs from SATA disks
- ◆ Array LUNs from storage arrays from different vendors
- ◆ Array LUNs from storage arrays from different families

Note

Storage arrays in the same *family* share the same characteristics. See the *V-Series Implementation Guide* for your vendor for information about how V-Series defines family members for the vendor.

- ◆ Array LUNs from storage arrays with 4-Gb HBAs and LUNs from storage arrays with 2-Gb HBAs

Example: The following command creates an aggregate called vs23b_vol1 with the same type of storage in it—four FCAL disks on systems running Data ONTAP 8.0 7-Mode and 7.x releases:

```
aggr create vs23b_vol1 -T FCAL -n 4
```

Considerations when preserving aggregates when moving a shelf

Data ONTAP does not support duplicate aggregate and volume names on the same system. If you want to preserve the data on a disk shelf you are moving to the V-Series system, check the aggregate and volume names on the disk shelf and V-Series system for duplicates before you move the disk shelf. If you discover

aggregate and volume names that are duplicated, you must change those aggregate and volume names before moving the shelf. See Chapter 5, “[Moving a Disk Shelf to a V-Series System](#),” on page 43 for more information.

SnapLock Compliance aggregates

SnapLock volumes are write once, read many (WORM) volumes that you create for data that you want to archive permanently. SnapLock enables you to create an aggregate with SnapLock as an attribute of that aggregate and then specify the type of SnapLock volume—either Compliance or Enterprise.

Note

SnapLock Compliance requires an additional license.

SnapLock Compliance is not supported in all Data ONTAP releases. See the *V-Series Support Matrix* to determine if it is supported with your release of Data ONTAP.

The following table summarizes SnapLock support with V-Series systems.

Type of SnapLock volume	How it is used	Data ONTAP support for
SnapLock Compliance volume	For strict regulatory environments, such as SEC 17a-4 compliant environments	Supported only for SnapLock aggregates composed of native disks. Data ONTAP prevents you from creating a SnapLock Compliance aggregate with array LUNs.
SnapLock Enterprise volume	For environments without regulatory restrictions	Supported for SnapLock aggregates composed of disks or array LUNs

Note

SnapLock Compliance prevents anyone, including the system administrator, from destroying or altering files, volumes, or aggregates before they reach their expiration date. SnapLock Compliance aggregates cannot be created for array LUNs because government regulations require that the WORM data be protected from alteration. Data ONTAP has no control over any changes to the array LUNs that the storage array administrators might want to make. For example, a storage array administrator could destroy an array LUN that Data ONTAP is using, which would result in noncompliance.

For more information about SnapLock Compliance, see the *Data ONTAP Archive and Compliance Management Guide*.

Where to find more information about creating aggregates

See the following additional guides for more information about planning for aggregates, including advantages and disadvantages of different approaches.

For information about...	See...
Planning for aggregates composed of disks	<i>Data ONTAP Storage Management Guide</i>
Planning for aggregates composed of array LUNs	<i>V-Series Installation Requirements and Reference Guide</i>

Considerations when planning for the root volume

Root volume defined

Each V-Series system has a special volume called a *root volume*. The root volume is required for the system to be able to boot and to function properly. The root volume contains system files and configuration information. The root volume can also contain data if you want it to. Core dump files, which are important for troubleshooting, are written to the root volume if there is enough space.

On a V-Series system with native disks, the root volume can be on a disk shelf or on an array LUN.

Note

Systems running Data ONTAP 8.0 Cluster-Mode must be configured with the root volume located on a disk shelf.

Options for the location of the root volume

The benefits of the root volume on a native disk shelf are as follows:

- ◆ For a new V-Series system, initial setup and configuration is easier because you have fewer steps to perform.
If you order a V-Series system with disk shelves, your system is delivered with the root volume and with three disks—a RAID-DP aggregate with one parity disk and two data disks. The root volume configured at the factory is a FlexVol volume with a block checksum type.
- ◆ Greater availability of core dump files
When a native disk shelf is installed on a V-Series system and the root volume is installed on a disk, core dumps are likely to be more available. If the root volume is on an array LUN, a core file cannot be saved if the array LUN becomes unavailable for some reason (for example, because of switch or storage array issues). Installing the root volume on a native disk eliminates the impact of network and network device failures on core file production.

Note

A core dump file is a valuable tool for technical support to use while troubleshooting why a V-Series system crashed. When a V-Series system crashes, a core dump file is saved to the root volume if adequate space is available on the root volume.

Root volumes on nodes in an HA pair

In an HA pair, it is highly recommended that the root volume be located on the same type of storage for both nodes—either on a disk shelf for both nodes or on an array LUN for both nodes. Although an asymmetrical approach is highly discouraged, it is not a requirement to follow the same strategy for both nodes.

Note

The root volume on a Data ONTAP 8.0 Cluster-Mode system must be located on a disk shelf.

Guidelines for V-Series port usage

Setting required for HBAs

All V-Series HBAs that are used to access disks or array LUNs must be set to initiator ports. The factory might configure some HBAs as target ports. Therefore, when you set up your system you need to check the HBA settings and, if necessary, change a target HBA setting to an initiator. See Appendix A, [“Configuring V-Series Ports as Initiators,”](#) on page 69 for information about how to check the setting of the HBAs and change them if necessary.

Number of ports required to connect to a storage array

V-Series requires two paths to each array LUN. Use two redundant initiator ports on the V-Series system to access each LUN group on the storage array.

Number of ports required to connect to a disk shelf

You need to use the following number of ports to connect the V-Series system to a disk shelf:

For a stand-alone system:

- ◆ If you are using one loop, use one FC initiator port
- ◆ If you are connecting two loops, use two FC initiator ports, one for each loop.

For an HA pair:

- ◆ If you are using Multipath Storage, use two FC initiator ports for each loop
- ◆ If you are not using Multipath Storage, use one FC initiator port for each controller for each loop in the configuration.

Restrictions for sharing initiator ports

Use a separate V-Series initiator port to connect each of the following:

- ◆ disk shelves
- ◆ array LUNs
- ◆ tape devices

If your V-Series model does not have enough internal ports for your needs, you must order an additional HBA.

General guidelines for determining port assignment

Consider redundancy and load balancing when selecting which initiator ports to use to connect to array LUNs and disk shelves: use redundant port pairs to access storage. For example, on a V30xx system, the redundant port pairs are as follows:

- ◆ 0a and 0c
- ◆ 0b and 0d

If you are using 0a and 0c as the redundant port pair to access array LUNs, 0b and 0d is the redundant port pair to access the disk shelf.

See the Data ONTAP *High Availability Configuration Guide* or Data ONTAP *Active/Active Configuration Guide* and the *System Configuration Guide* for additional guidelines for selecting ports for connections to disk shelves.

See the V-Series *Installation Requirements and Reference Guide* for information about which ports are redundant on specific V-Series models.

Limits on a V-Series system using both disks and array LUNs

How minimum and maximum limits apply

Minimum and maximum limits that are documented for V-Series are the limits for the disks and array LUNs *combined*. Plan carefully to ensure that your combination of disks and array LUNs does not exceed these limits.

Where to find more information about limits

The following documents provide details about the limits.

Limit type	Document
Per platform limits for minimum and maximum system capacity	<i>V-Series Support Matrix</i>
Maximum number of disks and array LUNs per V-Series platform	
Minimum and maximum array LUN size—vendor equivalent of V-Series limits	<i>V-Series Implementation Guide</i> for your storage array vendor
How array LUN capacity is calculated	<i>V-Series Installation Requirements and Reference Guide</i>
Disks supported, capacity, available disk capacity	Search on the NOW™ site for this information for FAS systems
Root volume and RAID group limits	<i>V-Series Support Matrix</i>

About this chapter

This chapter describes the tasks you need to perform to complete setup of a new V-Series system ordered with disks.

If you...	See instead...
Are adding a disk shelf to an existing V-Series system	Chapter 4, “ Adding a New Disk Shelf to a V-Series System ,” on page 35
Want to move a disk shelf from a FAS system to this V-Series system or from another V-Series system to this V-Series system	Chapter 5, “ Moving a Disk Shelf to a V-Series System ,” on page 43

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Factory default setup of a V-Series system with disk shelves](#)” on page 30
- ◆ “[Installing a V-Series system with disk shelves](#)” on page 32

Factory default setup of a V-Series system with disk shelves

Factory cabling of a stand-alone system	When you order a stand-alone system, the factory cables it according to your instructions. All internal adapters, such as networking adapters, Fibre Channel adapters, and other adapters, arrive preinstalled in the nodes.
Factory cabling of an HA pair	<p>For information about factory cabling of an HA pair and the cabling that you need to perform to complete the process, see the <i>Data ONTAP High Availability Configuration Guide</i> or the <i>Data ONTAP Active/Active Configuration Guide</i>.</p> <p>Attention —</p> <p>An HA pair is not cabled or configured for Multipath Storage by default. If you want to use Multipath Storage, you must set a specific environment variable for V-Series before cabling and configuring the nodes for Multipath Storage. See “Configuring the V-Series environment variable for Multipath Storage” on page 77 for information about the environment variable that you need to set. See the <i>Data ONTAP Configuration Guide</i> for cabling and configuration requirements for Multipath Storage.</p>

Default software setup of a V-Series system with disks The following table shows the factory default setup of a V-Series system ordered with disks.

For this...	The factory default is...
Data ONTAP software	Data ONTAP software is installed.
Root volume	Installed on the disk shelf. The root volume is a FlexVol volume named vol0. You can resize it later, if desired.
V-Series license	Installed
Protocol licenses	Licenses for protocols and features that you purchased might have been installed—for example, licenses for CIFS, NFS, and HTTP. After setting up your system, you can add any additional licenses that you need.

For this...	The factory default is...
CompactFlash card	Bootloader files and firmware updates, including primary and secondary BIOS images, are installed on the CompactFlash card that is shipped with the system.
Disks	Disks with the latest firmware are installed.
V-Series port settings	The V-Series ports might or might not be set as initiators. They must be set as initiators to connect to either a storage array port or a native disk shelf. Appendix A, “ Configuring V-Series Ports as Initiators ,” on page 69 describes how to check the port settings and change them if necessary.

Installing a V-Series system with disk shelves

Overview of the installation process

Because the root volume and Data ONTAP are already installed on a disk for new V-Series systems ordered with disk shelves, software installation is simpler. When you boot the V-Series system for the first time, Data ONTAP starts the setup script and prompts you to enter basic information about system configuration, network services, network addresses, and protocols.

After the setup script finishes and you reboot your V-Series system, set up your storage by assigning the disks and array LUNs that you want this V-Series system to own.

Prerequisites to booting the new V-Series system

Before booting the V-Series system for the first time, complete the following:

- ◆ If you are using third-party storage as well as disk shelves, prepare the storage on the third-party storage arrays for V-Series, and connect the storage arrays to the V-Series system.
See the *V-Series Implementation Guide* for your storage array vendor.
- ◆ Check the settings of the V-Series initiator ports, as described in Appendix A, “[Configuring V-Series Ports as Initiators](#),” on page 69.
- ◆ Gather the data you need for the setup script.

It will be faster to complete the prompts in the setup script if you gather information needed for the setup parameters before you boot your system. Use the setup parameters worksheet in the *Data ONTAP Software Setup Guide* to gather the information you need.

Booting the new V-Series system with disks for the first time

To boot your system for the first time, complete the following steps.

Step	Action
1	Check that the disk shelves are set up properly and that the switches are terminated, where applicable.
2	<p>First turn on the power to any switches, then to the disk shelves, and then to any tape backup devices.</p> <p>Attention _____ Make sure that the network and SAN switches are powered on and configured before you turn on your V-Series system.</p>
3	<p>Turn on the power to the V-Series system.</p> <p>Note _____ It takes the LEDs on the system power supplies a few seconds to illuminate.</p> <p>Result: The system begins to boot, and then stops at the first installation setup script question, which is displayed on the system console.</p>
4	Verify that all network and RLM LEDs are properly illuminated.
5	<p>Go to the system console and complete the setup script using the information you gathered and entered in the worksheet from the appropriate <i>Data ONTAP Software Setup Guide</i>.</p> <p>Note _____ If you configured the RLM, the setup script pauses for a short time while the RLM reboots and sets its configuration information.</p>
6	<p>Note _____ For Data ONTAP 8.0 Cluster-Mode systems, use the node shell command to enter the following 7-Mode commands: run -node node command</p> <p>Check the licenses on your system by entering the following command: license</p>

Step	Action
7	<p>Add any missing licenses by entering the following command for each missing license:</p> <pre>license add xxx</pre> <p>xxx is the license code for the product.</p>
8	<p>Use the <code>disk assign</code> command to specify this V-Series system as the owner of array LUNs and additional disks.</p> <p>See “Assigning disks and array LUNs to the V-Series system” on page 64 for information about the <code>disk assign</code> command.</p> <p>Note_____</p> <p>If the <code>disk.auto_assign</code> option is enabled (the default), Data ONTAP automatically assigns new disks to your system. However, you must explicitly assign array LUNs to your system.</p> <p>_____</p>
9	<p>Set up Data ONTAP features, as applicable, using guides in the V-Series and Data ONTAP library.</p>

About this chapter

This chapter provides information about how to add a new disk shelf (a disk shelf that has not been used on another system) to an existing V-Series system.

If you...	See instead...
Are setting up a new V-Series system with disks	Chapter 3, “ Installing a New V-Series System with Disks ,” on page 29
Want to move a disk shelf from a FAS system to this V-Series system or from another V-Series system to this V-Series system	Chapter 5, “ Moving a Disk Shelf to a V-Series System ,” on page 43

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[General requirements when adding a disk shelf](#)” on page 36
- ◆ “[Overview of adding a new disk shelf](#)” on page 38

General requirements when adding a disk shelf

Minimum Data ONTAP requirement

Before you add a disk shelf to your V-Series system, the system must be running Data ONTAP 7.3 or higher.

Loop speed requirements

If you are adding a disk shelf to an existing V-Series system, be sure that you understand the characteristics of the disk shelf—including its speed. The output from the `sysconfig -a` command provides information about the loop speed of the disk shelf.

Review the characteristics of the disk shelf models that are documented in the guide for your disk shelf.

Attention

A common error in setting up disk shelves is a loop speed mismatch between components on the same loop. The Gb-loop speed for all components—including the disk shelves—must be set to the lowest speed of any component in the loop.

When there is a loop speed mismatch, an error message might be generated, the V-Series system might panic, or the V-Series system might not be able to see the disks on the disk shelf.

Availability of the V-Series system while a disk shelf is being added

You can add new disk shelves to a V-Series system without powering down the system. (This is referred to as *hot-adding disk shelves*.) While the disk shelf is being added, the V-Series system can continue to access data on any existing disk shelves and on the storage arrays that provide storage (array LUNs) for the V-Series systems.

Attention

Hot-adding disk shelves enables you to add (but not swap) disk shelves without a service interruption. However, you cannot add more disk shelves to your system than it can support; for an HA pair, you cannot add more disk shelves than either the entire configuration or either node can support.

Hot-adding a disk shelf is different from hot-swapping a disk shelf. Hot-swapping a disk shelf, which means removing an existing disk shelf and installing a new one in its place, is not supported. Your system supports hot-adding of disk shelves only.

Note

If you are using the Multipath Storage feature to connect disk shelves to the V-Series system, you must configure the `fc-non-array-adapter-list` environment variable *before* you connect the first disk shelf on a new loop. See Appendix B, “[Setting up and Managing Multipath Storage](#),” on page 71 for more information.

Overview of adding a new disk shelf

Assumptions

This section is based on the assumption that the following basic V-Series system setup has already been completed and tested:

- ◆ You have completed the basic system setup (setup script).
- ◆ Data ONTAP software was installed.
- ◆ You verified that the paths to existing storage were set up correctly.
- ◆ If the system is an HA pair, you have verified that takeover works.

High-level steps

The following table lists the major steps for adding a new disk shelf and provides information about where you can find detailed procedures.

Step	Action	Details
1	Review the restrictions for adding disk shelves.	See the following guides: <ul style="list-style-type: none">◆ <i>Appropriate</i> disk shelf guide for your disk shelf model.◆ Data ONTAP <i>High Availability Configuration Guide</i> or Data ONTAP <i>Active/Active Configuration Guide</i>
2	If you plan to use Multipath Storage with an HA pair and this is the first disk shelf in the loop, set the environment variable required for V-Series with Multipath Storage. Note — This step must be performed before connecting the first disk shelf in a loop to the V-Series system.	See “ Configuring the V-Series environment variable for Multipath Storage ” on page 77.
3	Install the new disk shelf in your cabinet or equipment rack.	See the disk shelf guide for your disk shelf model.

Step	Action	Details
4	Cable the new disk shelf to the V-Series system.	<p>For a stand-alone system, see the hardware <i>Installation and Setup Overview</i> for your platform.</p> <p>For an HA pair, see the Data ONTAP <i>High Availability Configuration Guide</i> or Data ONTAP <i>Active/Active Configuration Guide</i>.</p>
5	Add the new disk shelf to a new loop or to an existing loop.	<p>For a stand-alone system, see the hardware disk shelf guide for your disk shelf.</p> <p>For an HA pair, see the sections in the Data ONTAP <i>High Availability Configuration Guide</i> or Data ONTAP <i>Active/Active Configuration Guide</i> about adding new loops and adding storage to an existing loop.</p>

Step	Action	Details
6	Assign the disks to your V-Series system.	<p>Note_____</p> <p>For Data ONTAP Cluster-Mode systems, use the node shell command to enter the following 7-Mode commands:</p> <p>run -node node command</p> <hr/> <ol style="list-style-type: none"> At the system console, enter the following command to determine the type of disks on the disk shelf: <p>storage show disk -T</p> <p>You need to know the disk ID and type of the disks on the disk shelf to be able to configure disk ownership.</p> Use the <code>disk assign</code> command to assign disks on the disk shelf to the V-Series system. <p>See “Assigning disks and array LUNs to the V-Series system” on page 64 for details about the <code>disk assign</code> command.</p> <p>Example:</p> <p><code>disk assign -n 4 -T FCAL</code> assigns four FCAL type disks to the V-Series system.</p> Enter the following command and review the list of disks to ensure that all disks on the disk shelf you moved are now owned by this V-Series system: <p>disk show -v</p>
7	Assign array LUNs to the V-Series system now or later.	<p>See “Assigning disks and array LUNs to the V-Series system” on page 64 for details about the <code>disk assign</code> command.</p>

Step	Action	Details
8	Configure aggregates on your system.	<p>Follow the guidelines in “Restrictions on mixing storage types in an aggregate” on page 21.</p> <p>Example: The following command creates an aggregate called vs23b3b_vol1 with four FCAL disks in it:</p> <pre>aggr create vs23b3b_vol1 -T FCAL -n 4</pre> <p>See the <i>Data ONTAP Storage Management Guide</i> for details about creating aggregates for disks.</p>

About this chapter

This chapter describes how to move a disk shelf from a FAS system to a V-Series system.

If you...	See instead...
Are setting up a new V-Series system with disks	Chapter 3, “ Installing a New V-Series System with Disks ,” on page 29
Are adding a new disk shelf to a new loop or an existing loop	Chapter 4, “ Adding a New Disk Shelf to a V-Series System ,” on page 35

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Planning for moving a disk shelf](#)” on page 44
- ◆ “[Tasks to perform on the FAS system before disconnecting the shelf](#)” on page 48
- ◆ “[Process overview for physically moving the disk shelf](#)” on page 55
- ◆ “[Optional configuration after moving the disk shelf](#)” on page 56
- ◆ “[Moving a disk shelf that is installed on a V-Series system](#)” on page 58

Planning for moving a disk shelf

Types of disk shelves that you can move

You can move any disk shelf that supports software disk ownership (SANOWN) from a FAS system to a V-Series system. (See “[Supported disk types and disk shelves](#)” on page 10 for a list of disk shelves that support software disk ownership.)

Note

You can also move disk shelves installed on a V-Series system to another V-Series system or to a FAS system if the destination system supports the disk shelf type.

Requirement to upgrade to Data ONTAP

If your existing system is running a version of Data ONTAP software prior to 7.3, you must upgrade to at least Data ONTAP 7.3 before you move the disk shelf to the V-Series system.

Loop speed requirements

If you are moving a disk shelf to an existing V-Series system, be sure that you understand the characteristics of the disk shelf you plan to move—including its speed. The output from the `sysconfig -a` command provides information about the loop speed of the disk shelf.

Review the characteristics of the disk shelf models that are documented in the guide for your disk shelf.

Attention

A common error in setting up disk shelves is a loop speed mismatch between components on the same loop. The Gb-loop speed for all components—including the disk shelves—must be set to the lowest of any component in the loop.

When there is a loop speed mismatch, an error message might be generated, the V-Series system might panic, or the V-Series system might not be able to see the disks on the disk shelf.

Requirement if the FAS system root volume is on the disk shelf

Only one root volume is permitted on a system. A root volume that was configured for the FAS system cannot be used on the V-Series system.

If the root volume for the FAS system is on the disk shelf that you plan to move to the V-Series system, the requirements are as follows.

If...	Then...
You still need the root volume for the FAS system	You must move the root volume to another disk shelf that will remain with the FAS system <i>before</i> you move the disk shelf.
You do not need the root volume for the FAS system	Move the disk shelf to the V-Series system. Data ONTAP prevents any conflict with the root volume already configured on the V-Series system; when you move the disk shelf with the FAS root volume to the V-Series system, Data ONTAP flags the FAS root volume as a foreign object. Data ONTAP renames the root volume from aggr0 to aggr0(1) and vol0 to vol0(1). Data ONTAP also removes the “root” flag from the disk set so that the V-Series system cannot boot from the FAS root volume.

Availability of data while the shelf is being moved

While you are installing the disk shelf on the V-Series system, the data on the disk shelf is not available until installation on the V-Series system is complete (assuming you are preserving the data on the disk shelf). However, access to other storage associated with the V-Series system is not interrupted. You do not need to shut down the V-Series system to be able to install the disk shelf on it. You can *hot-add* disk shelves on the V-Series system (that is, add disk shelves without shutting the system down.)

Impact on clients of moving a disk shelf

The following table shows the impact on clients of moving a disk shelf.

For users of this protocol...	The impact is...
CIFS	Shares to the moved data need to be created and the UNC path will likely change.
NFS	Volumes need to be exported on the new system with a new path. Users must remount to the new location.
iSCSI	igroups need to be created on the new system and the array LUNs mapped to them. New iSCSI sessions must be established from the hosts.

Importance of understanding aggregate layout

Before moving a disk shelf from a FAS system to a V-Series system you need to understand how the aggregates have been laid out, for the following reasons:

- ◆ You must know which disks are in which aggregate so that you know which disks to assign to your V-Series system.
- ◆ If an aggregate spans shelves, you cannot split the aggregate.

If you...	Then...
Want to preserve the data in the aggregate for use on the V-Series system	You must move all the shelves to the V-Series system that the aggregate spans. (See “Requirements if you want to preserve data on the disk shelf” on page 47.)
Want to preserve the data for use on the FAS system	Move the entire aggregate onto disk shelves that will remain on the FAS system
Do not want to preserve the data	Destroy the aggregate, then move the desired number of disk shelves to the V-Series system.

Requirements if you want to preserve data on the disk shelf

If you want to preserve the data on a disk shelf for use with the V-Series system, the requirements are as follows:

- ◆ Ensure that the combined total of aggregates and volumes on the disk shelf and on array LUNs does not exceed the limits.
- ◆ Turn off Data ONTAP automatic assignment of disks *before* you move the shelf to the V-Series system (see [“Turning off automatic disk assignment”](#) on page 52). If desired, you can turn automatic assignment of disks back on after setup of the disk shelf on the V-Series system is complete.

Note

With Data ONTAP automatic disk assignment (the default behavior), Data ONTAP automatically looks for any unassigned disks on a predetermined schedule and assigns them to the same system and pool as all other disks on the same loop.

- ◆ Resolve any duplicated names of aggregates and volumes on the disk shelf you are moving and the V-Series system.

All aggregate names and volume names on the same system must be unique. If you want to preserve the data on the disk shelf, you need to determine whether any aggregate and volume names are the same on the two systems before you move the disk shelf to the V-Series system. If there are duplicates, you need to rename the duplicated names on the disk shelf you are moving.

Attention

Data on the disk shelf being moved remains intact through the ownership transition if you have followed the steps described in this section about preserving data on the disk shelf.

Requirement to change disk ownership

The disks on the disk shelves you are moving have already have been assigned to a FAS system. A V-Series system cannot use disks that a FAS system owns. Therefore, you must remove FAS system ownership of the disks and assign them to a V-Series system before you move disk shelf from a FAS system to a V-Series system. See [“Reassigning ownership of FAS system disks to a V-Series system”](#) on page 53.

Tasks to perform on the FAS system before disconnecting the shelf

Task list

The following table shows the tasks to perform before disconnecting the disk shelf from the FAS system.

Stage	Task	Relevant procedure
1	Check the disks, aggregates, and volumes on the FAS system.	“Checking the disks, aggregates, and volumes on the FAS system” on page 49
	If you want to preserve your data when you move the disk shelf to the V-Series system, check for duplicate aggregate and volume names between the two systems. Change the aggregate and volume names on the disks on the disk shelf you are moving, as necessary.	<ul style="list-style-type: none">◆ “Renaming an aggregate” on page 50◆ “Renaming a volume” on page 52
2	If the root volume for the FAS system is on the disk shelf you are moving, and you still need the root volume for the FAS system, move the root volume to another disk shelf that is to remain with the FAS system.	<ul style="list-style-type: none">◆ If you need to move the root volume to another disk shelf on the FAS system, see the <i>Data ONTAP Storage Management Guide</i> for information about moving volumes.
3	If you want any other aggregates or volumes on the disk shelf you are moving to remain on the FAS system, move them to a disk shelf that will remain on the FAS system.	Information about moving aggregates and volumes in the <i>Data ONTAP Storage Management Guide</i>
4	If you want the data on the disk shelf to be preserved, turn off the automatic disk assignment functionality before moving the shelf.	“Turning off automatic disk assignment” on page 52

Stage	Task	Relevant procedure
5	<p>Change ownership of the disks on the disk shelf you are moving from the FAS system to the V-Series system.</p> <p>Note _____ Complete this step before you move the disk shelf.</p>	<p>“Reassigning ownership of FAS system disks to a V-Series system” on page 53.</p>
6	<p>If this is an HA pair and you want to use Multipath Storage with it, configure the boot environment variable required for V-Series with Multipath Storage. You must set this environment variable for the first disk shelf in each loop.</p>	<p>“Configuring the V-Series environment variable for Multipath Storage” on page 77</p>

Checking the disks, aggregates, and volumes on the FAS system

To compare the aggregate and volume names on the FAS system and the V-Series system.

Step	Action
1	<p>For Data ONTAP 8.0 7-Mode and 7.x releases Data ONTAP 8.0 7-Mode and 7.x releases, enter the following command:</p> <pre>sysconfig -r</pre> <p>For Data ONTAP 8.0 Cluster-Mode enter the following command:</p> <pre>run -node node sysconfig -r</pre> <p>The output from this command shows information about all aggregates, volumes, file system disks, spare disks, maintenance disks, and failed disks.</p>
2	<p>Record the names of the disks on the disk shelves that you plan to move and the aggregates and volumes on the disk shelves.</p> <p>Note _____ Note if any of the aggregates span onto disk shelves that you do not plan to move.</p>
<p>If you want to preserve the data on the disk shelf you are moving.</p>	

Step	Action	
3	On the V-Series system, enter the following command: sysconfig -r	
4	Record the names of the aggregates and volumes on the V-Series system.	
5	Compare the aggregates and volumes on the two systems to look for duplicate names.	
6	If...	Then...
	No aggregate or volume names are duplicated.	No renaming is necessary.
	There are duplicate aggregate names.	You need to rename the duplicate aggregate name on one of the systems. Go to “Renaming an aggregate” on page 50.
	There are duplicate volume names.	You need to rename the duplicate volume name on one of the systems. See “Renaming a volume” on page 52.

Renaming an aggregate

To rename an aggregate on the disk shelf you are moving, complete the following steps.

Step	Action
1	<p>For Data ONTAP 8.0 7-Mode and 7.x releases, enter the following command:</p> <pre>aggr rename <i>aggr_name</i> <i>new_name</i></pre> <p>For Data ONTAP 8.0 Cluster-Mode enter the following command:</p> <pre>storage aggregate rename -aggregate <i>aggr_name</i> -newname <i>new_name</i></pre> <p><i>aggr_name</i> is the name of the aggregate you want to rename. <i>new_name</i> is the new name of the aggregate.</p> <p>Result: The aggregate is renamed.</p>

Renaming a volume

To rename a volume on the disk shelf you are moving, complete the following steps.

Step	Action
1	<p>For Data ONTAP 8.0 7-Mode and 7.x releases, enter the following command:</p> <pre>vol rename vol_name new-name</pre> <p>For Data ONTAP 8.0 Cluster-Mode enter the following command:</p> <pre>volume rename -vservers vs1 -volume vol_name -newname new-name</pre> <p><i>vservers</i> is the name of the virtual server the volume is on.</p> <p><i>vol_name</i> is the name of the volume you want to rename.</p> <p><i>new-name</i> is the new name of the volume.</p> <p>Result: The following events occur:</p> <ul style="list-style-type: none">◆ The volume is renamed.◆ If NFS is in use and the <code>nfs.exports.auto-update</code> option is On, the <code>/etc/exports</code> file is updated to reflect the new volume name.◆ If CIFS is running, shares that refer to the volume are updated to reflect the new volume name.◆ The in-memory information about active exports gets updated automatically, and clients continue to access the exports without problems.
2	<p>If you access the V-Series system using NFS, add the appropriate mount point information to the <code>/etc/fstab</code> or <code>/etc/vfstab</code> file on clients that mount volumes from the V-Series system.</p>

Turning off automatic disk assignment

If you want to preserve the data on the disk shelf, you must turn off disk auto-assignment *before* you move the shelf.

To turn off automatic disk assignment, complete the following step.

Step	Action
1	<p>For Data ONTAP 8.0 7-Mode and 7.x releases, enter the following command:</p> <pre>options disk.auto_assign off</pre> <p>For Data ONTAP 8.0 Cluster-Mode, enter the following command:</p> <pre>run -node node options disk.auto_assign off</pre>

Reassigning ownership of FAS system disks to a V-Series system

You need to change ownership of the disks on the shelf to be moved to the V-Series system while the shelf is still attached to the FAS system. However, you must find out the system name or NVRAM ID of the V-Series system before you can start reassigning the disks on the FAS system. You must enter the V-Series system name as part of the `disk reassign` command.

Finding out the V-Series system ID: To obtain the V-Series system ID, complete the following steps.

Step	Action
1	<p>For Data ONTAP 8.0 7-Mode and 7.x releases, enter the following command, then check the Owner column for a LUN entry to find the system ID for the V-Series system:</p> <pre>sysconfig</pre> <p>For Data ONTAP 8.0 Cluster-Mode, enter the following command:</p> <pre>run -node node sysconfig</pre> <p>Note _____ You need the system ID, not the serial number of the system.</p>
2	Record the V-Series system ID.
3	Go to “Changing the disk owner to the V-Series system” on page 53.

Changing the disk owner to the V-Series system: To change the assignment of disks on the disk shelf from the FAS system to the V-Series system, complete the following steps.

Step	Action
1	<p>On the FAS system, reboot the system and interrupt the boot process by pressing Ctrl-C when you see the following message on the console:</p> <p>Press Ctrl-C for special boot menu</p>
2	Enter Maintenance mode from the boot menu.
3	<p>Enter the following command to list the disks that are visible to the FAS system:</p> <p>disk show -v</p> <p>Result: The output shows the disks, the system that owns the disk, the pool to which the disk is assigned, and the serial number of the disk.</p>
4	<p>Reassign disk ownership of all the disks on the disk shelf that you plan to move to the V-Series system by entering the following command:</p> <p>disk reassign -s <i>old_systemname</i> -d <i>new_system name</i></p> <p><i>old_systemname</i> is the system name or NVRAM ID associated with the FAS system.</p> <p><i>new_systemname</i> is the system name or NVRAM ID of the V-Series system to which you are moving the disk shelf.</p> <p>Result: Entering this command automatically removes ownership of the disks from the FAS system and assigns them to the V-Series system.</p>
5	<p>Enter the following command and review the list of disks to make sure that all disks on the disk shelf that you plan to move to the V-Series system are now owned by this V-Series system:</p> <p>disk show -v</p>
6	Disconnect the disk shelf from the FAS system, following the instructions in the guide for your disk shelf.

Process overview for physically moving the disk shelf

The following table provides an overview of the process for physically moving the shelf and tells you where to find the procedure for how to do each step.

Step	Action	Where to find instructions
1	<p>On the FAS system, disconnect the disk shelf.</p> <p>Note_____</p> <p>You must shut down the FAS system to remove the disk shelf from it.</p> _____	See the appropriate guide for your disk shelf.
2	<p>Connect the disk shelf to the V-Series controller, ground the Fibre Channel disk shelves and the controller, and connect the disk shelves to the power source.</p> <p>Note_____</p> <p>You do not need to shut down the V-Series system to install a disk shelf on it.</p> _____	<ul style="list-style-type: none">◆ For a stand-alone system, see the hardware <i>Installation and Setup Overview</i> for your platform and the disk shelf guide for your disk shelf.◆ For an HA pair, see the Data ONTAP <i>High Availability Configuration Guide</i> or Data ONTAP <i>Active/Active Configuration Guide</i>.
3	Set up the disk shelf	See the appropriate guide for your disk shelf.

Optional configuration after moving the disk shelf

Turning on automatic disk assignment

Automatic disk assignment is an optional feature that enables Data ONTAP to look for and automatically assign any unassigned disks to same system and pool as all other disks on the same loop.

On V-Series systems, automatic disk assignment is available only for disks on the disk shelves connected to it. Automatic disk assignment is not supported for assigning array LUNs to V-Series systems.

When to reenable automatic disk assignment: Automatic disk assignment is enabled by default. However, when you move a disk shelf from a FAS system to a V-Series system, you must disable automatic disk autoassignment if you want to preserve the data on the disk shelf. If you want the automatic disk assignment functionality after the shelf is moved, you must reenable it after you finish moving the disk shelf and then reassigning the disks on the disk shelf to the V-Series systems.

Reenabling automatic disk assignment: To turn automatic disk assignment back on, complete the following step.

Step	Action
1	<div>For Data ONTAP 8.0 7-Mode and 7.x releases, enter the following command: <code>options disk.auto_assign off</code> For Data ONTAP 8.0 Cluster-Mode, enter the following command: <code>run -node node options disk.auto_assign off</code></div>

Completing setup for Multipath Storage

If you are using Multipath Storage with an HA pair, you need to cable it and finish configuring it after you move the disk shelf to the V-Series system. See Appendix B, “[Setting up and Managing Multipath Storage](#),” on page 71 for instructions if you want to use Multipath Storage with a disk shelf on a V-Series system.

Attention

On the V-Series system, you must set the environment variable required for the V-Series for the first disk shelf in each loop before you connect the disk shelf to the V-Series system.

Moving a disk shelf that is installed on a V-Series system

You can move a disk shelf that is installed on a V-Series system to another V-Series system or to a FAS system if the destination system supports the disk shelf type.

About this chapter

This chapter helps you distinguish between data about disks and array LUNs in Data ONTAP command output and to inform you about any differences in managing disks and array LUNs.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Identifying disk and array LUN information in command output”](#) on page 60
- ◆ [“Assigning disks and array LUNs to the V-Series system”](#) on page 64

Where to find additional information about managing disks

For detailed information about managing disks on Data ONTAP systems, see the Data ONTAP *Storage Management Guide*.

Identifying disk and array LUN information in command output

Format of disk and array LUN addressing

Information about native disks and array LUNs is shown in output from some Data ONTAP commands—for example, in the output from the `storage show disk` command. The format of the “disk addressing” is different, depending on whether the information shown is an array LUN or a native disk.

Format for disk names: Names of disks on native disk shelves are in the following format: *HA.disk_id*.

In Data ONTAP 8.0 Cluster-Mode, the node name is specified:
nodex:*HA.disk_id*.

HA represents the host adapter number, which is the slot number on the V-Series system to which the host adapter is attached, as shown in the following examples:

- ◆ 0a—for a disk shelf attached to an onboard Fibre Channel port
- ◆ 7a—for a shelf attached to the A port of a host adapter installed in slot 7

Example: 0a.16

Format for array LUNs: Array LUN name format is as follows.

Direct-attached configuration	Fabric-attached configuration
<p>For Data ONTAP 8.0 7-Mode and 7.x releases:</p> <p><i>adapter.id.lun-id</i></p> <p><i>adapter</i> is the host adapter number on the V-Series system.</p> <p><i>id</i> is the adapter port on the storage array.</p> <p><i>lun-id</i> is the full LUN ID.</p> <p>For Data ONTAP 8.0 Cluster-Mode:</p> <p><i>node-name</i> is the name of the Cluster-Mode node. The node name is prepended to the LUN name so that the path-based name is unique within the cluster.</p>	<p>For Data ONTAP 8.0 7-Mode and 7.x releases:</p> <p><i>switch-name:port.id.lun-id</i></p> <p><i>switch-name</i> is the name of the switch.</p> <p><i>port</i> is the switch port.</p> <p><i>id</i> is the device ID.</p> <p><i>lun-id</i> is the LUN ID.</p> <p>Example: s2:4.12614 or s2:4.12714 (depending on the type of switch)</p> <p>For Data ONTAP 8.0 Cluster-Mode:</p> <p><i>node-name</i> is the name of the Cluster-Mode node. The node name is prepended to the LUN name so that the path-based name is unique within the cluster.</p> <p>Example: nodex:s2:4.12614 or s2:4.12714</p>

storage show disk -T output

You can use the `storage show disk` (use the node shell command on Data ONTAP Cluster-Mode systems: run `-node node`) command to display attribute information about all the disks and array LUNs connected to your V-Series system. When the `-T` option is used with the `storage show disk` command, the “disk” type is shown. (If your output from the `storage show disk -T` command is truncated, try adding the `-x` option either before or after `-T`—for example, `storage show disk -T -x`.)

In the following `storage show disk -T` command output, the topmost entries are for Fibre Channel disks and the bottom entries are for array LUNs.

DISK	SHELF	BAY	SERIAL	VENDOR	MODEL	REV	TYPE

a.16	1	0	V5Y6GGAA	NETAPP	X274_HPYTA146F10	NA03	FCAL
a.17	1	1	V5Y6E1RA	NETAPP	X274_HPYTA146F10	NA03	FCAL
a.18	1	2	3KS63VLY00009728WRQU	NETAPP	X274_S10K7146F10	NA07	FCAL
a.19	1	3	3HY8P9N6000075033RAG	NETAPP	X274_SCHT6146F10	NA05	FCAL
a.20	1	4	V5Y6GWAA	NETAPP	X274_HPYTA146F10	NA03	FCAL
a.21	1	5	V5Y6GPGA	NETAPP	X274_HPYTA146F10	NA03	FCAL
a.22	1	6	V5Y6L0BA	NETAPP	X274_HPYTA146F10	NA03	FCAL
a.23	1	7	3KS60HHB00009728VFEY	NETAPP	X274_S10K7146F10	NA07	FCAL
a.24	1	8	V5Y6DM3A	NETAPP	X274_HPYTA146F10	NA03	FCAL
a.25	1	9	V5Y694KA	NETAPP	X274_HPYTA146F10	NA03	FCAL
a.26	1	10	V5Y6GNXA	NETAPP	X274_HPYTA146F10	NA03	FCAL
a.27	1	11	V5Y6GRLA	NETAPP	X274_HPYTA146F10	NA03	FCAL
a.28	1	12	V5Y61XVA	NETAPP	X274_HPYTA146F10	NA03	FCAL
a.29	1	13	V5Y6E5YA	NETAPP	X274_HPYTA146F10	NA03	FCAL
c.16	1	0	V5Y6HA2A	NETAPP	X274_HPYTA146F10	NA03	FCAL
c.17	1	1	3KS60J7500009728WA4B	NETAPP	X274_S10K7146F10	NA07	FCAL
c.18	1	2	V5Y6E86A	NETAPP	X274_HPYTA146F10	NA03	FCAL
c.19	1	3	3KS638PR00009728VEUS	NETAPP	X274_S10K7146F10	NA07	FCAL
c.20	1	4	V5Y6A12A	NETAPP	X274_HPYTA146F10	NA03	FCAL
c.21	1	5	V5Y6HA3A	NETAPP	X274_HPYTA146F10	NA03	FCAL
c.22	1	6	3KS60JDL00009728WAF2	NETAPP	X274_S10K7146F10	NA07	FCAL
c.23	1	7	V5Y6E0ZA	NETAPP	X274_HPYTA146F10	NA03	FCAL
c.24	1	8	V5Y63XRA	NETAPP	X274_HPYTA146F10	NA03	FCAL
c.25	1	9	V5Y6DR0A	NETAPP	X274_HPYTA146F10	NA03	FCAL
c.26	1	10	3KS63SC500009728WSKZ	NETAPP	X274_S10K7146F10	NA07	FCAL
c.27	1	11	V5Y6EPXA	NETAPP	X274_HPYTA146F10	NA03	FCAL
c.28	1	12	3KS60HTJ00009728WBMR	NETAPP	X274_S10K7146F10	NA07	FCAL
c.29	1	13	3KS60GR100009728WBQ2	NETAPP	X274_S10K7146F10	NA07	FCAL
ss_b3800_2:0.126L0	-	-	751373110FE	IBM	2107900	0000	LUN
ss_b3800_1:0.126L1	-	-	751373110FF	IBM	2107900	0000	LUN
ss_b3800_1:0.126L2	-	-	75137311200	IBM	2107900	0000	LUN
ss_b3800_2:0.126L3	-	-	75137311201	IBM	2107900	0000	LUN
ss_b3800_2:0.126L4	-	-	75137311202	IBM	2107900	0000	LUN
ss_b3800_1:0.126L5	-	-	75137311203	IBM	2107900	0000	LUN
ss_b3800_1:0.126L6	-	-	75137311204	IBM	2107900	0000	LUN
ss_b3800_2:0.126L7	-	-	75137311205	IBM	2107900	0000	LUN
ss_b3800_2:0.126L8	-	-	75137311206	IBM	2107900	0000	LUN
ss_b3800_1:0.126L9	-	-	75137311207	IBM	2107900	0000	LUN

The following table describes the columns in the storage show disk -T output.

Column	Description
Disk	This is the disk or array LUN “address.” See “Format of disk and array LUN addressing” on page 60.
Shelf	These columns are relevant only for native disks.
Bay	
Serial	<ul style="list-style-type: none"> ◆ For a native disk: This is the disk serial number. ◆ For an array LUN: This is the LUN serial number. Typically the LUN serial number is a combination of a unique identifier for the storage array and the internal logical device number by which the storage array identifies the LUN.
Vendor	<ul style="list-style-type: none"> ◆ For disks: This column always shows NETAPP. ◆ For array LUNs: This column shows the product ID.
Model	<ul style="list-style-type: none"> ◆ For disks: This is the disk model number. ◆ For array LUNs: This is the product ID, which is supplied by Data ONTAP.
Rev	<ul style="list-style-type: none"> ◆ For disks: The disk firmware revision is shown. ◆ For array LUNs: Data ONTAP is not able to obtain the information about storage array firmware.
Type	<p>Shows the object type, as follows:</p> <ul style="list-style-type: none"> ◆ For disks: This column shows the disk type. ◆ For array LUNs: This column shows LUN.

Assigning disks and array LUNs to the V-Series system

Assigning ownership of disks, array LUNs, or both

Automatic disk assignment is enabled for disks by default. If you have disabled this feature, you must manually assign disk ownership for disks. Automatic disk ownership does not apply to array LUNS. To assign ownership of disks, array LUNs, or both to the V-Series system, complete the following steps.

Note

With the `disk assign` command, using the `-T` parameter in conjunction with the `-n count` parameter enables you to specify the number of disks of a particular type—or array LUNs—that you want to assign. Adding the `-x` parameter aligns the columns in the output—for example, `disk assign -T disktype -x`.

Step	Action
1	<p>Note</p> <p>For Data ONTAP 8.0 Cluster-Mode, use the node shell command to enter the following 7-Mode commands: run -node node</p> <hr/> <p>Use the following command:</p> <pre>disk assign {<i>disk_name</i> all -n <i>count</i> -T <i>disktype</i>} [-p <i>pool</i>] [-c <i>block zoned</i>]</pre> <p><i>disk_name</i> <i>all</i> <i>-n count</i>: Required option. Identifies the disks and array LUNs assigned to this V-Series system. Parameters are as follows:</p> <ul style="list-style-type: none"> ◆ <i>disk_name</i> specifies, by disk name or array LUN name, the disks or array LUNs to be assigned. In most cases, you identify the name of the specific disks or array LUNs that this V-Series system is to own in either of following formats: <pre>disk assign fullLUNname1 fullLUNname2 ... disk assign fulldiskname1 fulldiskname2 ...</pre> ◆ <i>all</i> causes all disks and array LUNs that were made available on the storage array to V-Series systems to be assigned to this V-Series system. ◆ <i>-n count</i> causes the number of unassigned disks and array LUNs specified by <i>count</i> to be assigned to this V-Series system. If, for example, you enter <code>disk assign -n 5</code>, Data ONTAP assigns three disks and two array LUNs. ◆ <i>-T</i>, when used with <i>-n count</i>, specifies the “disk type.” Valid types for the <i>-T</i> option are ATA, EATA, FCAL, LUN, SAS, SATA, SCSI, XATA, XSAS.

Step	Action
1 (cont)	<p>◆ -p <i>pool</i></p> <hr/> <p>Note</p> <p>For Data ONTAP 8.0 Cluster-Mode, the pool is not specified because SyncMirror is not supported.</p> <hr/> <p>Data ONTAP assigns pools automatically for disks. For array LUNs, the requirements for the -p parameter are as follows:</p> <ul style="list-style-type: none"> ❖ If you are <i>not</i> going to be deploying the SyncMirror feature, you do not need to enter the pool parameter (-p). The pool parameter defaults to 0 (zero). ❖ If you <i>are</i> going to be deploying SyncMirror, you must use the pool parameter with array LUNs, and you must enter it correctly. See the information about RAID mirroring in the <i>V-Series Installation Requirements and Reference Guide</i> for details. <p>Command example:</p> <pre>disk assign storeAlun1 storeAlun2 -p 0 disk assign storeBlun1 storeBlun2 -p 1</pre> <p>Typically, pool 0 is assigned to the local pool and pool 1 is assigned to the remote pool.</p>

Step	Action
1 (cont)	<p>◆ [-c block zoned]</p> <p>Checksum type can be specified only for array LUNs. The checksum type can be either block or zoned. Block is the default value.</p> <hr/> <p>Note</p> <p>If you use the <code>all</code> parameter, Data ONTAP ignores any checksum type specified. Only block checksum types are supported for disks and the <code>all</code> parameter picks both disks and array LUNs.</p> <hr/> <p>When a LUN from the storage array is mapped to be used by a V-Series system, the V-Series system treats it as a “raw,” unformatted disk (LUN). This “raw” array LUN needs to be formatted by the V-Series system in either of two ways:</p> <ul style="list-style-type: none"> ❖ Block checksum: Reserves 12.5 percent of the space in the array LUN for checksum ❖ Zoned checksum <p>An array LUN of type “block” is expected to be used with block checksum aggregates. An array LUN of type “zoned” is expected to be used with zoned checksum aggregates. See the V-Series <i>Installation Requirements and Reference Guide</i> for details about and planning for checksums.</p>
2	<p>Enter the following command:</p> <pre>disk show -v</pre> <p>Result: The Local System ID is shown and the disks and array LUNs exported on the port are shown.</p>
3	<p>Confirm that all disks and array LUNs that you expected to be assigned to this V-Series system (the local system) are shown with the ID of this V-Series system.</p>

Why you need to check the port settings

The factory might configure some V-Series HBAs as target ports. However, most V-Series systems do not have targets. Therefore, before you start setup, you need to ensure that the V-Series HBAs that you are using to access both storage array LUNs and disks are configured as initiators.

On V3xxx systems, for example, typical V-Series installations use 0a and 0c as a port pair or 0b and 0d as a port pair, for redundancy. The factory might not have set the initiator port pairs in this manner.

Changing target ports to initiator ports

Complete the following steps to ensure that V-Series HBAs are configured as initiators.

Step	Action
1	Power on the V-Series system and interrupt the boot process by pressing Ctrl-C when you see the following message on the console: <code>Starting Press CTRL-C for special boot menu</code>
2	Specify <code>Maintenance Mode</code> from the boot menu. Do not proceed any further with V-Series system installation and setup at this time.
3	Enter the following command to see which ports are configured as target ports: <code>fcadmin config</code>
4	For each target port that you want to change to an initiator port, enter the following command: <code>fcadmin config -t initiator port#</code>

About this appendix This appendix provides information specific to V-Series about how to set up and manage Multipath Storage for disk shelves connected to a V-Series HA pair. Use the Data ONTAP *High Availability Configuration Guide* or Data ONTAP *Active/Active Configuration Guide* for information about cabling and configuration for Multipath Storage that pertains to both FAS systems and V-Series systems.

Note

The Multipath Storage feature is not used with third-party storage arrays.

Topics in this appendix

This appendix contains the following topics:


- ◆ [“Description of Multipath Storage”](#) on page 72
- ◆ [“V-Series specific requirements for Multipath Storage”](#) on page 73
- ◆ [“Setup overview—Multipath Storage on an HA pair”](#) on page 75
- ◆ [“Configuring the V-Series environment variable for Multipath Storage”](#) on page 77

Description of Multipath Storage

Multipath Storage described

The optional Multipath Storage feature for HA pairs provides two paths from each controller to every disk shelf. Setting up Multipath Storage V-Series systems requires configuring an environment variable and adding cables. (See [“Setup overview—Multipath Storage on an HA pair”](#) on page 75.)

The following illustration shows two nodes using Multipath Storage cabled to a single disk shelf.



Advantages of Multipath Storage for HA pair

The advantages of Multipath Storage are as follows:

- ◆ The loss of a disk shelf module, connection, or host bus adapter (HBA) does not require a failover. The same node can continue to access the data using the redundant path.
- ◆ The loss of a single disk shelf module, connection, or HBA does not prevent a successful failover. The takeover node can access its partner's disks using the redundant path.
- ◆ You can replace modules without having to initiate a failover.

Connection types

See the Data ONTAP *High Availability Configuration Guide* or Data ONTAP *Active/Active Configuration Guide* for details about the types of connections with Multipath Storage.

V-Series specific requirements for Multipath Storage

System requirements

Multipath Storage is supported on all V-Series systems that support native disk shelves. See the V-Series *Support Matrix* for a list of the V-Series systems supported.

HA pair type requirements

Multipath Storage is available for the following types of V-Series HA pairs:

- ◆ Standard HA pairs
- ◆ Mirrored HA pairs

Note

V-Series does not support MetroCluster with native disk shelves.

Disk shelf requirements

Note

See the V-Series *Support Matrix* for a list of the combinations of disk shelves and modules supported in HA pairs with Multipath Storage. Only AT-FCX modules shipped in December 2005 or later support Multipath Storage for HA pairs. If you are unsure whether your module is the correct version, see the Data ONTAP *High Availability Configuration Guide* or Data ONTAP *Active/Active Configuration Guide* for the procedure for checking the version of the AT-FCX module.

General requirements

The following table summarizes the general requirements for V-Series with Multipath Storage.

For this...	The requirement is...
V-Series port requirements	<ul style="list-style-type: none">◆ Two Fibre Channel ports for each loop◆ Enough onboard Fibre Channel ports or HBAs on each node to accommodate the extra cables required for Multipath Storage.
V-Series boot environment variable setting	<p>Configure the <code>fc-non-array-adapter-list</code> environment variable for each new loop <i>before</i> you connect and configure the shelf for Multipath Storage.</p> <p>See “Setup overview—Multipath Storage on an HA pair” on page 75 for instructions.</p>
Number of loops to be cabled	<p>For maximum protection, it is recommended that all loops are cabled for Multipath Storage. If you have a mixed configuration in which some loops are cabled for Multipath Storage and some are not, the system displays a configuration error message when you boot the system or when a disk on a loop that is cabled for Multipath becomes single-pathed.</p>

Setup overview—Multipath Storage on an HA pair

The following table summarizes the high-level tasks necessary to use Multipath Storage with a V-Series system with disk shelves. Perform the tasks in the order shown in the table.

Step	Action	Location of detailed information
Prerequisites to setting up Multipath Storage		
1	<p>Install your HA pair and ensure that it is fully operational before performing any activities to set up Multipath Storage.</p> <hr/> <p>Note</p> <p>When testing your basic HA pair, be sure that failover and giveback is successful.</p> <hr/>	Data ONTAP <i>High Availability Configuration Guide</i> or Data ONTAP <i>Active/Active Configuration Guide</i>
2	<p>If the V-Series systems in the HA pair are not already running Data ONTAP 7.3 or higher, you must upgrade your software.</p>	Data ONTAP <i>Upgrade Guide</i>
V-Series specific Multipath Storage setup requirement		
3	<p>On each node, configure the <code>fc-non-array-adapter-list</code> environment variable for the first disk shelf on a new loop to the V-Series system.</p> <p>When you configure the <code>fc-non-array-adapter-list</code> environment variable, you explicitly enter the V-Series initiator port used for every path to the first disk shelf on the loop. If you do not perform this step, the V-Series systems in the HA pair panic.</p> <hr/> <p>Attention</p> <p>You must configure the <code>fc-non-array-adapter-list</code> environment variable <i>before</i> you connect the first disk shelf on a new loop to the V-Series system. This step is not necessary for FAS systems.</p> <hr/>	“Configuring the V-Series environment variable for Multipath Storage” on page 77
General instructions for completing Multipath Storage setup (these instructions pertain to V-Series systems and FAS systems)		

Step	Action	Location of detailed information
4	Connect the additional cables necessary for redundant connections.	Data ONTAP <i>High Availability Configuration Guide</i> or Data ONTAP <i>Active/Active Configuration Guide</i>
5	Configure Data ONTAP to add the new Multipath Storage.	Data ONTAP <i>High Availability Configuration Guide</i> or Data ONTAP <i>Active/Active Configuration Guide</i>
Completing your setup		
6	Assign the new disks to the V-Series system if you do not have automatic disk assignment enabled. Automatic disk assignment, which pertains only to native disks, is enabled by default.	“Assigning disks and array LUNs to the V-Series system” on page 64 (if you are manually assigning the disks)

Configuring the V-Series environment variable for Multipath Storage

When you need to set the environment variable

The following table summarizes when you need to set the V-Series environment variable for Multipath Storage.

Attention

If you need to set the environment variable for Multipath Storage, be sure to set the variable on each node before you connect the new disk shelf to it.

If you are...	Then...
Adding a disk shelf to a new loop	You must set this environment variable on both nodes before connecting the disk shelf to the V-Series initiator ports.
Adding a disk shelf to an existing loop (daisy-chaining)	You do not need to configure the V-Series environment variable for Multipath Storage again for the loop.

Setting the environment variable

This procedure describes how to set the V-Series required environment variable, called `fc-non-array-adapter-list`, to identify the V-Series initiator ports that you are connecting to the disk shelf. To set the environment variable for Multipath Storage for V-Series, complete the following steps.

Note

In the following procedure, assume that the nodes in the HA pair are named `vs1` and `vs2`.

Step	Action
1	<p>Verify that the FC ports on both vs1 and vs2 are set as initiators.</p> <p>See Appendix A, “Configuring V-Series Ports as Initiators,” on page 69 for instructions.</p>
To set the environment variable on vs2	
2	<p>On vs1, enter the following command to force vs1 to take over vs2:</p> <p>cf takeover</p>
3	<p>On vs2, interrupt the boot process by pressing Ctrl-C when you see the following message on the console:</p> <p>Starting AUTOBOOT press Ctrl-C to abort</p>
4	<p>On vs2, enter the following command to check which FC initiator ports are already specified:</p> <p>printenv fc-non-array-adapter-list</p> <p>If no FC initiator ports have been specified, the following appears on the screen:</p> <pre>Variable Name Value ----- fc-non-array-adapter-list ** Undefined ***</pre>

Step	Action
5	<p>On vs2, enter the following command:</p> <pre>setenv fc-non-array-adapter-list "adapter,adapter"</pre> <p>For example, if you are using V-Series initiator ports 0b and 0d to connect to the disk shelf, you would enter the command as follows:</p> <pre>setenv fc-non-array-adapter-list "0b,0d"</pre> <p>The following message appears:</p> <pre>*** Variable(s) modified but NOT saved until the OS is booted ***</pre> <p>The ports you specified are not saved until after you reboot your V-Series system.</p> <p>Note</p> <p>Be sure to put the quotes around the initiator ports. If you include a space between initiator port names and you have not enclosed the ports in quotes, then Data ONTAP ignores the initiator port names after the space.</p>
6	<p>On vs2, enter the following command to verify that the FC initiator ports you specified are set:</p> <pre>printenv fc-non-array-adapter-list</pre> <p>If you entered ports 0b and 0d, as shown in the previous step, the printenv output should be as follows:</p> <pre>Variable Name Value ----- fc-non-array-adapter 0b,0d</pre>
7	<p>Reboot vs2 using the appropriate boot command for your model—for example, <code>boot_ontap</code>.</p>
8	<p>Wait until vs2 is ready for giveback, then proceed to the next step.</p> <p>When the following message on the console of vs2 appears, you can continue to the next step:</p> <pre>Waiting for giveback...(Press Ctrl-C to abort wait)</pre>

Step	Action
9	On vs1, enter the following command to start giveback to vs2, then wait for vs2 to boot fully: cf giveback
To set the environment variable on vs1	
10	On vs2, enter the following command to force vs2 to take over vs1: cf takeover
11	Repeat the same steps you followed for setting the environment variable for vs2, but this time for vs1.

About this appendix This appendix provides information to help you troubleshoot V-Series configurations that use native disk shelves.

Topics in this appendix This appendix contains the following topics:

- ◆ “[Common problems](#)” on page 82
- ◆ “[Error messages](#)” on page 84
- ◆ “[Reinstalling the system](#)” on page 86

Determining which storage is affected by the problem If you encounter a storage-related issue, it is important to first determine whether the problem is with the native disk shelves or with the array LUNs.

Common problems

The following table shows how to resolve some common problems with a V-Series system that uses native disk shelves.

Problem	Explanation and solution
You are not aware that Data ONTAP automatically assigns new disks to a system by default.	<p>If you do not want the V-Series system to take ownership of the disks immediately, disable automatic disk assignment through the <code>disk.auto_assign</code> option.</p> <hr/> <p>Note— Disk autoassignment does not apply to array LUNs. You must explicitly assign array LUNs to your system.</p> <hr/>
Not all disks are visible.	<ul style="list-style-type: none">◆ Verify that the cables are plugged in.◆ Verify that the disk modules are properly seated.◆ Verify that there is not a mismatch between the disk shelf and components in the same loop. <p>For example, if the disk drives are limited to 2 Gb and you set the ESH4 module to 4 Gb, the port comes online but the V-Series system cannot see the disks.</p>
There is a loop speed mismatch between components on the same loop.	<p>The Gb-loop speed for all components in the loop must be set to the lowest of any component in the loop. Be sure that you confirm the speed of <i>all</i> components in the loop—including the disk shelves. See the guide for your disk shelf model for information about the characteristics of your disk shelf.</p>

Problem	Explanation and solution
Disk shelf IDs are duplicated.	<p>Each Fibre Channel and ATA disk shelf in a loop must have a unique ID between 1 and 6.</p> <p>The behavior when disk shelf IDs are duplicated depends on how many disk shelves are in the loop.</p> <p>If there are fewer than seven disk shelves in the loop, Data ONTAP assigns a soft address to each disk drive on the shelves whose disk shelf IDs are duplicated. Data ONTAP also issues the following message:</p> <pre>[fci.device.invalidate.soft.address:error]</pre> <p>See the guide for your disk shelf for information about how to change the disk shelf ID.</p>

Error messages

This section includes some error messages for problems commonly seen when users new to native disk shelves on V-Series systems set up their systems. This section is not meant to describe all possible reasons why a particular error message is generated. The Syslog Translator on the NOW™ site (<http://now.netapp.com>) can give you more information about error messages.

Message	Explanation
<pre>PANIC: sanown: received add for already known disk in process sanown_notify on release NetApp Release <release name> on <date and timestamp></pre>	<p>If you are using Multipath Storage for disks on a V-Series HA pair, and you do not set the V-Series required boot environment variable (<code>fc-non-array-adapter-list</code>), this message is issued and the V-Series system panics. See “Configuring the V-Series environment variable for Multipath Storage” on page 77 for instructions.</p> <p>Note _____ This message is issued in other circumstances that are not related to Multipath Storage on V-Series—for example, when array LUN IDs are duplicated.</p>
<pre>[cf.nm.nicTransitionDown:warning]: Cluster Interconnect link 0 is DOWN Tue Oct 16 17:25:28 GMT [cf.rv.notConnected:error]: Connection for cfo_rv failed Tue Oct 16 17:25:48 GMT [esh.auto.term.error:error]: AUTO TERMINATION ERROR: TERM switch set improperly on ESH A on channels 0b/0d disk shelf ID 1. Tue Oct 16 17:25:48 GMT [esh.speed.error:error]: Disk shelf previously connected to the OUT port of ESH A on channels 0b/0d disk shelf ID 1 is now inaccessible. (Check cables, speed selection.)</pre>	<p>You might see this error message when there is a loop speed mismatch for components on the same loop. The Gb-loop speed for all components must be set to the lowest of any component in the loop.</p>

Message	Explanation
<p>Tue Jan 8 18:54:59 GMT [fcid.device.invalidate.soft.address:error]: Fibre Channel adapter 0c is invalidating disk drive 0c.14 (0x0300000e) which appears to have taken a soft address. Expected hard address 26 (0xba), assigned soft address 14 (0xd1).</p>	<p>This error message is generated when there is a disk shelf ID conflict.</p> <p>See the guide for your disk shelf for information about how to change the disk shelf ID.</p>

Reinstalling the system

Default selection of the root volume location

If technical support advises you to completely reinstall your system because of a problem, you need to follow the appropriate procedure so that the root volume is installed where you want it. When no aggregates or volumes are on the native disk shelves attached to the V-Series system, Data ONTAP is optimized to select a disk for the root volume. Before you start reinstalling your system, you need to decide whether this default behavior for selection of the location of the root volume is the right approach for your system. See “[Choices for root volume installation](#)” on page 86 before you start installation.

Choices for root volume installation

The following summarizes the choices you have regarding root volume installation when you need to reinstall the system, and where to find the procedures that you need for each choice.

You need to decide how you want the disk for the root volume to be selected:

- ◆ If you want to specify the disk for the root volume, see “[Choosing the disk for the root volume](#)” on page 87
- ◆ If you want Data ONTAP to select the disk automatically, see “[Allowing Data ONTAP to choose the disk for the root volume](#)” on page 86

Allowing Data ONTAP to choose the disk for the root volume

If you do not have a preference for which disk to use for the root volume, you can let Data ONTAP automatically select the disk. Data ONTAP automatically assigns three disks and selects one of them for the root volume.

Note _____
Data ONTAP requires three disks for RAID-DP, which is the default RAID type for disks starting in Data ONTAP 7.3.

To let Data ONTAP automatically select the disk for the root volume, complete the following steps.

Step	Action
1	Gather information that you need for the setup script, using the setup parameters worksheet in the <i>Data ONTAP Software Setup Guide</i> .

Step	Action
2	Boot the V-Series system and interrupt the boot process by pressing Ctrl-C when you see the following message on the console: Press Ctrl-C for Boot Menu
3	Create the root volume by selecting the following menu item: Clear configuration and initialize root volume
4	Enter y when prompted as to whether you want to install a new file system.
5	Enter y to the following message: This will erase all the data on the disks, are you sure? Result: The V-Series system creates a root volume and begins prompting for system information.
6	Answer the prompts in the setup program, either through the command line or the Web interface, using the information that you recorded in the setup parameters worksheet in the <i>Data ONTAP Software Setup Guide</i> .
7	After the setup script is complete, do the following: <ul style="list-style-type: none"> ◆ Assign additional disks, array LUNs, or both to the V-Series system (see “Assigning disks and array LUNs to the V-Series system” on page 64). ◆ Install protocol and features licenses. ◆ Install and download Data ONTAP software, following the instructions in the <i>Data ONTAP Software Setup Guide</i>.
8	If you have an HA pair, repeat this same procedure on the other node.

Choosing the disk for the root volume

If you want to choose the disk for the root volume, complete the following steps.

Step	Action
1	Gather information that you need for the setup script, using the setup parameters worksheet in the <i>Data ONTAP Software Setup Guide</i> .
2	Boot the V-Series system and interrupt the boot process by pressing Ctrl-C when you see the following message on the console: Press Ctrl-C for Boot Menu

Step	Action
3	Enter <code>Maintenance Mode</code> from the boot menu.
4	<p>Enter the following command to see the disks on the disk shelf so that you can decide which disk you want to use for the root volume.</p> <pre>disk show -v</pre> <p>Result: A list of all disks and array LUNs that are visible to the V-Series system is displayed.</p>
5	<p>Use the <code>disk assign</code> command to assign at least three disks on the disk shelf that you want this V-Series system to own.</p> <hr/> <p>Note</p> <p>You must assign at least three disks to this V-Series system before you can create the root volume because Data ONTAP requires three disks for RAID-DP, the default RAID type.</p> <hr/> <p>See “Assigning disks and array LUNs to the V-Series system” on page 64 for information about the parameters that are available for the <code>disk assign</code> command.</p>
6	<p>Enter the following command:</p> <pre>disk show</pre>
7	Review the list of disks in the output to confirm that all the disks you expected to be assigned to this V-Series system are shown with the ID of this V-Series system.
8	Enter halt .
9	<p>Boot the V-Series system and interrupt the boot process by pressing Ctrl-C when you see the following message on the console:</p> <pre>Press Ctrl-C for Boot Menu</pre>
10	<p>Create the root volume by selecting the following menu item:</p> <pre>Clear configuration and initialize root volume</pre>
11	Enter y when prompted as to whether you want to install a new file system.

Step	Action
12	<p>Enter y to the following message:</p> <p>This will erase all the data on the disks, are you sure?</p> <p>Result: The V-Series system creates a root volume and begins prompting for system information.</p>
13	<p>Answer the prompts in the setup program, either through the command line or the Web interface, using the information that you recorded in the setup parameters worksheet in the <i>Data ONTAP Software Setup Guide</i>.</p>
14	<p>After the setup script is complete, do the following:</p> <ul style="list-style-type: none"> ◆ Assign additional disks, array LUNs, or both to the V-Series system (see “Assigning disks and array LUNs to the V-Series system” on page 64). ◆ Install protocol and features licenses. ◆ Install and download Data ONTAP software, following the instructions in the <i>Data ONTAP Software Setup Guide</i>.
15	<p>If you have an HA pair, repeat this same procedure on the other node.</p>

Index

A

- adding a disk shelf 35, 38
- aggregates
 - checking names of 49
 - creating
 - guidelines for 21
 - information sources 23
 - layout considerations 46
 - renaming 50
 - resolving duplicate names 47, 48, 50
 - restrictions on mixing storage types 21
 - SnapLock Compliance 22
- array LUNs
 - assigning ownership 64
 - checksum type with 16
 - defined 1
 - format of LUN ID 60
 - RAID types supported for 12
 - sizing 16
 - V-Series systems striping across 12
- automatic disk assignment
 - reenabling 56
 - troubleshooting 82
 - turning it off
 - how to 52
 - when required 47, 48
 - what supported for 34, 82
 - when it does not apply 64
 - when manual assignment is necessary 64

B

- block checksums
 - description 12
 - with array LUNs 16
 - with disk shelves 16
- boot environment variable
 - Multipath Storage 49, 74, 84
- booting the V-Series system, first time
 - how to 33
 - prerequisites 32
- Bootloader files

system firmware updates 31

C

- cabling of a new system
 - HA pair 30
 - stand-alone 30
- checksum type
 - assigning to array LUNs 67
 - block 12
 - with array LUNs 16
 - with native disks 16
 - zoned 12
- common problems 82
- CompactFlash card
 - Bootloader files on 31
 - firmware updates on 31
- connections
 - to a storage array 19
 - to disk shelves 19
 - with array LUNs 18
 - with disk shelves 11, 18
- core dump file
 - availability 6, 24
 - how it is used 24
 - in root volume 24

D

- Data ONTAP
 - default software setup 30
 - features for use with array LUNs 16
 - features for use with disk shelves 5, 16
 - minimum version with disk shelves on V-Series 8
- disk assignment
 - automatic
 - reenabling 56
 - turning it off 48, 52
 - when you must turn it off 47
- disk names, format 60
- disk ownership
 - assigning 64

- changing 49, 53
 - scheme defined 10
 - type supported for array LUNs 16
 - type supported for disk shelves 16
- disk reassign command 54
- disk sanitization
 - with array LUNs, not supported 16
 - with disk shelves 16
- disk shelf module, connection to 11
- disk shelves
 - adding
 - availability of system during 36
 - Data ONTAP version required 36
 - high-level steps 38
 - loop speed requirements 36, 44
 - restrictions 38
 - advantages of adding 6
 - checksum type with 16
 - connections on, described 11
 - hot-adding, defined 37
 - how to add to V-Series 5
 - loops on, described 11
 - moving 36, 44
 - aggregate layout considerations 46
 - availability of system during 45
 - Data ONTAP version required 44
 - from a V-Series system 58
 - impact on clients 46
 - optional configuration after 56
 - physically moving, overview 55
 - preserving data 21, 47
 - root volume considerations 45, 48
 - moving to a V-Series system 5
 - planning for 15
 - preserving data, requirements 47
 - RAID types supported for 12
 - root volume requirements 45
 - types you can move 44
 - understanding storage provided by 11
 - ways of using with V-Series systems 6
 - which are supported 10
- disk shelves with V-Series
 - minimum Data ONTAP version 8
- disk sizing, on disk shelves 16
- disk types, which are supported 10

- disk.auto_assign option
 - reenabling 56
 - turning it off 53
 - what happens if enabled 34, 82
- disks
 - assigning ownership 64
 - changing ownership of 49, 53
 - checking names of 49
 - format of disk ID 60
 - viewing information about 61
- duplicate disk shelf IDs
 - symptoms of 83
 - troubleshooting 83

E

- environment variable, required for Multipath Storage 84
- error messages 84

F

- family, defined 21
- fc-non-array-adapter-list environment variable
 - message if not set 84
 - requirement for 74, 75
 - setting 77
- firmware updates (system) 31
- front-end LUN, defined 1

H

- HA pair
 - Multipath Storage
 - advantages of 72
 - completing setup for 56
 - connection types 72
 - described 72
 - for path redundancy 11, 17, 72
 - required step for V-Series 38
 - root volume location recommendation 25
- HBAs (V-Series)
 - configuring as initiators 69
- hot-adding disk shelves, defined 37

I

I/O module 11

installing a V-Series system with disk shelves 32

installing root volume, where 18

installing, simplifying 6

L

license

- for V-Series 17

- on new system with disk shelves 30

- protocol, default installation 30

- requirement for native disk shelves 8, 17

- SnapLock Compliance 22

limits

- for array LUNs 17, 28

- for disks 17, 28

- on V-Series system with disk shelves 17, 28

location of data, planning for 20

loop speed

- common error 36, 44

- requirements 36, 44

loop speed mismatch

- error message generated for 84

loops on disk shelves, described 11

LUN ownership, defined 10, 16

LUN, defined 1

LUNs (array)

- assigning checksum type 67

- checksum type with 16

- name format 60

LUNs. See *array LUNs*, *LUNs (array)*, *front-end LUN*

M

MetroCluster, not supported 8, 17

models supporting disk shelves 8

moving a disk shelf 43

- Data ONTAP version required 44

- from a V-Series system 58

- impact on clients 46

- optional configuration after 56

- physically moving, overview 55

- preserving data 21, 47

- root volume requirements 45

- types you can move 44

Multipath Storage 17

- advantages of 72

- completing setup for 56

- connection types 72

- described 72

- disk shelf requirements 73

- environment variable required for V-Series 30

- fc-non-array-adapter-list environment variable
requirement for 74

- HA pair types supported 73

- number of loops for 74

- required environment variable 57, 84

- required step for V-Series 38

- setup overview 75

- V-Series requirements for 73, 74

N

naming

- disk names, format 60

- format for array LUN IDs 60

native disk shelves

- adding to V-Series system 35

- advantages of adding 6

- checksum type with 16

- connections on, described 11

- defined 2

- loops on, described 11

- minimum Data ONTAP version 8

- moving

 - optional configuration after 56

 - preserving data 21, 47

- moving to V-Series system 43

- planning for 15

- RAID types supported for 12

- understanding storage provided by 11

- ways of using with V-Series systems 6

- which are supported 10

native disks, defined 2

native storage, defined 2

P

panic, due to loop speed mismatch 36, 44

- pathing
 - Multipath Storage feature 17
 - requirements for array LUNs 17
 - requirements for disk shelves 17
 - to the V-Series controller 17
- planning
 - aggregates, guidelines for 21
 - for disk shelves with V-Series systems 15
 - location of data 20
 - storage requirements 20
- ports on V-Series system
 - guidelines for determining usage 27
 - required number to disk shelf 26
 - required number to storage array 26
 - required setting 19, 26
 - restrictions for sharing 26
 - setting as initiators 31, 69
- protocol licenses, default installation 30

R

RAID

- RAID0 17
- type used with array LUNs 17
- type used with disk shelves 17

RAID group size

- with array LUNs 18
- with disk shelves 18

RAID types

- RAID0, described 17
- supported 12
 - RAID4, described 12
 - RAID-DP, described 12, 17

RAID0

- use by V-Series systems 12, 17

RAID4, described 12

RAID-DP, described 12, 17

reassigning disk ownership 53

reinstalling the system 86

root volume

- configuring 87, 88
- considerations when moving a disk shelf 45, 48
- core dump files in 24

- defined 24

location

- choices for system reinstallation 86
- default 30
- on HA pair 25
- options for 24
- where to install 18

system reinstallation

- you choose disk 87

S

- SANOWN, defined 10, 16

simplifying installation 6

SnapLock Compliance

- aggregates 22
- how you can use it 6, 22
- license 22
- SnapLock volumes, defined 22
- when supported 18

SnapLock volumes, defined 22

SnapMover

- with array LUNs 18
- with disk shelves 18

software disk ownership

- requirement for 10, 16
- scheme defined 10

storage commands

- displaying information about disks 61

storage requirements, planning 20

storage show disk command

- output example 61

storage, expanding 6

striping, benefits with RAID0 12

SyncMirror, support for 8

T

target ports, changing to initiators 69

troubleshooting

- automatic disk assignment 82
- common problems 82
- not all disks are visible 82
- when disks are on V-Series system 81

V

volumes

- checking names of 49
- renaming 52
- resolving duplicate names 48, 52
- root volume
 - configuring 87, 88
 - default location 30
 - defined 24
 - location of 18
 - options for location of 24

V-Series neighborhood

- with array LUNs 19

- with disk shelves 19

V-Series system ports

- configuring as initiators 69
- guidelines for determining usage 27
- required number to disk shelf 26
- required number to storage array 26
- required setting 19, 26
- restrictions for sharing 26

Z

zoned checksum disks 12

