




V-Series Systems

Installation Requirements and Reference Guide



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1(408) 822-6000
Fax: +1(408) 822-4501
Support telephone: +1(888) 4-NETAPP
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 210-06053_A0
April 2013

Contents

V-Series technology overview	7
How a V-Series system uses different types of storage	7
Supported methods for connecting a V-Series system to a storage array	8
Number of storage arrays supported behind a V-Series system	8
Sharing storage arrays among hosts	9
Support for MetroCluster configurations with V-Series systems	9
V-Series planning overview	10
Where to find V-Series interoperability and limits information	10
Limits information for configurations with storage arrays	10
Interoperability information about support for storage arrays	11
32xx system configuration restrictions	12
Stages of implementation for a system using array LUNs	12
Planning summary for a V-Series system using native disks	13
Planning summary for a V-Series system using array LUNs	14
Planning for RAID implementation	17
RAID protection for array LUNs	17
Implications of LUN size and number for Data ONTAP RAID groups	17
Planning for Data ONTAP use of array LUNs	19
How array LUNs are made available for host use	19
How array LUNs become available for Data ONTAP storage use	20
Considerations when planning for disk ownership	20
Array LUN assignment changes	21
Array LUN types supported by Data ONTAP	21
Factors that impact the number and size of array LUNs you need	21
Factors impacting the number of array LUNs needed	21
Minimum number of array LUNs required per V-Series system	22
Spare core array LUN requirement for core dumps	23
Minimum and maximum array LUN sizes supported by Data ONTAP	23
Minimum array LUN size for the root volume	24
Elements that reduce the usable space in an array LUN	24
Considering checksum type when planning array LUN size and number	25
Location of the root volume	29

Planning for LUN security on the storage arrays	31
Available LUN security methods	31
Planning for paths to array LUNs	33
Requirements for redundant setup of components in a path	33
When to check for redundant paths to array LUNs	34
Required number of paths to an array LUN	34
Advantages of four paths to an array LUN (clustered Data ONTAP 8.1 and later)	35
Using LUN groups to partition the load over V-Series connections	35
Implementation requirements for a multiple LUN group configuration	36
Example of a configuration with multiple LUN groups	37
How paths are reflected in array LUN names	38
Array LUN name format	39
How the array LUN name changes in Data ONTAP displays	41
Valid pathing: stand-alone system with a single 2-port array LUN group	42
Valid pathing: a single 4-port array LUN group in a fabric-attached configuration	43
Planning a port-to-port connectivity scheme	45
How FC initiator ports are labeled	45
Requirements for V-Series FC initiator port usage	46
Rules for sharing an FC initiator port with multiple target ports	47
Example configuration: shared FC initiator ports	48
Checking the number of array LUNs visible on an FC initiator port (clustered Data ONTAP)	52
What to include in a port-to-port plan for connecting to storage arrays	53
What to include in a port-to-port plan for connecting to native disk shelves	54
Zoning guidelines	56
Zoning requirements	56
Zoning recommendation for a configuration with storage arrays	57
Example of zoning in a configuration with storage arrays	58
Determining the array LUNs for specific aggregates	60
Rules for mixing storage in aggregates for V-Series systems	60
How the checksum type is determined for aggregates with array LUNs	61
Checking the checksum type of spare array LUNs	62
Aggregate rules when the storage arrays are from the same family	63
Aggregate rules when the storage arrays are from different vendors or families	64

Determining whether to use neighborhoods (7-Mode in 8.x)	67
What a V-Series neighborhood is	67
Requirements for V-Series neighborhoods	68
Planning the number of array LUNs and disks in a neighborhood	68
How to establish a neighborhood	70
Calculation of the visible limit for both array LUNs and disks	71
Connecting a V-Series system to a storage array	74
Commands for checking back-end configuration (clustered Data	
ONTAP 8.1 and later)	78
Back-end configuration errors detected by the storage errors show command	
(clustered Data ONTAP 8.1 and later)	80
Situations not identified by the storage errors show command	81
Verifying installation with storage arrays (clustered Data ONTAP	
8.1 and later)	83
Checking for back-end configuration errors preventing system operation	
(clustered Data ONTAP 8.1 and later)	85
The storage errors show messages and their resolution (clustered Data	
ONTAP 8.1 and later)	86
Array LUN is too small or too large	87
LUN IDs for the same LDEV do not match	88
Volume Set Addressing is inconsistent	91
Duplicate LUN IDs on a target port	93
Fewer than two paths to an array LUN	96
An access control LUN was presented	99
All paths to an array LUN are on the same storage array controller	99
Array LUNs are configured with conflicting failover modes (clustered	
Data ONTAP 8.2 and later)	102
Checking whether the configuration matches your intentions (clustered Data	
ONTAP 8.1 and later)	104
Reasons for fewer array LUN groups than expected	106
Reasons for more array LUN groups than expected	107
Reasons for more paths to an array LUN than expected	109
Reasons for the incorrect number of LUNs in array LUN groups	110
Reasons storage arrays are missing from command output	111
Verifying an installation with storage arrays (7-Mode in 8.x)	112
Checking the number of paths (7-Mode in 8.x)	112

Example output showing correct and incorrect pathing (7-Mode in 8.x)	113
Troubleshooting	116
Getting started with troubleshooting	116
Invalid path setup examples	117
Invalid path setup: too many paths to an array LUN (7-Mode)	117
Invalid path setup: alternate paths are not configured	118
What happens when a link failure occurs	120
Relationship between zoning and host group configuration	120
Dependency between zone and host group definitions	120
Example of cascading zoning and host group configuration errors	122
Installation quick start (7-Mode and storage arrays only)	125
Sample configuration for the installation quick start (7-Mode and storage arrays)	125
Performing preinstallation tasks on the storage array	126
Installing the V-Series system	127
Setting up the switches	128
Setting up LUN security	128
Assigning an array LUN to a V-Series system and creating the root volume	129
Installing Data ONTAP and licenses	130
Testing your setup	131
Additional setup	133
Obtaining WWNs manually	134
Settings for connecting to an ASCII terminal console	135
Target queue depth customization	137
Guidelines for specifying the appropriate target queue depth	137
Setting the target queue depth (clustered Data ONTAP prior to 8.2 and 7-Mode) .	138
Setting the target queue depth (clustered Data ONTAP 8.2)	138
Displaying target queue depth statistics (clustered Data ONTAP 8.2)	139
Terminology comparison between storage array vendors	141
Glossary	145
Copyright information	147
Trademark information	148
How to send your comments	149
Index	150

V-Series technology overview

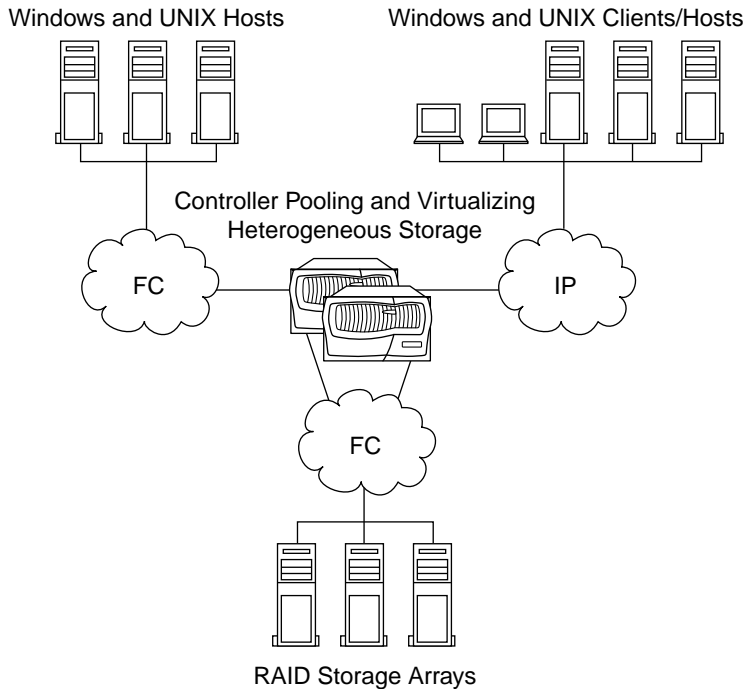
A V-Series system is an open storage controller that virtualizes storage from storage array vendors, native disks, or both into a single heterogeneous storage pool.

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.

How a V-Series system uses different types of storage

A V-Series system can use LUNs from storage arrays only, native disks only, or both.

A V-Series system that uses array LUNs pools the LUNs from storage arrays and serves data to Windows and UNIX hosts and clients.



A V-Series system presents storage to clients either in the form of Data ONTAP file system volumes, which you manage on the system by using Data ONTAP management features, or as a SCSI target that creates LUNs for use by clients. In both cases (file system clients and LUN clients), on the V-Series system you combine the array LUNs into one or more Data ONTAP volumes for presentation

to the clients as files or as LUNs served by Data ONTAP. Likewise, you can combine native disks into one or more Data ONTAP volumes for presentation to clients.

Supported methods for connecting a V-Series system to a storage array

You can incorporate a fabric-attached configuration with storage arrays into an existing FC SAN infrastructure. Fabric-attached configurations are supported for both stand-alone systems and HA pairs. Direct-attached configurations are not supported for new deployments.

For older deployments, direct-attached configurations are limited to some storage arrays and some Data ONTAP releases. The Interoperability Matrix at support.netapp.com contains information about the connection methods supported for specific storage arrays and platforms running Data ONTAP.

Number of storage arrays supported behind a V-Series system

For most storage arrays, you can connect a stand-alone V-Series system or the nodes in an HA pair to multiple storage arrays. For a few storage arrays, you are limited to one storage array behind the V-Series system.

The *Interoperability Matrix* identifies the storage arrays for which only one storage array is supported behind the V-Series system.

If multiple storage arrays behind a V-Series system are supported for your storage array, the rules are as follows:

- There is no limit to the number of storage arrays you can deploy behind your system.
- The storage arrays can be from the same vendor, either all from the same family or from different families.

Note: Storage arrays in the same family share the same performance and failover characteristics. For example, members of the same family all perform active-active failover, or they all perform active-passive failover. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same. The *V-Series Implementation Guide for Third-Party Storage* contains information about the storage array families for each vendor.

- The storage arrays can be from different vendors.

Related concepts

[Rules for sharing an FC initiator port with multiple target ports](#) on page 47

[Rules for mixing storage in aggregates for V-Series systems](#) on page 60

Sharing storage arrays among hosts

A typical storage array provides storage for hosts from different vendors. However, Data ONTAP requires some storage arrays to be dedicated to V-Series systems.

To determine whether your vendor's storage array must be dedicated to V-Series systems, see the *Interoperability Matrix*.

Support for MetroCluster configurations with V-Series systems

The type of support for MetroCluster configurations with V-Series systems differs according to Data ONTAP release.

In Data ONTAP releases prior to 8.2, support for MetroCluster configurations with V-Series systems is as follows:

- The V-Series systems must use *only* LUNs on storage arrays (no native disks).
- The V-Series systems must be 7-Mode systems (not clustered systems).
- The storage arrays and V-Series models must be identified in the *Interoperability Matrix* as being supported in MetroCluster configurations.

Starting in Data ONTAP 8.2, support for MetroCluster configurations with V-Series systems is as follows:

- A MetroCluster configuration with V-Series systems can include native disks only, LUNs on storage arrays only, or both.
Setting up a MetroCluster configuration with V-Series systems requires FibreBridge 6500N bridges and the native disks supported by this bridge. These requirements are the same for fabric-attached and stretch MetroCluster configurations.
- The V-Series systems must be 7-Mode systems (clustered Data ONTAP is not supported with MetroCluster configurations).
- The storage arrays and V-Series models must be identified in the *Interoperability Matrix* as being supported in MetroCluster configurations.

See the *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode* for information about how to plan for and set up a MetroCluster configuration with storage arrays. If your MetroCluster configuration with V-Series systems will include disks, also see the document *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges*.

V-Series planning overview

Successful implementation of a V-Series deployment requires careful planning and verifying proper installation and configuration of all devices in your deployment.

Related concepts

- [Planning for Data ONTAP use of array LUNs](#)* on page 19
- [Planning for LUN security on the storage arrays](#)* on page 31
- [Planning for paths to array LUNs](#)* on page 33
- [Determining the array LUNs for specific aggregates](#)* on page 60
- [Planning a port-to-port connectivity scheme](#)* on page 45

Where to find V-Series interoperability and limits information

When planning your configuration, you need to check sources in addition to product documentation for interoperability details and limits information.

Tools available on the NetApp Support Site provide, in a central location, specific information about which features, configurations, and storage array models are supported in particular releases.

Related concepts

- [Limits information for configurations with storage arrays](#)* on page 10
- [Interoperability information about support for storage arrays](#)* on page 11

Limits information for configurations with storage arrays

The *Hardware Universe* (formerly the *System Configuration Guide*) at support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml contains information about limits that you need to consider when planning for a configuration with storage arrays.

The *Hardware Universe* includes the following limits that do not apply to native disks:

- Minimum and maximum array LUN size that Data ONTAP supports
- Minimum size for the array LUN for the root volume
- Spare core array LUN minimum size
- Limits for RAID groups with array LUNs
- Minimum aggregate size for an aggregate of array LUNs
- Maximum number of array LUNs and disks combined, per platform

The *Hardware Universe* also contains storage limits that are the same for native disks and array LUNs.

Note: Limits for V-Series systems were previously provided in the *V-Series Limits Reference for Third-Party Storage* and the V-Series Support Matrix.

Related concepts

[Interoperability information about support for storage arrays](#) on page 11

Interoperability information about support for storage arrays

Not all Data ONTAP releases support the same features, configurations, storage array models, and V-Series models. During your deployment planning, you need to check Data ONTAP support information to ensure that your deployment conforms to Data ONTAP hardware and software requirements for all systems in the deployment.

Previously all support information used to set up deployments with storage arrays was included in the V-Series Support Matrix. As of November 2012, the information is divided into two different tools, as shown in the following table:

For information about...	You should look here...
<p>Data ONTAP working with devices, including the following:</p> <ul style="list-style-type: none"> Supported storage arrays and storage array firmware Supported switches and switch firmware Whether your storage array supports nondisruptive (live) upgrade of the storage array firmware Whether a MetroCluster configuration is supported with your storage array 	<p>Interoperability Matrix at support.netapp.com</p> <p>Note: The <i>Interoperability Matrix</i> shows the Brocade and Cisco switches that are supported. You can find guides for configuring specific vendor switches on the NetApp Support Site at support.netapp.com.</p>
<p>Data ONTAP limits for releases and platforms, including the following:</p> <ul style="list-style-type: none"> Minimum and maximum array LUN sizes, including the minimum array LUN size for the root volume and spare core array LUNs Minimum aggregate size for aggregates with array LUNs Supported block size Minimum and maximum capacity Neighborhood limits 	<p><i>Hardware Universe</i> (formerly the <i>System Configuration Guide</i>) at support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml</p>

For information about...	You should look here...
Information about what is supported for specific arrays, including supported configurations.	<i>V-Series Implementation Guide for Third-Party Storage</i>

Related concepts

[*Limits information for configurations with storage arrays*](#) on page 10

32xx system configuration restrictions

There are some restrictions for 32xx systems that do not apply to other models.

32xx system on-board storage port guidelines

The two onboard FC ports, labeled 0c and 0d, are not on independent busses. Therefore, they do not provide storage redundancy. Some port failures can cause the system to panic.

To configure redundant port pairs, you need to use an FC HBA in an available expansion slot.

V3210 configuration restrictions

The following restrictions apply to V3210 configurations without an added FC HBA:

- Only a single storage array can be connected to a V3210 configuration.
- The storage array must present all LUNs in a single LUN group.

Stages of implementation for a system using array LUNs

Implementing a configuration in which your system uses LUNs from a storage array has two stages: a back-end implementation and a front-end implementation. It is helpful when planning your configuration to understand the high-level tasks in each stage.

Stage 1: back-end implementation

Setting up the back-end implementation includes all tasks that are required to set up the V-Series system with a storage array, up to the point where you can install Data ONTAP software.

Tasks to set up the back-end implementation include the following:

1. Creating and formatting array LUNs
2. Assigning ports
3. Cabling
4. Zoning switches (if applicable)
5. In Data ONTAP, assigning specific array LUNs to a V-Series system
6. In Data ONTAP, providing information to set up a V-Series system on the network

This process is similar to FAS system setup.

7. Installing Data ONTAP software

If a V-Series system is ordered with disk shelves, the Data ONTAP software is installed by the factory. In such a configuration, you do not need to create the root volume and install licenses and Data ONTAP software. The *Data ONTAP Software Setup Guide for 7-Mode* and the *Clustered Data ONTAP Software Setup Guide* provide instructions for how to set up your system in Data ONTAP, depending on whether your system is shipped with disk shelves.

Stage 2: front-end implementation

Tasks to set up the front-end implementation are the same as for a system using disks, including the following:

- Configuring the V-Series system for all protocols (NAS, FC, or both)
- Configuring Data ONTAP features such as SnapVault, SnapMirror, SnapValidator, Snapshot copies, and so on
- Creating volumes and aggregates
- Setting up data protection, including NDMP dumps to tapes

Planning summary for a V-Series system using native disks

Native disk shelves can be installed on new or existing V-Series systems. Setting up native disk shelves on V-Series systems and FAS systems is essentially the same. There are just a few additional considerations for V-Series systems.

Additional planning considerations for basic setup on V-Series systems with disks

You need to consider the following when determining basic setup and installation requirements for your system that uses native disks:

- If your system is ordered with disk shelves, the factory configures the root volume and installs licenses and Data ONTAP software (just as it does for FAS systems).
Clustered V-Series systems must be ordered with native disks.
- If your system is not ordered with disk shelves, you need to plan to install the Data ONTAP software and appropriate licenses.
- Data ONTAP automatically assigns ownership to native disks attached to your system.

The *Data ONTAP Software Setup Guide for 7-Mode* contains basic setup and installation instructions. Information in other Data ONTAP documents related to disks is relevant for V-Series systems unless otherwise noted in the document, *Hardware Universe*, or other support matrices.

Additional planning considerations if your system uses both disks and array LUNs

The following table summarizes additional planning considerations and the location of information to help you with each task.

Planning consideration	Where to find guidelines
Location of the root volume	<i>V-Series Installation Requirements and Reference Guide</i>
How many disks and array LUNs combined can be assigned to your system without exceeding the supported maximum assigned device limit for your system.	<i>Hardware Universe</i> (formerly the <i>System Configuration Guide</i>) at support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml
FC initiator port usage	<i>V-Series Installation Requirements and Reference Guide</i>
Which data should go on disks and which data should go on array LUNs	The system administrator should evaluate the type of data that needs to be handled. Data that needs to be accelerated, for example, flash aggregates and Flash Pools, should be on disks.

Planning considerations for moving a disk shelf

If you are moving a disk shelf from a FAS system to a V-Series system, you need to determine whether you want to preserve the data on a disk shelf.

Note: The `bootarg.storageencryption.support` variable must be set to `true` to be able to move an FDE disk shelf from a FAS system to a V-Series system. If it is not, all disks are marked as failed.

Related concepts

[Location of the root volume](#) on page 29

[Requirements for V-Series FC initiator port usage](#) on page 46

Planning summary for a V-Series system using array LUNs

Planning for Data ONTAP use of array LUNs involves considering what needs to be done on the storage array, what needs to be done on switches, and what needs to be done on the V-Series systems.

When your V-Series systems use array LUNs, you must communicate with the storage array and switch administrators to ensure that the back-end devices are configured to work with V-Series systems.

The following table summarizes major planning tasks and the location of information to help you with each task.

Planning task...	Where to find information...
<p>Determining requirements for setting up your storage array to work with Data ONTAP</p>	<p>See the <i>V-Series Implementation Guide for Third-Party Storage</i> for information about the following:</p> <ul style="list-style-type: none"> • Configuration settings on the storage array that are required for Data ONTAP to work with the storage array • Supported configurations for your vendor <p>See the Interoperability Matrix at support.netapp.com for information about the following:</p> <ul style="list-style-type: none"> • Environment requirements, for example, which storage array, storage array firmware, and switch firmware are supported
<p>Guidelines for array LUN use with Data ONTAP</p>	<p>See the <i>V-Series Installation Requirements and Reference Guide</i>.</p>
<p>Determining Data ONTAP limits regarding array LUNs</p>	<p>See the <i>Hardware Universe</i> (formerly the <i>System Configuration Guide</i>) at support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml for limits such as the following:</p> <ul style="list-style-type: none"> • Minimum and maximum array LUN sizes for data LUNs, the root volume array LUN, and spare core array LUNs • Number of devices supported on a V-Series model
<p>Determining a LUN security scheme setting access controls on the storage array and, if switches are deployed, setting zoning on switches</p>	<p>See the <i>V-Series Installation Requirements and Reference Guide</i>.</p>

Planning task...	Where to find information...
Determining a port-to-port connectivity scheme between the V-Series systems and the storage array	<p>See the <i>V-Series Installation Requirements and Reference Guide</i> for guidelines about the following:</p> <ul style="list-style-type: none"> • FC initiator port usage • Cabling of redundant paths between the V-Series system and storage array • Zoning of switches (for a fabric-attached configuration) • Mapping (exporting) array LUNs to the ports to which the V-Series systems are connected <p>See the <i>Interoperability Matrix</i> for supported connectivity between V-Series systems and your storage arrays.</p>
Determining which V-Series system is to “own” which array LUN (disk ownership)	<p>See the <i>V-Series Installation Requirements and Reference Guide</i> for information about the following:</p> <ul style="list-style-type: none"> • Data ONTAP disk ownership requirements • Guidelines for desired workload on specific V-Series systems <p>For procedures for configuring disk ownership, see the <i>Data ONTAP Storage Management Guide</i> for your systems' Data ONTAP release and operating mode.</p>
If your 7-Mode configuration includes multiple V-Series systems, if you want to use neighborhoods	<p>See the <i>V-Series Installation Requirements and Reference Guide</i> for a discussion of neighborhoods.</p> <p>See the <i>Hardware Universe</i> and the <i>Interoperability Matrix</i> for limits related to neighborhoods.</p>

Related concepts

[Planning for Data ONTAP use of array LUNs](#) on page 19

[Planning for LUN security on the storage arrays](#) on page 31

[Requirements for V-Series FC initiator port usage](#) on page 46

[How array LUNs become available for Data ONTAP storage use](#) on page 20

[Determining whether to use neighborhoods \(7-Mode in 8.x\)](#) on page 67

Planning for RAID implementation

You need to plan the size of and number of LUNs in the storage array RAID groups and decide whether you want to share the RAID group among hosts.

RAID protection for array LUNs

Storage arrays provide the RAID protection for the array LUNs that they make available to Data ONTAP; Data ONTAP does not provide the RAID protection.

Data ONTAP uses RAID0 (striping) for array LUNs. Data ONTAP supports a variety of RAID types on the storage arrays, except RAID0 because RAID0 does not provide storage protection.

When creating *RAID groups* on storage arrays, you need to follow the best practices of the storage array vendor to ensure that there is an adequate level of protection on the storage array so that disk failure does not result in loss of data or loss of access to data.

Note: A *RAID group* on a storage array is the arrangement of disks that together form the defined RAID level. Each RAID group supports only one RAID type. The number of disks that you select for a RAID group determines the RAID type that a particular RAID group supports. Different storage array vendors use different terms to describe this entity—RAID groups, parity groups, disk groups, Parity RAID groups, and other terms.

Data ONTAP supports RAID4 and RAID-DP on the native disk shelves connected to a V-Series system but does not support RAID4 and RAID-DP with array LUNs.

Implications of LUN size and number for Data ONTAP RAID groups

Part of planning for aggregates is to plan the size and number of Data ONTAP RAID groups you need for those aggregates, and the size and number of array LUNs for the Data ONTAP RAID groups. Setting up Data ONTAP RAID groups for array LUNs requires planning and coordination with the storage array administrator.

Planning for Data ONTAP RAID groups involves the following:

1. Planning the size of the aggregate that best meets your data needs.
2. Planning the number and size of the RAID groups that you need for the size of the aggregate. RAID groups in the same aggregate should be the same size, with the same number of array LUNs in each RAID group. Use the default RAID group size if possible.
3. Planning the size of the array LUNs that you need in your Data ONTAP RAID groups.
 - To avoid a performance penalty, all array LUNs in a particular Data ONTAP RAID group should be the same size.

- The array LUNs should be the same size in all RAID groups in the same aggregate.
4. Communicating with the storage array administrator to create the number of array LUNs of the size you need for the aggregate.
The array LUNs should be optimized for performance, according to the instructions in the storage array vendor documentation.

For more recommendations about setting up Data ONTAP RAID groups for use with storage arrays, including minimum and maximum RAID group size, see the appropriate Data ONTAP *Storage Management Guide*.

Related concepts

[*Determining the array LUNs for specific aggregates*](#) on page 60

Planning for Data ONTAP use of array LUNs

For Data ONTAP to use array LUNs, a storage array administrator must first create LUNs on the storage array and make them available to Data ONTAP. Then the Data ONTAP administrator must configure Data ONTAP to use the array LUNs that the storage array made available.

Planning how to provision array LUNs for Data ONTAP use includes the following considerations:

- The types of array LUNs that Data ONTAP supports
- Data ONTAP minimum and maximum array LUN sizes
- The number of array LUNs you need

Note: Data ONTAP considers an array LUN to be a virtual disk.

How array LUNs are made available for host use

A storage array administrator must create array LUNs and make them available to specified FC initiator ports of V-Series systems.

The process to make LUNs available to hosts and the terminology to describe it varies among storage array vendors. The basic process that the storage array administrator follows to make LUNs available for host use is as follows:

1. Creates logical devices (LDEVs).

Note: *LDEV* is a term used by some vendors and this guide to describe a piece of logical RAID storage configured from disks.

2. Creates a host group (or vendor equivalent).

The host group includes the WWPNs of the initiator ports of the hosts that are allowed to see the LDEV.

Note: To simplify management, most storage arrays enable you to define one or more *host groups*. You can define specific WWPNs (ports) and WWNs (hosts) to be members of the same group. You then associate specific array LUNs with a host group. Hosts in the host group can access the LUNs associated with the host group; hosts that are not in that host group cannot access those LUNs. Different vendors use different terms to describe this concept. The process of creating a host group differs among vendors.

3. Maps the LDEVs to host groups as LUNs.

Related concepts

[*How array LUNs become available for Data ONTAP storage use*](#) on page 20

How array LUNs become available for Data ONTAP storage use

A V-Series system cannot use an array LUN presented to it until after Data ONTAP has been configured to use the array LUN.

Although the storage array administrator makes an array LUN accessible to Data ONTAP, Data ONTAP cannot use the array LUN for storage until both of the following tasks are completed:

1. One V-Series system must be assigned to be the *owner* of the array LUN.
2. The array LUN must be added to an aggregate.

When you assign an array LUN to a V-Series system, Data ONTAP writes data to the array LUN to identify the assigned system as the *owner* of the array LUN. This logical relationship is referred to as *disk ownership*.

When you assign an array LUN to a V-Series system, it becomes a spare LUN owned by that system and it is no longer available to any other V-Series system.

A spare array LUN cannot be used for storage until you add it to an aggregate. Thereafter, Data ONTAP ensures that only the owner of the array LUN can write data to and read data from the LUN.

In an HA pair, both nodes must be able to see the same storage, but only one node in the pair is the owner of the array LUN. The partner node takes over read/write access to an array LUN in case of a failure of the owning node. The original owning node resumes ownership after the problem that caused unavailability of the node is fixed.

Related concepts

[*How array LUNs are made available for host use*](#) on page 19

[*Considerations when planning for disk ownership*](#) on page 20

[*Determining the array LUNs for specific aggregates*](#) on page 60

Considerations when planning for disk ownership

If you are deploying multiple V-Series systems, you must determine which system will *own* which array LUNs.

You should consider the following when planning which system will own which array LUNs:

- The maximum assigned device limit supported by your platform
The *Hardware Universe* shows the maximum assigned device limit that is supported for different platforms. This is a hard-coded limit. If your system uses both array LUNs and disks, this maximum limit is the maximum of disks and array LUNs combined. You must account for both types of storage when determining how many array LUNs and disks you can assign to a system.
- The amount of load that you expect to be generated by different applications used in your environment

Some types of applications are likely to generate a lot of requests, whereas other applications (for example, archival applications) generate fewer requests. You might want to consider weighing ownership assignments based on expected load from specific applications.

Related concepts

How array LUNs become available for Data ONTAP storage use on page 20

Array LUN assignment changes

You can change assignment of a *spare* array LUN from one V-Series system to another. You might want to change ownership for load balancing over nodes.

Information about changing the ownership of an array LUN is provided in the *Data ONTAP Storage Management Guide* for your operating mode.

Array LUN types supported by Data ONTAP

You can map only storage array LUNS to Data ONTAP.

Some storage arrays have a non-storage *command* LUN. You cannot map a command type LUN to a system running Data ONTAP.

Starting in Data ONTAP 8.1, you can map LUN 0 to Data ONTAP if it is a storage type LUN.

See the *V-Series Implementation Guide for Third-Party Storage* for a discussion about array LUN types used by particular vendors that are not supported by Data ONTAP.

Factors that impact the number and size of array LUNs you need

You must consider a number of factors, including usable space in a LUN, when determining how many array LUNs you need and their size.

Factors impacting the number of array LUNs needed

Factors such as array LUN size, Data ONTAP overhead, and checksum type impact the number of array LUNs that you need.

You should consider the following when determining the number of array LUNs that you need:

- The smaller the array LUNs, the more LUNs you need for the storage that you want. Ideally, creating one large array LUN from a given storage array RAID group is recommended.
- Device limits define the maximum number of disks and array LUNS that can be assigned to a V-Series system.
The *Hardware Universe* contains information about device limits.
- The more usable space in an array LUN, the fewer array LUNs are needed.

The amount of usable space in an array LUN is determined by the space that Data ONTAP requires, checksum type, and additional factors such as space required for optional Snapshot reserves.

- Different applications generate different loads.
When determining the assignment of array LUNs to V-Series systems, you should consider what the storage will be used for and the number of requests that are likely to be generated by different applications.
- If you plan to share an FC initiator port with multiple target ports, you need to be aware of the requirements to do so and the maximum number of array LUNs from all target ports combined that can be visible on an FC initiator port.

Starting in Data ONTAP 8.2, the supported number of visible array LUNs on an FC initiator increased from previous releases as part of the shared FC initiator support introduced in 8.2. (Prior to Data ONTAP 8.2, support is limited to one FC initiator connecting to just one storage array target port, and fewer visible array LUNs than in 8.2.)

The V-Series *Best Practice Guide* documents the supported number of visible array LUNs on an FC initiator port for different Data ONTAP releases. The *V-Series Installation Requirements and Reference Guide* contains information about the requirements for shared FC initiator ports.

Related concepts

[Elements that reduce the usable space in an array LUN](#) on page 24

[Rules for sharing an FC initiator port with multiple target ports](#) on page 47

[Example configuration: shared FC initiator ports](#) on page 48

Minimum number of array LUNs required per V-Series system

The number of array LUNs that you need per V-Series system depends on the location of the root volume.

If the root volume is on a storage array, each stand-alone V-Series system and each node in an HA pair must own at least one array LUN. If the root volume is on a native disk, the only array LUNs needed are those for data storage.

For a MetroCluster configuration with V-Series systems, two array LUNs are required (one LUN from each site) if the root volume is on a storage array. The two LUNs are necessary so that the root volume can be mirrored.

Note: Data ONTAP 8.2 and later supports both disks and array LUNs in a MetroCluster configuration with V-Series systems. This means that you can choose whether the root volume will be on disks or array LUNs. Prior to Data ONTAP 8.2, only array LUNs are supported in a MetroCluster configuration with V-Series systems.

Note: MetroCluster configurations are not supported for clustered Data ONTAP.

Related concepts

[Factors impacting the number of array LUNs needed](#) on page 21

Spare core array LUN requirement for core dumps

For both stand-alone V-Series systems and nodes in HA pairs, you need to create a spare core array LUN of adequate size to hold core dumps if no spare disk is available.

On a system that uses both disks and array LUNs, a spare array LUN is not required for a core dump if a spare disk is available. If neither a spare array LUN or spare disk is available, there is no place for the core to be dumped.

The *Hardware Universe* contains the minimum spare core array LUN size for each platform.

A core dump contains the contents of memory and NVRAM. During a system panic, Data ONTAP dumps core to a spare array LUN or spare disk, if a spare exists. Upon reboot, the core is read from the spare and saved to a core dump on the root filesystem. Technical support can then use the core dump to help troubleshoot the problem.

Related concepts

[Limits information for configurations with storage arrays](#) on page 10

Minimum and maximum array LUN sizes supported by Data ONTAP

The array LUNs presented to Data ONTAP must fall within the minimum and maximum size that Data ONTAP requires for data LUNs, according to the way that Data ONTAP calculates size. Data ONTAP issues an error message identifying an array LUN that does not adhere to the minimum or maximum array LUN size requirements.

The Data ONTAP minimum and maximum array LUN sizes are calculated according to the way that Data ONTAP defines units of measure. The Data ONTAP definition of a GB and TB is as follows:

One...	Equals...
GB	1000 x 1024 x 1024 bytes (1000 MB)
TB	1000 x 1000 x 1024 x 1024 bytes (1000 GB)

Different vendors use different formulas for calculating units of measurement. You must use your vendor's units of measure to calculate the minimum and maximum array LUN sizes that are equivalent to the minimum and maximum array LUN sizes that Data ONTAP supports.

The maximum LUN size that Data ONTAP supports differs according to Data ONTAP release. For information about the minimum and maximum array LUN sizes, see the *Hardware Universe*. The *Hardware Universe* contains the maximum array LUN size, calculated according the mathematical formulas used by most vendors. The maximum array LUN size shown is the size that will work for all storage array vendors.

Note: The minimum array LUN size for a data (storage) LUN is different from the minimum array LUN size for the root volume. See the *Hardware Universe* for more information.

Related concepts

[Limits information for configurations with storage arrays](#) on page 10

[Minimum array LUN size for the root volume](#) on page 24

Related references

[Array LUN is too small or too large](#) on page 87

Minimum array LUN size for the root volume

The array LUN used for the root volume must be larger than the minimum size required for other array LUNs. When planning for your root volume, you need to determine the minimum root volume size for the Data ONTAP release that is running on your system.

It is strongly recommended that you do not set the size of a root volume below the minimum size array LUN for the root volume that is shown in the *Hardware Universe*. The reason is that you want to ensure that there is sufficient space in the root volume for system files, log files, and core files. You need to provide these files to technical support if a system problem occurs.

Note: Both the minimum array LUN size for the root volume and the minimum array LUN size for non-root volumes are shown in the *Hardware Universe*. The minimum array LUN size for a non-root volume is considerably smaller than for the root volume.

Related concepts

[Limits information for configurations with storage arrays](#) on page 10

[Minimum and maximum array LUN sizes supported by Data ONTAP](#) on page 23

[Location of the root volume](#) on page 29

Elements that reduce the usable space in an array LUN

Not all capacity in an array LUN is available for storage. The usable space in an array LUN is impacted by a number of fixed and optional elements.

When planning the number and size of the array LUNs you need, you must consider the usable space in the array LUN according to the checksum type you are using and the optional elements you might configure.

When calculating the usable space in an array LUN, you must consider the following factors that decrease the usable space of the LUN:

- 10%—Reserved for use by Data ONTAP
- 0.1% for LUNs less than or equal to 2 TiB, and 0.2% for LUNs greater than 2 TiB—Core dump
- 5%—Volume-level Snapshot reserve (default, configurable)
- 0%—Aggregate-level Snapshot copy (default, configurable)
- Checksum type (you assign one type):
 - 12.5%—Block checksum (BCS)

- 1.56 %—Advanced Zoned checksum (AZCS)

Note: The percentages shown are for Data ONTAP 8.1.1 and later and might differ from the percentages for previous Data ONTAP releases.

Space for aggregate-level and volume-level Snapshot reserves are optional and changeable through Data ONTAP. See the *Data ONTAP Data Protection Guide* for more information about Snapshot copies.

Related concepts

[Characteristics of checksum types that Data ONTAP supports](#) on page 25

[Considering checksum type when planning array LUN size and number](#) on page 25

[Formulas for calculating array LUN size, considering the checksum type](#) on page 26

Considering checksum type when planning array LUN size and number

When planning the number and size of array LUNs that you need for Data ONTAP, you must consider the impact of the checksum type on the amount of usable space in the array LUN. A checksum type must be specified for each array LUN assigned to a V-Series system.

When an array LUN on the storage array is mapped to be used by a V-Series system, Data ONTAP treats the array LUN as a *raw*, unformatted disk. When you assign an array LUN to a V-Series system you specify the checksum type, which tells Data ONTAP how to format the raw array LUN. The impact of the checksum type on usable space depends on the checksum type you specify for the LUN.

Related concepts

[Elements that reduce the usable space in an array LUN](#) on page 24

[Characteristics of checksum types that Data ONTAP supports](#) on page 25

[Formulas for calculating array LUN size, considering the checksum type](#) on page 26

Characteristics of checksum types that Data ONTAP supports

Starting in Data ONTAP 8.1.1, Data ONTAP supports block checksum type (BCS) and advanced zoned checksum type (AZCS) for array LUNs, disks, and aggregates. Zoned checksum (ZCS) is supported in Data ONTAP 8.1.1 and later as legacy support only.

The type of checksum assigned to an array LUN in Data ONTAP can impact performance or the usable space of an array LUN. Therefore, the number and size of array LUNs you need can be impacted depending on the checksum type you assign to array LUNs.

The following table describes the characteristics of the checksum types:

Checksum type	Description
BCS	<p>BCS is the default and recommended checksum type for array LUNs. BCS provides better performance for array LUNs than AZCS.</p> <p>BCS has a greater impact on the usable space in an array LUN than AZCS. BCS uses 12.5 percent of the usable space in an array LUN.</p>
AZCS (advanced_zoned)	<p>Starting in Data ONTAP 8.1.1, AZCS is an alternative to BCS. The impact of AZCS on usable space in an array LUN is less than with BCS; AZCS uses 1.56 percent of the device capacity. However, you must weigh the need for more usable space against performance. AZCS can sometimes cause performance problems for array LUNs.</p> <p>AZCS is not recommended for array LUNs for high-performance random workloads. However, you can use AZCS with array LUNs for DR, archive, or similar workloads.</p> <p>Note: This caution about the possible performance impact of AZCS does not apply to native disks.</p>
ZCS (zoned)	<p>Prior to Data ONTAP 8.1.1, zoned checksum array LUNs were used with ZCS type aggregates. Starting in 8.1.1, any new aggregates created with zoned checksum array LUNs are AZCS aggregates, and the array LUNs become advanced_zoned LUNs after they are added to the aggregate. You can add zoned checksum array LUNs to existing ZCS aggregates.</p> <p>Like AZCS, ZCS has a performance penalty for high-performance random workloads. However, ZCS does not provide as much functionality as AZCS. For example, unlike ZCS, AZCS supports deduplication and compression.</p>

Note: Different guidelines apply to assigning the checksum type for disks. Those guidelines differ according to disk size and type. See *TR3838 Storage Subsystem Configuration Guide* at support.netapp.com for more information.

Related concepts

Elements that reduce the usable space in an array LUN on page 24

Formulas for calculating array LUN size, considering the checksum type on page 26

Formulas for calculating array LUN size, considering the checksum type

A number of elements, including checksum type, impact the usable capacity of an array LUN. You can use a formula to calculate how much usable capacity there would be in a given size array LUN, or to calculate how large an array LUN needs to be to provide the amount of storage that you want.

A number of elements, including checksum type, impact the size of the array LUN you need for the amount of *usable capacity*. Usable capacity is the amount of space that is available for storage.

The following table shows the ways of calculating the array LUN size you need:

If you know...	You want to find out...
How large your array LUNs are	How much capacity is available for storage (usable capacity). You need to consider the amount of space required for all elements.
How much storage that you want in the array LUN	How large an array LUN you need. You need to take into account your desired amount of storage and space required for other elements.

Note: 2 TB in these formulas represents 2 TiB, or 2199023255552 bytes, which is 2097.152 GnaB or 2.097 TnaB according to the way that Data ONTAP calculates measurements.

Formula for calculating the usable capacity

When you know how large your array LUNs are, you can use the following formula to determine the usable capacity for storage in an array LUN. This formula takes into account the Snapshot reserve.

- Y is the usable capacity for storage.
- N is the total capacity of the array LUN.
- Snapshot reserve.

Checksum type	Formula
BCS—array LUNs less than 2 TB	$N \times \{0.875 \times 0.9 \times 0.99 \times \text{Snapshot reserve}\} = Y$
BCS—array LUNs greater than 2 TB	$N \times \{0.875 \times 0.9 \times 0.998 \times \text{Snapshot reserve}\} = Y$
AZCS—array LUNs less than 2 TB	$N \times \{0.984 \times 0.9 \times 0.99 \times \text{Snapshot reserve}\} = Y$
AZCS—array LUNs greater than 2 TB	$N \times \{0.984 \times 0.9 \times 0.998 \times \text{Snapshot reserve}\} = Y$

Example 1: calculations *with* a Snapshot reserve

In the following example, the total capacity of the array LUN is 4 GB, with a volume Snapshot reserve set at the default for Data ONTAP 8.1.1 (5 percent).

The following examples are for an array LUN less than 2 TB:

Checksum type	Formula
BCS (array LUN less than 2 TB)	$4 \times \{0.875 \times 0.9 \times 0.99 \times 0.95\} = 2.96 \text{ GB usable space for storage}$

Checksum type	Formula
AZCS (array LUN less than 2 TB)	$4 \times \{0.984 \times 0.9 \times 0.99 \times 0.95\} = 3.33 \text{ GB usable space for storage}$

Formula for calculating maximum array LUN size needed

When you know the array LUN capacity needed to obtain the storage capacity you want, you can use the following formula to determine the total array LUN size you need, considering elements that require space in the LUN.

- Y is the exact amount of space in the array LUN that you want.
- If you are using Snapshot copies, the Snapshot reserve is taken into account.

The following examples are for an array LUN less than 2 TB:

Checksum type	Formula
BCS (array LUN less than 2 TB)	$Y \div \{0.875 \times 0.9 \times 0.99 \times \text{Snapshot reserve}\} = \text{Actual capacity needed}$
AZCS (array LUN less than 2 TB)	$Y \div \{0.984 \times 0.9 \times 0.99 \times \text{Snapshot reserve}\} = \text{Actual capacity needed}$

Example 2: calculations *with* Snapshot reserves

In this example, the volume Snapshot reserve is the default setting for Data ONTAP 8.1.1 (5 percent).

The following examples are for an array LUN less than 2 TB:

Checksum type	Formula
BCS (array LUN less than 2 TB)	$10 \text{ GB} \div \{0.875 \times 0.9 \times 0.99 \times 0.95\} = 13.5 \text{ GB actual capacity needed}$
AZCS (array LUN less than 2 TB)	$10 \text{ GB} \div \{0.984 \times 0.9 \times 0.99 \times 0.95\} = 12.05 \text{ GB actual capacity needed}$

Example 3: calculations *without* Snapshot reserves

You want 10 GB of usable capacity for storage. The following example shows calculating actual array LUN size when you are not using Snapshot copies.

The following examples are for an array LUN less than 2 TB:

Checksum type	Formula
BCS (array LUN less than 2 TB)	$10 \text{ GB} \div \{0.875 \times 0.9 \times 0.99\} = 12.8 \text{ GB actual capacity needed}$
AZCS (array LUN less than 2 TB)	$10 \text{ GB} \div \{0.984 \times 0.9 \times 0.99\} = 11.41 \text{ GB actual capacity needed}$

Related concepts

Elements that reduce the usable space in an array LUN on page 24

Characteristics of checksum types that Data ONTAP supports on page 25

Location of the root volume

For systems running clustered Data ONTAP, the root volume must be installed on a native disk shelf. For systems running Data ONTAP operating in 7-Mode, you might need to decide where the root volume should be located.

For systems running Data ONTAP operating in 7-Mode

For 7-Mode systems, follow these guidelines:

- The root volume can be on a storage array or on a native disk shelf.
However, installing the root volume on a native disk is recommended.
If you order your 7-Mode system with disks, the factory installs the root volume on a native disk.
- In a 7-Mode HA pair, it is highly recommended that the root volume be located on the same type of storage for both nodes—either on a native disk shelf for both nodes or on a storage array for both nodes.
Although an asymmetrical approach is highly discouraged, it is not a requirement to follow the same strategy for both nodes.
- For MetroCluster configurations with both disks and array LUNs (supported with V-Series systems starting in Data ONTAP 8.2), it is recommended that you create the root volume on a disk if you are setting up a new configuration.
If you are adding disks to an existing MetroCluster configuration with array LUNs, you can leave the root volume on an array LUN.

For clustered Data ONTAP systems

Clustered Data ONTAP systems must be ordered with native disks, and the root volume must be on a native disk shelf. The factory installs the root volume on a native disk shelf.

Related concepts

Minimum array LUN size for the root volume on page 24

Planning for LUN security on the storage arrays

If you are using your V-Series system with storage arrays, you must use a LUN security method to eliminate the possibility of a non V-Series system overwriting array LUNs owned by a V-Series system, or vice versa.

LUN security is a method for isolating the hosts that can access particular array LUNs. LUN security is similar to switch zoning in concept, but it is performed on the storage array. *LUN security* and *LUN masking* are equivalent terms to describe this functionality.

Attention: The Data ONTAP disk ownership scheme prevents one V-Series system from overwriting an array LUN owned by another V-Series system. However, it does not prevent a V-Series system from overwriting an array LUN accessible by a non V-Series host. Likewise, without a method of preventing overwriting, a non V-Series host could overwrite an array LUN used by a V-Series system.

Available LUN security methods

With LUN security, you can mask array LUNs for viewing by only certain hosts, present LUNs only for a specific host on a port, or dedicate a storage array to a particular host.

You should use both zoning and LUN security for added protection and redundancy for the V-Series system. If, for example, you do not have LUN security configured and you have to replace a SAN switch, the V-Series system could panic before you can configure the zoning on the new switch because the switch is wide open.

In addition to reading about the LUN security methods described here, you should also see the *V-Series Implementation Guide for Third-Party Storage* for any additional details regarding LUN security for your vendor's storage arrays. Some storage arrays must be dedicated for V-Series use.

Method 1: Port-level security

You can use port-level security to present only the array LUNs for a particular host. That port then becomes dedicated to that host.

Note: Not all storage arrays support port-level security. Some storage arrays present all LUNs on all ports by default, and they do not provide a way to restrict the visibility of LUNs to particular hosts. For these arrays you must use either a LUN security product or dedicate the storage array to the V-Series system. You should check your storage array documentation to determine whether your storage array supports port-level security.

Method 2: LUN security products

You can use a LUN security product to control hosts that are zoned to the same port so that they can see specific array LUNs over that port. This prevents other hosts from accessing those same array LUNs by masking them from the other hosts.

Method 3: Dedicate the storage array for V-Series use

You can dedicate the storage array to V-Series use. In this case, no hosts other than V-Series systems are connected to the storage array.

Planning for paths to array LUNs

Paths are the physical connections between the V-Series system and the storage array. Redundant paths are required to eliminate any single point of failure (SPOF) between the V-Series system and the storage array.

Requirements for redundant setup of components in a path

V-Series systems must connect to the storage array through a redundant Fibre Channel (FC) network. Two FC networks or fabric zones are required to protect against a connection failing and so that fabric ports or switches can be taken offline for upgrades and replacements without impacting the V-Series systems.

V-Series system redundancy requirements

- You must attach each connection to a different FC initiator port in the port pair on the V-Series system.
- Each V-Series FC initiator port in the same initiator port pair must be on a different bus.

FC switch redundancy requirements

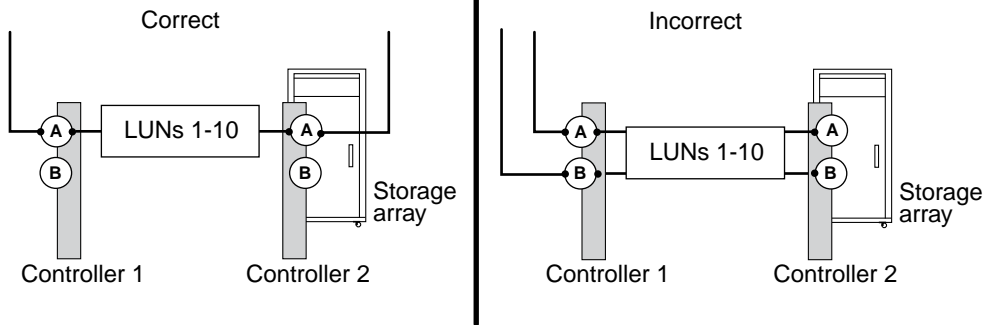
- You must use redundant switches.
- You must use redundant ports on the FC switches.

Storage array redundancy requirements

Be sure that the ports on the storage array that you select to access a given LUN are from different components that could represent a single point of failure, for example, from alternate controllers, clusters, or enclosures. The reason is that you do not want all access to an array LUN to be lost if one component fails.

Note: A given array LUN is accessed through only one port at a time.

The following illustration shows correct and incorrect storage array port selection for redundancy. The path setup in the example on the left is correct because the paths to the array LUN are redundant—each connection is to a port on a different controller on the storage array.



When to check for redundant paths to array LUNs

You need to check for redundant paths to an array LUN after installation and during fabric maintenance activities.

You should recheck for path redundancy when performing the following activities:

- Initial installation
- Fabric maintenance, for example:
 - Before, during, and after an infrastructure upgrade
 - Before and after taking a switch out of service for maintenance

Be sure that the paths were configured as redundant paths before you remove a switch between the V-Series systems and the storage array so that access to the array LUNs is not interrupted.
 - Before and after maintaining hardware on a storage array

For example, you should recheck for path redundancy when maintaining the hardware component on which host adapters and ports are located. (The name of this component varies on different storage array models).

Required number of paths to an array LUN

The required number of paths to an array LUN varies depending on Data ONTAP release and mode.

Data ONTAP release	Number of paths supported
8.1 and later clustered systems	2 or, with an active-active storage array, 4
8.1 and later systems operating in 7-Mode	2
Releases prior to 8.1 (for both clustered systems and systems operating in 7-Mode)	2

For all releases, Data ONTAP expects and requires that a storage array provide access to a specific array LUN on two redundant storage array ports; that is, through two redundant paths. A given array LUN is accessed through only one port at a time.

Be sure that the ports on the storage array that you select to access a given LUN are from different components that could represent a single point of failure, for example, from alternate controllers, clusters, or enclosures. The reason is that you do not want all access to an array LUN to be lost if one component fails.

Related concepts

[Advantages of four paths to an array LUN \(clustered Data ONTAP 8.1 and later\)](#) on page 35

Advantages of four paths to an array LUN (clustered Data ONTAP 8.1 and later)

When planning the number of paths to an array LUN for clustered Data ONTAP 8.1 and later, you need to consider whether you want to set up two or four paths.

The advantages of setting up four paths to an array LUN include the following:

- If a switch fails, both storage array controllers are still available.
- If a storage array controller fails, both switches are still available.
- Performance can be improved because load balancing is over four paths instead of two.

Note: Only two paths to an array LUN are supported for Data ONTAP 8.1 and later operating in 7-Mode and releases prior to Data ONTAP 8.1.

Related concepts

[Required number of paths to an array LUN](#) on page 34

Using LUN groups to partition the load over V-Series connections

Using multiple LUN groups allows for additional capacity, as well as potentially improving system performance by spreading the workload across more target ports. Use of multiple LUN groups is not supported for all storage arrays.

A *LUN group* is set of logical devices on the storage array that a V-Series system accesses over the same paths. The storage array administrator configures a set of logical devices as a group to define which host WWPNs can access them. Data ONTAP refers to this set of devices as a *LUN group*.

Advantages of using multiple LUN groups are as follows:

- There are limits on the number of LUNs that a given FC initiator port pair can support. For large storage arrays in particular, the needed capacity might exceed what a single LUN group can provide.

Note: The supported number of visible array LUNs on an FC initiator port varies according to Data ONTAP release. See the V-Series *Best Practices Guide* for information about this limit for the release running on your system.

- You can partition the load of array LUNs over the FC initiator port pairs.

The number of paths to a LUN group varies according to release and operating mode.

See the *Interoperability Matrix* to determine whether a configuration using multiple LUN groups is supported for your storage array.

Related concepts

[*Implementation requirements for a multiple LUN group configuration*](#) on page 36

Related tasks

[*Checking whether the configuration matches your intentions \(clustered Data ONTAP 8.1 and later\)*](#) on page 104

Related references

[*Reasons for fewer array LUN groups than expected*](#) on page 106

[*Reasons for more array LUN groups than expected*](#) on page 107

Implementation requirements for a multiple LUN group configuration

Implementing a multiple LUN group configuration requires setup on the V-Series systems and on the storage arrays.

The following setup is done on the *storage array* to implement a multiple LUN group configuration:

- As many ports as possible are used to provide access to the array LUNs you allocated for the V-Series system.
- Host groups (or your vendor's equivalent) are used to define which array LUN groups are presented to each V-Series FC initiator port.

The following setup is done on the *V-Series system* to implement a multiple LUN group configuration:

- One initiator port pair is used for each array LUN group.
Each FC initiator port pair accesses a different LUN group on the storage array through redundant paths.
- One large aggregate is created (in the Data ONTAP configuration), and array LUNs from multiple RAID groups (parity groups) are added to the aggregate.
By doing so, the I/O is spread across more disks. The combination of spreading I/O across the RAID groups and creating one large aggregate results in a significant performance boost.

The following is done on the *switch* to implement a multiple LUN group configuration:

- Switch zoning is configured to define which target ports the V-Series initiator ports are to use to access each array LUN group.
- For active-passive storage arrays, both ports of the storage array controller are cabled to the same switch fabric.

Related concepts

[Example of a configuration with multiple LUN groups](#) on page 37

Related references

[Reasons for fewer array LUN groups than expected](#) on page 106


[Reasons for more array LUN groups than expected](#) on page 107

Example of a configuration with multiple LUN groups

Using multiple LUN groups enables you to partition the load over V-Series connections. This configuration cannot be used with all storage arrays.

The following illustration shows how one V-Series system FC initiator port pair (0c and 0f) accesses one LUN group over one storage array port pair, and a second FC initiator port pair (0a and 0h) accesses a second LUN group on the same storage array over a different storage array port pair.

Note: This configuration is referred to as *stand-alone with two 2-port array LUN groups*. A multiple LUN group configuration could have an HA pair instead of a stand-alone system.



This multiple LUN group configuration enables you to spread the I/O across the RAID groups (parity groups) on the storage array. You set up your configuration so that different V-Series FC initiator

port pairs access different groups of LUNs on the storage array. The V-Series system sees any given array LUN over only two paths because a given LDEV (logical device) is mapped to only two redundant ports on the storage array. Each LUN group is accessed through a different target port pair.

Each LDEV is identified externally by a LUN ID. The LDEV must be mapped to the same LUN ID on all storage array ports over which it will be visible to V-Series systems.

Note: The same LUN ID cannot refer to two different LDEVs, even if the LUNs using the same ID are in different host groups on a target port. Although LUN ID reuse is not supported on the same target port, LUN ID reuse is supported on a storage array if the LUNs are mapped to different storage array ports.

The following table summarizes the zoning for this example. Single-initiator zoning is the recommended zoning strategy.

Zone	V-Series system FC initiator port	Storage array
Switch 1		
z1	Port 0a	Controller 1 Port B
z2	Port 0c	Controller 1 Port A
Switch 2		
z3	Port 0f	Controller 2 Port A
z4	Port 0h	Controller 2 Port B

See the *Interoperability Matrix* to determine whether multiple array LUN groups are supported for your storage array.

Related concepts

[Using LUN groups to partition the load over V-Series connections](#) on page 35

[Implementation requirements for a multiple LUN group configuration](#) on page 36

Related references

[Reasons for fewer array LUN groups than expected](#) on page 106

[Reasons for more array LUN groups than expected](#) on page 107

How paths are reflected in array LUN names

The array LUN name is a path-based name that includes the devices in the path between the V-Series system and the storage array.

By looking at the array LUN name as it is displayed in Data ONTAP output, you can identify the following:

- Devices in the path between the V-Series system and the storage array
- Ports used
- The LUN identifier that the storage array presents externally for mapping to hosts

The format of the array LUN differs depending on the type of configuration and the Data ONTAP operating mode that the system is running.

Related concepts

[Array LUN name format](#) on page 39

Array LUN name format

The array LUN name is a path-based name that includes the devices in the path between the V-Series system and the storage array, ports used, and the SCSI LUN ID on the path that the storage array presents externally for mapping to hosts.

On a V-Series system operating in 7-Mode, there are two names for each array LUN because there are two paths to each LUN, for example, *brocade3:6.126L1* and *brocade15:6.126L1*.

On a clustered Data ONTAP 8.0.x V-Series system, there are two names for each array LUN because there are two paths to each LUN.

Array LUN format for systems operating in 7-Mode and releases prior to 8.0

Configuration	Array LUN name format	Component descriptions
Direct-attached	<i>adapter.idlun-id</i>	<p><i>adapter</i> is the adapter number on the V-Series system.</p> <p><i>id</i> is the channel adapter port on the storage array.</p> <p><i>lun-id</i> is the array LUN number that the storage array presents to hosts.</p> <p>Example:</p> <p><i>0a.0L1</i></p>

Configuration	Array LUN name format	Component descriptions
Fabric-attached	<i>switch-name:port.idlun-id</i>	<p><i>switch-name</i> is the name of the switch.</p> <p><i>port</i> is the switch port that is connected to the target port (the end point).</p> <p><i>id</i> is the device ID.</p> <p><i>lun-id</i> is the array LUN number that the storage array presents to hosts.</p> <p>Example:</p> <p><i>brocade3:6.126L1</i></p> <p><i>brocade3:6.126</i> is the path component and <i>L1</i> is the SCSI LUN ID.</p>

Array LUN name format for clustered Data ONTAP systems

Configuration	Array LUN name format	Component descriptions
Direct-attached	<i>node-name.adapter.idlun-id</i>	<p><i>node-name</i> is the name of the clustered node. With clustered Data ONTAP, the node name is prepended to the LUN name so that the path-based name is unique within the cluster.</p> <p><i>adapter</i> is the adapter number on the V-Series system.</p> <p><i>id</i> is the channel adapter port on the storage array.</p> <p><i>lun-id</i> is the array LUN number that the storage array presents to hosts.</p> <p>Example: <i>node1.0a.0L1</i></p>

Configuration	Array LUN name format	Component descriptions
Fabric-attached	<i>node-name:switch-name:port.idlun-id</i>	<p><i>node-name</i> is the name of the node. With clustered Data ONTAP, the node name is prepended to the LUN name so that the path-based name is unique within the cluster.</p> <p><i>switch-name</i> is the name of the switch.</p> <p><i>port</i> is the switch port that is connected to the target port (the end point).</p> <p><i>id</i> is the device ID.</p> <p><i>lun-id</i> is the array LUN number that the storage array presents to hosts.</p> <p>Example: <i>node1:brocade3:6.126L1</i></p>

Related concepts

[How the array LUN name changes in Data ONTAP displays](#) on page 41

[How paths are reflected in array LUN names](#) on page 38

How the array LUN name changes in Data ONTAP displays

For array LUN names shown in Data ONTAP displays, the paths shown are from the perspective of the V-Series system.

- On each V-Series system there are multiple valid names for a single array LUN, one per path. The name for a given LUN that is displayed in Data ONTAP can change depending on which path is active at a given time. For example, if the primary path becomes unavailable and Data ONTAP switches to the alternate path, the LUN name that is displayed changes.
- For clustered Data ONTAP 8.xe, the node name is prepended to the array LUN name to provide a unique name for each array LUN.
- Each node in a cluster typically accesses a given LUN through a different storage array port to limit contention on a single port.

Related concepts


[Array LUN name format](#) on page 39

Valid pathing: stand-alone system with a single 2-port array LUN group

A fabric-attached stand-alone system with a single 2-port array LUN group is supported with most storage arrays for all Data ONTAP releases.

Note: Different storage arrays, even those from the same vendor, might label the ports differently from those shown in the example. On your storage array, you need to ensure that the ports you select are on alternate controllers.

The following illustration shows a single 2-port array LUN group with a stand-alone V-Series system:



When validating your installation, it is helpful to check your command output against command output in this guide to verify that the number of LUN groups is what you intended and that there are redundant paths.

For clustered Data ONTAP, the following example shows the expected `storage array config show` output for the illustrated configuration—a single LUN group (LUN group 0) with two redundant paths to each array LUN. (The redundant paths in the output that match the illustration are shown in the array target port names `201A00a0b80fee04` and `202A00a0b80fee0420`.)

```
vs1::> storage array config show
      LUN  LUN
Node      Group Count      Array Name      Array Target Port      Initiator
```

```

-----
vs1          0    50          IBM_1742_1      201A00a0b80fee04    0a
                                           202A00a0b80fee04    0h
    
```

For 7-Mode, the following example shows the expected `storage array show-config` output for the illustrated configuration—a single LUN group (LUN group 0) with two redundant paths to each array LUN.

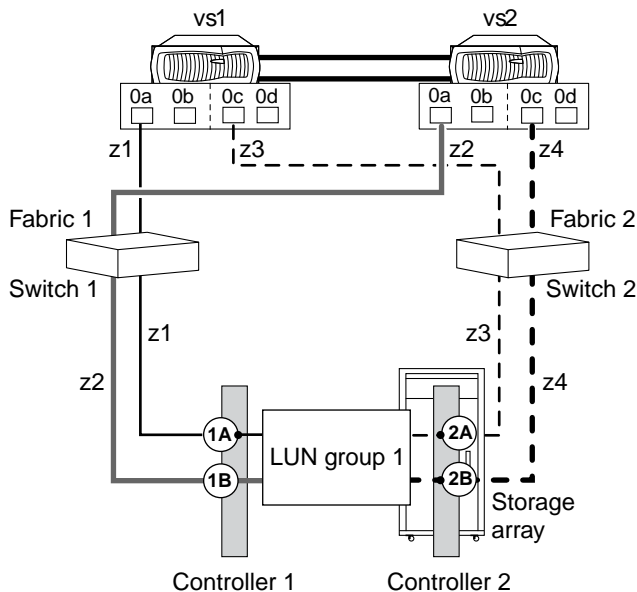
```

LUN Group      Array Name    Array Target Ports      Switch Port  Initiator
Group 0 (4 LUNS) HP_HSV210_1 50:00:1f:e1:50:0a:86:6d vnmc4300s35:11    0a
                                           50:00:1f:e1:50:0a:86:69 vnbr4100s31:15    0h
    
```

Valid pathing: a single 4-port array LUN group in a fabric-attached configuration

A single 4-port array LUN group configuration works with all storage arrays for all Data ONTAP releases.

The following illustration shows pathing in a configuration with a single 4-port array LUN group:



In this configuration with a single 4-port LUN group, array LUNs are mapped to four ports on the storage array. The array LUN group is presented to both nodes in the HA pair on different array target ports. However, each V-Series system can see an array LUN, end-to-end, through only two paths. Zoning is configured so that each FC initiator port on a V-Series system can access only a single target array port.

For clustered Data ONTAP, it is helpful to compare your output against valid `storage array config show` output when you are checking that the number of LUN groups you expected were set up. The following `storage array config show` output example shows the expected output for this configuration—a single array LUN group:

```
vs::> storage array config show
Node      LUN      LUN
          Group Count  Array Name  Array Target Port  Initiator
-----
vs1        1      10   IBM_2107900_1  50050763030301241A  0a
          50050763031301242A  0c
vs2        1      10   IBM_2107900_1  50050763030881241B  0a
          50050763031881242B  0c

4 entries were displayed.
```

For 7-Mode, you can check the `storage array show-config` output to see how many array LUN groups were created.

Planning a port-to-port connectivity scheme

Planning connectivity between the V-Series FC initiator ports and storage array ports includes determining how to achieve redundancy and meeting requirements for the number of paths to an array LUN.

Related concepts

How FC initiator ports are labeled on page 45

Requirements for V-Series FC initiator port usage on page 46

Rules for sharing an FC initiator port with multiple target ports on page 47

Example configuration: shared FC initiator ports on page 48

What to include in a port-to-port plan for connecting to storage arrays on page 53

What to include in a port-to-port plan for connecting to native disk shelves on page 54

Zoning recommendation for a configuration with storage arrays on page 57

Related tasks

Checking the number of array LUNs visible on an FC initiator port (clustered Data ONTAP) on page 52

How FC initiator ports are labeled

All FC initiator ports on V-Series systems are identified by a number and a letter. Labeling differs depending on whether the ports are on the motherboard or cards in expansion slots.

- Port numbering on the motherboard
Ports are numbered 0a, 0b, 0c, 0d...
- Port numbering on expansion cards
Ports are numbered according to the slot in which the expansion card is installed. A card in slot 3 yields ports 3A and 3B.
The FC initiator ports are labeled 1 and 2. However, the software refers to them as A and B. You see these labels in the user interface and system messages displayed on the console.

Related concepts

Requirements for V-Series FC initiator port usage on page 46

Requirements for V-Series FC initiator port usage

When you plan V-Series FC initiator port usage, you need to ensure that you are using redundant FC initiator ports and that you meet the configuration requirements for FC initiator ports.

The requirements for FC port usage are as follows:

For...	The requirements are...
FC initiator port pair redundancy	Redundant FC initiator port pairs are required to connect a V-Series system to array LUNs.
FC initiator port setting for HBAs	All V-Series HBAs that are used to access disks or array LUNs must be set to <i>initiator</i> ports. You need to check the port settings on new systems because the factory might not always set the FC ports to initiators.
FC initiator ports connecting to target ports	<p>For releases prior to Data ONTAP 8.2, sharing an FC initiator port with multiple target ports is not supported. You need to use a different FC initiator port for every target port set on the same storage array and for every storage array.</p> <p>For Data ONTAP 8.2 and later releases, the same initiator port can connect to multiple target ports if the configuration is set up according to the rules for doing so. See Rules for sharing an FC initiator port with multiple target ports on page 47 for details.</p>
FC initiator port usage for different storage and devices	<p>A separate V-Series FC initiator port must be used to connect each of the following:</p> <ul style="list-style-type: none"> • Disk shelves • Array LUNs • Tape devices <p>If your V-Series model does not have enough internal ports for your needs, you must order an additional HBA.</p>

Related concepts

[How FC initiator ports are labeled](#) on page 45

Rules for sharing an FC initiator port with multiple target ports

Starting in Data ONTAP 8.2, connecting a V-Series FC initiator port to up to four target ports on *separate* storage arrays is supported. Connecting a V-Series FC initiator port to up to four target ports on the *same* storage array is also supported; however, it is possible to do so only on storage arrays that can present different sets of logical devices to an FC initiator based on the target port being accessed.

Being able to share a single FC initiator port with multiple target ports is useful for organizations that want to minimize the number of FC initiator ports used. For example, an organization that expects low I/O usage over an FC initiator port might prefer to connect an FC initiator port to multiple target ports instead of dedicating each FC initiator port to a single target port.

Note: Connecting a V-Series FC initiator port to multiple storage array target ports is not supported in Data ONTAP releases prior to 8.2.

There are limits on how many array LUNs can be visible over an FC initiator port. These limits vary according to release. If you think the total number of visible LUNs from all target ports sharing an FC initiator port might exceed 100, you should check the V-Series *Best Practice Guide* to determine the supported maximum number of array LUNs per FC initiator for the Data ONTAP release you are running. Then you should check the number of LUNs that are visible on the FC initiator port.

Rules: FC initiator port connected to multiple target ports on *separate* storage arrays

The rules for this configuration are as follows:

- All storage arrays must be in the same vendor model family.
Storage arrays in the same family share the same performance and failover characteristics. For example, members of the same family all perform active-active failover, or they all perform active-passive failover. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same. See the *V-Series Implementation Guide for Third-Party Storage* for details about the models in the storage array families for your vendor.
- Connecting a single V-Series FC initiator port to multiple target ports is not supported for MetroCluster configurations.
- A single FC initiator port on a V-Series system can connect to up to four target ports on multiple storage arrays.
- The best practice zoning recommendation is to put each FC initiator-target port pair in a separate zone (1:1), even if the same FC initiator is talking to multiple target ports.

Rules: FC initiator port connected to multiple targets ports on the *same* storage array

This configuration can be used only with storage arrays whose LUN masking, presentation, or host group capability allows for different LUN group presentations to the same FC initiator based on the target port being accessed.

Some storage arrays can present different sets of logical devices to an FC initiator based on the target port being accessed. These types of storage arrays allow the same FC initiator to be in multiple host groups. On storage arrays with this capability, it is possible for each V-Series FC initiator to access multiple array target ports on the same storage array, with each target port presenting a different LUN group to the FC initiator. See your storage array documentation to determine whether your storage array allows the same FC initiator to be in multiple host groups.

The rules for this configuration are as follows:

- A single FC initiator port on a V-Series system can connect to up to four target ports on a storage array.
- Connecting a single V-Series FC initiator port to multiple target ports is not supported for MetroCluster configurations.
- The best practice zoning recommendation is to put each FC initiator-target port pair in a separate zone (1:1), even if the same FC initiator is talking to multiple target ports.

Related concepts

Example configuration: shared FC initiator ports on page 48

Related tasks

Checking the number of array LUNs visible on an FC initiator port (clustered Data ONTAP) on page 52

Example configuration: shared FC initiator ports


Starting in Data ONTAP 8.2, one FC initiator port on a V-Series system can be connected to up to four target ports on separate storage arrays or, for some storage arrays, to four target ports on the same storage arrays.

The best practice zoning recommendation is to put each FC initiator-target port pair in a separate zone (1:1), even if the same FC initiator is talking to multiple target ports.

Shared FC initiator port attached to multiple targets ports on *separate* storage arrays

The following illustration shows connections and zoning for sharing an FC initiator port with target ports on *different* storage arrays.

The solid lines in the following illustration show the connections from the FC initiator ports on system vs1 and the dashed lines show the connections from the FC initiator ports on system vs2.



The following table shows the 1:1 zoning definitions for the example of one FC initiator port sharing multiple target ports on different storage arrays.

Zone	V-Series system	Storage array
Switch 1		
z1	vs1:0a	Storage array 1: Controller 1 Port 1A
z2	vs1:0a	Storage array 2: Controller 1 Port 1A
z3	vs2:0a	Storage array 1: Controller 1 Port 1B
z4	vs2:0a	Storage array 2: Controller 1 Port 1B
Switch 2		

Zone	V-Series system	Storage array
z5	vs1:0c	Storage array 1: Controller 2 Port 2A
z6	vs1:0c	Storage array 2: Controller 2 Port 2A
z7	vs2:0c	Storage array 1: Controller 2 Port 2B
z8	vs2:0c	Storage array 2: Controller 2 Port 2B


Shared FC initiator port attached to multiple target ports on the *same* storage array

This configuration can be used only with storage arrays whose LUN masking, presentation, or host group capability allows for different LUN group presentations to the same FC initiator based on the target port being accessed.

Some storage arrays can present different sets of logical devices to an FC initiator based on the target port being accessed. These types of storage arrays allow the same FC initiator to be in multiple host groups. On storage arrays with this capability, it is possible for each V-Series FC initiator to access multiple array target ports on the same storage array, with each target port presenting a different LUN group to the FC initiator. Check your storage array documentation to determine whether your storage array allows the same FC initiator to be in multiple host groups.

The following illustration shows connections and zoning for sharing an FC initiator port with multiple target ports on the *same* storage array. In this example, zoning definitions are configured as 1:1, that is, one FC initiator to one target port.

The solid lines in the following illustration show the connections from the FC initiator ports on system vs1 and the dashed lines show the connections from the FC initiator ports on system vs2. Two LUN groups are required for this configuration.



The following table shows the 1:1 zoning definitions for the example of one FC initiator port sharing multiple target ports on the same storage array.

Zone	V-Series system and port	Storage array and port
Switch 1		
z1	vs1:0a	Controller 1 Port 1A
z2	vs1:0a	Controller 1 Port 1C
z3	vs2:0a	Controller 1 Port 1B
z4	vs2:0a	Controller 1 Port 1D
Switch 2		
z5	vs1:0c	Controller 2 Port 2A
z6	vs1:0c	Controller 2 Port 2C
z7	vs2:0c	Controller 2 Port 2B
z8	vs2:0c	Controller 2 Port 2D

Related concepts

[Rules for sharing an FC initiator port with multiple target ports](#) on page 47

Checking the number of array LUNs visible on an FC initiator port (clustered Data ONTAP)

The supported number of array LUNs that can be visible on an FC initiator port varies for different Data ONTAP releases.

About this task

If you think that the total number of visible LUNs from all target ports sharing the FC initiator port might exceed 100, you should check how many array LUNs per FC initiator are supported for your Data ONTAP release. Data ONTAP supports 100 array LUNs per FC initiator for all Data ONTAP releases (which is why you do not have to check if you do not exceed 100). Data ONTAP supports significantly more array LUNs per FC initiator for some releases; you need to make sure that you do not exceed the number supported for the release you are running.

Steps

1. Check the V-Series *Best Practice Guide* to determine the supported maximum number of array LUNs per FC initiator for the Data ONTAP release you are running.
2. Enter the following command to check the visible number on the FC initiator port:
`storage array config show -initiator initiator_number`

Example

```
storage array config show -initiator 0a
```

3. If there is more than one array LUN group for a node, add together the number of array LUNs for all the LUN groups for that node to determine the combined total of array LUNs visible to that node's specified FC initiator.

Example

The following example shows the output for FC initiator 0a for all nodes. To determine the number of array LUNs visible over a specific FC initiator for a *particular* node, you need to look at the entries for that node from all the target ports shown for that node. For example, to find the number of array LUNs seen over vgv3070f51-01:0a, you would add the LUN count of 24 for LUN group 1 (HP) to the LUN count of 1 for LUN group 2 (IBM_1722_1), for a total of 25 array LUNs visible over vgv3070f51-01:0a.

You would follow the same process to determine the number of array LUNs visible over vgv3070f51-02:0a for LUN group 0 and LUN group 2, which is also 25.

```
vgv3070f51:> storage array config show -initiator 0a
Node      LUN      LUN
          Group Count   Array Name   Array Target Port Initiator
```

vgv3070f51-01	1	24	HP	50014380025d1508	0a
	2	1	IBM_1722_1	200600a0b819e16f	0a
vgv3070f51-02	0	24	HP	50014380025d1508	0a
	2	1	IBM_1722_1	200600a0b819e16f	0a

Related concepts

[Rules for sharing an FC initiator port with multiple target ports](#) on page 47

[Example configuration: shared FC initiator ports](#) on page 48

What to include in a port-to-port plan for connecting to storage arrays

When planning how to connect your V-Series system to a storage array, your port-to-port connectivity plan should address redundancy, pathing, and other guidelines.

The requirements to set up connections are as follows:

- Each connection in a redundant port pair on the storage array must be attached to a different FC initiator port on the V-Series system.
- The ports used on the FC switches must be redundant.
- Connectivity must be set up to avoid a SPOF.

Be sure that the ports on the storage array that you select to access a given LUN are from different components that could represent a single point of failure, for example, from alternate controllers, clusters, or enclosures. The reason is that you do not want all access to an array LUN to be lost if one component fails.

- The number of paths cannot exceed the number of paths supported for your Data ONTAP release and operating mode.

Clustered Data ONTAP 8.1 and later supports two or four paths to an array LUN. Data ONTAP 8.1 and later operating in 7-Mode, and releases prior to Data ONTAP 8.1, support only two paths to an array LUN.

- If you want to share an FC initiator port with multiple target ports, you must follow the rules discussed in [Rules for sharing an FC initiator port with multiple target ports](#) on page 47.
- If your storage array supports fewer LUNs per host group per port than the number of LUNs that the V-Series systems will be using, you need to add additional cables between the V-Series system and the storage array.

See the *V-Series Implementation Guide for Third-Party Storage* for information about any unique port-to-port connectivity requirements for your storage array type.

Related tasks

[Connecting a V-Series system to a storage array](#) on page 74

What to include in a port-to-port plan for connecting to native disk shelves

When planning how to connect your V-Series system to native disk shelves, your port-to-port connectivity plan should address redundancy and other guidelines.

Number of FC initiator ports needed for disks

The number of FC initiator ports required to connect a V-Series system to a disk shelf depends on whether your V-Series system is a stand-alone system or in an HA pair

Configuration	Number of FC initiator ports
Stand-alone system	<ul style="list-style-type: none">• If you are using one loop: one FC initiator port.• If you are connecting two loops: two FC initiator ports, one for each loop.
In an HA pair	<ul style="list-style-type: none">• If you are using Multipath Storage, two FC initiator ports for each loop.• If you are not using Multipath Storage, one FC initiator port for each controller for each loop in the configuration.

Note: Starting in Data ONTAP 8.1, the `fc-non-array-adapter-list` environment variable is no longer required for V-Series systems with native disk shelves that use Multipath Storage.

Connections between a V-Series system and disks

You follow the same processes to cable a V-Series system to a native disk shelf as you would to cable a FAS system to a native disk shelf. When creating your port-to-port connectivity scheme, this guide and the Data ONTAP and hardware guides in the following table provide information about setup and management of disks and disk shelves.

For information about...	See...
Disk support, including supported disk speeds, and disk capacity	The NetApp Support Site for information about disk installation and management and FAS system management.
Installing a 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.

For information about...	See...
Connecting a disk shelf to a stand-alone V-Series system	The <i>Hardware Overview</i> for your platform.
Connecting an HA pair to a disk shelf	The <i>Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode</i> or the <i>Clustered Data ONTAP High-Availability Configuration Guide</i> , as appropriate.
Adding a disk shelf	The appropriate guide for your disk shelf type.
Moving a disk shelf	<p>The appropriate guide for your disk shelf type.</p> <p>Note: The <code>bootarg.storageencryption.support</code> variable must be set to <code>true</code> to be able to move an FDE disk shelf from a FAS system to a V-Series system. If it is not, all disks are marked as failed.</p>
Disk management	The <i>Storage Management Guide</i> for the mode and release of Data ONTAP that your system is running.

Zoning guidelines

A common error when installing a configuration with storage arrays is to misconfigure zoning. For successful zoning, you should follow the guidelines provided for setting up zoning between a V-Series system and a storage array. The Interoperability Matrix at support.netapp.com contains information about specific switch guidelines and potential issues. The *V-Series Best Practice Guide* also contains information about zoning.

Zoning requirements

Configuring zoning on a Fibre Channel (FC) switch enables you to define paths between connected nodes, restricting visibility and connectivity between devices connected to a common FC SAN.

Zoning enables a storage administrator to restrict the array LUNs that a particular V-Series system can see. Data ONTAP requires that an array LUN be visible on only one target port for each initiator port.

The requirements for configuring zoning in a V-Series deployment are as follows:

- The *Interoperability Matrix* must identify a switch and the switch firmware as supported.
- Zoning must be configured to restrict each initiator port to a single target port on each storage array.
- On the switch, ports on the V-Series system and ports on the storage array must be assigned to the same zone.

This enables the V-Series systems to see the LUNs on the storage arrays.

- When storage array ports are shared across heterogeneous systems, array LUNs from the V-Series system cannot be exposed to other systems, and vice versa.
LUN security or array LUN masking must be used to ensure that only the array LUNs that are for Data ONTAP storage are visible to the V-Series systems.
- Do not put a host configuration port (such as the HP EVA management port) in the same zone as a target port.

Zoning can be configured by specifying WWNs (worldwide names) or ports.

You can find guides for configuring specific vendor switches on the NetApp Support Site at support.netapp.com.

Related concepts

[*Zoning recommendation for a configuration with storage arrays*](#) on page 57

[*Example of zoning in a configuration with storage arrays*](#) on page 58

[*Rules for sharing an FC initiator port with multiple target ports*](#) on page 47

[*Example configuration: shared FC initiator ports*](#) on page 48

[*Dependency between zone and host group definitions*](#) on page 120

Related references

[*Relationship between zoning and host group configuration*](#) on page 120

Zoning recommendation for a configuration with storage arrays

The recommended type of zoning for a configuration with storage arrays is 1:1 zoning. With 1:1 zoning, each zone contains a single V-Series FC initiator port and a single storage array target port.

The benefits of creating 1:1 zoning are as follows:

- You limit the number of ports over which a specific array LUN can be accessed.
- There are discovery and boot time improvements because the V-Series FC initiators do not attempt to discover each other.

Related concepts

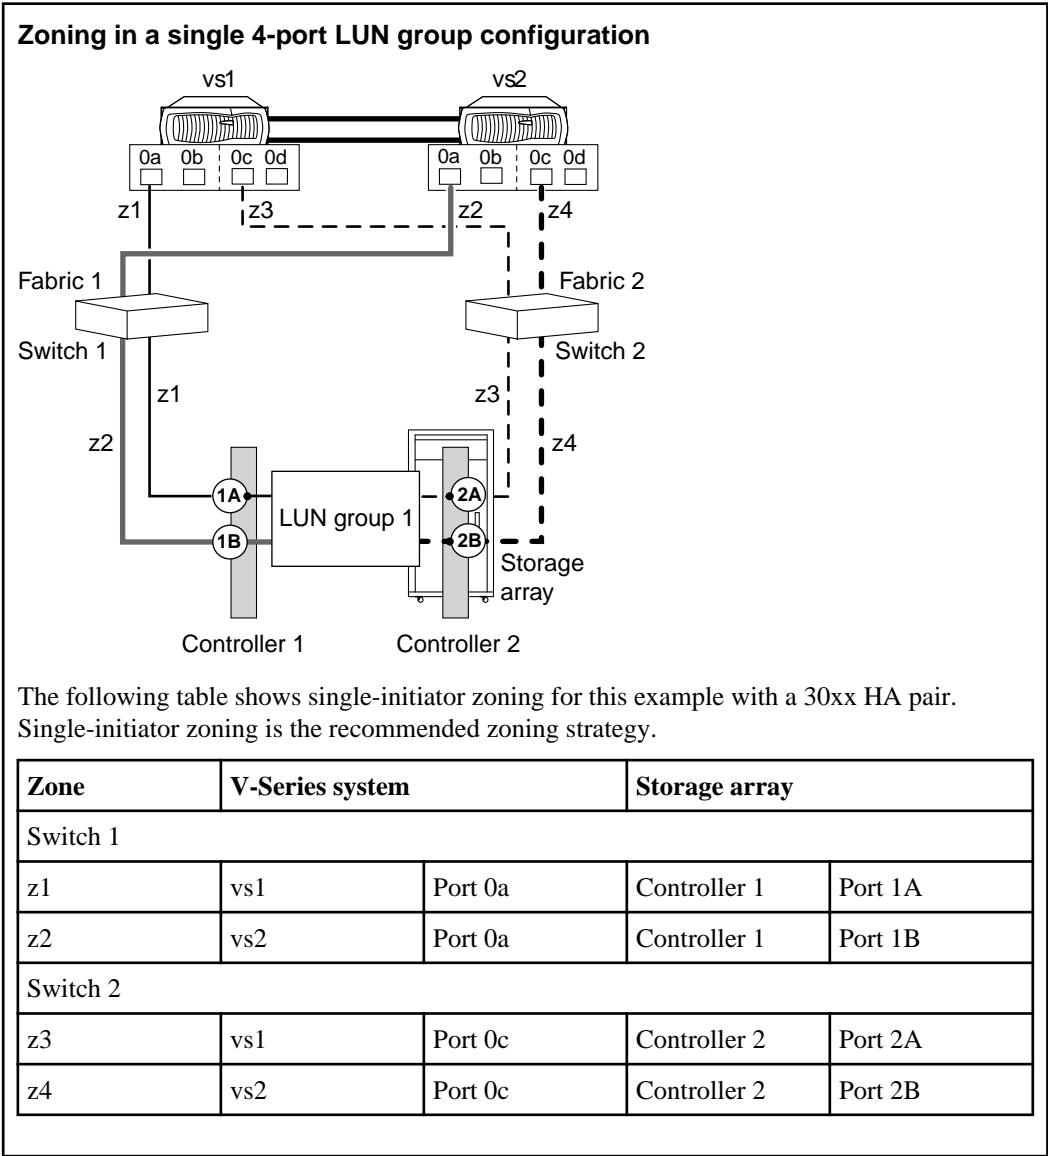
[*Zoning requirements*](#) on page 56

[*Example of zoning in a configuration with storage arrays*](#) on page 58

[*Example configuration: shared FC initiator ports*](#) on page 48

Example of zoning in a configuration with storage arrays

Using LUN security when configuring the switches for zoning ensures that different hosts do not see LUNs mapped to another host.



Related concepts

Zoning recommendation for a configuration with storage arrays on page 57

Dependency between zone and host group definitions on page 120

Related references

Relationship between zoning and host group configuration on page 120

Determining the array LUNs for specific aggregates

There are a number of rules about mixing different types of storage in aggregates that are unique to V-Series systems that use array LUNs. You must understand these rules when planning which array LUNs and disks to add to which aggregates.

Rules for mixing storage in aggregates for V-Series systems

When planning for aggregates, you must consider the rules for mixing storage in aggregates. You cannot mix different storage types or array LUNs from different vendors or vendor families in the same aggregate.

Adding the following to the same aggregate is not supported:

- Array LUNs and disks
- Array LUNs with different checksum types
- Array LUNs from different drive types (for example, FC and SATA) or different speeds
- Array LUNs from different storage array vendors
- Array LUNs from different storage array model families

Note: Storage arrays in the same family share the same performance and failover characteristics. For example, members of the same family all perform active-active failover, or they all perform active-passive failover. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same.

Related concepts

[*How the checksum type is determined for aggregates with array LUNs*](#) on page 61

[*Aggregate rules when the storage arrays are from the same family*](#) on page 63

[*Aggregate rules when the storage arrays are from different vendors or families*](#) on page 64

How the checksum type is determined for aggregates with array LUNs

Each Data ONTAP aggregate has a checksum type associated with it. The aggregate checksum type is determined by the checksum type of the array LUNs that are added to it.

The checksum type of an aggregate is determined by the checksum type of the first array LUN that is added to the aggregate. The checksum type applies to an entire aggregate (that is, to all volumes in the aggregate). Mixing array LUNs of different checksum types in an aggregate is not supported.

- An array LUN of type *block* must be used with block checksum type aggregates.
- An array LUN of type *zoned* must be used with advanced zoned checksum (AZCS or advanced_zoned) type aggregates.

Note: Prior to Data ONTAP 8.1.1, zoned checksum array LUNs were used with ZCS (zoned) type aggregates. Starting in 8.1.1, any new aggregates created with zoned checksum array LUNs are AZCS aggregates. You can, however, add zoned checksum array LUNs to existing ZCS aggregates.

Before you add array LUNs to an aggregate, you must know the checksum type of the LUNs you want to add, for the following reasons:

- You cannot add array LUNs of different checksum types to the same aggregate.
- You cannot convert an aggregate from one checksum type to the other.

When you create an aggregate you can specify the number of array LUNs to be added, or you can specify the names of the LUNs to be added. If you want to specify a number of array LUNs to be added to the aggregate, the same number or more array LUNs of that checksum type must be available.

Related concepts

[*Rules for mixing storage in aggregates for V-Series systems*](#) on page 60

[*Aggregate rules when the storage arrays are from the same family*](#) on page 63

[*Aggregate rules when the storage arrays are from different vendors or families*](#) on page 64

Checking the checksum type of spare array LUNs

If you plan to add a spare array LUN to an aggregate by specifying its name, you need to make sure that the checksum type of the array LUN you want to add is the same as the aggregate checksum type.

About this task

You cannot mix array LUNs of different checksum types in an array LUN aggregate. The checksum type of the aggregate and the checksum type of the array LUNs added to it must be the same.

If you specify a number of spare array LUNs to be added to an aggregate, by default Data ONTAP selects array LUNs of the same checksum type as the aggregate.

Note: Data ONTAP 8.1.1 and later supports a new checksum scheme called *advanced zoned checksum* (AZCS). Existing zoned checksum aggregates are still supported. The checksum type of all newly created aggregates using zoned checksum array LUNS is AZCS, which provides more functionality than the “version 1” zoned checksum type that was supported in previous releases and continues to be supported for existing zoned aggregates. Zoned checksum spare array LUNs added to an existing zoned checksum aggregate continue to be zoned checksum array LUNs. Zoned checksum spare array LUNs added to an AZCS checksum type aggregate use the AZCS checksum scheme for managing checksums.

Step

1. Check the checksum type of the spare array LUNs by entering the following command:

For...	The command is...
7-Mode	<pre>aggr status -s</pre> <p>The output shows information about the spare disks or array LUNs on the system, including the checksum type of each. You can add a block checksum array LUN to a block checksum aggregate and a zoned array LUN to either a zoned aggregate or an AZCS checksum aggregate.</p>
Clustered Data ONTAP	<pre>storage disk show -fields checksum-compatibility -container-type spare</pre> <p>You can add a block checksum array LUN to a block checksum aggregate and a zoned array LUN to either a zoned checksum aggregate or an AZCS checksum aggregate.</p>

Aggregate rules when the storage arrays are from the same family

Specific rules apply to how you can lay out array LUNs in aggregates when the storage arrays are from the same storage array vendor and model family.

If your storage arrays are from the same vendor, the rules for adding array LUNs to aggregates are as follows:

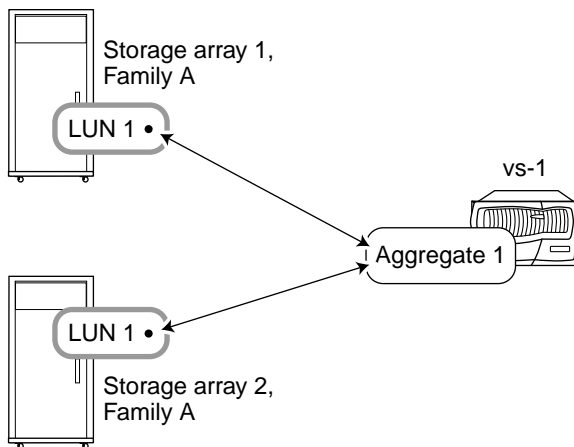
- You can mix array LUNs from the storage arrays in the same aggregate if the storage arrays are in the same family.
- You can separate the array LUNs into different aggregates.

The following examples show some options for laying out array LUNs in aggregates when the storage arrays behind a V-Series system are in *the same vendor family*.

Note: For simplicity, the illustrations show only two storage arrays; your deployment can include more storage arrays.

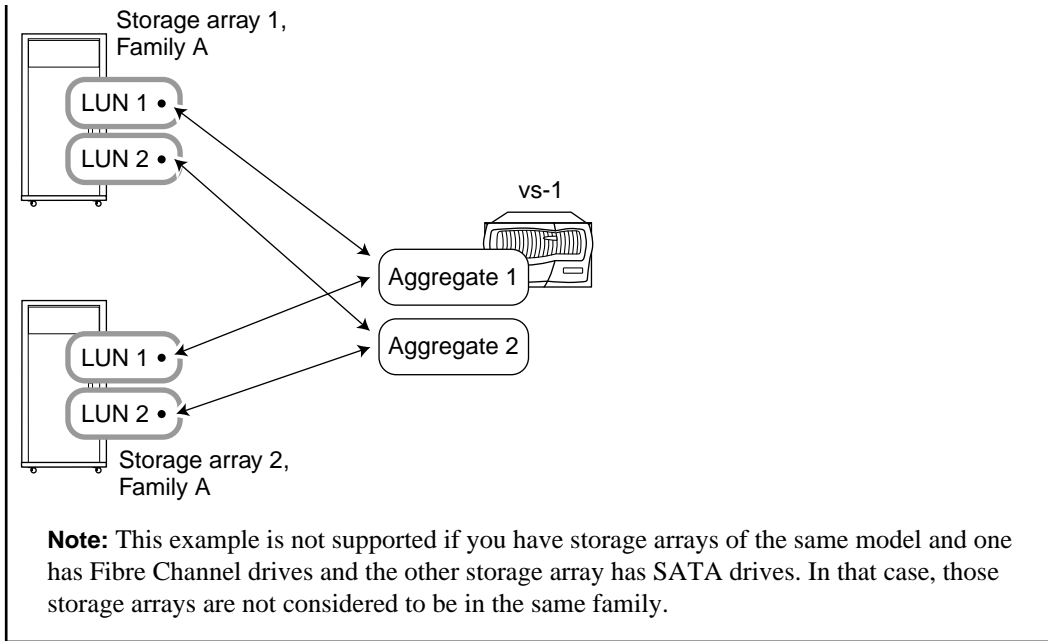
Example 1: Add LUNs from all storage arrays to a single aggregate

As shown in the following illustration, you can create one aggregate, then add all LUNs from all the storage arrays in the same family to the same aggregate:



Example 2: Distribute and mix LUNs from the storage arrays over multiple aggregates

As shown in the following illustration, you can create multiple aggregates, then distribute and mix the array LUNs from the different storage arrays in the same family over the aggregates:



Related concepts

[Rules for mixing storage in aggregates for V-Series systems](#) on page 60

[Aggregate rules when the storage arrays are from different vendors or families](#) on page 64

Aggregate rules when the storage arrays are from different vendors or families

Specific rules apply to how you can lay out array LUNs in aggregates when the storage arrays are from different vendors or from different storage array families from the same vendor.

The following rules apply if your storage arrays are from different vendors or different families from the same vendor:

- You cannot mix array LUNs from storage arrays from different vendors, or from different families of the same vendor, in the same aggregate.
- You can associate the aggregate containing the root volume with any of the storage arrays, regardless of the family type of the storage array.

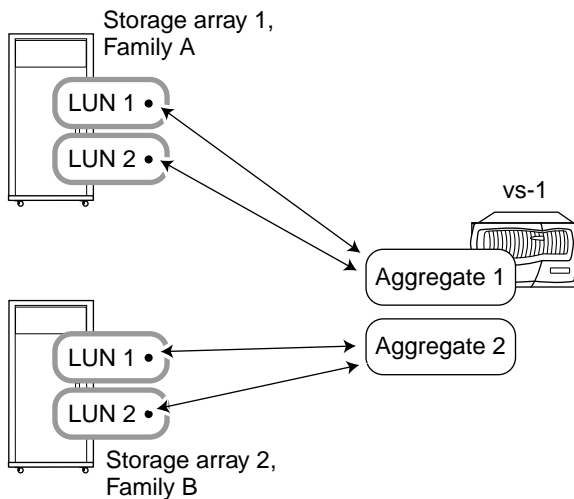
Note: When you create your aggregate, be sure that you explicitly specify the IDs of the array LUNs that you want to add to the aggregate. Do not use the parameters for specifying the number and size of array LUNs to be picked up because the system might automatically pick up LUNs from a different family or from a different vendor's storage array. After array LUNs

from different families or vendors are in the same aggregate, the only way to fix the problem of mixed array LUNs in an aggregate is to destroy the aggregate and re-create it.

The following examples show options for how to lay out array LUNs in aggregates when the storage arrays are from *different vendors* or *from different families from the same vendor*.

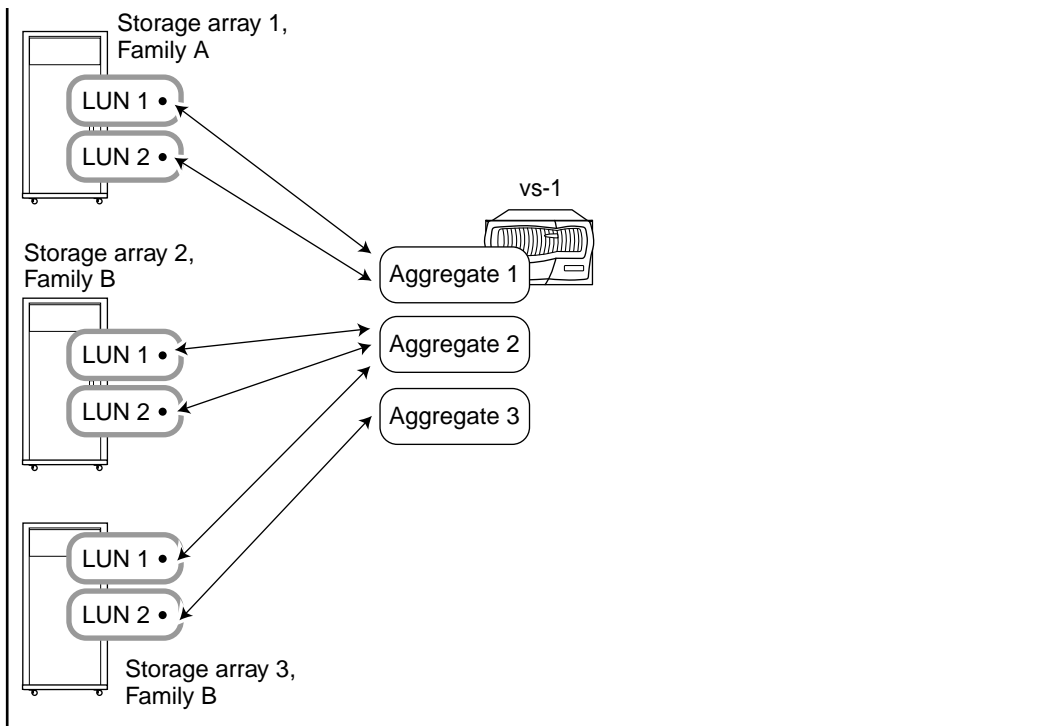
Example 1: LUNs from the two storage arrays are in different aggregates

In this example, some LUNs for Data ONTAP are from Storage array 1, Family A, while the other LUNs for Data ONTAP are from Storage array 2, Family B. The LUNs from the two storage arrays cannot be added to the same aggregate because the two storage arrays are from different families from the same vendor. The same would be true if the two storage arrays were from different vendors.



Example 2: Some LUNs can be mixed in the same aggregate and some cannot

In this example, one storage array is from Family A and the other two storage arrays are from Family B. The LUNs from the Family A storage array cannot be added to the same aggregate as the LUNs from a Family B storage array because the storage arrays are from different families. However, LUN 1 of storage array 3 can be assigned to aggregate 2, which also contains LUNs from storage array 2, because the two storage arrays are in the same family.



Related concepts

[Rules for mixing storage in aggregates for V-Series systems](#) on page 60

[Aggregate rules when the storage arrays are from the same family](#) on page 63

Determining whether to use neighborhoods (7-Mode in 8.x)

A neighborhood is a logical entity that enables V-Series systems to see the same array LUNs. Ordinarily, systems that are not part of an HA pair have no relationship with each other; they cannot see the same array LUNs and you cannot load balance between them. A neighborhood makes it possible to do these things.

Neighborhoods do not pertain to configurations with only one stand-alone system or only one HA pair and no other systems.

Note: Nodes in clustered Data ONTAP provide like functionality to neighborhoods, which are supported only for 7-Mode V-Series systems.

What a V-Series neighborhood is

A V-Series neighborhood is a logical relationship between up to six V-Series systems that enables you to easily reassign ownership of array LUNs from one neighborhood member to another through Data ONTAP. The members of a neighborhood can be stand-alone systems or HA pairs. Neighborhood functionality is limited to storage arrays.

With a neighborhood, you can also transparently load balance data service among the V-Series systems in a neighborhood by moving vFiler units among neighborhood members. The physical storage always remains on the storage array.

Although neighborhood members see the same array LUNs, the systems outside of an HA pair cannot see each other's disks.

The neighborhood relationship does not provide any failover between neighborhood members if a member becomes unavailable. Failover of services is a function of the relationship between two nodes in an HA pair, and can occur only between the two nodes in an HA pair.

Requirements for V-Series neighborhoods

If you are thinking about using neighborhoods, you need to check to be sure they are supported for your storage array and for the Data ONTAP release running on your V-Series systems.

For...	This is supported...
Data ONTAP releases	<ul style="list-style-type: none"> Releases 8.0 and later operating in 7-Mode Data ONTAP releases earlier than 8.x (which support only 7-Mode) <p>Mixing Data ONTAP releases from different release families is not supported. For example, mixing Data ONTAP release 7.3.x with an 8.x release is not supported.</p>
Storage arrays	Storage arrays that the <i>Interoperability Matrix</i> does not indicate are not supported with neighborhoods.
Platforms	All platform types. Mixed platform types are supported.
Configurations	<p>Any combination of stand-alone systems and HA pairs.</p> <p>Neighborhoods are not supported for a MetroCluster configuration.</p>
Number of V-Series systems	A maximum of six systems.
Number of LUNs and disks supported	No more than the limits for the maximum devices in a neighborhood and the neighborhood maximum LUN limit.

Related concepts

[Planning the number of array LUNs and disks in a neighborhood](#) on page 68

Planning the number of array LUNs and disks in a neighborhood

When determining the number of LUNs that storage arrays can present to the V-Series systems in your neighborhood, you must consider both the *platform maximum assigned device limit* and the

neighborhood maximum LUN limit. You cannot exceed the limits for the maximum devices in a neighborhood and the neighborhood maximum LUN limit.

The relationship between the following two limits determines the limits for your neighborhood. You should use the limits information in the *Hardware Universe* (formerly the *System Configuration Guide*) at support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml to calculate the limits for your neighborhood.

Type of limit...	What it is...
Platform maximum assigned device limit	<p>Maximum number of array LUNs and disks combined that can be <i>assigned</i> to a stand-alone system or nodes in an HA pair.</p> <p>You cannot exceed the platform maximum assigned device limit.</p> <p>Furthermore, the combined total of array LUNs and disks assigned to a system cannot exceed the <i>neighborhood maximum LUN limit</i>.</p>
Neighborhood maximum LUN limit (the <i>visible</i> limit)	<p>Maximum number of array LUNs that <i>all</i> systems in the neighborhood combined can <i>see</i>, and the maximum number of array LUNs that can be assigned to all systems in the neighborhood combined.</p>

Platform maximum assigned device limit

For every platform, Data ONTAP hard codes the maximum number of devices (disks and array LUNs combined) that can be *assigned* to a stand-alone system or an HA pair (through the Data ONTAP disk ownership feature).

The platform maximum assigned device limit does not limit the number of disks and array LUNs that the systems in a V-Series neighborhood can *see* (that is, the visible limit); it limits only the number of disks and LUNs that you can assign to the platform. The visible limit is determined by the neighborhood maximum LUN limit.

Note: The platform maximum assigned device limit applies whether or not a V-Series system is in a neighborhood. The platform maximum assigned device limit is the same for a stand-alone system and an HA pair because each node in the pair must be able to handle its storage and its partner's storage if the partner becomes unavailable.

Neighborhood maximum LUN limit

The neighborhood maximum LUN limit has the following two components:

- It is the maximum *visible* limit for the neighborhood.

This limit is the maximum number of the same array LUNs that V-Series systems in a neighborhood are allowed to *see*. (All members of the neighborhood see all the same array LUNs but a LUN can be assigned to only one of them.)

Individual V-Series systems in the neighborhood that have disks attached cannot see more array LUNs and disks combined than the maximum visible limit.

- It is the *maximum assigned LUN limit for all the systems in the neighborhood combined*. The platform maximum assigned device limit, not the neighborhood maximum LUN limit, dictates the maximum number of disks and array LUNs that can be *assigned* to a stand-alone system or HA pair. However, in the context of the neighborhood, you might not be able to assign the maximum number of devices that the platform can support because the combined total of assigned devices must not exceed the neighborhood limit.

Neighborhood members can never see more array LUNs than the neighborhood maximum LUN limit for the platform type. However, the following factors can reduce the maximum LUN limit from that number:

- If the systems in the neighborhood are mixed platform types
- If there are neighborhood restrictions for the storage arrays that are presenting LUNs to the neighborhood systems
- If disks are connected to the V-Series systems in the neighborhood

Note: The lowest limit based on any factor becomes the maximum LUN limit for your neighborhood.

Example of how the limits work together

Assume that you have two stand-alone systems in the neighborhood, and the *maximum assigned limit* for the platform is 600 devices. If the *neighborhood LUN limit* is 1,000 array LUNs, you cannot assign 600 array LUNs to each system because the total assigned LUNs for the two systems would be 1,200, which is 200 LUNs more than the 1,000 neighborhood maximum LUN limit.

Related concepts

[Requirements for V-Series neighborhoods](#) on page 68

How to establish a neighborhood

A neighborhood exists only when a V-Series system can see array LUNs that belong to another V-Series system that is not its partner in an HA pair. To establish a neighborhood, the storage arrays

and switches must be configured to enable all V-Series systems in the same neighborhood to see the same array LUNs.

Data ONTAP configuration to establish a neighborhood

No explicit configuration is required on a V-Series system to support neighborhoods. The underlying functionality that enables the V-Series systems to operate as a neighborhood is the Data ONTAP disk ownership feature (assigning array LUNs to a specific V-Series system).

Note: After you determine the maximum number of LUNs that the storage arrays can present to your neighborhood, be sure to communicate that information to the storage array administrators.

Storage array configuration to establish a neighborhood

The storage array administrator must configure one or more storage arrays to present the same LUNs to the V-Series systems that you want to be in the neighborhood.

How a storage array administrator creates and presents LUNs to hosts varies on different storage arrays. Typically the storage array administrator specifies the FC initiator ports of a number of V-Series systems to be in the same *host group* on the storage array. This host group configuration enables all the systems to see the same LUNs.

The storage array administrator must also set the storage array access controls so that all the V-Series systems can see the same array LUNs.

Switch configuration to establish a neighborhood

If your configuration is fabric-attached, you must zone switch ports that connect to your V-Series system FC initiator ports. This ensures that all V-Series systems in the neighborhood can see the same array LUNs.

Note: It is recommended that you use single-initiator zoning, which limits each zone to a single V-Series system FC initiator port.

Calculation of the visible limit for both array LUNs and disks

Calculating the visible limit for a neighborhood requires that you consider the total number of disks and array LUNs combined.

Assume that in this neighborhood example there are two HA pairs and one stand-alone system, with disks connected to each system. The following table shows the maximum value used in this example for each limit.

Limit type	Maximum allowed
Neighborhood maximum LUN limit (<i>visible</i> device limit)	1,200

Limit type	Maximum allowed
Platform maximum assigned limit for the platform type (disks and array LUNs <i>assigned</i> combined)	500

The following table shows the calculations to determine whether the planned visible limit of array LUNs and disks in this neighborhood is within the supported limits. For simplicity, assume that there are no restrictions related to the storage arrays that present LUNs to the neighborhood members. (This table does not show the number of disks and LUNs assigned to each system.)

System	Number of array LUNs seen	Number of disks connected to it	Total disks and array LUNs	Within visible limits?
HA pair 1				
vs-X	1,000	200	1,200	Yes
vs-Y	1,000	150	1,150	Yes
HA pair 2				
vs-A	1,000	200	1,200	Yes
vs-B	1,000	350	1,350	No (The maximum visible device limit is 1,200.)
Stand-alone system				
vs-standalone	1,000	100	1,100	Yes

All neighborhood members see all the same array LUNs. Therefore, for every system in this example, the 1,000 LUNs is counted toward the maximum visible limit of 1,200. (When disks and array LUNs are both used for storage, you can think of this as the device limit).

Disks are counted on a system-by-system basis toward the visible device limit. Therefore, the disks are counted toward the limit for the system to which the disks are connected, and not toward the visible device limit of any other systems.

For four systems in this example neighborhood, the number of visible devices is the same as or less than the visible device limit. However, if 1,000 array LUNs were presented to the neighborhood, you would not be able to connect 350 disks to system vs-B in HA pair 2. The reason is that the total of array LUNs and disks seen by system vs-B would be 1,350, which is more than the visible device limit of 1,200.

No more than 500 array LUNs and disks combined can be assigned to any one system in this example (so that the platform maximum assigned limit is not exceeded). However, you cannot assign 500 devices to each of the systems in the neighborhood because the total assigned devices would be

2,500. The combined assigned devices for the neighborhood cannot exceed the neighborhood LUN limit, so the total assigned devices must be 1,200 or less.

Connecting a V-Series system to a storage array

Connecting your V-Series system to a storage array involves cabling the V-Series system, switches, and storage arrays together and connecting additional devices such as tape backup devices.

About this task

This procedure describes how to connect a V-Series system to a storage array through two paths, with each initiator dedicated to one target port. Starting in Data ONTAP you can connect one FC initiator port to multiple target ports, if you follow the guidelines for doing so.

Steps

1. Identify the onboard ports and expansion adapter ports for your V-Series model.
2. For an HA pair, do the following:
 - a) Check the *Hardware Universe* to ensure that the HA interconnect adapter is in the correct slot for your system in an HA pair.
 - b) Plug one end of the optical cable into one of the local node’s HA adapter ports, then plug the other end into the partner node’s corresponding HA adapter port.

Note: Do not cross-cable the HA interconnect adapter. Cable the local node ports only to the identical ports on the partner node.
 - c) Repeat the previous step for the two remaining ports on the HA adapter.
3. Locate the ports on the storage array that you want to use to connect the V-Series system to the storage array.
4. Connect the V-Series system to the switches as shown in the following table:

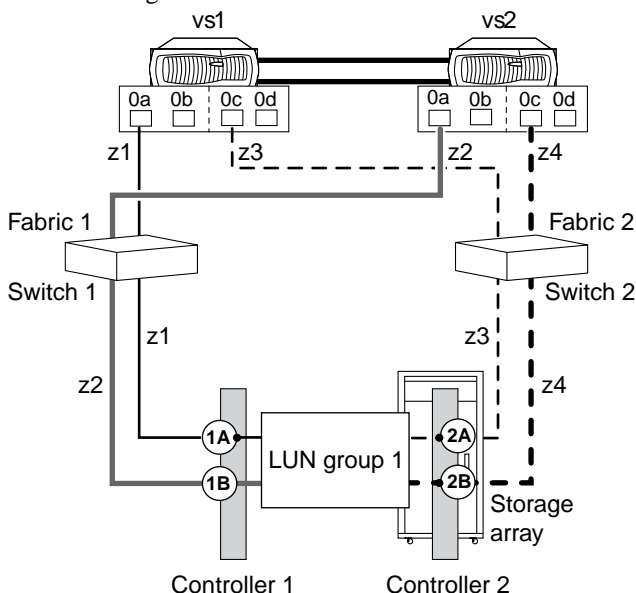
For...	Follow these steps...
A stand-alone system	<ol style="list-style-type: none">a. Connect one cable from one FC initiator port on the V-Series system to a port on Switch 1.b. Connect another cable from a redundant FC initiator port to a port on Switch 2.

For...	Follow these steps...
An HA pair	<ol style="list-style-type: none"> On the first node in the HA pair, connect one cable from one FC initiator port to a port on Switch 1. Connect another cable from a redundant FC initiator port on the same node to a port on Switch 2. On the second node in the HA pair, connect one cable from one FC initiator port to a port on Switch 1. Connect another cable from a redundant FC initiator port on the same node to a port on Switch 2.

5. Connect the switches to the storage array by using the instructions in the following table and, for an HA pair, the illustration following the table:

For a stand-alone system...	For an HA pair...
<ol style="list-style-type: none"> Connect Switch 1 to the storage array controller 1, port 1A. Connect Switch 1 to the storage array controller 2, port 1A. 	<ol style="list-style-type: none"> Connect Switch 1 to the storage array controller 1, port 1A. Connect Switch 2 to the storage array controller 2, port 1A. Connect Switch 1 to the storage array controller 1, port 2B. Connect Switch 1 to the storage array controller 1, port 2B.

The following illustration shows the connections for an HA pair.



6. Optional: Connect the V-Series system to a tape backup device through a separate FC initiator port or SCSI tape adapter.
7. Connect a console cable to the console port on each V-Series system to the RJ-45 to DB-9 adapter that is included with your system.
8. Install the cable management tray as follows:
 - a) Pinch the arms of the tray and fit the holes in the arms through the motherboard tray pins.
 - b) Push the cables into the cable holders, thread the adapter cables through the top rows of the cable holders, and thread the port cables through the lower cable holders.
9. Connect the V-Series system to an Ethernet network by plugging the network cable into the networking port.

If you are connecting more than one network cable to the network, you need to connect to the ports sequentially. Use the cable management tray to direct all the cabling from your system.
10. Optional: Connect a remote management device from the back of the V-Series system to the network by using an Ethernet cable.

Note: The network switch port for the RLM connection must negotiate down to 10/100 or autonegotiate.

11. Verify that the storage array is configured and connected properly, and that it is powered on.

Note: Your configured and connected storage array must be powered on before you power on your V-Series system. See your storage array documentation for how to power on the storage array. The V-Series system expects these units to be ready for input/output when it powers on and performs its reset and self-test.

12. If your deployment includes switches, make sure that all switch IDs are set, then turn on each switch 10 minutes apart from one another.
13. If applicable, turn on any tape backup devices.
14. For each power supply on the V-Series system, do the following:
 - a) Ensure that the power switch is in the Off (O) position.
 - b) Connect the socket end of the power cord to the power plug on the power supply.
 - c) Secure the power cord with the retaining adjustable clip on the power supply.
 - d) Plug the other end of the power cord into a grounded electrical outlet.

Note: To obtain power supply redundancy, you must connect the second power supply to a separate AC circuit.

15. Start a communications program.

You must use some form of communications program to be able to perform initial network setup and V-Series configuration. You can start a communications program through a remote management device after connecting to the serial port.

16. Turn the power switch on the V-Series system to the On (I) position.

The system verifies the hardware and loads the operating system.

17. If the storage array does not automatically discover V-Series system WWNs after you connect the V-Series system to the storage array, you must obtain the WWNs manually.

After you finish

Continue with the appropriate setup of Data ONTAP management features and to work with the storage arrays.

Related concepts

What to include in a port-to-port plan for connecting to native disk shelves on page 54

What to include in a port-to-port plan for connecting to storage arrays on page 53

Settings for connecting to an ASCII terminal console on page 135

Related tasks

Obtaining WWNs manually on page 134

Commands for checking back-end configuration (clustered Data ONTAP 8.1 and later)

Starting in Data ONTAP 8.1, several clustered Data ONTAP commands provide information about the storage array configuration, including back-end configuration errors. These commands are particularly useful during installation verification and troubleshooting.

The `storage array config show` command is the first command to use during installation verification. It is also the first command to use if you notice that your system is not working as you expect, or if you receive an error message.

Commands that are particularly useful for installation verification and troubleshooting are shown in the following table:

Command	Description
<code>storage array config show</code>	<p>Provides information, at the storage array level, about the configuration of the back-end devices in a V-Series deployment. This command shows how storage arrays connect to the cluster.</p> <p>If Data ONTAP detects an issue that would prevent the V-Series systems and storage arrays from operating properly together, <code>storage array config show</code> instructs you to run <code>storage errors show</code> to obtain details about the error.</p> <p>This command is also useful for verifying that the configuration is set up as you intended. For example, you can look at the output to confirm that the number of array LUN groups that you intended were created.</p>
<code>storage array path quiesce</code>	<p>Temporarily suspends I/O to a specific array LUN on a specific path. The path becomes active again on reboot or by running <code>storage array path resume</code>.</p> <p>Some storage arrays require ceasing I/O for a period of time to remove or move an array LUN.</p> <p>Note: The <code>storage array path quiesce</code> command cannot be used with IBM DS storage arrays.</p>

Command	Description
<code>storage array path resume</code>	<p>Allows I/O to start flowing again; this is the inverse of quiesce. The <code>storage array path resume</code> command is used primarily for hardware maintenance (for example, cable or GBIC pulls) or after an accidental quiescence of a path to an array LUN occurs.</p> <p>It is not always necessary to run this command after quiescing a path. For example, Data ONTAP can discover a newly mapped array LUN.</p>
<code>storage array show</code>	<p>Displays information about the storage arrays that are visible to the cluster, for example, name, vendor, model, and failover type.</p>
<code>storage disk show</code>	<p>Entering <code>storage disk show</code> without parameters shows the following for all disks and array LUNs: the name, usable size, container type, position, aggregate, and owner.</p> <p>Entering <code>storage disk show <i>diskname</i></code> or <code>storage disk show <i>arrayLUNname</i></code> shows details about an individual disk or array LUN, for example, the status (assigned or unassigned), the owner, and the paths for the array LUN. The output is broken into three sections: information about the array LUN, information about the paths to the array LUN, and any errors associated with the array LUN.</p> <p>For example, if <code>storage errors show</code> output indicates that there are multiple IDs for the same LUN, you can use <code>storage disk show <i>arrayLUNname</i></code> to see the LUN IDs that are used on the two paths.</p>
<code>storage errors show</code>	<p>Provides details, at the array LUN level, about back-end configuration errors that prevent the V-Series system and the storage array from operating together. You must fix errors identified by <code>storage errors show</code> before you can configure Data ONTAP to work with storage arrays.</p>

Back-end configuration errors detected by the storage errors show command (clustered Data ONTAP 8.1 and later)

Starting in clustered Data ONTAP 8.1, `storage errors show` provides details, at the array LUN level, about common back-end configuration errors. You must fix errors identified by `storage errors show` before you can configure Data ONTAP to work with array LUNs.

When there is a back-end configuration error that would prevent the devices in your configuration from working together, the `storage array config show` command instructs you to run `storage errors show` to obtain the details of the error.

Back-end misconfiguration identified by storage errors show

The `storage errors show` command identifies the following back-end configuration errors:

- There are fewer than two paths to an array LUN.
- All paths to an array LUN are to the same storage array controller.
- Two array LUNs are presented with the same LUN ID.
- The LUN IDs for the same LDEV do not match on all target ports over which the LDEV will be visible.
- The array LUN exceeds the Data ONTAP maximum LUN size.
- The array LUN does not meet the Data ONTAP minimum LUN size.
- The block size of an array LUN is invalid.
- An Engenio access LUN was presented to Data ONTAP.

Output example for the storage errors show command

The `storage errors show` output is grouped by storage array (if there is more than one storage array behind the V-Series system). The name and unique identifier (UID) of an array LUN are shown, when applicable.

The following output example shows one type of error—only one path to an array LUN. This is an error because Data ONTAP requires two paths to an array LUN.

Note: Starting in Data ONTAP 8.1, four paths to an array LUN are supported for clustered configurations.

```
systemf47ab:> storage errors show
Disk: systemf47a:vgbr300s181:5.126L32
UID:
60060160:3C542400:60F91317:DE87E111:00000000:00000000:00000000:00000
000:00000000:00000000
-----
vgbr300s181:5.126L32 (600601603c54240060f91317de87e111): This array
LUN is only available on one path. Proper configuration requires
```

```

two paths.

Disk: systemf47a:vgbr300s181:5.126L33
UID:
60060160:3C542400:62F91317:DE87E111:00000000:00000000:00000000:00000
000:00000000:00000000
-----
vgbr300s181:5.126L33 (600601603c54240062f91317de87e111): This array
LUN is only available on one path. Proper configuration requires
two paths.

Disk: systemf47a:vgbr300s181:5.126L34
UID:
60060160:3C542400:64F91317:DE87E111:00000000:00000000:00000000:00000
000:00000000:00000000
-----
vgbr300s181:5.126L34 (600601603c54240064f91317de87e111): This array
LUN is only available on one path. Proper configuration requires
two paths.

```

Related concepts

[Situations not identified by the storage errors show command](#) on page 81

[Required number of paths to an array LUN](#) on page 34

Related tasks

[Checking for back-end configuration errors preventing system operation \(clustered Data ONTAP 8.1 and later\)](#) on page 85

Related references

[LUN IDs for the same LDEV do not match](#) on page 88

[Fewer than two paths to an array LUN](#) on page 96

Situations not identified by the storage errors show command

There might be situations that you consider to be an issue but are not errors from the perspective of Data ONTAP because the situation does not prevent the system from operating. The `storage errors show` command does not identify configurations that do not prevent system operation.

The `storage errors show` command does not alert you to situations such as the following:

- Configurations that do not conform to best practice recommendations; that is, they are not required
- Conditions that can occur during transitional states

For example, you might see more LUN groups than you intended in `storage array config show` output until migration of LUNs from one LUN group to another is complete.

- Conditions that do not match your intentions

For example, if you wanted multiple LUN groups and only one was configured, Data ONTAP does not identify that as an error because it cannot determine your intentions, and a single LUN group is supported.

Related concepts

Back-end configuration errors detected by the `storage errors show` command (clustered Data ONTAP 8.1 and later) on page 80

Verifying installation with storage arrays (clustered Data ONTAP 8.1 and later)

It is important to detect and resolve any back-end configuration errors before you deploy your system in a production environment.

The two stages for verifying back-end configuration are as follows:

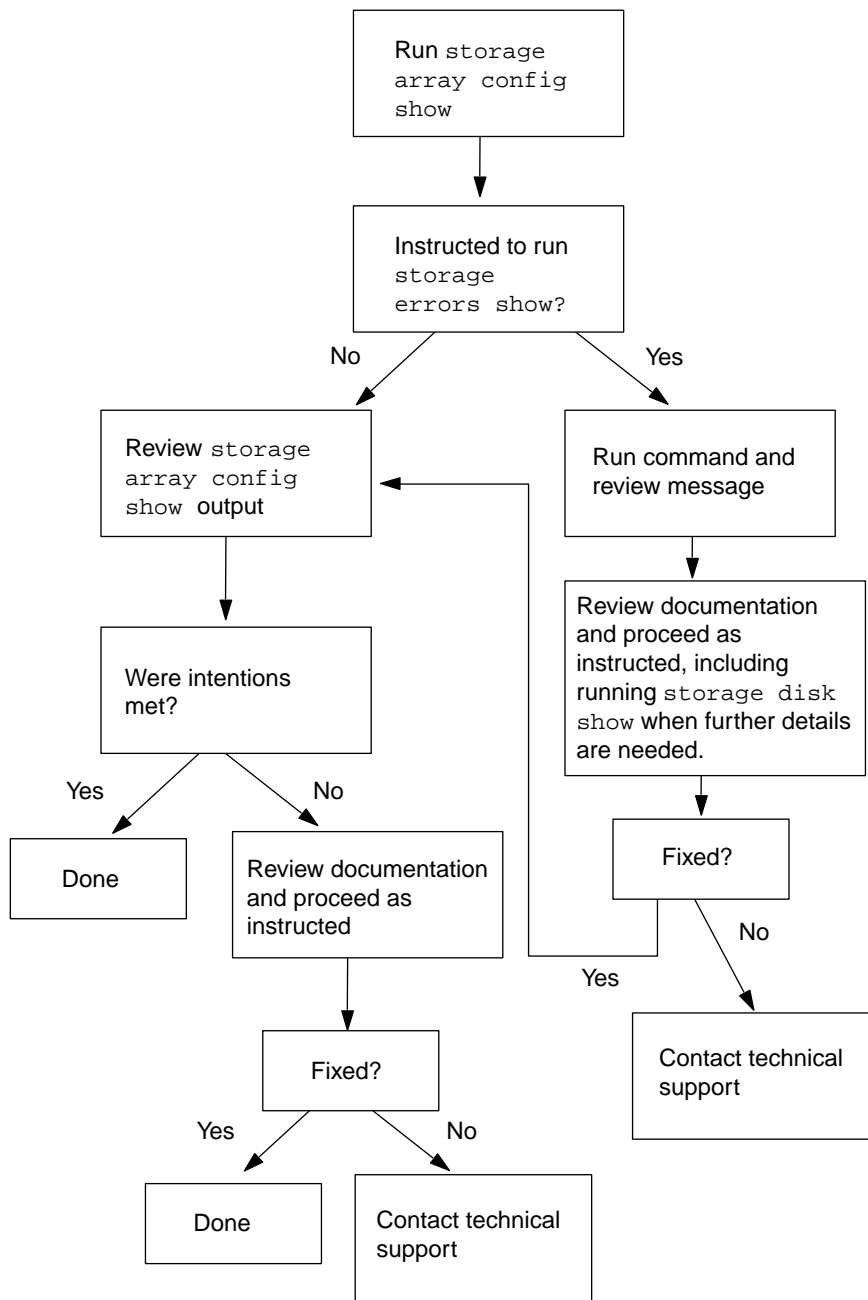
1. Checking for any back-end configuration errors that prevent Data ONTAP from operating with the storage array.

These are the errors that are flagged by `storage errors show`. You must fix these errors.

2. Checking to ensure that the configuration is as you intended.

There are a number of situations that are not errors from the system perspective but might not be what you intended. For example, the `storage array config show` output shows two LUN groups but you only intended one LUN group. This document refers to such situations as situations that do not meet your “intentions”.

The following illustration shows the workflow in which you first verify that there are no configuration errors from the system perspective and then you verify that installation is as you intended.



Checking for back-end configuration errors preventing system operation (clustered Data ONTAP 8.1 and later)

Data ONTAP requires you to fix common back-end configuration errors that would prevent the storage array and V-Series system from operating together normally.

Steps

1. Enter the following command:

```
storage array config show
```

The first step in installation verification (and troubleshooting) is to run the `storage array config show` command. If Data ONTAP detects an error in the back-end configuration, the following message is displayed at the bottom of the `storage array config show` output:

```
Warning: Configuration errors were detected. Use 'storage
errors show' for detailed information.
```

2. Take the appropriate action, as follows:

If...	Then...
<code>storage array config show</code> instructs you to run <code>storage errors show</code>	Go to Step 3.
<code>storage array config show</code> does NOT instruct you to run <code>storage errors show</code>	Review the <code>storage array config show</code> output to make sure that the output reflects the configuration you intended. See Checking whether the configuration matches your intentions (clustered Data ONTAP 8.1 and later) on page 104. (You do not need to continue with the next steps in this procedure.)

3. Enter the following command:

```
storage errors show
```

The `storage errors show` command enables you to see details of the problem at the array LUN level.

4. Review the error message and fix any errors shown.

You must fix any errors shown by `storage errors show`. See [The storage errors show messages and their resolution \(clustered Data ONTAP 8.1 and later\)](#) on page 86 to learn about the cause of each problem detected by `storage errors show` and how to fix it.

5. After fixing the problem, run `storage errors show` again to confirm that the error was fixed.

If `storage errors show` continues to show the problem, review the documentation again for more information about what to do, or contact technical support.

6. After you resolve the back-end configuration error, run `storage array config show` again so that you can review the output to ensure that the configuration meets your intentions.

Related concepts

Back-end configuration errors detected by the storage errors show command (clustered Data ONTAP 8.1 and later) on page 80

Situations not identified by the storage errors show command on page 81

Related tasks

Checking whether the configuration matches your intentions (clustered Data ONTAP 8.1 and later) on page 104

The storage errors show messages and their resolution (clustered Data ONTAP 8.1 and later)

When `storage errors show` flags an error condition, you need to determine why the error occurred and how to fix it.

The following table lists the back-end configuration errors detected by `storage errors show` and refers you to detailed information about the causes of each error and its resolution.

storage errors show message	For more information about this message see...
NAME (Serial #): All paths to this array LUN are connected to the same fault domain. This is a single point of failure.	<i>All paths to an array LUN are on the same storage array controller</i> on page 99
NAME (Serial #), port WWPN1: LUN 1 occurs more than once. LUNs cannot be reused on the same array target port.	<i>Duplicated LUN IDs on a target port</i> on page 93
NAME (Serial #): This array LUN is an access control LUN. It is not supported and should be masked off or disabled.	<i>Access control LUN was presented</i> on page 99
NAME (Serial #) This array LUN is configured with conflicting failover modes. Each path to this LUN must use the same mode.	<i>Array LUNs are configured with conflicting failover modes (clustered Data ONTAP 8.2 and later)</i> on page 102

storage errors show message	For more information about this message see...
NAME (Serial #): This Array LUN is only available on one path. Proper configuration requires two paths.	<i>Fewer than two paths to an array LUN</i> on page 96
NAME (Serial #): This array LUN is too large and is not usable. The maximum array LUN size supported is xTB.	<i>Array LUN is too small or too large</i> on page 87
NAME (Serial #): This array LUN is too small and is not usable. The minimum array LUN size supported is 1GB.	<i>Array LUN is too small or too large</i> on page 87
NAME (Serial #): This Array LUN is using multiple LUN IDs. Only one LUN ID per serial number is supported.	<i>LUN IDs for the same LDEV do not match</i> on page 88 Volume Set Addressing is inconsistent

Array LUN is too small or too large

When planning array LUN sizes, you must conform to Data ONTAP minimum and maximum array LUN size limits. These limits vary according to Data ONTAP release. The `storage errors show` output identifies array LUNs that do not meet size requirements.

Data ONTAP does not allow you to assign array LUNs with size issues to a V-Series system.

Storage errors show message

```
NAME (Serial #): This array LUN is too large and is not usable. The
maximum array LUN size supported is xTB
```

OR

```
NAME (Serial #): This array LUN is too small and is not usable. The
minimum array LUN size supported is 1GB.
```

Explanation

This message is generated when the array LUN is smaller than the minimum array LUN size allowed by Data ONTAP or exceeds the maximum size allowed. Primary reasons why an array LUN might be too small or too large are as follows:

- The storage array administrator did not convert the Data ONTAP array LUN size limits into the equivalent limits according to the vendor's definition of units of measure.

The Data ONTAP minimum and maximum array LUN size limits are calculated according to the way that Data ONTAP defines units of measure. Many vendors calculate measurements differently than Data ONTAP.

The *Hardware Universe* (formerly the *System Configuration Guide*) at support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml contains a maximum array LUN size calculated with the formula used by most storage array vendors. The maximum array LUN size shown works for all storage array vendors.

- The array LUNs are for another host whose size limits are different from Data ONTAP limits. In an open SAN, Data ONTAP can see array LUNs intended for other hosts if those array LUNs have not been masked off. Data ONTAP generates an error message about size issues for any array LUN it can see.

Troubleshooting and problem resolution

1. Review the `storage errors show` output so that you can tell the storage array administrator which array LUN has a size problem.
 - If the array LUN with the size problem is for Data ONTAP, the storage array administrator must resize the array LUN to meet Data ONTAP requirements, then present it again to Data ONTAP.
 - If the array LUN with the size problem is for another host, the storage array administrator must mask off the array LUN so that Data ONTAP cannot see it.
2. After the problem is fixed, run `storage array config show` again to confirm that it no longer indicates that there is an error.

LUN IDs for the same LDEV do not match

A logical device (LDEV) must be mapped to the same LUN ID on all storage array ports over which it will be visible to V-Series systems. The `storage errors show` output identifies the LDEVs whose LUN IDs do not match.

Data ONTAP does not allow you to assign array LUNs to a V-Series system if the LUN IDs do not match.

Storage errors show message

```
vgci9148s76:1-9.126L3 (4849544143484920443630303035323430303132): This Array LUN
is using multiple LUN IDs. Only one LUN ID per serial number is supported.
```

Explanation

One of the following errors was made during storage array configuration:

- The LDEV was presented to the same V-Series FC initiator port from multiple target ports and the LUN IDs are not consistent.

- The LUN IDs of two LDEVs were swapped.
In this case, an error is reported for each array LUN.
- Different LUN IDs for the same LDEV were used when mapping the LDEV to storage array ports that are presenting the LDEV to the V-Series system.

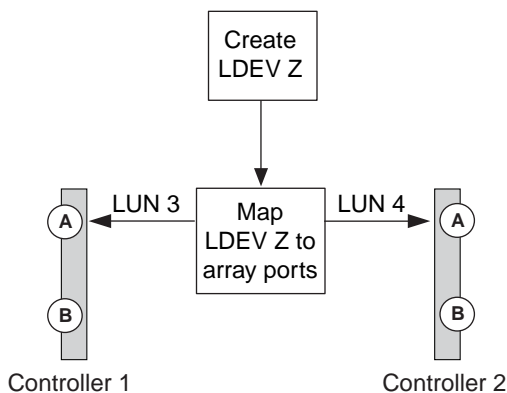
Note: This error is more likely to occur on storage arrays on which each port is configured separately, for example, on Hitachi storage arrays. On some storage arrays, for example IBM storage arrays, ports are not configured separately.

- The Volume Set Addressing setting is inconsistent on the ports on which the LUN is mapped.
On an EMC Symmetrix storage array, the problem would be that the Volume Set Addressing setting varies on the channel director ports.

Problem scenario

This scenario discusses the case of inconsistent LUN IDs as it applies to most storage arrays. See [Volume Set Addressing is inconsistent](#) on page 91 for a discussion of this same error message in the context of misconfigured Volume Set Addressing.

Assume that the storage array administrator creates a new LDEV Z. The LUN ID for LDEV Z is supposed to be LUN 3. However, the administrator presents LDEV Z as LUN 3 on storage array controller port 1A and as LUN 4 on storage array controller port 2A, as the following illustration shows:



To fix this problem, the same LUN ID must be assigned to an LDEV on all ports to which the LDEV is mapped. In this example, the LDEV should be presented as LUN ID 3 on both ports.

Troubleshooting and problem resolution

To fix the problem, the storage array administrator must remap the LUN, using the correct LUN ID. You can use Data ONTAP commands to obtain the details that you need to provide information about the problem to the storage administrator.

1. Review the `storage errors show` output to identify the array LUN whose LUN IDs do not match.

When the LUN IDs for the same LDEV do not match, the output identifies the serial number of the LDEV with the problem. For example:

```
mysystemla:> storage errors show
Disk: mysystemla:vgci9148s76:1-9.126L3
UID: 48495441:43484920:44363030:30353234:30303132:00000000:...
-----
HITACHI_DF600F_1
-----
vgci9148s76:1-9.126L3 (4849544143484920443630303035323430303132): This Array
LUN is using multiple LUN IDs. Only one LUN ID per serial number is supported.
```

Note: THE UID in this example is

48495441:43484920:44363030:30353234:30303132:00000000:00000000:00000000:00000000. It is truncated in the example because of space.

2. Obtain details about which LUN IDs are being used for the same LDEV by entering the following command:

storage disk show arrayLUNname

The storage disk show output for this example shows the following:

```
mysystemla:> storage disk show -disk mysystemla:vgci9148s76:1-9.126L3
Disk: mysystemla:vgci9148s76:1-9.126L3
Container Type: unassigned
Owner/Home: - / -
DR Home: -
Array: HITACHI_DF600F_1
Vendor: HITACHI
Model: DF600F
Serial Number: D600020C000C
UID: 48495441:43484920:44363030:30353234:30303132:00000000:...
BPS: 512
Physical Size: -
Position: present
Checksum Compatibility: block
Aggregate: -
Plex: -

Paths:
Controller Initiator ID LUN Initiator Side Target Side
TPGN... Switch Port Switch Port Acc Use Target Port
-----
mysystemla 0c 4 vgci9148s76:1-2 vgci9148s76:1-9 AO INU 50060e80004291c1
1
mysystemla 0a 3 vgbr300s89:1 vgbr300s89:9 S RDY 50060e80004291c0
2
mysystemlb 0c 4 vgci9148s76:1-4 vgci9148s76:1-9 AO INU 50060e80004291c1
1
mysystemlb 0a 3 vgbr300s89:3 vgbr300s89:10 S RDY 50060e80004291c2
2

Errors:
vgci9148s76:1-9.126L3 (4849544143484920443630303035323430303132): This Array LUN is using
multiple LUN IDs. Only one LUN ID per serial number is supported.
```

Note: THE UID in this example is

48495441:43484920:44363030:30353234:30303132:00000000:00000000:00000000:00000000. It is truncated in the example because of space.

By looking at the LUN IDs in the `Paths` section of the `storage disk show` output, you can see that LUN IDs 3 and 4 are both being used for this LDEV.

3. Determine which LUN ID is not correct for the LDEV.

LUN ID 4 is the incorrect LUN ID in this example.

4. In Data ONTAP, use the `storage path quiesce` command to quiesce the incorrect path for the array LUN.

The following example shows the options to add to the `storage path quiesce` command for the path that is being quiesced—LUN ID 4 on initiator 0c.

```
storage path quiesce -node mysystemla -initiator 0c -target-wwpn
50060e80004291c1 -lun-number 4
```

The `storage array path quiesce` command temporarily suspends I/O to a specific array LUN on a specific path. Some storage arrays require ceasing I/O for a period of time when an array LUN is to be removed or moved.

After the path is quiesced, Data ONTAP can no longer see that LUN.

5. Wait 1 minute for the storage array's activity timer to expire.

Although not all storage arrays require ceasing I/O for a period of time, it is good practice to do so.

6. On the storage array, remap the LUN to the target port by using the correct LUN ID, LUN ID 3 in this scenario.

The next time the Data ONTAP discovery process runs, it discovers the new array LUN. Discovery runs every minute.

7. After Data ONTAP discovery is complete, run `storage array config show` in Data ONTAP again to confirm that there is no longer an error.

Volume Set Addressing is inconsistent

For all storage arrays, Data ONTAP detects inconsistent LUN IDs across a path set. For storage arrays on which Volume Set Addressing is set, a mismatch of the settings on the ports to which the LUN is mapped is one problem that causes a mismatch of LUN IDs.

On EMC Symmetrix storage arrays, for example, inconsistent setting of the Volume Set Addressing parameter on the channel director ports to which a LUN is mapped triggers a LUN mismatch error.

Storage errors show message

```
vgci9148s76:1-9.126L3 (4849544143484920443630303035323430303132): This Array LUN
```

```
is using multiple LUN IDs. Only one LUN ID per serial number is supported.
```

Explanation

There are a number of configuration errors that can cause this “inconsistent LUN IDs” message. This explanation is about the display of this message when Volume Set Addressing is set inconsistently.

Data ONTAP explicitly looks for the problem of inconsistent setting of the Volume Set Addressing parameter on ports to which a LUN is mapped. If the settings are different, Data ONTAP reports this as a LUN ID mismatch in `storage errors show` output and in EMS messages.

Attention: See the *V-Series Implementation Guide for Third-Party Storage* for the setting that Data ONTAP requires for Volume Set Addressing for EMC Symmetrix storage arrays. Data ONTAP does not alert you if Volume Set Addressing is not configured as it expects; it alerts you only if the configuration is inconsistent.

Troubleshooting and problem resolution

If the `storage errors show` command displays the “inconsistent LUN IDs” message and your storage array is an EMC Symmetrix, doing one of the following helps you identify whether the problem is due to a Volume Set Addressing inconsistency:

- In Data ONTAP, run

```
storage disk show -disk
```

for the identified array LUN.
This command shows all the paths to the array LUN and the LUN ID assigned to each path.
- On the storage array, check the Volume Set Addressing settings for the channel director ports to which the identified LUN is mapped.

If you determine that the settings are inconsistent, fix the configuration problem on the storage array, making sure that you are setting the parameter on both channel director ports to the setting required by Data ONTAP (as described in the *V-Series Implementation Guide for Third-Party Storage*).

Related references

[LUN IDs for the same LDEV do not match](#) on page 88

Duplicate LUN IDs on a target port

Each array LUN on the same storage array target port must have a unique LUN ID. The `storage errors show` output identifies the LUNs that are presented with the same LUN ID on the same target port.

Storage errors show message

```
NAME (UID), port WWPNx: LUN x occurs more than once. LUNs cannot be
reused on the same array target port.
```

Explanation

The usual cause of duplicate LUN IDs on a target port is a zoning error. An administrator puts the V-Series FC initiators into different host groups to create multiple LUN groups on a storage array, but then makes a zoning mistake that allows the initiators in the different host groups to access the same target port.

When this type of zoning error is made, `storage array config show` output shows two LUN groups with the same target ports.

Problem scenario

The administrator wants to map four LDEVs (a, b, c, and d) for Data ONTAP use, two LDEVs in each of two LUN groups. For this scenario, assume that the storage array presents LDEVs to initiator ports without considering the target port by which the initiator accesses the storage arrays; that is, the host groups are *not* specific to a target port. Zoning must be used to create LUN groups by controlling which target ports are accessed by each initiator.

Note: For some storage arrays, such as HP EVA, IBM DS4xxx, and IBM DS5xxx, host groups are the same for all target ports. For other storage arrays, such as Hitachi, host groups are specific to a target port.

The administrator correctly sets up two host groups, as follows, so that there will be two LUN groups:

Host group	V-Series initiators in the host group	LDEVs and associated LUN IDs
1	0a 0c	LDEV a/LUN 1 LDEV b/LUN 2
2	0b 0d	LDEV c/LUN 1 LDEV d/LUN 2

The zoning should be set up as follows:


- The initiators in host group 1, 0a and 0c, should be zoned to the target port pair 1A and 2A.
- The initiators in host group 2, 0b and 0d, should be zoned to the target port pair 1B and 2B.

Notice in the preceding table that LDEV a and LDEV c both have the same LUN ID (L1). Likewise, LDEV b and LDEV d both have the same LUN ID (L2). If the zoning is set up correctly, this duplication of LUN IDs is not a problem because reuse of LUN IDs on different target ports is supported.

The problem in this scenario is that some initiators were placed in the wrong zone when the zoning was configured, as shown in the following table:

Zone	V-Series system		Storage array	
Switch vnbr200es25				
z1	vs1	Port 0a	Controller 1	Port 1A
z2	vs1	Port 0b	Controller 1	Port 1A (instead of 1B)
Switch vnci9124s53				
z3	vs1	Port 0c	Controller 2	Port 2A
z4	vs1	Port 0d	Controller 2	Port 2A (instead of 2B)

The following illustration shows the result of the zoning error:



Initiator 0a sees LDEV 'a' at vnbr200es25:5.126L1

Initiator 0b sees LDEV 'c' at vnbr200es25:5.126L1

Initiator 0c sees LDEV 'a' at vnci9124s53:6.126L1

Initiator 0d sees LDEV 'c' at vnci9124s53:6.126L1

The same pattern is true for LDEV b and LDEV d.

As you can see from the illustration, two LUN groups were created. However, because of the zoning error, LUN group 0 and LUN group 1 are on the same target port pair (1A and 2A), instead of one LUN group being on each target port pair.

Two LUN groups on the same target port pair is not supported. However, Data ONTAP can tolerate such a configuration. The problem in this scenario is that LDEV a and LDEV b have the same LUN ID, LDEV c and LDEV d have the same LUN ID, and all the initiators are accessing the same target port pair. Array LUN names are based on the switch port on the storage array side and the LUN ID. Therefore, these LDEV pairs have the same name, making them difficult to manage.

The following storage array config show output for this example shows two LUN groups. The problem is that the two LUN groups have the same target ports.

```
vs1::> storage array config show
```

Node	LUN Group	LUN Count	Array Name	Array Target Ports	Switch Port	Initiator
vs1	0	2	IBM_1742_1	20:1A:00:a0:b8:0f:ee:04	vnbr200es25:5	0a
				20:2A:00:a0:b8:0f:ee:04	vnci9124s53:6	0c
	1	2	IBM_1742_1	20:1A:00:a0:b8:0f:ee:04	vnbr200es25:5	0b
				20:2A:00:a0:b8:0f:ee:04	vnci9124s53:6	0d

Warning: Configuration were errors detected. Use 'storage errors show' for detailed information.

The following storage errors show output for this example identifies the LUNs with the problem:

```
vs1::> storage errors show
```

Disk: vs1:vnbr200es25:5.126L1

UID: UID-a

vnbr200es25:5.126L1 (UID-a), port WWP1: LUN 1 occurs more than once. LUNs cannot be reused on the same array target port.

Disk: vs1:vnbr200es25:5.126L2

UID: UID-b

vnbr200es25:5.126L2 (UID-b), port WWP1: LUN 2 occurs more than once. LUNs cannot be reused on the same array target port.

Disk: vs1:vnbr200es25:5.126L1

UID: UID-c

vnbr200es25:5.126L1 (UID-c), port WWP2: LUN 1 occurs more than once. LUNs

```
cannot be reused on the same array target port.
```

```
Disk: vs1:vnbr200es25:5.126L2
UID: UID-d
-----
```

```
vnbr200es25:5.126L2 (UID-d), port WWP2: LUN 2 occurs more than once. LUNs
cannot be reused on the same array target port.
```

From this `storage errors show` example you can see that the UIDs of all four LDEVs are shown, but there are only two unique LUN *names*, `vnbr200es25:5.126L1` and `vnbr200es25:5.126L2`, instead of four.

Troubleshooting and problem resolution

The storage array administrator must fix the zoning so that the initiators in different host groups do not have access to the same target port.

1. In `storage array config` output, look for initiators that are talking to the same target port.
2. Enter the following command to view the details of the error: `storage errors show`
3. Determine the LDEV for which the LUN IDs are duplicated.
4. For each target port on controller 1 that has multiple initiators from the same V-Series system mapped to it, change the zoning so that the two initiators are *not* talking to the same target port. You are performing this step because initiators in different host groups should not be in the same zone. You need to perform this step on one initiator at a time so that there is always a path to the array LUN.
5. Repeat the procedure on controller 2.
6. Enter `storage errors show` in Data ONTAP and confirm that the error has been fixed.

Related concepts

[Dependency between zone and host group definitions](#) on page 120

[Example of cascading zoning and host group configuration errors](#) on page 122

Related references

[Relationship between zoning and host group configuration](#) on page 120

Fewer than two paths to an array LUN

Common reasons for fewer than two paths to an array LUN are a mapping error, a zoning error, or a cable dropping out. The `storage errors show` output identifies array LUNs with only a single path.

Data ONTAP requires redundant paths to an array LUN so that access to the LUN is maintained if a device fails. Two paths must exist to each array LUN.

Storage errors show message

NAME (UID): This Array LUN is only available on one path. Proper configuration requires two paths.

Explanation

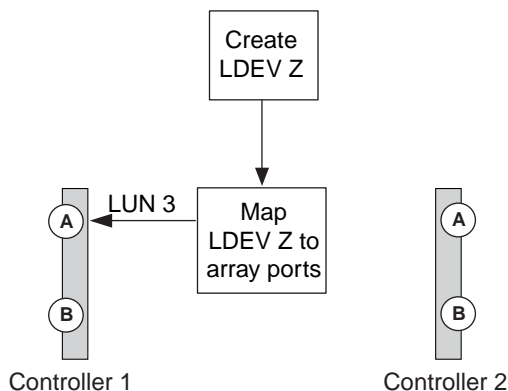
Reasons you would see fewer than two paths to an array LUN include the following:

- The LDEV was mapped on only one storage array port.
- The second path to the array LUN was not zoned.
- There is a problem with the host group mapping.
- There is a problem with switch connections.
- The cable dropped out.
- SFPs failed on the adapter.

Note: If a path drops out on a running system, an EMS message is generated.

Problem scenario

For this example of a mapping error, assume that the storage administrator created a new LDEV Z. The administrator mapped LDEV Z as LUN ID 3 to target port 1A. However, the administrator did not map the LDEV to target port 2A, as the following illustration shows. The result is only one path to the array LUN.



When this error is made, the `storage array config show` output shows only one path to the LUN , as the following example shows.

```
mysystem1::> storage array config show
```

[illegible]

```
mssystemla 0 1 IBM_1742_1 20:1A:00:a0:b8:0f:ee:04 vnbr20es25:5 0a
Warning: Configuration errors were detected. Use 'storage errors show' for detailed
information.
```

The `storage errors show` command provides the details you need to determine which LUN has fewer than two paths.

```
mssystemla:> storage errors show
Disk: mssystemla:vnbr20es25:5.126L1
UID: 600508B4:000B6314:00008000:00200000:00000000:00000000:00000000:...
-----
vnbr20es25:5.126L1 (600508b4000b63140000800000200000): This array LUN is only
available on one path. Proper configuration requires two paths.
```

Note: The UID for this example is 600508B4:000B6314:00008000:00200000:00000000:00000000:00000000:00000000:00000000. It is truncated because of space.

Troubleshooting and problem resolution

Looking at both the `storage array config show` output and the `storage errors show` output is helpful when troubleshooting fewer than two paths to an array LUN.

- 1. Review the `storage errors show` output to obtain the serial number of the array LUN that is available on only one path.
- 2. Review the `storage array config show` output to try to isolate the cause of the problem.

If the storage array config show output shows...	The cause is most likely...
Other array LUNs	A mapping error
No other array LUNs	A cabling error, zoning error, or hardware issue

- 3. If the cause is a mapping error, have the storage array administrator map the identified array LUN to two redundant storage array ports.
- 4. If the cause seems to be a problem other than mapping, check zoning, host group mapping, cabling, and connectivity.
- 5. After you fix the problem, run `storage array config show` again to confirm that the error is fixed.

Related concepts

Dependency between zone and host group definitions on page 120
Example of cascading zoning and host group configuration errors on page 122

Related references

[Relationship between zoning and host group configuration](#) on page 120

An access control LUN was presented

Data ONTAP does not support access control array LUNs. The `storage errors show` output alerts you if an Engenio access control LUN is being presented.

Storage errors show message

```
NAME (UID): This array LUN is an access control LUN. It is not supported
and should be masked off or disabled.
```

Explanation

Data ONTAP supports only storage array LUNs. When an Engenio access control LUN is presented to Data ONTAP, the `storage array config show` output looks normal; it shows the access control LUN in a LUN group, as the following example shows. The warning message at the bottom of the screen is your clue that there is a problem. You need to run `storage errors show` to find out that the problem is that an access control LUN was presented and which LUN it is.

```
mysystem1:> storage array config show
```

Node	LUN Group	LUN Count	Array Name	Array Target Port	Initiator
mysystem1	0	1	IBM_1742_1	20:1A:00:a0:b8:0f:ee:04 20:2A:00:a0:b8:0f:ee:04	0a 0c

```
Warning: Configuration errors were detected. Use 'storage errors show' for
detailed information.
```

Troubleshooting and problem resolution

1. On the storage array, mask off the access control LUN.
2. In Data ONTAP, run `storage errors show` again to confirm that the access control LUN is no longer being presented to Data ONTAP.

All paths to an array LUN are on the same storage array controller

Data ONTAP does not support configuring all paths to the same storage array controller because doing so sets up a configuration with a single point of failure (SPOF). The `storage errors show` command identifies any array LUN whose paths are set up to go to the same storage array controller.

Data ONTAP does not allow you to assign array LUNs to a V-Series system until after you fix this error.

Storage errors show message


```
NAME (UID): All paths to this array LUN are connected to the same fault domain. This is a single point of failure
```

Explanation

This error occurs because the paths to an array LUN are set up to go to the same storage array controller (on an active-passive array) or FRU (on an active-active storage array).

Note: Using four paths to an array LUN, a storage array with FRUs with multiple directors (such as an EMC Symmetrix or HDS USP), or a storage array with dual controllers (such as an IBM, EMC CX, or HP EVA) are good methods for achieving redundancy. However, if you set up the paths to go through a single storage array controller or FRU, you are setting up your configuration with a SPOF, even with such features. On an active-active storage array, the entire FRU is considered to be one fault domain. An EMC Symmetrix storage array, for example, has multiple channel directors on the same FEBC board. A FEBC board is considered to be one fault domain because if all paths go through the same FEBC board, you lose all paths if you must replace the board.

The following illustration shows correct and incorrect storage array port selection for setting up redundant paths to an array LUN so that you do not have a single fault domain. The path setup in the example on the left is correct because the paths to the array LUN are redundant—each connection is to a port on a different controller on the storage array. In the example on the right, both paths to the array LUN go to the same controller, which sets up a SPOF.



The storage errors show command shows the array LUN that is in the same fault domain. You can also see this problem in the storage disk show output if you look at the TPGN column (Target port group number). A different TPGN should be shown for each initiator in an initiator port pair. If the TPGN is the same for both initiators in the pair, both initiators are in the same fault domain.

The following storage disk show example shows TPGN 1 for LUN 30, which is accessed over initiators 0a and 0c. If the paths are redundant, each initiator shows a different TPGN.

```

mysystem1:> storage disk show mysystem1:vgbr300s70:9.126L30
      Disk: mysystem1:vgbr300s70:9.126L30
      Container Type: unassigned
      Owner/Home: - / -
      DR Home: -
      Array: HP_HSV300_1
      Vendor: HP
      Model: HSV300
      Serial Number: 600508B4000B63140000800001660000
      UID: 600508B4:000B6314:00008000:01660000:00000000:...
      BPS: 512
      Physical Size: -
      Position: present
      Checksum Compatibility: block
      Aggregate: -
      Plex: -

Paths:

```

Controller	Initiator	LUN ID	Initiator Side Switch Port	Target Side Switch Port	Acc Use	Target Port	TPGN...
mysystem1	0c	30	vgci9148s75:1-12	vgci9148s75:1-9	AO RDY	50014380025d1509	1
mysystem1	0a	30	vgbr300s70:12	vgbr300s70:9	AO INU	50014380025d1508	1

```

Errors:
vgbr300s70:9.126L30 (600508b4000b63140000800001660000): All paths to this array LUN are
connected to the same fault domain. This is a single point of failure.

```

Note: The full UID in this example is 600508B4:000B6314:00008000:01660000:00000000:00000000:00000000:00000000:00000000:00000000. It is truncated in the example because of space.

Troubleshooting and problem resolution

The paths to the array LUN must be reconfigured so that they go to redundant storage array controllers or FRUs.

1. Add a cable to the redundant target port on the other controller.
You should maintain redundancy while fixing this problem by adding a cable to the alternate controller *before* you remove a cable from the controller with the SPOF. Redundancy is maintained in this case because you are increasing the number of paths to three paths temporarily instead of decreasing the number of paths to one while you are fixing the problem.
2. Remove one cable from the controller that was set up with the SPOF.
You now have two redundant paths to the array LUN.
3. From the Data ONTAP command line, enter the following command again and confirm that the error has been fixed:

```
storage errors show
```

Array LUNs are configured with conflicting failover modes (clustered Data ONTAP 8.2 and later)

Data ONTAP requires that array LUNs that are visible to a particular V-Series system be configured with the same failover mode. On some storage arrays it is possible to configure inconsistent failover modes on different paths to an array LUN.

Storage errors show message

```
NAME (serial #) This array LUN is configured with conflicting failover
modes. Each path to this LUN must use the same mode.
```

Explanation

On some storage arrays, for example, EMS CLARiiON and Fujitsu storage arrays, the failover mode can be set by FC initiator port. On such storage arrays it is possible to set inconsistent failover modes for array LUNs visible to the FC initiators on the same V-Series system. Data ONTAP does not support inconsistent failover modes for paths to an array LUN from a particular V-Series system.

If your storage array allows setting the failover mode for an array LUN by FC initiator, part of your installation validation process should include checking to ensure that there are no problems with the failover mode settings for the array LUNs visible to the V-Series system. The `storage errors` show command alerts you about inconsistent failover mode settings for array LUNs and generates an EMS message.

Although your system can operate with inconsistent array LUN failover mode settings, you need to fix this problem as soon as possible. Otherwise, if a path fails, the V-Series system might not operate properly, failover might not occur, or the system might panic.

Note: Data ONTAP supports different failover mode settings between V-Series nodes. For example, node A can use Active/Passive mode for the paths to an array LUN and node B can use ALUA for the paths to the same array LUN.

Troubleshooting and problem resolution

The failover mode of the first path that Data ONTAP discovers during LUN initialization is the failover mode that Data ONTAP expects for all paths to the LUN from a particular V-Series system. If the failover mode of subsequent discovered paths does not match the failover mode of the first path, Data ONTAP issues an error message.

In the following `storage errors show` example, Data ONTAP tells you that the failover mode for LUN `vgbr300s89:20.126L4`, which is visible over `mysystem1` FC initiator `0a`, is *Proprietary*, and that the failover mode is different from the failover mode that Data ONTAP discovered on the first path for that array LUN.

```
mysystem1:> storage errors show
vgbr300s89:20.126L4 (60060160e1b0220008071baf6046e211): hba 0a port
500601603ce014de mode Proprietary: This array LUN is configured with
conflicting failover modes. Each path to this LUN must use the same mode.

Disk: mysystem1a:vgbr300s89:20.126L0
UID: 60060160:E1B02200:1C65EB20:BFF7E111:00000000:00000000:00000000:...
```

You need to fix the failover mismatch problem on the storage array. However, the entire procedure for fixing the mismatch depends on whether the failover mode that Data ONTAP detected on the first path is the failover mode that you want to be used for all paths on that V-Series system to the array LUN.

1. Enter `storage errors show` if you have not done so already as part of your installation verification process.

Note: The `storage array config` command tells you to run `storage error show` if there is a problem that you need to fix.

2. Review the `storage errors show` output to determine the failover mode setting for the array LUN that is not consistent with the failover mode that Data ONTAP is expecting.

If the failover mode that the system detected on the first path is...	Example	You need to...
What you want	You want a failover mode of ALUA and ALUA is the failover mode that Data ONTAP detected for the first path.	Change, on the storage array, the failover mode for the initiator that Data ONTAP identified in the error message. Go to Step 3.
Not what you want	You want a failover mode of Active/Passive but ALUA is the failover mode that Data ONTAP detected for the first path.	Remove the array LUN from the view of the V-Series system. Go to Step 4.

3. If you need to change the failover mode for the initiator, proceed as follows to fix the mismatch. You would use this step if the failover mode that the system detected on the first path *is* what you want.
 - a. In Data ONTAP, take the second path offline.
 - b. On the storage array, change the failover mode for the initiator that Data ONTAP identified in the error message.
 - c. In Data ONTAP, bring the second path back online.

4. If you need to remove the array LUN from the view of the V-Series system to fix the mismatch, select one of the following methods, depending on whether the array LUNs are spares or in an aggregate.

You would use one of these methods if the failover mode that the system detected on the first path *is not* what you want.

Method 1: Affected array LUNs are spares (not part of an aggregate)	Method 2: Affected LUNs are in an aggregate
<p>With this method, the V-Series system does not have to be rebooted.</p> <ol style="list-style-type: none">1. In Data ONTAP, run the following command for each affected spare LUN: <code>disk remove_ownership LUNfullname</code>2. On the storage array, mask each affected array LUN on all paths to the V-Series system.3. Wait for about a minute, then confirm that the array LUNs are no longer visible to the V-Series system.4. Set the same failover mode for each of the FC initiators on the V-Series system.5. Present all the affected array LUNs to the V-Series system again. Data ONTAP should discover the LUNs when it next runs LUN discovery6. Run <code>storage errors show</code> to confirm that there is no longer a failover mode error.	<p>With this method, the V-Series system must be rebooted.</p> <ol style="list-style-type: none">1. Reboot the V-Series system and hold it at the LOADER prompt.2. On the storage array, review the failover mode settings on the FC initiators for this system and update them as necessary to the failover mode you want.3. Reboot the V-Series system.4. Run <code>storage errors show</code> to confirm that there is no longer a failover mode error.

Checking whether the configuration matches your intentions (clustered Data ONTAP 8.1 and later)

After resolving any back-end configuration errors that were detected by `storage errors show`, you next need to check whether the back-end configuration matches your intentions. For example,

you want to check that the number of array LUN groups you intended were configured and that each LUN group has the correct number of LUNs.

Steps

1. Enter the following command:

```
storage array config show
```

The `storage array config show` output groups information about LUN groups, LUN counts, and paths by storage array, as the following example for an HA pair shows:

Example

```
mysystem1::> storage array config show
Node          LUN  LUN
              Group Count  Array Name      Array Target Port  Initiator
-----
mysystem1a    0    10    IBM_2107900_1  5005076303030124  1a
               5005076303088124  1b
               5005076303130124  1c
               5005076303188124  1d
mysystem1b    0    10    IBM_2107900_1  5005076303030124  1a
               5005076303088124  1b
               5005076303130124  1c
               5005076303188124  1d
8 entries were displayed.
```

2. Check the `storage array config show` output for each item in the following table.

If you do not see what you intended in the output, read the referenced information and fix the problem. Then continue to the next item.

Check this...	If your intentions are not met, see this information...
Are all the storage arrays listed that you expected?	<i>Reasons you might not see all storage arrays you expected</i> on page 111
Are there more array LUN groups than you expected?	<i>Reasons for more array LUN groups than expected</i> on page 107 Note: The <i>Interoperability Matrix</i> contains information about the storage arrays for which multiple LUN groups are not supported.
Are there fewer array LUN groups than you expected?	<i>Reasons for fewer array LUN groups than expected</i> on page 106

Check this...	If your intentions are not met, see this information...
Does each array LUN group show the number of LUNs you expected?	Reasons for the incorrect number of LUNs in array LUN groups on page 110
Are there fewer paths than expected?	Fewer than two paths to an array LUN on page 96
Are there more paths than expected?	Reasons for more paths to an array LUN than expected on page 109

Related tasks

[Checking for back-end configuration errors preventing system operation \(clustered Data ONTAP 8.1 and later\)](#) on page 85

Reasons for fewer array LUN groups than expected

When validating the back-end configuration, you need to check the `storage array config show` output to determine whether the number of array LUN groups in the output is what you intended.

Explanation

The most likely cause for fewer LUN groups than expected is that the LDEV-to-LUN mapping is the same for both V-Series FC initiator port pairs. If the LDEV-to-LUN mapping is the same for both V-Series FC initiator port pairs, the `storage array config show` output shows one fewer LUN group.

The following `storage array config show` output shows only one array LUN group due to both V-Series FC initiator port pairs being mapped to the same target port pair:

```
mysystem1:> storage array config show
```

Node	LUN Group	LUN Count	Array Name	Array Target Ports	Switch Port	Initiator
-----	-----	-----	-----	-----	-----	-----
mysystem1	0	2	IBM_1742_1	20:1A:00:a0:b8:0f:ee:04	vnbr200es25:5	0a
				20:2A:00:a0:b8:0f:ee:04	vnci9124s53:6	0c
				20:1A:00:a0:b8:0f:ee:04	vnbr200es25:5	0b
				20:2A:00:a0:b8:0f:ee:04	vnci9124s53:6	0d

By looking at the Array Target Ports column, you can see the same target port more than once within the LUN group, and each occurrence has a different initiator.

- Initiators 0a and 0b both have access to storage array port 1A.
- Initiators 0c and 0d both have access to storage array port 2A.

Resolving the issue

Data ONTAP does not flag this as an error because a V-Series system can operate when the LDEV-to-LUN mapping is the same for both V-Series FC initiator port pairs. However, multiple V-Series FC initiators to the same target port is not supported. You should fix the mapping so that your V-Series system follows a supported configuration, and so that you have the number of LUN groups that you intended.

If you want to fix this issue, do the following:

1. On the storage array, fix the mapping so that V-Series FC initiator port pair mapping is no longer the same for both V-Series FC initiator port pairs.
2. On the V-Series system, run `storage array config show` again and confirm that the number of LUN groups that you expected are shown and that V-Series FC initiator port pairs are not accessing the same target ports.

Reasons for more array LUN groups than expected

When validating your configuration, you should check the `storage array config show` output to ensure that the number of LUN groups is what you intended. The most common reason for too many array LUN groups is that the LDEV is spanning target port pairs.

Explanation

Common reasons why you would see more LUN groups than expected are as follows:

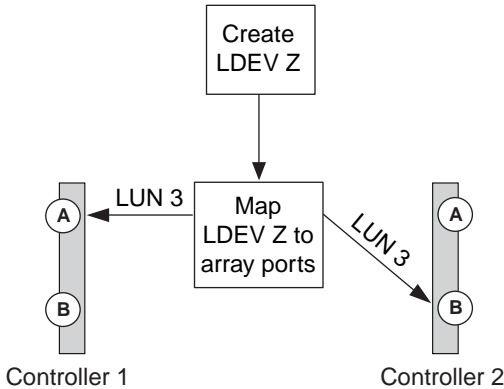
- An LDEV is not mapped to a redundant target port pair.
The most common cause of an extra LUN group is that the LDEV is spanning path pairs. An LDEV is said to be *spanning path pairs* if matching LUN IDs for the LDEV are not presented to a redundant target port pair on the storage array.
Note: On many storage arrays, it is not possible to present an LDEV to target ports that are not redundant if you follow the conventional process for creating and mapping LDEVs.
- Data ONTAP is trying to balance the load over the target ports during migration of array LUNs from one LUN group to another.
In this case, the LDEV spanning path pairs is a transitional state.
- Too many connections were set up.

Data ONTAP can run with a single array LUN group or multiple array LUN groups. Therefore, Data ONTAP does not consider any number of LUN groups to be a misconfiguration. However, multiple LUN groups are not supported for all storage arrays in all Data ONTAP releases. During configuration planning, you should check the information in the Interoperability Matrix at support.netapp.com about the storage arrays for which multiple array LUN groups are supported.

Problem scenario

Assume that there are two existing LUN groups, LUN group 0 and LUN group 1. LUN group 0 has two array LUNs and LUN group 1 has three array LUNs. The administrator created a new LDEV Z, intending to add LUN 3 to LUN group 0. The administrator mapped LDEV Z as LUN 3 on two

storage array ports. However, the ports to which LUN 3 are mapped, storage array ports 1A and 2B, are not a redundant target port pair, as the following illustration shows.



The result of this error is that a third, unintended LUN group was created (Group 2), as the following example shows:

```
mysystem1:> storage array config show
```

Node	LUN Group	LUN Count	Array Name	Array Target Ports	Switch Port	Initiator
mysystem1	0	3	IBM_1742_1	20:1A:00:a0:b8:0f:ee:04	vnbr20es25:5	0a
				20:2A:00:a0:b8:0f:ee:04	vncl9124s53:6	0c
	1	3	IBM_1742_1	20:1B:00:a0:b8:0f:ee:04	vnbr20es25:7	0b
				20:2B:00:a0:b8:0f:ee:04	vncl9124s53:8	0d
	2	1	IBM_1742_1	20:1A:00:a0:b8:0f:ee:04	vnbr20es25:5	0a
				20:2B:00:a0:b8:0f:ee:04	vncl9124s53:8	0d

If you look at the Array Target Ports column for Groups 0 and 1, you can see that the paths to the array target ports are redundant. Group 0 goes to target ports 1A and 2A, a target port pair. Group 1 goes to 1B and 2B, a different target port pair.

However, if you look at the Array Target Ports column for Group 2, you can see that the paths are not redundant. One path goes to target port 1A and the other goes to target port 2B. This is not a redundant target port pair; the array LUN is spanning LUN groups. The array LUN should have been mapped to either 1A and 2A or 1B and 2B.

Because the V-Series system can run with an LDEV spanning path pairs, there is no message at the bottom of `storage array config show` output telling you to run `storage errors show`. However, this is not a best practice configuration.

Resolving the issue

1. Wait 1 minute, then run `storage array config show` again to see whether the extra LUN group is still shown in the `storage array config show` output.

- If the extra LUN group is no longer in the output, you can conclude that the issue was transitional.
- If the extra LUN group still appears in the output, the storage array administrator must remap the LDEV, as documented in the next steps.

You need to maintain redundancy while fixing this problem. This procedure instructs you to map the LDEV to the correct target port *before* removing the mapping to the incorrect target port. Redundancy is maintained in this case because you are increasing the number of paths to three, temporarily, instead of decreasing the number of paths to one while you are fixing the problem.

2. Decide which of the inconsistent target ports should be remapped.
3. On the storage array, map the LDEV to the new (correct) target port.
4. In Data ONTAP, run `storage array config show` to confirm that three paths are shown.
5. Remove the incorrect mapping.
6. Wait 1 minute while Data ONTAP discovers the LUN.
7. In Data ONTAP, run `storage array config show` again to make sure that the extra LUN group is gone.

Reasons for more paths to an array LUN than expected

The primary reasons for an unexpected additional path to an array LUN are zoning issues and too many cables. Three paths to an array LUN is not best practice, but you are not required to fix this issue.

Example of storage array config show output

The FC initiator port pair on the V-Series system is 0a and 0c. The following example shows an extra path, 0b, in the `storage array config show` output:

```
mysystem1:> storage array config show
      LUN      LUN
Node      Group Count   Array Name      Array Target Port      Initiator
-----
mysystem1a    1      3    HITACHI_DF600F_1  50060e80004291c0      0a
                                     50060e80004291c1      0b
                                              0c

3 entries were displayed.
```

Note: If an initiator in the `storage array config show` output is not preceded by an array target port, the initiator is connecting to the same array target port as the initiator above it.

Explanation

Three paths within a LUN group indicates that there is an extra path. For clustered Data ONTAP configurations, the best practice is two or four paths.

Reasons there could be more paths than expected include the following:

- More cables were connected than are needed.
- A zoning configuration error caused an extra path.

Resolving the issue

Having an extra path is not a best practice, but it is not incorrect from the system perspective. You do not have to fix this issue. If you want to fix the issue to align with your intended configuration, complete the following steps:

1. Check your cabling and the zoning configuration for the cause, and then fix the issue that is causing the extra path.
2. After the issue is fixed, run `storage array config show` again to confirm that the extra path is gone.

Reasons for the incorrect number of LUNs in array LUN groups

When validating your configuration, you should check the `storage array config show` output to ensure that the number of LUNs in each LUN group is what you intended. The most likely cause for the incorrect number of LUNs in a LUN group is that the array LUN was not mapped to a V-Series system.

Example of storage array config show output

The number of array LUNs in each LUN group appears in the `storage array config show` output, as the following example shows:

```
mysystem1:> storage array config show
```

Node	LUN Group	LUN Count	Array Name	Array Target Port	Initiator
mysystem1	0	50	IBM_1742_1	201A00a0b80fee04 202A00a0b80fee04	0a 0c

Explanation

The most likely reasons that an array LUN you expected to be in a LUN group is missing are as follows:

- The array LUN was not mapped to the V-Series system.
- A mapping error was made that resulted in the array LUN being in the wrong LUN group. For example, the host group configuration might be incorrect.
- The storage array is still in the process of initializing and making the array LUNs available (transitional state).
- The Data ONTAP LUN scanner has not yet discovered the LUNs (transitional state).

Resolving the issue

1. If the array LUN has not been mapped to the V-Series system, the storage array administrator needs to map it.
The process for mapping array LUNs to hosts varies among storage arrays.
2. If the array LUN has been mapped to the V-Series system, check zoning and the host group configuration.
3. After the issue is fixed, run `storage array config show` again to confirm that the issue was really fixed.

Related concepts

[*Dependency between zone and host group definitions*](#) on page 120

[*Example of cascading zoning and host group configuration errors*](#) on page 122

Related references

[*Relationship between zoning and host group configuration*](#) on page 120

Reasons storage arrays are missing from command output

A storage array that is not connected to the V-Series system does not appear in the `storage array config show` output. Problems with cabling, zoning, and host group configuration can prevent a connection between the two devices.

Resolving the issue

1. Check cabling, host group configuration, and zoning as follows:
 - Check that cables are connected.
 - Check that the WWPNs for the V-Series initiators are in the host group.
 - Check that both the storage array and the V-Series initiator are in the same zone.
2. After the problem is fixed, run `storage array config show` in Data ONTAP to confirm that the issue was fixed.

Related concepts

[*Dependency between zone and host group definitions*](#) on page 120

[*Example of cascading zoning and host group configuration errors*](#) on page 122

Related references

[*Relationship between zoning and host group configuration*](#) on page 120

Verifying an installation with storage arrays (7-Mode in 8.x)

It is important to detect and resolve any configuration errors before you bring the configuration online in a production environment.

Checking the number of paths (7-Mode in 8.x)

For systems running 7-Mode in Data ONTAP 8.x, you use the `storage array show-config` command to list the array LUNs that the storage system can access, and to check that each array LUN is visible through both paths.

Steps

1. Enter the following command to show the array LUNs in your configuration:

```
storage array show-config
```

Example

The following output shows correctly configured array LUNs. Each LUN group contains two paths.

LUN Group	Array Name	Array Target Ports	Switch Port	Initiator
Group 0 (4 LUNS)	HP_HSV210_1	50:00:1f:e1:50:0a:86:6d	vnmc4300s35:11	0b
		50:00:1f:e1:50:0a:86:69	vnbr4100s31:15	0c
Group 1 (4 LUNS)	HP_HSV210_1	50:00:1f:e1:50:0a:86:68	vnbr4100s31:1	0a
		50:00:1f:e1:50:0a:86:6c	vnmc4300s35:6	0d
Group 2 (50 LUNS)	HP_HSV200_1	50:00:1f:e1:50:0d:14:6d	vnbr4100s31:5	0a
		50:00:1f:e1:50:0d:14:68	vnmc4300s35:3	0d

Example

The `storage array show-config` output shows array LUNs that are not configured with two paths as LUNs with a single path. In the following example output, the incorrectly configured LUNs are the 20 LUNs not belonging to a group and showing only a single path.

LUN Group	Array Name	Array Target Ports	Switch Port	Initiator
Group 2 (50 LUNS)	HP_HSV200_1	50:00:1f:e1:50:0d:14:68	vnmc4300s35:3	0d
		50:00:1f:e1:50:0d:14:6d	vnbr4100s31:5	0a
		(20 LUNS)	vnmc4300s35:2	0e

2. If you see array LUNs in the output from `storage array show-config` with only one path, enter the following command to show information about each array LUN on the storage array:

```
storage array show-luns
```

Example

The following shows output from `storage array show-luns`:

```

Name                                WWPNs
vnm4300s35:3.127L1 50:00:1f:e1:50:0d:14:68, 50:00:1f:e1:50:0d:14:6d
vnm4300s35:3.127L2 50:00:1f:e1:50:0d:14:68, 50:00:1f:e1:50:0d:14:6d
vnm4300s35:3.127L3 50:00:1f:e1:50:0d:14:68, 50:00:1f:e1:50:0d:14:6d
.
.
.
vnbr4100s31:5.126L49 50:00:1f:e1:50:0d:14:6d, 50:00:1f:e1:50:0d:14:68
vnm4300s35:3.127L50 50:00:1f:e1:50:0d:14:68, 50:00:1f:e1:50:0d:14:6d
vnm4300s35:3.127L51 50:00:1f:e1:50:0d:14:69,
vnm4300s35:3.127L52 50:00:1f:e1:50:0d:14:69,
.
.
.
vnbr4100s31:5.126L53 50:00:1f:e1:50:0d:14:69,
vnbr4100s31:5.126L70 50:00:1f:e1:50:0d:14:69,

```

LUNs 1 through 50 make up Group 2, the array LUNs configured with two ports as shown with the `storage array show-config` command. LUNs 51 through 70 are the 20 LUNs that contain a single path connected only to Port 50:00:1f:e1:50:0d:14:69 of the storage array.

A single path indicates a configuration problem. The second path can have a problem at the V-Series FC initiator port, the switch, or the storage array port.

Example output showing correct and incorrect pathing (7-Mode in 8.x)

By learning to interpret Data ONTAP command output, you can determine whether there are sufficient paths to an array LUN.

For systems running 7-Mode in Data ONTAP 8.x, you use the commands in the following table to check pathing:

Run this command...	To...
<code>storage array show-config</code>	<ul style="list-style-type: none"> List the array LUNs that the V-Series system can access Check that each array LUN is visible through both paths
<code>storage array show-luns</code>	<ul style="list-style-type: none"> Obtain details about an individual LUN

A single path indicates a configuration problem. The second path can have a problem at the V-Series port, the switch, or the storage array port.

Output of storage array show-config showing two paths

The following example shows output from a V-Series system connected to two storage arrays:

```
> storage array show-config
LUN Group      Array Name      Array Target Ports      Switch Port      Initiator
Group 0 (4 LUNS) HP_HSV210_1  50:00:1f:e1:50:0a:86:6d  vnm4300s35:11    0b
                    50:00:1f:e1:50:0a:86:69  vnbr4100s31:15    0c
Group 1 (4 LUNS) HP_HSV210_1  50:00:1f:e1:50:0a:86:68  vnbr4100s31:1    0a
                    50:00:1f:e1:50:0a:86:6c  vnm4300s35:6      0d
Group 2 (50 LUNS) HP_HSV200_1  50:00:1f:e1:50:0d:14:6d  vnbr4100s31:5    0a
                    50:00:1f:e1:50:0d:14:68  vnm4300s35:3      0d
```

In this valid example, each LUN group comprises LUNs that share the same two paths. Groups 0 and 1 contain a total of 8 LUNs on the HP_HSV210_1 array and Group 2 contains 50 LUNs on the HP_HSV200_1 array.

Output of storage array show-config if there are not two paths

Array LUNs that are not configured with two paths are shown as one or more LUNs with a single path, similar to the following example. The incorrectly configured LUNs are the 20 LUNs not belonging to a group and showing only a single path.

```
LUN Group      Array Name      Array Target Ports      Switch Port      Initiator
Group 2 (50 LUNS) HP_HSV200_1  50:00:1f:e1:50:0d:14:68  vnm4300s35:3    0d
                    50:00:1f:e1:50:0d:14:6d  vnbr4100s31:5    0a
(20 LUNS)      HP_HSV200_1  50:00:1f:e1:50:0d:14:69  vnm4300s35:2    0e
```

Output of storage array show-luns

If you see array LUNs in the output from the `storage array show-config` command with only one path, you need to use the `storage array show-luns` command to show information about each array LUN on the storage array. This information enables you to determine which array LUNs are members of groups and which are incorrectly configured. The output from the `storage array show-luns HP_HSV200_1` command produces output similar to the following (the output is abbreviated):

```
Name                                WWPNS
vnm4300s35:3.127L1  50:00:1f:e1:50:0d:14:68, 50:00:1f:e1:50:0d:14:6d
vnm4300s35:3.127L2  50:00:1f:e1:50:0d:14:68, 50:00:1f:e1:50:0d:14:6d
vnm4300s35:3.127L3  50:00:1f:e1:50:0d:14:68, 50:00:1f:e1:50:0d:14:6d
.
.
.
vnbr4100s31:5.126L49 50:00:1f:e1:50:0d:14:6d, 50:00:1f:e1:50:0d:14:68
vnm4300s35:3.127L50 50:00:1f:e1:50:0d:14:68, 50:00:1f:e1:50:0d:14:6d
vnm4300s35:3.127L51 50:00:1f:e1:50:0d:14:69,
vnm4300s35:3.127L52 50:00:1f:e1:50:0d:14:69,
```

```
.
.
.
vnbr4100s31:5.126L53 50:00:1f:e1:50:0d:14:69,
vnbr4100s31:5.126L70 50:00:1f:e1:50:0d:14:69,
```

LUNs 1 through 50 make up Group 2, the array LUNs configured with two ports as shown with the `storage array show-config` command. LUNs 51 through 70 are the 20 LUNs that contain a single path connected only to Port 50:00:1f:e1:50:0d:14:69 of the storage array.

Related concepts

Invalid path setup examples on page 117

Troubleshooting

You should validate your configuration during initial installation so you can resolve issues before your configuration is put into a production environment.

Getting started with troubleshooting

Before escalating a problem, there are a number of things you can check to figure out the problem yourself. You should check for the obvious and easiest problems first.

About this task

This procedure suggests an order for approaching troubleshooting.

Note: As you troubleshoot, you should save any information you gather. If you need to escalate the problem, you need to provide that data.

Steps

1. Determine whether the problem is unique to the V-Series system.

You need to determine whether the problem is a *front-end problem* (a Data ONTAP problem that affects all platforms running Data ONTAP) or a *back-end problem* (a problem with the switch or storage array configuration). For example, if you are trying to use a Data ONTAP feature and it is not working as you expected, the problem is likely to be a front-end problem.

2. If the issue is a front-end problem, proceed with troubleshooting the Data ONTAP feature by following troubleshooting instructions in the Data ONTAP guides.
3. If the issue is a back-end problem, check the Interoperability Matrix at support.netapp.com to ensure that the following is supported: the configuration, storage array, storage array firmware, switch, and switch firmware.

As of November 2012, the information that was in the *Interoperability Matrix* about these supported items was moved to the *Interoperability Matrix*.

4. If your system is running 7-Mode or a clustered Data ONTAP release prior to 8.1, check EMS messages for clues to the problem.
5. If your system is running clustered Data ONTAP 8.1 or later, use the `storage array show config` command to check whether there are any common back-end configuration errors that the system can detect.

If Data ONTAP detects a back-end configuration error, it tells you to run the `storage errors show` command to find details about the error.

When you troubleshoot your system using the `storage array show config` and `storage errors show` commands, you check for the same issues as you did during installation

verification of your system. You can check both errors that the system detects and differences in what you expected to see in your configuration.

6. If the cause of the problem is not yet apparent, check the following sources to ensure that your system is complying with requirements for working with storage arrays:
 - *V-Series Installation Requirements and Reference Guide*
 - *V-Series Systems Implementation Guide for Third-Party Storage*
 - *Interoperability Matrix*
 - *Hardware Universe* (formerly the *System Configuration Guide*) at support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml
7. If you still need help resolving the problem, contact technical support.

Related concepts

Verifying installation with storage arrays (clustered Data ONTAP 8.1 and later) on page 83

Verifying an installation with storage arrays (7-Mode in 8.x) on page 112

Invalid path setup examples

Path setup can be invalid because paths to an array LUN are not redundant or the number of paths to an array LUN does not meet Data ONTAP requirements.

Related concepts


Requirements for redundant setup of components in a path on page 33

Required number of paths to an array LUN on page 34

Invalid path setup: too many paths to an array LUN (7-Mode)

Systems operating in 7-Mode require two paths to an array LUN; more than two paths to an array LUN is not supported.

The path setup in the following example is invalid because the same array LUN would be accessed over four paths instead of only two paths:



In this invalid scenario, Connection 3 from 0a through Switch 1 is incorrectly zoned to both Connection 4 and Connection 5. Likewise, Connection 3 from 0c through Switch 2 is incorrectly zoned to both Connection 4 and Connection 5. The result is that array LUN 1 is seen over more than two paths.

For this configuration to be correct for 7-Mode, FC initiator port 0a must see either Controller 1 Port A or Port B on the storage array, but not both. Likewise, FC initiator port 0c must see either port Controller 2 Port A or Port B on the storage array, but not both.


Related concepts

[Required number of paths to an array LUN](#) on page 34

Invalid path setup: alternate paths are not configured

It is important to set up alternate paths to array LUNs to provide access to all V-Series LUNs from both V-Series FC initiators and to avoid a single point of failure (SPOF).

The following configuration is invalid because it does not provide alternate paths from each FC initiator port on the V-Series systems to each LUN on the storage array. Both FC initiator ports from the same V-Series system are connected to the storage array through the same switch.



Assume that the following zoning is in place in this invalid example:

- For vs1:
 - 0a is zoned to see Controller 1 Port A.
 - 0c is zoned to see Controller 1 Port B.
- For vs2:
 - 0a is zoned to see Controller 2 Port A.
 - 0c is zoned to see Controller 2 Port B.

In this sample configuration, each switch becomes a SPOF.

To make this a valid configuration, the following changes must be made:

- vs1's FC initiator port 0c must be connected to Switch 2.
- vs2's FC initiator port 0a must be connected to Switch 1.
- Appropriate zoning must be configured.

If you are using multiple ports on a storage array that supports configuring a specific set of LUNs on a selected set of ports, a given FC initiator port must be able to see all array LUNs presented on the fabric.

Related concepts

[Requirements for redundant setup of components in a path](#) on page 33

What happens when a link failure occurs

Data ONTAP monitors a link’s usage periodically. The Data ONTAP response to a link failure differs depending on where the failure occurs.

The following table shows what occurs if there is a failure in a fabric-attached configuration:

If a failure occurs in the link between the...	Then...
V-Series system and the switch	Data ONTAP receives notification immediately and sends traffic to the other path immediately.
Switch and the storage array	Data ONTAP is not immediately aware that there is a link failure because the link is still established between the V-Series system and the switch. Data ONTAP becomes aware that there is a failure when the I/O times out. Data ONTAP retries three times to send the traffic on the original path, then it fails over the traffic to the other path.

Relationship between zoning and host group configuration

When you fix zoning configuration errors you sometimes have to change host group configuration also, and the reverse.

Related concepts

[Dependency between zone and host group definitions](#) on page 120

[Example of cascading zoning and host group configuration errors](#) on page 122

[Zoning guidelines](#) on page 56

Dependency between zone and host group definitions

Errors made in zone definitions can require reconfiguration of host group definitions and the reverse.

When a zone definition is constructed, two ports are specified: the WWPN of the V-Series initiator port and the storage array port WWPN or WWNN for that zone. Likewise, when the host group for V-Series is configured on the storage array, the WWPNs of the V-Series initiator ports that you want to be members of the host group are specified.

The typical ordering of configuration is as follows:

1. Construct a zone definition.
2. Construct the host group on the storage array (by picking the WWPN of the V-Series initiator from the picklist).
3. Present array LUNs to the ports.

However, host groups are sometimes configured before zone definitions, which requires manually entering WWPNs in the host group configuration on the storage array.

Common errors

In Data ONTAP output, the V-Series initiator ports are identified by adapter number—for example, 0a, 0b, 0c, 0d, and so on for models with onboard ports. WWPNs are shown in the switch GUI and the storage array GUI. Because WWPNs are long and in hexadecimal format, the following errors are common:

How WWPNs are specified	Common error
The administrator types in WWPNs	A typing mistake is made.
WWPNs are automatically discovered by the switch	The wrong V-Series initiator port WWPN is selected from the picklist.

Note: When the V-Series systems, switches, and storage array are cabled together, the switch automatically discovers the WWPNs of the V-Series systems and storage array ports. The WWPNs are then available in picklists in the switch GUI, enabling selection of the WWPN of each zone member rather than typing it. To eliminate the possibility of typing errors, it is recommended that the switch discover WWPNs.

Cascading effect of errors

An obvious first step when troubleshooting problems with a fabric-attached configuration is to check whether zoning was configured correctly. Considering the relationship between the host group and zone definitions is also important. Fixing a problem might require reconfiguring both the zone definition and the host group definition, depending on where the error was made during the configuration process.

If the switch is automatically discovering WWPNs and zone definitions are configured first, the WWPNs of the V-Series initiator ports that will be used to access LUNs on the storage array are automatically propagated to the host group configuration picklists in the storage array GUI. Therefore, any zoning errors are also propagated to the storage array host group picklists. The picklists show the long, hexadecimal WWPNs instead of the short V-Series initiator port labels that are visible on the V-Series system (for example, 0a, 0b, and so on). Therefore, it is not easy to see that the WWPN you expected to be listed is not there.

The following table shows the effects of certain errors:

Zone definition on the switch	Host group configuration on the storage array	Symptom in Data ONTAP output
The V-Series initiator port in the zone definition is incorrect. This caused the incorrect V-Series initiator port WWPN to be propagated to the host group configuration.	The WWPN of the V-Series initiator port shown in the picklist was selected, not the WWPN you intended.	Array LUNs are not visible over the V-Series initiator port on which the expected LUNs would be visible.
The zone definition includes the correct V-Series initiator port.	<p>The WWPN in the host group definition is incorrect because of either of the following:</p> <ul style="list-style-type: none"> • The wrong WWPN was selected. • Host groups were configured manually before the zone definition was configured and a typing error was made when typing in the WWPN of the V-Series initiator port. 	

Related concepts

[Example of cascading zoning and host group configuration errors](#) on page 122

[Zoning requirements](#) on page 56

Example of cascading zoning and host group configuration errors

Errors made in zone definitions can impact host group definitions, and vice versa. When LUNs are not visible over a path, you need to check for both zoning and host group configuration errors.

Assume that your configuration sequence is as follows:

1. The zone definition was created on the switch.
The WWPN for V-Series initiator port 0a was put in the zone definition. However, the intention was that the WWPN for V-Series initiator port 0c was to be put into the zone definition.
2. The host group was created on the storage array.
The WWPN for V-Series initiator port 0a was selected (because that was the only WWPN available and it was not obvious that it was the WWPN for 0a and not 0c).
3. In Data ONTAP, you looked at array LUNs over the V-Series initiator ports, expecting to see array LUNs over 0c.
However, there were no array LUNs over 0c because both the zone definition and the host group definition incorrectly include the WWPN for V-Series initiator port 0a.

Note: You used the `sysconfig -v` command for 7-Mode or the `storage array show config` command for clustered Data ONTAP.

4. You start troubleshooting because you cannot see LUNs over the initiator over which you expected to see them.

You need to check both the zoning and host group configuration but it does not matter which of the following procedures you start with first. You might see different messages, depending on whether you start fixing things from the host group first or the zoning first.

Troubleshooting by checking the zoning first

1. Check the zone definitions for the V-Series system.

You realize that you have two zones with the WWPN for V-Series initiator port 0a in it and no zones with the WWPN for 0c in it.

2. Fix the incorrect zone definitions and activate them.

Note: For 7-Mode, if you were to run `sysconfig -v` on the V-Series system now to check the array LUNs over the V-Series initiator ports, you would still not be able to see the array LUNs over 0c because the host group still has the wrong initiator in it (0a is still in the host group). For clustered Data ONTAP, you would not see the array LUNs over the initiator ports when running `storage array config show`.

3. Go to the array and reconfigure the host group to include the WWPN for V-Series initiator port 0c.

Now that the WWPN for 0c is in a zone definition that has been activated, the WWPN for 0c shows up in the picklist in the host group configuration on the storage array.

4. On the V-Series system, run `sysconfig -v` for 7-Mode or `storage array config show` for clustered Data ONTAP to check the array LUNs over the V-Series initiator ports to confirm that array LUNs are shown over 0c.

Troubleshooting by checking the host group first

1. From the V-Series console, run `storage show adapter adapter#`, then write down the WWPN of the adapter that is missing—0c in this example.
2. Go to the storage array and compare the WWPN you wrote down to the WWPNs shown in the host group picklist to see whether the WWPN of the V-Series initiator port you expected is listed. If the WWPN you expected does not appear, then the initiator you intended is not in the zone definition.
3. If the storage array allows you to modify WWPNs in the host group, you could modify the WWPN shown to be the WWPN that you wrote down.

Note: If the storage array does not let you modify WWPNs in the host group, you need to modify the host group definition after modifying the zone definition.

You still cannot see LUNs over the initiator you were intending because the zoning has not been fixed yet.

4. Go to the switch and replace the incorrect WWPN with the correct V-Series port initiator, and then activate the zone definition.
5. If you could not correct the WWPN in the host group definition earlier in the process, go to the storage array and reconfigure the host group to include the WWPN for V-Series initiator port 0c. Now that the WWPN for 0c is in a zone definition that has been activated, the WWPN for 0c shows up in the picklist in the host group configuration on the storage array.
6. On the V-Series system, run `sysconfig -v` for 7-Mode or `storage array config show` for clustered Data ONTAP to check the array LUNs over the V-Series initiator ports to confirm that array LUNs are shown over 0c.
You should now see the LUNs you expected over the initiator you expected to see them.

Related concepts

[Dependency between zone and host group definitions](#) on page 120

[Zoning requirements](#) on page 56

Installation quick start (7-Mode and storage arrays only)

If you are familiar with setting up a system running Data ONTAP, quick start instructions might be sufficient to help you set up your V-Series system to work with a storage array.

The quick start installation instructions are for a V-Series 7-Mode configuration that uses only storage arrays.

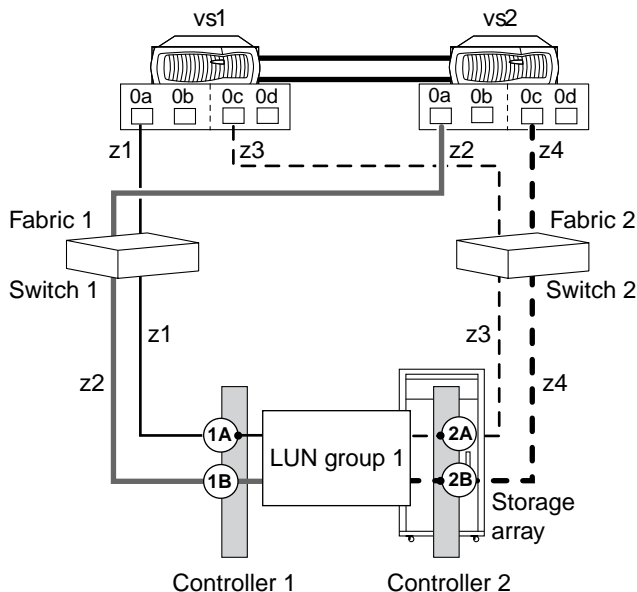
Steps

1. [Sample configuration for the installation quick start \(7-Mode and storage arrays\)](#) on page 125
2. [Performing preinstallation tasks on the storage array](#) on page 126
3. [Installing the V-Series system](#) on page 127
4. [Setting up the switches](#) on page 128
5. [Setting up LUN security](#) on page 128
6. [Assigning an array LUN to a V-Series system and creating the root volume](#) on page 129
7. [Installing Data ONTAP and licenses](#) on page 130
8. [Testing your setup](#) on page 131
9. [Additional setup](#) on page 133

Sample configuration for the installation quick start (7-Mode and storage arrays)

Setting up connectivity between a V-Series system and a storage array is basically the same for any configuration.

Refer to this *single 4-port array LUN group* example as you are using the quick start. This is the recommended configuration. It is supported with all V-Series models, all switches, and all storage arrays, regardless of vendor.



In this configuration with a single 4-port LUN group, array LUNs are mapped to four ports on the storage array. The array LUN group is presented to both nodes in the HA pair on different array target ports. However, each V-Series system can see an array LUN, end-to-end, through only two paths. Zoning is configured so that each FC initiator port on a V-Series system can access only a single target array port.

Note: V-Series FC initiator port names and storage array port names vary depending on the V-Series model and storage array model. In the illustration, *Controller 1* and *Controller 2* are the hardware components on which the ports are located. Different vendors and different storage array models use different terminology to represent these hardware components (for example, *cluster* or *controller* for Hitachi, HP XP, and IBM; *Storage Processor (SP)* for EMC CLARiiON; and *Controller Module* for Fujitsu ETERNUS).

Performing preinstallation tasks on the storage array

Before you can begin installing a V-Series configuration, the storage array administrator must prepare storage for Data ONTAP to use.

Steps

1. Ensure compliance with supported storage array models, firmware, switches, Data ONTAP version, and so on.

See the Interoperability Matrix at support.netapp.com for interoperability information and the *V-Series Limits Reference for Third-Party Storage* for information about limits such as the minimum array LUN size for the root volume.

2. Ask your storage administrator to create at least four LUNs on the storage array for the V-Series system.

Each node in the HA pair requires an array LUN for the root volume and an array LUN for core dumps.

3. Ask your storage array administrator to configure any parameters on the storage array that are required to work with Data ONTAP.

See the *V-Series Implementation Guide for Third-Party Storage* for information about the parameters that must be set to work with Data ONTAP.

4. Obtain appropriate Data ONTAP software.

Related concepts

[Zoning recommendation for a configuration with storage arrays](#) on page 57

[Minimum array LUN size for the root volume](#) on page 24

[Spare core array LUN requirement for core dumps](#) on page 23

Installing the V-Series system

After the storage administrator makes storage available to Data ONTAP, you are ready to install the V-Series system.

Steps

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:

```
Starting Press CTRL-C for special boot menu
```

2. Select “Maintenance mode boot” on the menu.

Do not proceed any further with V-Series system installation and setup at this time.

3. Check the settings of the V-Series HBAs to ensure that they are configured as initiators.

- a) Determine which ports are configured as target ports:

```
fcadmin config
```

- b) Configure the required ports as initiator ports:

```
fcadmin config -t initiator port#
```

4. Install the Fibre Channel cables connecting the V-Series system to switches and switches to the storage array.

Setting up the switches

Switch configuration is typically done by the storage or SAN administrator. Some customers use hard zoning and other customers use soft zoning.

About this task

The switches need to be zoned so that the V-Series systems and the storage arrays can see each other. You need to use single-initiator zoning so that the V-Series FC initiator ports do not see each other.

Step

1. Zone the switches:
 - a) Log in to the storage array and obtain the WWPNs of the FC adapters of the storage array.
 - b) Use the Fibre Channel switch commands to zone each switch so that the storage array and the V-Series system see each other's WWPNs.

In the example configuration, the zones are as follows for soft zoning:

Zone	V-Series system and port	Storage array controller and port
Switch 1		
z1	vs1 0a	Controller 1 1A
z2	vs2 0a	Controller 1 1B
Switch 2		
z3	vs1 0c	Controller 2 2A
z4	vs2 0c	Controller 2 2B

Related concepts

[Zoning guidelines](#) on page 56

Setting up LUN security

The storage array administrator must configure the storage array so that other hosts cannot access the array LUNs intended for use by Data ONTAP.

About this task

The concept of LUN security is similar to zoning except that LUN security is set up on the storage array. LUN security keeps different servers from using each other's storage on the SAN. LUN security might also be referred to as *LUN masking*.

Steps

1. Set up LUN security on the storage array.
2. Create host groups, or the equivalent, for the V-Series system.

The term *host group* is used on some storage arrays to describe a configuration parameter that enables you to specify host access to specific ports on the storage array. Different storage arrays use different terms to describe this configuration parameter. Each storage array vendor has its own process for creating a host group or the equivalent.

Related concepts

[Planning for LUN security on the storage arrays](#) on page 31

Assigning an array LUN to a V-Series system and creating the root volume

For each V-Series system, you must assign an array LUN to it and create its root volume before you can install Data ONTAP.

About this task

At this point it is easiest if you assign only one array LUN to each V-Series system. You can add additional array LUNs after you install Data ONTAP and verify your configuration.

For Data ONTAP 8.x, the root volume must be a FlexVol volume. You should assign only one array LUN to the aggregate with the root volume. The FlexVol volume can then use all of the space for the root volume.

Steps

1. Return to the V-Series console.

The system should still be in Maintenance mode.

2. Enter the following command to confirm that you can see the array LUNs created by the storage administrator:

```
disk show -v
```

If you do not see the array LUNs, reboot the V-Series system into Maintenance mode, and then double-check that the array LUNs exist, that the host groups were created correctly, that zoning is correct, and that cabling is correct.

3. Enter the following command to assign the first array LUN to the V-Series system:

```
disk assign {disk_name}
```

For example, on V-Series system 1, enter **disk assign L1**, and on V-Series system 2, enter **disk assign L2**.

Note: For Data ONTAP 8.1, use block checksum type (BCS), the default. For Data ONTAP 8.1.1 and later, advanced zoned checksum type (AZCS) is also available but BCS type is recommended.

4. Confirm that the system ID of the V-Series system is shown as the owner of the array LUN:

```
disk show -v
```

If the system ID of a V-Series system is shown, the array LUN was assigned to the V-Series system.

5. Exit Maintenance mode:

```
halt
```

6. At the boot environment prompt, enter the following command:

```
bye
```

7. Press Ctrl-c to interrupt the boot process and to display the boot options menu.
8. Select the **Clean configuration and initialize all disks** option to create the root volume with one of the array LUNs that you assigned to this storage system.
9. Enter **y** when the system prompts you whether you want to install a new file system.
10. Enter **y** when the system displays the This will erase all the data on the disks, are you sure? message.

The storage system creates a FlexVol root volume named `vol0` in an aggregate named `aggr0` (the system automatically creates the aggregate). After these are created on one of the assigned array LUNs, the system prompts for setup information.

Related concepts

Planning for Data ONTAP use of array LUNs on page 19

Characteristics of checksum types that Data ONTAP supports on page 25

Installing Data ONTAP and licenses

When you order a V-Series system without disk shelves, you must install Data ONTAP and the required V-Series license.

About this task

You can perform Data ONTAP software installation using either of the following methods:

- Map the V-Series C\$ share to your laptop as a CIFS share.
- On an HTTP server, use the `software install http://ipaddr/file.zip` command to install the software.

This procedure describes using the CIFS share method.

Steps

1. Install the CIFS license.
2. Run CIFS setup in workgroup mode.
3. Map the V-Series C\$ share to your laptop.
4. Make an `/etc/software` directory (or enter the `software list` command to create the `/etc/software` directory).
5. Copy the Data ONTAP executable to the `/etc/software` directory.
6. Enter the following command to run the software executable:
`software install release_setup.exe`
7. Download the Data ONTAP software.
8. Install the license for the V-Series system to use LUNs on storage arrays.
Starting in Data ONTAP 8.2, this license package is called `V_StorageAttach`.
9. Reboot the V-Series system.
10. If your V-Series system is a node in an HA pair, repeat the setup and Data ONTAP software installation steps on the partner node before validating your setup.

Testing your setup

Before putting a V-Series system into a production environment, you must test your setup to ensure that everything works.

About this task

Testing for proper setup of a V-Series system with a storage array includes the following:

- Checking your system to ensure that the configuration is as expected
- Verifying that there are two paths to storage
- Testing normal controller failover
- Testing path failover and controller failover

Steps

1. Use the following commands to confirm that the results of configuration are what you expect:

Use this command...	To...
<code>disk show -v</code>	Check whether all array LUNs are visible.
<code>sysconfig -v</code>	Check which V-Series FC initiator ports, switch ports, and array LUNs are used.

Use this command...	To...
<code>storage array show-config</code>	Display connectivity to back-end storage arrays.
<code>sysconfig -r</code>	Check the aggregate configuration and ensure that spare array LUNs are available.

2. Ensure that there are two paths to each array LUN so that the V-Series system can continue to operate when running on a single path:

- a) Enter the following command:

```
storage array show-config
```

- b) Check whether two paths to the array LUNs are shown.

If you do not see two paths to the array LUNs, check zoning, host group configuration, and cabling.

- c) Look at the adapters shown to see whether all paths are on a single adapter.

If you see both paths through only one V-Series system FC initiator port (the V-Series system's 0c port, for example) this is an indication that the back-end zoning is redundantly crossed. This is not a supported configuration.

Note: Do not continue with testing until you see two paths.

3. Test normal controller failover.

- a) On the partner node (vs2 in the example configuration), enter the following commands to perform a cf takeover; a takeover of vs1 should occur:

```
vs2> cf status
vs2> cf takeover (vs1 goes down; vs2 takes over)
vs1/vs2> df (vs2 looks at vs1)
vs1/vs2> partner (to switch to the partner)
vs2> sysconfig
vs2 (takeover)> cf status
vs2 (takeover)> cf giveback (to return to normal operation)
```

- b) Repeat the same commands on the local V-Series system (vs1 in the example configuration).

4. Test path failover and cluster failover:

- a) On the local V-Series system (vs1 in the example configuration), enter the following commands:

```
vs1> fcadmin offline 0a
vs1> storage show disk -p (you should see only one path) or enter
storage array show-config
vs1> fcadmin offline
vs1> fcadmin offline HA pair (takeover should occur)
```

- b) On the partner node (vs2 in the example configuration), enter the following commands:

```
cf takeover
```

```
cf giveback
```

- c) After both the local node (vs1) and the partner node (vs2) are back online, go to the partner node (vs2 in the example) and repeat the procedure.

Related concepts

[*Verifying an installation with storage arrays \(7-Mode in 8.x\)*](#) on page 112

Additional setup

After initial installation and testing, you can assign additional array LUNs to your V-Series systems and set up various Data ONTAP features on your systems.

Tasks after initial installation and testing include the following:

- Assign additional array LUNs to the V-Series systems as required.
After the basic V-Series setup is complete, you can have the storage administrator create additional array LUNs for the V-Series systems, as needed.
- Create Data ONTAP aggregates and volumes as desired.
- Set up additional Data ONTAP features on your V-Series system, for example, features for backup and recovery.

Related concepts

[*Planning for Data ONTAP use of array LUNs*](#) on page 19

[*Determining the array LUNs for specific aggregates*](#) on page 60

Obtaining WWNs manually

If the V-Series system is not connected to the SAN switch, you need to obtain the World Wide Port Names (WWPNs) of the V-Series FC initiator ports that will be used to connect the V-Series system to the switch.

About this task

Having the switch automatically discover WWPNs is the preferred method of obtaining WWPNs because you can avoid potential errors resulting from typing the WWPNs into the switch configuration.

Steps

1. Connect the V-Series system console connection to a laptop computer.
2. Power on your V-Series system.

Interrupt the boot process by pressing Ctrl-c when you see the following message on the console:

```
Starting Press CTRL-c for floppy boot menu
```

3. Select the Maintenance Mode option on the boot options menu.
4. Enter the following command to list the WWPNs of the V-Series system FC initiator ports:

```
storage show adapter
```

To list a specific adapter WWPN, add the adapter name, for example, `storage show adapter 0a`.

5. Record the WWPNs that will be used and leave the V-Series system in Maintenance mode.

Related tasks

[Connecting a V-Series system to a storage array](#) on page 74

Settings for connecting to an ASCII terminal console

You can attach an ASCII terminal console through the serial port on the back of your V-Series system if you want to do local system administration.

The ASCII terminal console enables you to monitor the boot process and helps you configure your V-Series system after it boots.

The ASCII terminal console is connected to your V-Series system with a DB-9 serial adapter, attached to an RJ-45 converter cable. The DB-9 adapter connects into the DB-9 serial port on the back of your V-Series system.

The following table shows how the DB-9 serial cable is wired. Input indicates data flow from the ASCII terminal to your V-Series system and output indicates data flow from your V-Series system to the ASCII terminal.

Pin number	Signal	Data flow direction	Description
1	DCD	Input	Data carrier detect
2	SIN	Input	Serial input
3	SOUT	Output	Serial input
4	DTR	Output	Data terminal ready
5	GND	N/A	Signal ground
6	DSR	Input	Data set ready
7	RTS	Output	Request to send
8	CTS	Input	Clear to send
9	RI	Input	Ring indicator

The following table shows the communications parameters for connecting an ASCII terminal console to a V-Series system. You need to set the following communications parameters to the same values for both your V-Series system and the ASCII terminal.

Parameter	Setting
Baud	9600
Data bit	8
Parity	None

Parameter	Setting
Stop bits	1
Flow control	None

Note: See your terminal documentation for information about changing the ASCII console terminal settings.

Related tasks

[Connecting a V-Series system to a storage array](#) on page 74

Target queue depth customization

The target queue depth defines the number of Data ONTAP commands that can be queued (outstanding) on a storage array target port. Data ONTAP supplies a default value. For most deployments, the default target queue depth is appropriate. But you can change it to correct performance issues.

The default target queue depth differs with different releases of Data ONTAP:

- For Data ONTAP 8.2, the default is 512.
- For all releases prior to Data ONTAP 8.2, the default is 256

When a storage array is configured with multiple initiators sharing target ports, you do not want the outstanding commands in the queue buffer from all initiators together to exceed what the storage array can handle. Otherwise, the performance of all hosts can suffer. Storage arrays differ in the number of commands they can handle in the queue buffer.

Note: Target queue depth might also be referred to as target queue length, Q-Depth, or Max Throttle.

Guidelines for specifying the appropriate target queue depth

You need to consider the impact of all the initiators on the storage array port when you are planning the configuration for a specific V-Series system or a specific non V-Series host.

If your deployment includes more than one initiator on a target port, you need to consider the total number of commands sent to a target port by all initiators when setting the target queue depth.

Guidelines for specifying the appropriate target queue depth are as follows:

- Do not configure a value of 0 (zero).
A value of 0 means that there is no limit on the outstanding commands.
- Consider the volume of commands that specific initiators would be likely to send to the target port.
You could then configure higher values for initiators likely to send a greater number of requests and a lower value for initiators likely to send a lesser number of requests.
- Configure non V-Series hosts according to the guidelines provided for those hosts.
- Consider setting the target queue depth on a per-target-port basis when workloads differ across ports (clustered Data ONTAP 8.2 and later).

Related tasks

[*Setting the target queue depth \(clustered Data ONTAP prior to 8.2 and 7-Mode\)*](#) on page 138

[*Setting the target queue depth \(clustered Data ONTAP 8.2\)*](#) on page 138

Setting the target queue depth (clustered Data ONTAP prior to 8.2 and 7-Mode)

The default target queue depth is acceptable for most implementations, but you can change it if you want to.

About this task

This setting is per V-Series system and it applies to all target ports on all storage arrays. For clustered Data ONTAP systems running Data ONTAP prior to 8.2, this option must be used (via nodeshell).

Step

1. Use the following option to set the target queue depth:

```
options disk.target_port.cmd_queue_depth value
```

Related concepts

[Guidelines for specifying the appropriate target queue depth](#) on page 137

Setting the target queue depth (clustered Data ONTAP 8.2)

The default target queue depth is acceptable for most implementations, but can be changed if performance issues are encountered.

About this task

In clustered Data ONTAP 8.2, you can set the target queue depth by storage array or on a per-target port basis.

Step

1. Use one of the following commands to set the target port queue depth on all target ports or on a specific target port of a storage array.

If you want to...	Use this command sequence...
Set the target port queue depth on all target ports for a storage array	<pre>set advanced storage array port modify -name array_name -max-queue-depth value</pre>

If you want to...	Use this command sequence...
Set the target port queue depth on a specific target port on a storage array	<pre>set advanced storage array port modify -name array_name -wwnn value -wwpn value -max- queue-depth value</pre>

Related concepts

Guidelines for specifying the appropriate target queue depth on page 137

Related tasks

Displaying target queue depth statistics (clustered Data ONTAP 8.2) on page 139

Displaying target queue depth statistics (clustered Data ONTAP 8.2)

If you suspect that a target queue depth setting is causing performance issues on your storage array, you should check the value that is set for the queue depth and check the state of requests on the FC initiator ports.

About this task

There are different levels of detail that you can access to determine whether there are problems processing requests on target ports. The following steps describe how to determine the current setting for the target port queue depth, determine if there are requests waiting on the ports, and display detailed port statistics to help you understand the workload on the port.

Steps

1. Use the `storage array show` command with the `-instance` parameter to show the current value of the target port queue depth.

```
>set advanced
>storage array show -instance

Name: HP2
      Prefix: HP-2
      Vendor: HP
      Model: HSV300
      options:
        Serial Number: 50014380025d1500
Target Port Queue Depth: 512
      LUN Queue Depth: 32
      Upgrade Pending: false
      Optimization Policy: eALUA
```

```
Affinity: aaa
Error Text: -
```

2. Use the storage array port show -fields max-queue-depth to show the queue depth setting for each port on the storage array.

```
>set advanced
> storage array port show -fields max-queue-depth
```

name	wwnn	wwpn	max-queue-depth
EMC_SYMMETRIX_1	50060480000001a0	50060480000001a0	-
EMC_SYMMETRIX_1	50060480000001a1	50060480000001a1	-
EMC_SYMMETRIX_1	50060480000001b0	50060480000001b0	-
EMC_SYMMETRIX_1	50060480000001b1	50060480000001b1	256

A value of '-' for Max Queue Depth indicates that the port does not have a specific max queue depth setting and is using the value set at the storage array level.

3. Use the storage array port show command to display performance information about storage array target ports.

The results of this command helps you determine if there are performance problems related to the ports. The %busy and %waiting values provide a high-level view of the performance on a port. If these values show a high percentage of requests waiting to be processed or show that the port is busy for a great percentage of time, then you might want to investigate further into the state of the port.

```
vgv3070f51::*> storage array port show
```

```
Array Name: HP2
WWNN: 50014380025d1500
WWPN: 50014380025d1508
Connection Type: fabric
Switch Port: vgbr300s70:9
Link Speed: 4 GB/s
Max Queue Depth: -
```

Node	Initiator	Count	LUN IOPS	KB/s	%busy	%waiting	Link Errs
vgv51-02	0a	21	2	53	0	0	0
vgv51-01	0a	21	2	48	1	0	0

4. You can obtain more detailed information about ports using the storage array port show -fields command with the average-latency-per-iop, average-pending, average-waiting, max-pending, or max-waiting fields

Terminology comparison between storage array vendors

Different vendors sometimes use different terms to describe the same things.

The following table provides a mapping between some common vendor terms.

Term	Vendor	Definition
host group	Hitachi, Sun	A configuration entity that enables you to specify host access to ports on the storage array. You identify the FC initiator port WWNs for the V-Series systems that you want to access the LUNs; the process differs according to vendor and sometimes differs for different storage array models of the same vendor.
	IBM DS4xxx/DS5xxx	
	EMC DMX	
	HP XP	
volume group	IBM DS8xxx	
Storage Group	EMC CX	
cluster	IBM XIV	
host affinity group	Fujitsu ETERNUS4000, ETERNUS6000, ETERNUS8000, ETERNUS DX8000, ETERNUS DX400	
host definition	3PAR	
host	3PAR, HP EVA	
—	IBM ESS	No concept of <i>host group</i> . You must create a host in the ESS user interface for each V-Series FC initiator port that you plan to connect to the storage array and map each host to a port.
parity group	IBM DS8xxx, IBM ESS, Hitachi, HP XP, Sun	The arrangement of disks in the back-end that together form the defined RAID level.
RAID group	Data ONTAP, EMC CX, Fujitsu ETERNUS	
array, RAID set	IBM DS4xxx/DS5xxx	

Term	Vendor	Definition
Parity RAID, Parity RAID group	EMC DMX	A DMX feature that provides parity data protection on the disk device level using physical parity volumes.
disk group	HP EVA	A set of physical disks that form storage pools from which you can create virtual disks.
parity set, RAID set	3PAR	A group of parity-protected <i>chunklets</i> . (A chunklet is a 256-MB block of contiguous space on a physical disk.)
cluster	Data ONTAP	In clustered Data ONTAP 8.x, a grouping of nodes that enables multiple nodes to pool their resources into a large virtual server and to distribute work across the cluster.
	Hitachi, HP XP, Sun	A hardware component on the storage arrays that contains the ports to which hosts attach.
	IBM XIV	An entity that groups multiple hosts together and assigns the same mapping to all the hosts.
controller	Data ONTAP	The component of a V-Series system that runs the Data ONTAP operating system and interacts with back-end storage arrays. Controllers are also sometimes called <i>heads</i> or <i>CPU modules</i> .
	Hitachi, HP EVA, HP XP, IBM	Hardware on the storage array on which the target ports are located.

Term	Vendor	Definition
interface module	IBM XIV	A hardware component on the storage arrays that contains the ports to which hosts attach.
node	3-PAR	
FEBE Board	EMC Symmetrix	
Storage processor (SP)	EMC CLARiiON	
Controller Module	Fujitsu ETERNUS	
LUN	Many storage arrays	A grouping of one or more disks or disk partitions into one span of disk storage space. In the Data ONTAP documentation, this is referred to as <i>array LUN</i> .
	Data ONTAP	The V-Series system can virtualize the storage attached to it and serve the storage up as LUNs to applications and clients outside the V-Series system (for example, through iSCSI and FC). Clients are unaware of where a front-end LUN is stored.
LUN, virtual disk	HP EVA	A virtual disk (called a “Vdisk” in the user interface) is a simulated disk drive created in a disk group. You can assign a combination of characteristics to a virtual disk, such as a name, redundancy level, and size. Presenting a virtual disk offers its storage to a host.
array LUN	Data ONTAP documentation, Data ONTAP storage management tools	The Data ONTAP documentation uses the term <i>array LUN</i> to distinguish LUNs on the storage arrays from front-end LUNs (Data ONTAP LUNs).

Term	Vendor	Definition
vLUN	3PAR	(volume-LUN) A pairing between a virtual volume and a logical unit number (LUN). For a host to see a virtual volume, the volume must be exported as a LUN by creating vLUNs on the storage array.
volume	IBM, IBM XIV	Equivalent to what other storage array vendors call a LUN.
	Data ONTAP	A logical entity that holds user data that is accessible through one or more of the access protocols supported by Data ONTAP, including Network File System (NFS), Common Internet File System (CIFS), HyperText Transfer Protocol (HTTP), Fibre Channel (FC), and Internet SCSI (iSCSI). The V-Series system treats an IBM volume as a disk.
	EMC DMX	A general term referring to a storage device. A physical volume corresponds to a single disk device.
virtual volume	3PAR	A virtual storage unit created by mapping data from one or more logical disks.

Glossary

- array LUN** The storage that third-party storage arrays provide to storage systems running Data ONTAP software. One array LUN is the equivalent of one disk on a native disk shelf.
- checksum** A form of redundancy check, a simple measure for protecting the integrity of data through error detection. A checksum is used mainly in data storage and networking protocols. It adds up the basic components of a message, typically the bytes, and stores the resulting value. Later, the checksum can verify that the message was not corrupted by doing the same operation on the data, and checking the sum.
- controller**
- For NetApp storage systems, the component that runs the Data ONTAP operating system and controls its disk subsystem.
Controllers are also sometimes called *storage controllers*, *storage appliances*, *appliances*, *storage engines*, *heads*, *CPU modules*, or *controller modules*.
 - For V-Series systems, the component that runs the Data ONTAP operating system and interacts with back-end storage arrays.
Controllers are also sometimes called V-Series systems, *heads*, or *CPU modules*.
 - For some storage array vendors, *controller* refers to the hardware component on which host adapters and ports are located.
Some vendors refer to this component as a *cluster*.
- family** Storage arrays in the same family share the same performance and failover characteristics. For example, members of the same family all perform active-active failover, or they all perform active-passive failover. More than one factor might be used to determine storage array families. For example, storage arrays with different architectures would be in different families even though other characteristics might be the same.
- host group** A management entity on a storage array that establishes connectivity between a host and storage (array LUNs) through port definitions. Hosts in the host group can access the LUNs associated with the host group; hosts that are not in the host group cannot access those LUNs. Different vendors use different terms to describe this concept. The process of creating a host group differs among vendors.
- LDEV** A piece of logical RAID storage configured from disks. Each LDEV has an internal number that is unique to the storage array. When an LDEV is presented out of a port on the storage array, the hosts see it as a LUN. The LUN ID—the external number seen by the hosts—must match on each of the two ports over which they are presented. LUN IDs do not have to be unique on the storage array but must be unique on a port. Different storage array vendors use different terms to describe this concept.

LUN group	A set of logical devices on the storage array that a V-Series system accesses over the same paths. The storage array administrator configures a set of logical devices as a group to define which host WWPNs can access them. Data ONTAP refers to this set of devices as a LUN group. The number of paths to a LUN group varies according to release and operating mode.
LUN security	A method for designating which hosts can access particular array LUNs. LUN security is similar to <i>switch zoning</i> in concept, but it is performed on the storage array. LUN security and <i>LUN masking</i> are equivalent terms to describe this functionality.
open target port	<p>A port that Data ONTAP can see on the fabric, but the port presents no array LUNs to Data ONTAP.</p> <p>An open target port can occur for a variety of reasons, and the reasons might vary among storage arrays. For example, an open target port would occur if the V-Series system is zoned into the storage array target port, but the initiator's WWPN is not in the host group for that target port. It occurs if the host group is empty.</p> <p>A V-Series system continues to function normally with an open target port. However, Data ONTAP cannot access any array LUNs that are not associated with a host group.</p>
target group port number	The identifier (TPGN) for the group of target ports on a specific storage array controller or field-replaceable unit (FRU), which is called a <i>target port group</i> . Data ONTAP uses the TPGN to detect whether paths to an array LUN are redundant. Some Data ONTAP command output shows the TPGN.

Copyright information

Copyright © 1994–2013 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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 NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

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 www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States 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 United States and/or other countries.

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

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

32xx

configuration restrictions [12](#)

A

access control LUNs

troubleshooting presentation to Data ONTAP [99](#)

advanced zoned checksums

Data ONTAP support for [25](#)

formula for calculating usable capacity with [26](#)

usable space in an array LUN [24](#)

aggregates

determination of checksum type of array LUN [61](#)

rules about mixing storage types in [60](#), [63](#), [64](#)

rules for storage array families [60](#)

array LUN groups

See LUN groups

array LUNs

See LUNs (array)

ASCII terminal

settings for connecting to [135](#)

AZCS

See advanced zoned checksums

B

back-end configurations

commands for troubleshooting [78](#)

detecting errors in [85](#)

troubleshooting errors with [80](#)

validating in clustered Data ONTAP [85](#)

block checksums

Data ONTAP support for [25](#)

formula for calculating usable capacity with [26](#)

impact on usable space in LUN [24](#)

C

checksums

advanced zoned

defined [25](#)

block

defined [25](#)

characteristics of supported types [25](#)

checking the type [62](#)

defined [145](#)

determining size of array LUNs with [26](#)

formulas for usable capacity and size [26](#)

impact on array LUN size and number [25](#)

impact on usable space in a LUN [24](#)

rules for aggregates [61](#)

zoned

defined [25](#)

commands

for checking back-end configuration [78](#)

storage array config show

validating back-end configuration with [85](#)

storage array path quiesce [78](#)

storage array path resume [78](#)

storage array show [78](#)

storage array show-config [112](#), [113](#)

storage array show-luns [112](#), [113](#)

storage disk show [78](#)

storage errors show

validating back-end configuration with [85](#)

configuration restrictions

V32xx [12](#)

configuration types

direct-attached [8](#)

fabric-attached [8](#)

MetroCluster support [9](#)

connecting FC initiator ports

usage requirements [46](#)

connecting to a native disk shelf

location of instructions for [54](#)

connecting to storage arrays

guidelines for [53](#)

how to [74](#)

redundancy requirement [53](#)

supported methods [8](#)

connections

redundancy requirements for [53](#)

controller

defined [145](#)

core dumps

contents of [23](#)

impact on usable space in an array LUN [24](#)

space requirement in array LUNs [23](#)

D

- Data ONTAP RAID groups
 - See* RAID groups (Data ONTAP)
- desired capacity of array LUNs
 - calculating [26](#)
- disk ownership
 - changing [21](#)
 - considerations when planning for [20](#)
 - defined [20](#)
 - how array LUN becomes available [20](#)
- disk shelves
 - guidelines for connecting FC initiators to [54](#)
- disks
 - FC initiator port usage when connecting to [46](#)
 - planning for use with V-Series systems [13](#)
 - with V-Series systems [7](#)

E

- EMC Symmetrix storage array
 - troubleshooting
 - Volume Set Addressing mismatches [91](#)
- Engenio access control LUNs
 - troubleshooting presentation to Data ONTAP [99](#)
- errors
 - back-end configuration troubleshooting [80](#)
 - storage errors show [86](#)

F

- fabric-attached configurations with array LUNs
 - support for [8](#)
- family
 - defined [60](#), [145](#)
- FC initiator ports
 - checking the number of LUNs visible on [52](#)
 - labeling guidelines [45](#)
 - sharing target ports
 - example [48](#)
 - rules for [47](#)
 - usage requirements for [46](#)

G

- glossary [145](#)

H

- host groups
 - defined [145](#)
 - impact of zoning errors on
 - example [122](#)
 - what to include in [19](#)

I

- implementation
 - stages with storage arrays [12](#)
- initiator ports
 - checking array LUNs visible on [52](#)
 - checking LUNs visible on [52](#)
 - example for sharing with multiple target ports
 - example [48](#)
 - FC labeling guidelines [45](#)
 - sharing one with multiple target ports
 - rules for [47](#)
 - usage requirements [46](#)
- installation
 - additional setup possibilities [133](#)
 - assigning array LUN [129](#)
 - creating root volume [129](#)
 - quick start for Data ONTAP in 7-Mode
 - software and licenses [130](#)
 - quickstart for Data ONTAP 8.x in 7-Mode [125](#)
 - setting up LUN security [128](#)
 - storage array preinstallation tasks [126](#)
 - switches [128](#)
 - testing [83](#), [112](#), [131](#)
 - validating for 8.x 7-mode [112](#)
 - validating for clustered Data ONTAP 8.1 and later
 - [83](#), [104](#)
- Interoperability Matrix
 - support information in [11](#)

L

- labeling guidelines
 - FC initiator ports [45](#)
- LDEVs
 - defined [145](#)
 - spanning path pairs [107](#)
 - troubleshooting
 - inconsistent LUN IDs detected [80](#)
 - LUN ID mismatches [88](#)
 - when created [19](#)
- limits for storage arrays

- where to find information needed for planning [10](#)
- link failures
 - Data ONTAP response to [120](#)
- load partitioning
 - using LUN groups for [35](#)
- location
 - of the root volume [29](#)
- LUN groups
 - configuration example with multiple [37](#)
 - configuration requirements for multiple [36](#)
 - defined [35](#)
 - load partitioning with [35](#)
 - paths to
 - single 2-port in stand-alone system [42](#)
 - single 4-port array LUN group example [43](#)
 - troubleshooting
 - fewer than expected [106](#)
 - incorrect number of LUNs in [110](#)
 - more than expected [107](#)
- LUN IDs
 - requirements for [93](#)
 - troubleshooting
 - duplicated on same target port [93](#)
 - mismatch in LDEVs [88, 91](#)
- LUN limits
 - maximum for neighborhoods [68](#)
 - where to find information needed for planning [10](#)
- LUN ownership
 - See disk ownership
- LUN security
 - defined [31](#)
 - methods and requirements [31](#)
 - planning for on storage arrays [31](#)
 - quick start setup for Data ONTAP in 7-Mode [128](#)
- LUNs (array)
 - assignment changes [21](#)
 - availability for Data ONTAP storage [20](#)
 - checking the checksum type of [62](#)
 - conflicting failover modes
 - troubleshooting [102](#)
 - defined [145](#)
 - FC initiator port usage when connecting to [46](#)
 - how they become available for host use [19](#)
 - impact of checksum on number [25](#)
 - impact of checksum on size [25](#)
 - names
 - changes in Data ONTAP displays [41](#)
 - format of [39](#)
 - paths reflected in [38](#)
 - number

- factors impacting [21](#)
 - per V-Series system, minimum [22](#)
- number on an FC initiator port
 - checking [52](#)
- partitioning the load over connections [35](#)
- paths to
 - advantages of four [35](#)
 - how reflected in names [38](#)
 - invalid examples [117, 118](#)
 - number required [34](#)
 - troubleshooting non-redundant [99](#)
 - when to check for redundancy [34](#)
- planning for Data ONTAP use of [14, 19](#)
- RAID protection for [17](#)
- result of link failure [120](#)
- rules about mixing storage types in aggregates [60, 63, 64](#)
- size
 - calculating desired capacity [26](#)
 - calculating usable capacity [26](#)
 - determining when using checksums [26](#)
 - maximum [23](#)
 - minimum [23](#)
 - root volume minimum [24](#)
 - troubleshooting [87](#)
- types supported by Data ONTAP [21](#)
- usable space in [24](#)

M

- MetroCluster configurations
 - type of storage supported for [9](#)
- minimum array LUNs per system [22](#)

N

- names of array LUNs
 - changes in Data ONTAP displays [41](#)
 - format of [39](#)
 - how paths are reflected in [38](#)
- native disks on a V-Series system
 - connecting [54](#)
 - description [7](#)
 - planning port-to-port connectivity for [54](#)
- neighborhood maximum LUN limits
 - defined [68](#)
- neighborhoods
 - described [67](#)
 - determining whether to use [67](#)
 - establishing, requirements for [70](#)

limits

- calculating visible limit [71](#)
- maximum number of devices in [68](#)
- requirements [68](#)

O

open target port

- defined [145](#)

owning array LUNs

- considerations when planning for [20](#)
- defined [20](#)

P

paths to array LUNs

- advantages of four [35](#)
- examples
 - invalid, no alternate paths [118](#)
 - invalid, too many paths [117](#)
 - valid, single 2-port array LUN group [42](#)
 - valid, single 4-port array LUN group [43](#)

how reflected in names [38](#)

invalid setup examples

- alternate paths are not configured [118](#)
- too many paths to an array [117](#)

number required [34](#)

redundancy requirements for [33](#), [34](#)

troubleshooting [96](#), [109](#)

troubleshooting inconsistent [102](#)

ports

FC initiator

- guidelines for connecting to native disk shelves [54](#)
- requirements for [46](#)
- rules for sharing with multiple target ports [47](#)

Q

quick start installation for Data ONTAP in 7-Mode

- adding additional LUNs [133](#)
- assigning array LUNs [129](#)
- installing the V-Series system [127](#)
- LUN security setup [128](#)
- preinstallation tasks [126](#)
- sample configuration [125](#)
- switch setup [128](#)
- testing [131](#)

quiescing a path [78](#)

R

RAID groups (Data ONTAP)

- LUN number implications [17](#)
- LUN size implications [17](#)

RAID protection

- for array LUNs [17](#)

RAID0

- aggregate checksum type for array LUNs [61](#)
- how Data ONTAP uses for array LUNs [17](#)
- use by Data ONTAP [17](#)

redundancy requirements

- FC initiator port usage [46](#)
- to array LUNs [33](#)

requirements

- multiple LUN group configuration [36](#)
- neighborhoods [68](#)
- redundancy to array LUN [33](#)

restrictions

- V32xx [12](#)

root volumes

- array LUN size minimum [24](#)
- location of [29](#)

S

shared FC initiator ports

- checking array LUNs visible on [52](#)
- example [48](#)
- rules for [47](#)

size

- minimum for the root volume [24](#)
- supported for array LUNs [23](#)

Snapshot copies

- impact on array LUN usable space [24](#)

Snapshot reserve

- impact on array LUN usable space [24](#)

spare array LUNs

- checking the type [62](#)

spare core array LUNs

- minimum size for [23](#)
- when needed [23](#)

storage array config show command

- described [78](#)
- interpreting output [104](#)
- troubleshooting with [85](#)
- validating back-end configuration with [85](#)

storage array path quiesce command

- described [78](#)

storage array path resume command

- described [78](#)
- storage array show command
 - described [78](#)
- storage array show-config command
 - using to check number of paths [112](#)
- storage array show-luns command
 - using to check number of paths [112](#)
- storage arrays
 - guidelines for connecting system to [53](#)
 - methods for connecting to [8](#)
 - number behind a V-Series system [8](#)
 - planning LUN security for [31](#)
 - preinstallation tasks [126](#)
 - requirement for dedicated [9](#)
 - rules about mixing in aggregates [60](#)
 - supported, where to find information about [11](#)
 - troubleshooting when missing in command output [111](#)
 - zoning recommendation for [57](#)
- storage disk show command
 - described [78](#)
- storage errors show command
 - described [78](#)
 - errors detected by [80](#)
 - index to error details [86](#)
 - issues it does not identify [81](#)
 - troubleshooting with [85](#)
 - validating back-end configuration with [85](#)
- Support Matrix
 - See* Interoperability Matrix
- switches
 - setting up [128](#)
 - supported with back-end storage arrays
 - finding information about [11](#)
 - zoning in configurations with storage arrays [56](#)

T

- target port group number
 - defined [145](#)
- target port queue depth
 - displaying value of [139](#)
- target queue depth
 - defined [137](#)
 - displaying statistics for [139](#)
 - displaying value of target port que depth [139](#)
 - guidelines for specifying [137](#)
 - setting for 7-Mode [138](#)
 - setting in 8.2 clustered Data ONTAP [138](#)
 - setting prior to Data ONTAP 8.2 [138](#)

- technology overview for V-Series systems [7](#)
- terminology
 - checksum [25](#)
 - core dump [23](#)
 - family [60](#)
 - host group [19](#)
 - LDEV [19](#)
 - LUN group [35](#)
 - open target port [145](#)
 - storage array vendor [141](#)
 - target port group [145](#)
- troubleshooting
 - access control array LUN presentation [99](#)
 - array LUN size [87](#)
 - back-end configuration errors [80](#)
 - checking for intended configuration [104](#)
 - commands for [78](#), [85](#)
 - conflicting failover mode of array LUNs [102](#)
 - duplicated LUN IDs on same target port [93](#)
 - getting started with [116](#)
 - invalid pathing
 - no alternate paths [118](#)
 - too many paths [117](#)
 - LUN groups
 - fewer than expected [106](#)
 - incorrect number of LUNs in [110](#)
 - more than expected [107](#)
 - missing storage arrays in command output [111](#)
 - non-redundant paths to an array LUN [99](#)
 - path failover mode conflicts [102](#)
 - paths to array LUNs [96](#), [109](#)
 - Volume Set Addressing mismatch [91](#)
- troubleshooting V-Series configurations [116](#)

U

- usable capacity in array LUNs
 - calculating [26](#)
 - elements impacting [24](#)

V

- V-Series implementation
 - planning tasks for [10](#)
- V-Series MetroCluster configurations
 - See* MetroCluster configurations
- V-Series neighborhoods
 - See* neighborhoods
- V-Series Support Matrix
 - See* Interoperability Matrix

V-Series systems

- connecting to a storage array [74](#)
- how they use storage [7](#)
- minimum number of array LUNs for [22](#)
- number of storage arrays behind [8](#)

V3210

- configuration restrictions [12](#)

V32xx

- configuration restrictions [12](#)

validating a clustered Data ONTAP installation [83](#)

validating a Data ONTAP 8.x 7-Mode installation [112](#)

visible array LUNs on an FC initiator

- determining number of [52](#)

visible device limit

- calculating [71](#)

Volume Set Addressing

- mismatches [91](#)

W**WWNs**

- obtaining manually [134](#)

Z**zoning**

errors impacting host groups

- example [122](#)

examples with shared FC initiator ports [48](#)

examples with storage arrays [58](#)

guidelines with storage arrays [56](#)

recommendation for storage arrays [57](#)

requirements in a V-Series configuration [56](#)