



Installing FlexArray

ONTAP FlexArray

NetApp

February 11, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap-flexarray/install/concept_systems_that_can_use_array_luns_on_storage_arrays.html on February 11, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- FlexArray® Virtualization Installation Requirements and Reference 1
 - FlexArray Virtualization technology overview—using array LUNs for storage 1
 - Zoning in a configuration with storage arrays 3
 - Planning a configuration with array LUNs 5
 - Planning for RAID implementation 10
 - Planning for ONTAP use of array LUNs 11
 - Planning for LUN security on the storage arrays 19
 - Planning for paths to array LUNs 20
 - Planning a port-to-port connectivity scheme 33
 - Determining the array LUNs for specific aggregates 45
 - Preparing a storage array for use with ONTAP systems 49
 - Connecting an ONTAP system to a storage array 50
 - Setting up the switches 51
 - Configuring LUN security 52
 - Setting up ONTAP to work with array LUNs 53
 - Commands for checking back-end configuration 61
 - Verifying installation with storage arrays 65
 - Managing array LUNs using ONTAP 95
 - Troubleshooting configurations with storage arrays 99
 - Additional tasks after installing and testing an ONTAP configuration with array LUNs 105
 - Obtaining WWPNs manually 106
 - Target queue depth customization 106
 - Terminology comparison between storage array vendors 111

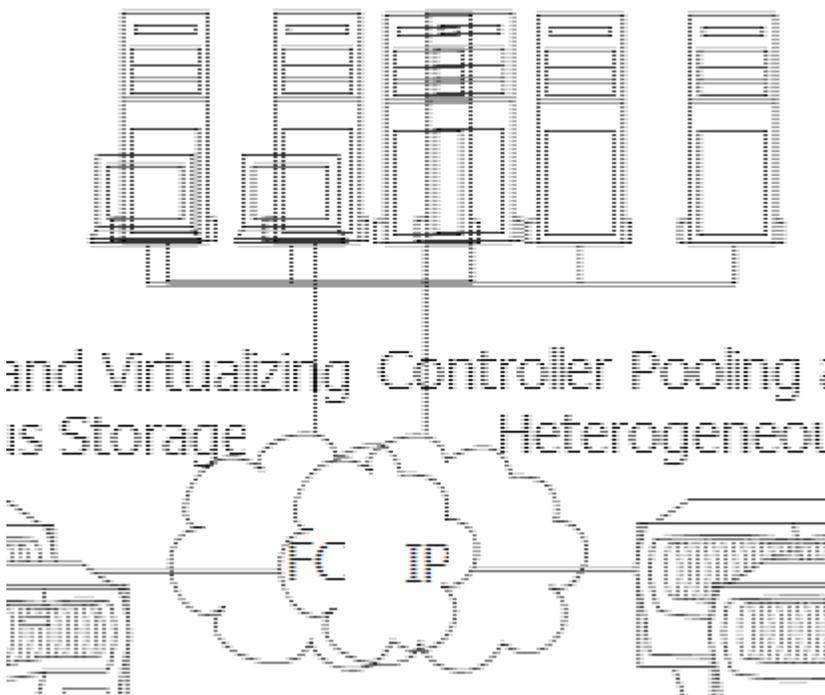
FlexArray® Virtualization Installation Requirements and Reference

FlexArray Virtualization technology overview—using array LUNs for storage

The 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 require it, without disruption. This functionality is provided by *FlexArray Virtualization Software*, formerly known as V-Series.

The following illustration shows a configuration in which ONTAP systems that are licensed to attach to storage arrays are pooling LUNs from the storage arrays and presenting that storage to clients.

Windows and Linux Hosts/Hosts



An ONTAP system presents storage to clients either in the form of ONTAP file system volumes, which you manage on the system by using 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 systems that can use array LUNs, you combine the array LUNs into one or more array LUN aggregates. In an ONTAP environment, you can associate these array LUN aggregates with storage virtual machines (SVMs), and create ONTAP volumes for presentation to the clients as files or as LUNs served by ONTAP.

ONTAP systems that can use array LUNs on storage arrays

You can use supported FAS and V-Series systems with array LUNs.

The NetApp Interoperability Matrix Tool lists the supported combinations of hardware and software.

Related information

[NetApp Interoperability Matrix Tool](#)

Considerations when connecting multiple storage arrays to an ONTAP system

Depending on the storage array models supported with ONTAP systems, you can connect either multiple storage arrays or a single storage array to the ONTAP system.



The Interoperability Matrix identifies any storage array model for which only one storage array is supported with an ONTAP system.

Consider the following when connecting multiple storage arrays to an ONTAP system:

- If you can connect multiple arrays of the same supported storage array model to the ONTAP system, there is no limit on the number of storage arrays that you can deploy.
- The storage arrays can either be from the same supported vendor or from different vendors.

If the storage arrays are from the same vendor, then they can be all from the same family or from different families.



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 information

[FlexArray virtualization implementation for third-party storage](#)

[NetApp Interoperability Matrix Tool](#)

Supported methods for connecting an ONTAP system to a storage array

You can connect ONTAP system in a fabric-attached configuration with storage arrays. Fabric-attached configurations are supported for both stand-alone systems and HA pairs. Direct-attached configurations are limited to some storage arrays and some ONTAP releases.

The Interoperability Matrix contains information about the connection methods supported for specific storage arrays and platforms running ONTAP.

Related information

[NetApp Interoperability Matrix Tool](#)

Sharing storage arrays among hosts

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

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

Related information

[NetApp Interoperability Matrix Tool](#)

Requirements for MetroCluster configurations with array LUNs

Before you set up a MetroCluster configuration with array LUNs, you must ensure that the ONTAP systems and the storage used in the configuration follow basic supportability requirements.

The basic requirements for supporting MetroCluster configurations with array LUNs are as follows:

- You can use both V-Series systems and FAS systems with support for array LUNs in MetroCluster configurations.

However you must ensure that all the ONTAP systems used in a MetroCluster configuration are of the same model.

- The ONTAP systems can use native disks only, LUNs on storage arrays only, or both.
- If you are setting up a MetroCluster configuration with both native disks and array LUNs, you must use FC-to-SAS bridges and native disks supported by those bridges.
- The ONTAP systems and storage arrays must be identified in the *Interoperability Matrix* as being supported in MetroCluster configurations.

Related information

[NetApp Interoperability Matrix Tool](#)

[Fabric-attached MetroCluster installation and configuration](#)

[Stretch MetroCluster installation and configuration](#)

Zoning in a configuration with storage arrays

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

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.

Requirements for zoning in a configuration with storage arrays

You must follow the zoning requirements in your ONTAP configuration with storage arrays to ensure that the ONTAP system can access the correct set of LUNs.

- The *Interoperability Matrix* must identify a switch and the switch firmware as supported for the ONTAP configuration.

- Zoning must be configured to restrict each initiator port to a single target port on each storage array.
- On the switch, ports on the ONTAP system and ports on the storage array must be assigned to the same zone.

This enables the ONTAP systems to access the LUNs on the storage arrays.

- When storage array ports are shared across heterogeneous systems, array LUNs from the ONTAP system cannot be exposed to other systems.

LUN security or array LUN masking must be used to ensure that the array LUNs for ONTAP storage are visible only to the ONTAP systems.

- A host configuration port must not be included in the same zone as a target port.

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 FC initiator port and a single storage array target port.

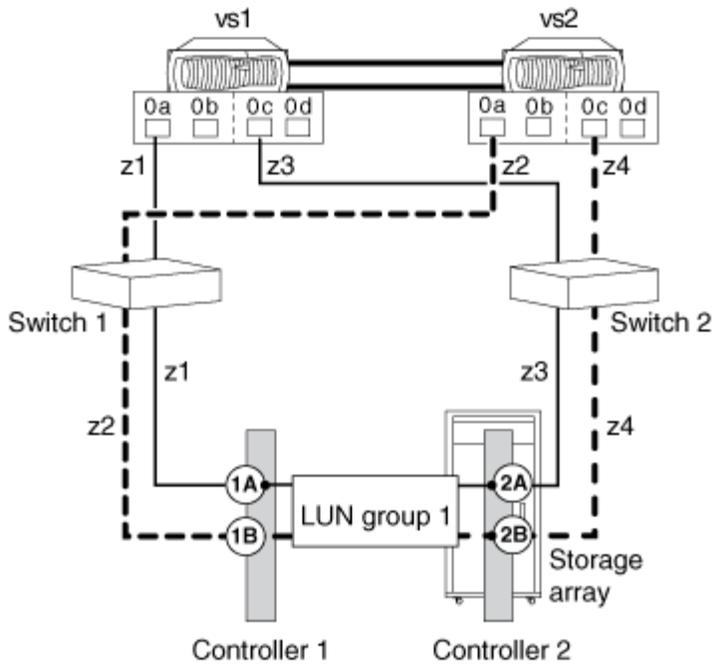
The benefits of creating 1:1 zoning are as follows:

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

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.

Zoning in a single 4-port LUN group configuration



The following table shows zoning for this example configuration of ONTAP systems in an HA pair. Single-initiator zoning is the recommended zoning strategy.

Zone	ONTAP system		Storage array	
Switch 1				
z1	vs1	Port 0a	Controller 1	Port 1A
z2	vs2	Port 0a	Controller 1	Port 1B
Switch 2				
z3	vs1	Port 0c	Controller 2	Port 2A
z4	vs2	Port 0c	Controller 2	Port 2B

Planning a configuration with array LUNs

For planning an ONTAP configuration with array LUNs, you must verify if the storage array is supported for a particular ONTAP release. In addition, you must check the interoperability and support information to ensure that all the systems meet the requirements for a configuration with array LUNs.

Related information

[NetApp Support](#)

[NetApp Interoperability Matrix Tool](#)

Where to find information for configurations with storage arrays

When planning your configuration for using ONTAP systems with storage arrays, you should check various sources for information about array LUN configuration in addition to the product documentation.

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 information

[NetApp Support](#)

Limit types for configurations with storage arrays

You must consider certain storage array limits when planning for an ONTAP configuration.

The *Hardware Universe* contains specific limit values for storage arrays and native disks.

The following types of limits apply only to storage arrays and not to native disks:

- Minimum and maximum array LUN size that 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

Related information

[NetApp Hardware Universe](#)

Where to find information about ONTAP support for storage arrays

Not all ONTAP releases support the same features, configurations, system models, and storage array models. During your deployment planning, you must check ONTAP support information to verify that your deployment conforms to ONTAP hardware and software requirements for all systems in the deployment.

The following table lists the information sources that contain the details of the hardware and software requirements associated with ONTAP systems:

For information about...	You should look here...
<p>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>NetApp Interoperability Matrix Tool</p>
<p>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>NetApp Hardware Universe</p>
<p>Setting up E-Series storage arrays, including the following:</p> <ul style="list-style-type: none"> • Site preparation requirements • Cabling instructions • SANtricity software installation and configuration instructions 	<p>The following E-Series documentation:</p> <ul style="list-style-type: none"> • <i>E-Series Storage Systems Site Preparation Guide</i> • <i>E-Series Storage Systems Hardware Cabling Guide</i> • <i>SANtricity ES Storage Manager documentation</i> <p>You can access these documents from the NetApp Support site.</p> <p>NetApp Support</p>
<p>What is supported for specific storage arrays, including supported configurations</p>	<ul style="list-style-type: none"> • FlexArray virtualization implementation for third-party storage • FlexArray virtualization implementation for NetApp E-Series storage

32xx system configuration restrictions

There are some restrictions for 32xx systems that do not apply to other models. You must be aware of these while configuring the systems.

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.

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 ONTAP system with a storage array, up to the point where you can install 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 ONTAP, assigning specific array LUNs to an ONTAP system
6. In ONTAP, providing information to set up an ONTAP system on the network
7. Installing ONTAP software

If an ONTAP system is ordered with disk shelves, the ONTAP software is installed by the factory. In such a configuration, you do not need to create the root volume and install licenses and ONTAP software.

If an ONTAP system is ordered without disk shelves, you need to configure ONTAP before setting up the cluster.

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

Planning summary for a V-Series system using native disks

You can install native disk shelves on new or existing V-Series systems. However, you must consider a few additional factors if you plan to install native disk shelves on V-Series systems, compared to installing the disk shelves on FAS systems.

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

You must consider the following when determining basic setup and installation requirements for your V-Series system that uses native disks:

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

Additional planning considerations if the V-Series 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	Location of the root volume
How many disks and array LUNs combined can be assigned without exceeding the supported maximum limit for the V-Series system	NetApp Hardware Universe
FC initiator port usage	Requirements for FC initiator port usage
The type of data that should reside on disks and the type of data that should reside on array LUNs	You must evaluate the type of data that needs to be managed, and then decide whether the data can reside on native disks or array LUNs.

Planning summary for ONTAP systems using array LUNs

When you plan to use array LUNs with ONTAP systems, you must communicate with the storage array and switch administrators so that the back-end devices are configured to work with ONTAP systems.

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

Planning task	Where to find information
Determining requirements for setting up your storage array to work with ONTAP	<ul style="list-style-type: none">• FlexArray virtualization implementation for third-party storage• FlexArray virtualization implementation for NetApp E-Series storage• NetApp Interoperability Matrix Tool
Guidelines for array LUN use with ONTAP	Planning for ONTAP use of array LUNs

Planning task	Where to find information
Determining ONTAP limits regarding array LUNs	NetApp Hardware Universe
Determining a LUN security scheme, setting access controls on the storage array and, if switches are deployed, setting zoning on switches	<ul style="list-style-type: none"> • Planning for LUN security on storage arrays • Zoning requirements
Determining a port-to-port connectivity scheme between the ONTAP systems and the storage array	<ul style="list-style-type: none"> • Planning a port-to-port connectivity scheme • NetApp Interoperability Matrix Tool
Determining which ONTAP system is to “own” which array LUN (disk ownership)	<ul style="list-style-type: none"> • How array LUNs become available for ONTAP storage use • Assigning ownership of array LUNs

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

ONTAP uses RAID0 (striping) for array LUNs. 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.



- 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.

- ONTAP supports RAID4 and RAID-DP on native disk shelves, but supports only RAID0 on array LUNs.

Planning considerations for ONTAP RAID groups

Before you set up ONTAP RAID groups for array LUNs, you must plan the size of aggregates, the number and size of RAID groups, and the size of array LUNs. RAID groups that best meet your data requirements ensure adequate protection on the array

for data and its availability.

Planning for ONTAP RAID groups involves the following tasks:

1. Planning the size of the aggregate that best meets your data requirements.
2. Planning the number and size of RAID groups required for the size of the aggregate.
3. Planning the size of the array LUNs required in the ONTAP RAID groups:
 - To avoid a performance penalty, all array LUNs in a particular 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 ONTAP RAID groups for use with storage arrays, see [Disks and Aggregates management](#).

Planning for ONTAP use of array LUNs

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

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

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



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 ONTAP 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).



LDEV is a term used by some vendors and this content 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.



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.

How array LUNs become available for ONTAP storage use

An ONTAP system cannot use an array LUN presented to it until ONTAP has been configured to use the array LUN.

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

1. One ONTAP system (licensed to use array LUNs) 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 an ONTAP system, 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 an ONTAP system, it becomes a spare LUN owned by that system and it is no longer available to any other ONTAP system.

A spare array LUN cannot be used for storage until you add it to an aggregate. Thereafter, 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.

Considerations when planning for disk ownership

If you are deploying multiple ONTAP systems for use with array LUNs, you must determine which system *owns* which array LUNs. Disk ownership ensures that only the ONTAP system that owns a particular array LUN can read data from and write data to the array LUN.

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 information

[NetApp Hardware Universe](#)

Array LUN assignment changes

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

Array LUN type supported by ONTAP

You can map only storage array LUNs to ONTAP. You can map LUN 0 to ONTAP if it is a storage type LUN.

Some storage arrays have a nonstorage *command* LUN. You cannot map a command type LUN to an ONTAP system.

Related information

[FlexArray virtualization implementation for third-party storage](#)

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 required number of array LUNs

While planning for ONTAP use of array LUNs, you must consider factors such as array LUN size, ONTAP overhead, and checksum type that impact the number of array LUNs required in your ONTAP environment.

You should consider the following when determining the required number of array LUNs:

- The smaller the array LUNs, the more LUNs you require for the storage.

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 an ONTAP system.

The *Hardware Universe* contains information about device limits.

- The more usable space in an array LUN, the fewer array LUNs are required.

The amount of usable space in an array LUN is determined by the space that 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 ONTAP systems, you must consider what the storage will be used for and the number of requests that are likely to be generated by different applications.

Minimum number of array LUNs required per ONTAP system

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

The root volume can be either on a disk or on array LUNs. The location of the root volume then determines the minimum number of array LUNs you need. If the root volume is on a storage array, each stand-alone ONTAP 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 that is using array LUNs, 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.

Spare array LUN requirement for core dumps

For both stand-alone ONTAP systems and nodes in HA pairs, you must create a spare 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 nor a spare disk is available, there is no place for the core to be dumped.

A core dump contains the contents of memory and NVRAM. During a system panic, 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.

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

Related information

[NetApp Hardware Universe](#)

Minimum and maximum array LUN sizes supported for ONTAP configurations

The array LUNs presented from the storage arrays must fall within the minimum and maximum size requirements for an ONTAP configuration with array LUNs. ONTAP issues an error message identifying an array LUN that does not adhere to the minimum or maximum size requirements.

The minimum and maximum array LUN sizes are calculated according to the way that ONTAP defines units of measure. The 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 storage array vendors use different formulas for calculating units of measurement. You must use your vendor's units of measure to calculate the array LUN sizes equivalent to the minimum and maximum sizes supported with an ONTAP configuration.

The maximum LUN size that ONTAP supports differs with ONTAP releases. For information about the minimum and maximum array LUN sizes, see the *Hardware Universe*.



The minimum array LUN size for a data (storage) LUN is different from the minimum array LUN size for the root volume.

Related information

[NetApp Hardware Universe](#)

Array LUN size required for the root volume

You must set the size of the root volume to be greater than the supported minimum array LUN size, to ensure that there is sufficient space in the root volume for system files, log files, and core files. You must provide these files to technical support if a system issue occurs.

The *Hardware Universe* lists the minimum array LUN size for a root volume.

Related information

[NetApp Hardware Universe](#)

Elements that reduce the usable space in an array LUN

Several factors impact the usable space in an array LUN. When planning for the required number and size of the array LUNs, you must consider the usable space in the array LUN according to the checksum type you are using, and the elements you might configure.

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

- Space reserved for use by ONTAP
- Space for core dump
- Volume-level Snapshot reserve
- Aggregate-level Snapshot copy
- Checksum type (you assign one type):
 - Block checksum (BCS)

- Advanced Zoned checksum (AZCS)

Considering checksum type when planning array LUN size and number

When planning the number and size of array LUNs that you need for 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 an ONTAP system.

When an array LUN on the storage array is mapped to be used by an ONTAP system, ONTAP treats the array LUN as a raw, unformatted disk. When you assign an array LUN to an ONTAP system you specify the checksum type, which tells 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.

Characteristics of checksum types that ONTAP supports

ONTAP supports block checksum type (BCS) and advanced zoned checksum type (AZCS) for array LUNs, disks, and aggregates.

The type of checksum assigned to an array LUN in 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.

Block checksum (BCS)

BCS is the default and recommended checksum type for array LUNs. BCS provides better performance for array LUNs than AZCS.

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.

Advanced zoned checksum (AZCS)

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.

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.

There is no AZCS performance impact on native disks.

The guidelines for checksum types differ according to disk size and type. See *TR3838 Storage Subsystem Configuration Guide* for more information.

Related information

[NetApp Technical Report 3838: Storage Subsystem Configuration Guide](#)

Formulas for calculating array LUN size based on 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 has 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 require 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 require:

If you know...	You want to find out...
How large your array LUNs are	How much capacity is available for storage (usable capacity). You have 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 require. You have to take into account your required amount of storage and space required for other elements.



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 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.

Checksum type	Formula
BCS—array LUNs less than 2 TB	$N \times \{0.875 \times 0.9 \times 0.99 \times (1 - \text{Snapshot reserve})\} = Y$
BCS—array LUNs greater than 2 TB	$N \times \{0.875 \times 0.9 \times 0.998 \times (1 - \text{Snapshot reserve})\} = Y$
AZCS—array LUNs less than 2 TB	$N \times \{0.984 \times 0.9 \times 0.99 \times (1 - \text{Snapshot reserve})\} = Y$
AZCS—array LUNs greater than 2 TB	$N \times \{0.984 \times 0.9 \times 0.998 \times (1 - \text{Snapshot reserve})\} = Y$

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)	$4 \times \{0.875 \times 0.9 \times 0.99 \times 0.95\} = 2.96$ GB usable space for storage
AZCS (array LUN less than 2 TB)	$4 \times \{0.984 \times 0.9 \times 0.99 \times 0.95\} = 3.33$ 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)	$Y \div \{0.875 \times 0.9 \times 0.99 \times (1 - \text{Snapshot reserve})\} =$ Actual capacity required
AZCS (array LUN less than 2 TB)	$Y \div \{0.984 \times 0.9 \times 0.99 \times (1 - \text{Snapshot reserve})\} =$ Actual capacity required

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

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

Location of the root volume

The location of the root volume depends on whether the ONTAP system is pre-ordered with native disks or whether you are adding disks to an ONTAP system that is configured for array LUNs.

Follow these guidelines for determining the location of the root volume on the ONTAP system:

- The root volume can be on a storage array or on a native disk shelf.

However, you should install the root volume on a native disk if the ONTAP system has both native disks and array LUNs.

If you order your storage system with disks, the factory installs the root volume on a native disk.

- In an HA pair, the best practice is for the root volume to be located on the same type of storage for both the nodes—either on a native disk shelf or on a storage array for both the nodes.
- For MetroCluster configurations with both disks and array LUNs, you should 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.

Planning for LUN security on the storage arrays

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



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

Available LUN security methods

Various LUN security methods help you designate which hosts can access particular array LUNs. You can use port-level security or LUN security products, or dedicate a storage for use by ONTAP systems.

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.

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 either use a LUN security product or dedicate the storage array to the ONTAP system. You should check your storage array documentation to determine whether your storage array supports port-level security.

LUN security products

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

Dedicate the storage array for ONTAP use

You can dedicate the storage array for use by ONTAP systems. In this case, no hosts other than the ONTAP systems are connected to the storage array.

You should use both zoning and LUN security to achieve added protection and redundancy for the ONTAP systems.

In addition to following the LUN security methods, you should also check for any additional details regarding LUN security for your vendor's storage arrays. Some storage arrays must be dedicated for use by ONTAP systems.

Related information

[FlexArray virtualization implementation for third-party storage](#)

[FlexArray virtualization implementation for NetApp E-Series storage](#)

Planning for paths to array LUNs

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

Requirements for redundant setup of components in a path

ONTAP systems must connect to the storage array through a redundant Fibre Channel (FC) network. Two FC networks 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 ONTAP systems.

ONTAP systems redundancy requirements

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

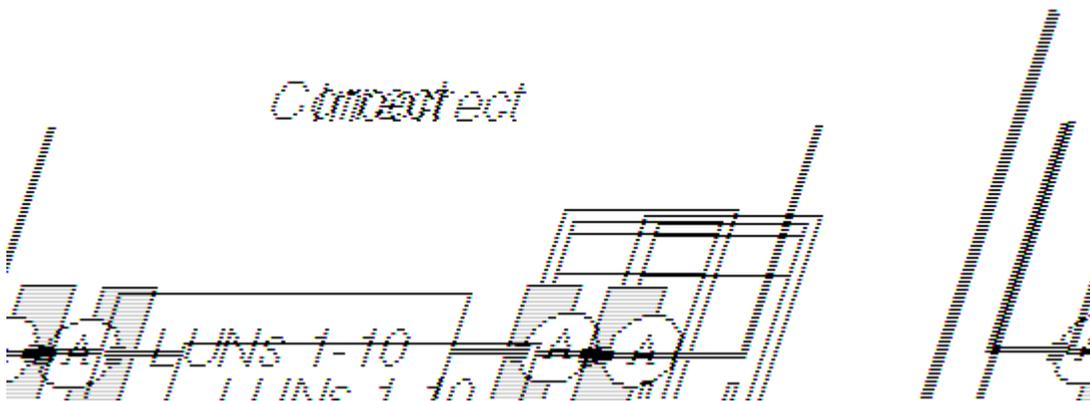
FC switch redundancy requirements

- You must use redundant 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, so as to avoid 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 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 ONTAP 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

ONTAP supports either four or two paths to array LUNs.

ONTAP expects and requires that a storage array provide access to a specific array LUN on at least two redundant storage array ports; that is, through a minimum two redundant paths.

Be sure that the ports on the storage array that you select to access a given LUN are from different components, so as to avoid 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.

Advantages of four paths to an array LUN

When planning the number of paths to an array LUN for ONTAP, 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.

Multiple paths to an array LUN can be used to distribute load

I/O requests for a given LUN can be distributed over all the available optimized paths to the LUN. This is unlike in previous releases, where in spite of multiple paths being available, I/O requests for a given LUN were sent over a single active optimized path only.

Distribution of I/O requests for a given LUN over multiple paths results in the following benefits:

- Improved efficiency due to maximum utilization of all available and optimized paths
- Improved performance due to load balancing over multiple paths

For example, on an active-active array, I/O requests for a given LUN can be distributed over all the four target ports available for that LUN. For an asymmetric active-active array LUN, I/O requests can be distributed over all the optimized paths for a given LUN.

Commands to view load balancing over multiple paths for a given LUN

You can run the following commands to view the distribution of load for a given LUN across multiple paths:

- `storage disk show -disk <LUN name>` displays the distribution of I/O load across the available paths for a given array LUN.
- `storage path show-by-initiator -array-name <array name>` displays the distribution of I/O load across all the initiator ports on the ONTAP system that is connected to a given storage array.
- `storage path show -by-target -array-name <array name>` displays the distribution of I/O load across all the target ports on a given storage array.

Benefits of using multiple LUN groups

You can use multiple LUN groups in your storage configuration for additional capacity and for improving the system performance by spreading the workload across more target ports.

A *LUN group* is set of logical devices on the storage array that an ONTAP 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. ONTAP refers to this set of devices as a *LUN group*.

Benefits 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 required capacity might exceed what a single LUN group can provide. Therefore, using multiple LUN groups can be beneficial.

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



Using multiple LUN groups is not supported for all storage arrays. See the *Interoperability Matrix* to determine whether a configuration using multiple LUN groups is supported for your storage array.

Related information

[NetApp Interoperability Matrix Tool](#)

Implementation requirements for a multiple LUN group configuration

You can improve the system performance by implementing a multiple LUN group configuration in your storage environment. Certain setup tasks must be performed on the ONTAP systems and the storage arrays for implementing this configuration.

A multiple LUN group configuration is supported for most storage arrays. See the *Interoperability Matrix* to confirm that this configuration is supported with your particular storage array.

You must work with the storage array administrator to set up the following on the *storage array* for a multiple LUN group configuration:

- Use as many ports as possible, to provide access to the array LUNs allocated for the ONTAP system.
- Use host groups (or your vendor's equivalent) to define which array LUN groups are presented to each FC initiator port on an ONTAP system.

You can set up the following on the ONTAP system to implement a multiple LUN group configuration:

- Use one FC initiator port pair for each array LUN group.

Each FC initiator port pair accesses a different LUN group on the storage array through redundant paths.

- Create a large aggregate in the ONTAP configuration, and add array LUNs from multiple RAID groups (parity groups) 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.

You must set up the following on the *switch* to implement a multiple LUN group configuration:

- Configure switch zoning to define which target ports the FC initiator ports on the ONTAP system should use to access each array LUN group.

Related information

[NetApp Interoperability Matrix Tool](#)

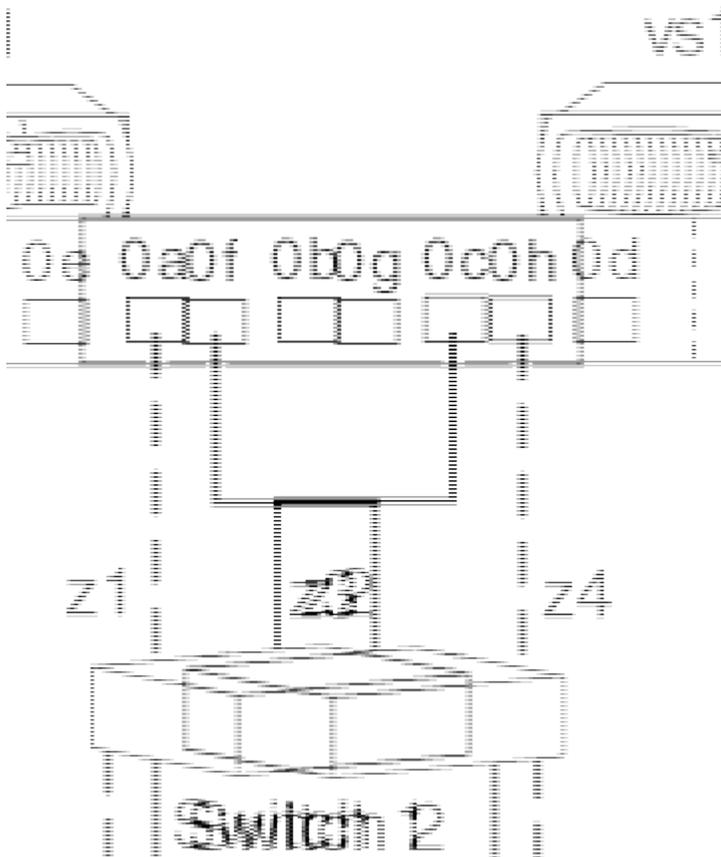
Example of a configuration with multiple LUN groups

You can use a multiple LUN group configuration to improve system performance by spreading the workload over multiple target ports.

A multiple LUN group configuration is supported for most storage arrays. See the *Interoperability Matrix* to confirm this configuration is supported with your storage array.

The following illustration shows how one FC initiator port pair (0c and 0f) on an ONTAP system 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.

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 FC initiator port pairs access different

groups of LUNs on the storage array. The ONTAP 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 ONTAP systems.



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	FC initiator port on the ONTAP system	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

Related information

[NetApp Interoperability Matrix Tool](#)

Array LUN name format

The name assigned to an array LUN has a new format to ensure that the name is unique across a cluster.

The array LUN name consists of two components and looks like the following:

<array_prefix>.<offset>, for example EMC-1.1.

- The `array_prefix` is a unique prefix that ONTAP assigns by default to each storage array.

This field is composed of <array_name-array_instance> (EMC-1 in this case).

`array_name` can be denoted by the first three letters of the name of the vendor.

If there is more than one array from the same vendor, the value of `array_instance` proceeds in an

ascending order.

- The offset is the ascending virtual disk number that ONTAP assigns to each LUN. It is independent of the LUN ID of the host.

You can modify the `<array_prefix>` field by using the `storage array modify -name -prefix` command.

Pre-cluster array LUN name format

Before a node joins to a cluster or when the system is in the maintenance mode, the array LUN name follows a format used before Data ONTAP 8.3, the *pre-cluster* format.

In this format, the array LUN name is a path-based name that includes the devices in the path between the ONTAP 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 an ONTAP system that supports array LUNs, each array LUN can have multiple names because there are multiple paths to each LUN.

Array LUN name format for ONTAP systems

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

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

How the array LUN name is displayed in ONTAP

A cluster-wide unique name is assigned to each array LUN, even though there are multiple paths to the array LUN. This is unlike in earlier releases, where each array LUN had multiple names based on the path accessing the LUN at a given time.

You can view the old names mapped to each array LUN by running the `storage disk show --disk <disk name> -fields diskpathnames` command.

Output example for the storage disk show --disk <disk name> -fields diskpathnames command

```

vgv3270f47ab::*> storage disk show -type LUN
          Usable          Disk      Container  Container
Disk      Size Shelf Bay Type      Type      Name      Owner
-----
-----
EMC-1.7          8.66GB      -   - LUN      spare     Pool0
vgv3270f47a
EMC-1.8          8.66GB      -   - LUN      spare     Pool0
vgv3270f47a
EMC-1.9          8.66GB      -   - LUN      spare     Pool0
vgv3270f47a
EMC-1.10         8.66GB      -   - LUN      spare     Pool0
vgv3270f47a

vgv3270f47ab::*> storage disk show -disk EMC-1.10 -fields diskpathnames
disk      diskpathnames
-----
-----
EMC-1.10
vgv3270f47a:vgbr300s181:5.126L9,vgv3270f47a:vgbr300s139:5.126L9,vgv3270f47
b:vgbr300s181:5.126L9,vgv3270f47b:vgbr300s139:5.126L9

```

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 ONTAP releases.



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 ONTAP system:

When validating your installation, you can check your command output against the example below to verify that the number of LUN groups is what you intended and that there are redundant paths.

Example

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`.)

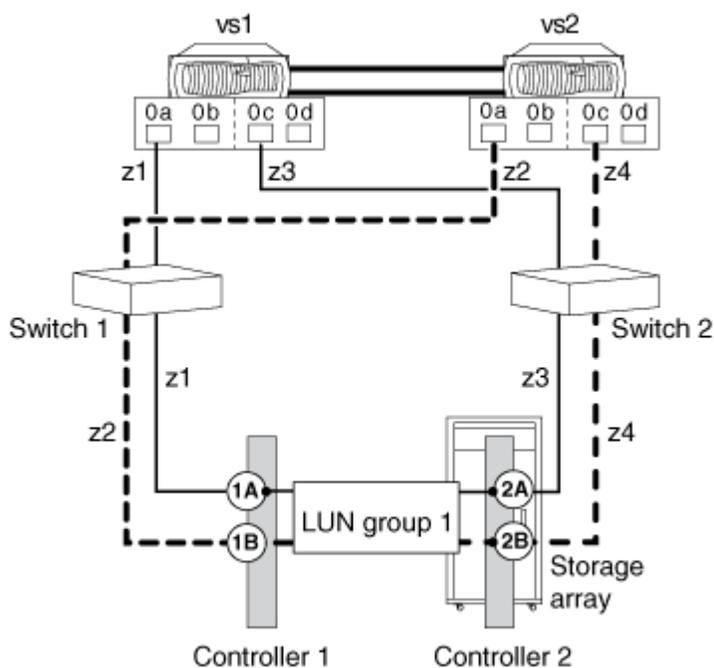
```
vs1::> storage array config show
```

Node	LUN Group	LUN Count	Array Name	Array Target Port	Initiator
vs1	0	50	DGC_RAID5_1	201A00a0b80fee04 202A00a0b80fee04	0a 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 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 node can see an array LUN, end-to-end, through only two paths. Zoning is configured so that each FC initiator port on a node can access only a single target array port.

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 Group	LUN Count	Array Name	Array Target Port	Initiator
vs1	1	10	DGC_RAID5_1	50050763030301241A 50050763031301242A	0a 0c
vs2	1	10	DGC_RAID5_1	50050763030881241B 50050763031881242B	0a 0c

```
4 entries were displayed.
```

Valid pathing: eight-port array LUN group configuration

You can use an eight-port LUN group configuration to connect storage arrays with ONTAP systems in large clustered deployments that require greater path redundancy and load balancing than possible with fewer ports per LUN group.

You can set up this configuration with the back-end connections crossed or uncrossed.

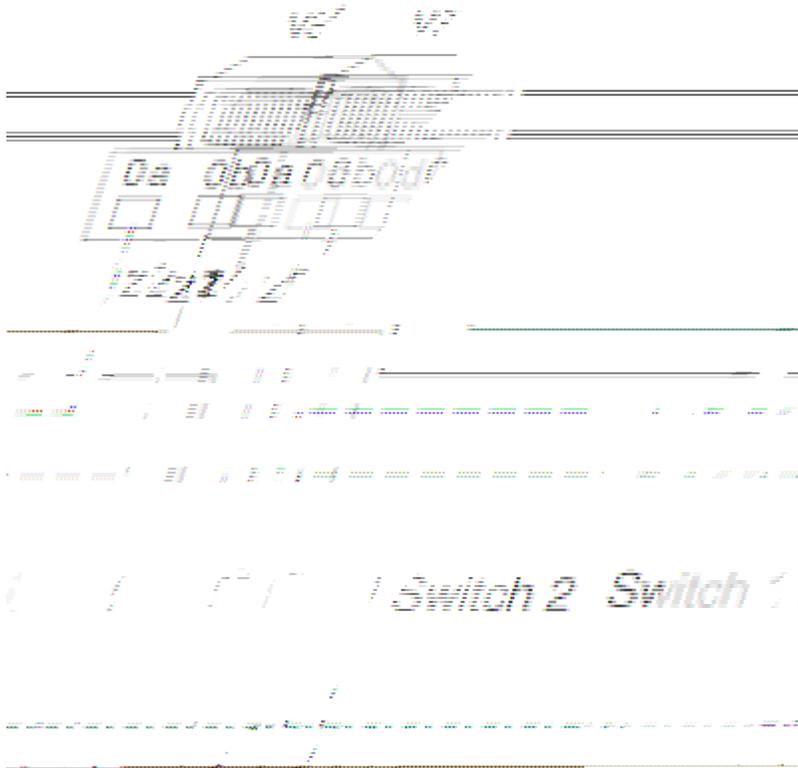
Variation in which back-end connections are crossed

In a configuration with the back-end connections crossed, the FC connections from the same storage array controller go to both fabric switches (redundant).

This connection scheme makes better use of both switch ports and storage array ports than if the back-end connections are not crossed, which reduces the impact of a switch or storage array controller failure.

For storage arrays with only two controllers, a crossed eight-port LUN group configuration is preferred over an eight-port array LUN group configuration that is not crossed.

You can only cross the eight-port array LUN group when there are dedicated paths from each node (one FC initiator-to-one-target-port zoning per path).



In this illustration of crossed back-end connections, note how the ONTAP systems are attached to the switches and to the storage array. Vs1 uses switch 1 when attaching to the storage array Controller 1 port 1A and Controller 2 port 2C, and uses switch 2 when attaching to storage array Controller 2 ports 2A and Controller 1 port 1C.

The following table summarizes the zoning for an eight-port array LUN group with crossed back-end connections. Single-initiator zoning is the recommended zoning strategy.

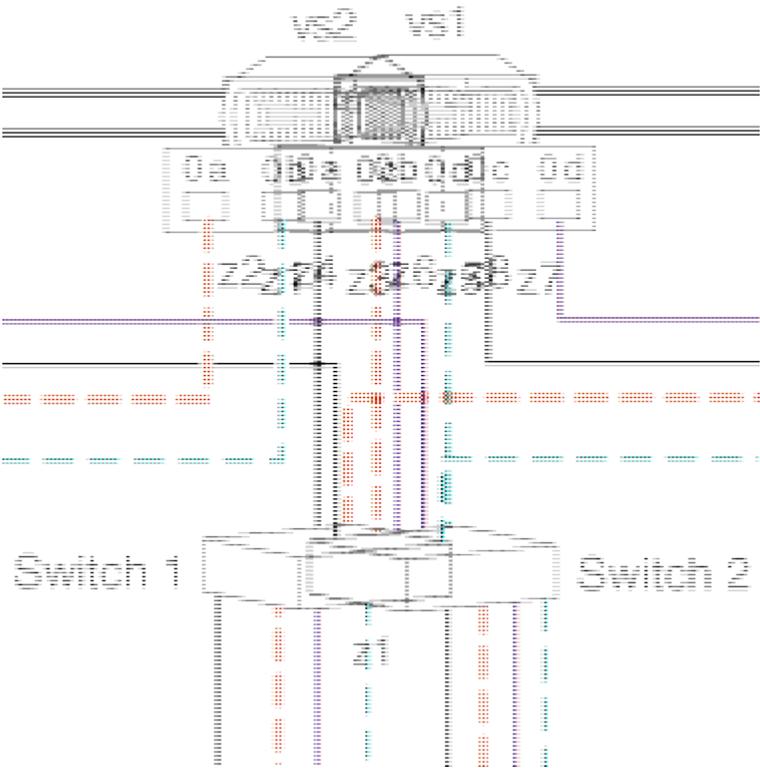
Zone	FC initiator port on the ONTAP system	Storage array
Switch 1		
z1	vs1, Port 0a	Controller 1, Port 1A
z2	vs2, Port 0a	Controller 1, Port 1B
z3	vs1, Port 0b	Controller 2, Port 2C
z4	vs2, Port 0b	Controller 2, Port 2D
Switch 2		
z5	vs1, Port 0c	Controller 2, Port 2A
z6	vs2, Port 0c	Controller 2, Port 2B

Zone	FC initiator port on the ONTAP system	Storage array
z7	vs1, Port 0d	Controller 1, Port 1C
z8	vs2, Port 0d	Controller 1, Port 1D

Variation in which back-end connections are *not* crossed

In a configuration in which the back-end connections are not crossed, the FC connections from the same storage array controller go to only one fabric switch.

The following illustration shows pathing in a configuration with an eight-port array LUN group in which the back-end connections are not crossed:



The following table summarizes the zoning for an eight-port array LUN group when the back-end connections are not crossed. Single-initiator zoning is the recommended zoning strategy.

Zone	FC initiator port on the ONTAP system	Storage array
Switch 1		
z1	vs1, Port 0a	Controller 1, Port 1A
z2	vs2, Port 0a	Controller 1, Port 1B
z3	vs1, Port 0b	Controller 1, Port 1C

Zone	FC initiator port on the ONTAP system	Storage array
z4	vs2, Port 0b	Controller 1, Port 1D
Switch 2		
z5	vs1, Port 0c	Controller 2, Port 2A
z6	vs2, Port 0c	Controller 2, Port 2B
z7	vs1, Port 0d	Controller 2, Port 2C
z8	vs2, Port 0d	Controller 2, Port 2D

Considerations for maximum number of array LUNs per FC initiator

When setting up a configuration with an eight-port array LUN group, you cannot exceed the number of array LUNs that ONTAP supports per FC initiator port.

Planning a port-to-port connectivity scheme

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

Requirements for using FC initiator ports

If you want to use FC initiator ports in an ONTAP configuration with array LUNs, you must follow specific requirements for port pair redundancy, port settings for HBAs, connecting to target ports, array LUN limits, and connecting to different storage and devices.

For this feature...	The requirements are...
Port pair redundancy	You must use redundant FC initiator port pairs while connecting an ONTAP system to array LUNs.
Port setting for HBAs	You must set all the HBAs used to access disks or array LUNs as <i>initiator</i> ports.
Connecting to storage array target ports	You can connect the same FC initiator port to multiple target ports on the storage array. You can connect a maximum of two FC initiator ports to a single target port.

For this feature...	The requirements are...
Array LUN limits	There are limits on how many array LUNs can be visible over an FC initiator port. These limits vary with ONTAP releases.
Connecting to different storage and devices	<p>You must use a separate FC initiator port to connect the ONTAP system to each of the following:</p> <ul style="list-style-type: none"> • Disk shelves • Array LUNs • Tape devices <p>If the ONTAP system does not have enough internal ports for your requirement, you must order an additional HBA.</p>

How FC initiator ports are labeled

All FC initiator ports on ONTAP 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.

Configuring FC ports as initiators

You can configure individual FC ports as initiators on ONTAP systems. The initiator mode enables the ports to connect to storage arrays.

Steps

1. **Optional:** If the adapter port already has LIFs configured, delete all the LIFs by using the `network interface delete` command.

If the LIF is in a port set, you must remove the LIF from port set before you can delete the LIF.

The following example shows how you can delete LIFs on an SVM vs3: `network interface delete -vserver vs3 -lif lif2,lif0`

2. Take the port offline by using the `network fcp adapter modify` command.

The following example shows how you can take the port 0c offline for the node sysnode: `network fcp adapter modify -node sysnode1 -adapter 0c -state down`

3. Use the `system hardware unified-connect modify` command to change the offline port from target to initiator.

The following example shows how you can change the port type for 0c from target to initiator: `system node hardware unified-connect modify -node sysnode1 -adapter 0c -type initiator`

4. Reboot the node hosting the adapter you changed.
5. Use the `system hardware unified-connect show` command to verify that the FC ports are configured correctly for your configuration.

The following example shows the change in port type for 0c:

```
system node hardware unified-connect show -node sysnode1
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Status
sysnode1	0a	fc	target	-	-	online
sysnode1	0b	fc	target	-	-	online
sysnode1	0c	fc	initiator	-	-	offline
sysnode1	0d	fc	target	-	-	online

6. Use the `storage enable adapter` command to bring the offline port back online.

The following example shows how you can make the port 0c online: `node run -node sysnode1 -command storage enable adapter -e 0c`

Rules for sharing an FC initiator port with multiple target ports

You can connect an FC initiator port on an ONTAP system to a maximum of four target ports on *separate* storage arrays. Sharing an FC initiator port with multiple targets is useful when you want to minimize the number of FC initiator ports used.

You can also connect an FC initiator port to a maximum of four target ports on the *same* storage array if the array can present different sets of logical devices to the FC initiator based on the target port being accessed.

There are limits on how many array LUNs can be visible over an FC initiator port. These limits vary according to release.

Rules when the FC initiator port is connected to multiple target ports on *separate* storage arrays

The rules for this configuration are as follows:

- All storage arrays must be from 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.

- Connecting a single FC initiator port to multiple target ports is supported for MetroCluster configurations.
- A single FC initiator port can connect to a maximum of four target ports on multiple storage arrays.
- You must have each FC initiator-target port pair in a separate zone (1:1), even if the same FC initiator port is accessing multiple target ports.

Rules when the FC initiator port is 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 FC initiator port 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 following are the rules for this configuration:

- A single FC initiator port can connect to up to four target ports on a storage array.
- Connecting a single FC initiator port to multiple target ports is supported for MetroCluster configurations.
- You must have each FC initiator-target port pair in a separate zone (1:1), even if the same FC initiator is accessing multiple target ports.

Related information

[FlexArray virtualization implementation for third-party storage](#)

[FlexArray virtualization implementation for NetApp E-Series storage](#)

[Fabric-attached MetroCluster installation and configuration](#)

Example configuration: shared FC initiator ports

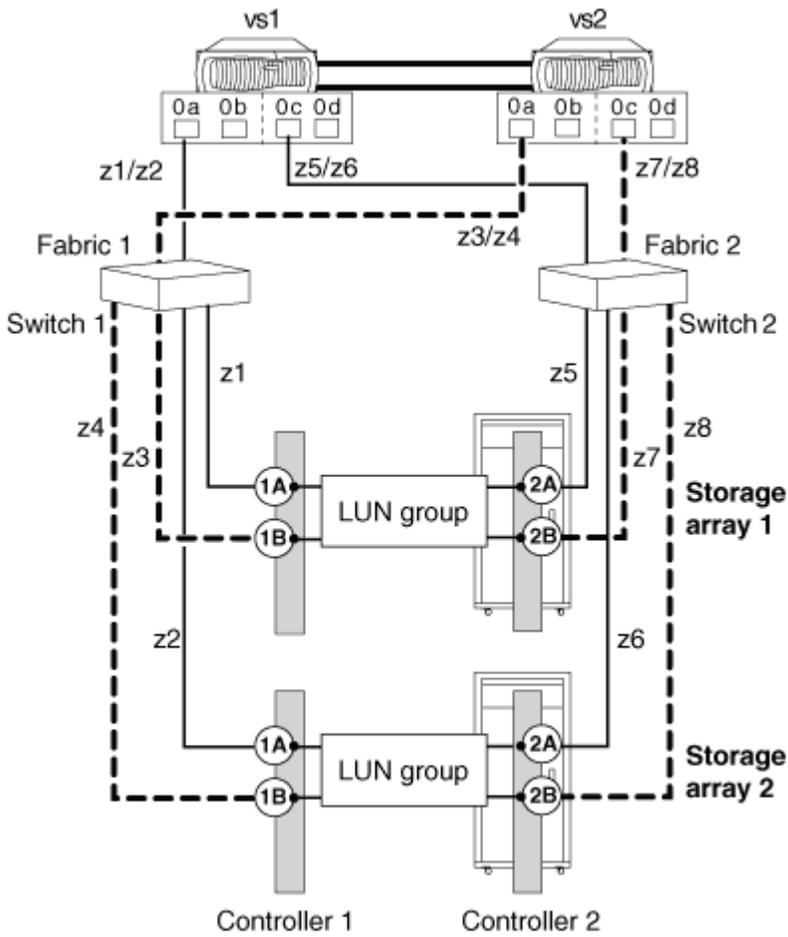
One FC initiator port on an ONTAP system can be connected to up to four target ports on separate storage arrays or, for some storage arrays, to four target ports on the same storage arrays.

The best practice zoning recommendation is to have 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	ONTAP system and FC initiator port	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	ONTAP system and FC initiator port	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

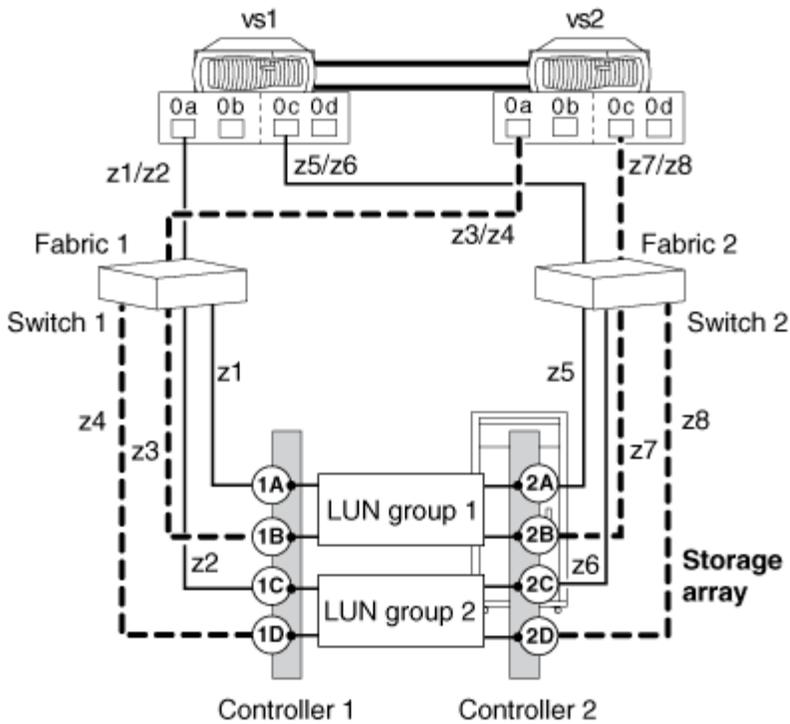
Shared FC initiator port attached to multiple targets ports on the *same* storage array

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 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	ONTAP system and FC initiator 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

Rules for sharing a target port with multiple FC initiator ports

Connecting a maximum of two ONTAP FC initiator ports to a single target port on the storage array is supported. Each target port is zoned into two FC initiator ports, one from each cluster node. Sharing a target port connected to ONTAP systems with any other host is not supported.

Sharing a target port with multiple initiators helps in optimizing the use of storage array ports for connectivity with ONTAP systems.

The rules for this configuration are as follows:

- When the ONTAP systems are in an HA pair, each node can share at the most one FC initiator port with the same target port.
- All storage arrays must be from the same vendor and model family.
- Connecting a single target port to multiple FC initiator ports is supported for MetroCluster configurations.
- The best practice for zoning is to have each FC initiator-target port pair in a separate zone (1:1).

Related information

[Fabric-attached MetroCluster installation and configuration](#)

Example configuration: shared target ports

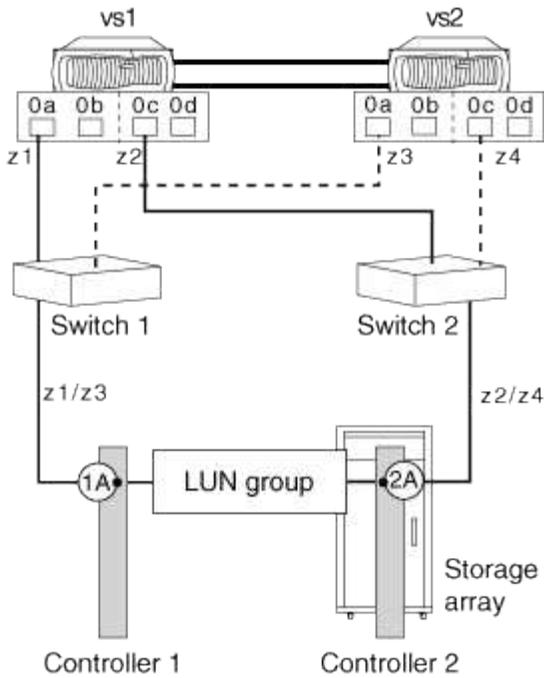
A maximum of two ONTAP FC initiator ports can be connected to a single target port on the storage array.

The best practice for zoning is to have each FC initiator-target port pair in a separate zone (1:1).

Shared target port attached to multiple initiator ports

The following illustration shows connections and zoning for sharing a target port with multiple FC initiator ports on different ONTAP systems.

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 target port sharing two FC initiator ports from controllers in an HA pair:

Zone	ONTAP system and FC initiator port	Storage array
Switch 1		
z1	vs1:0a	Controller 1: Port 1A
z3	vs2:0a	
Switch 2		
z2	vs1:0c	Controller 2: Port 2A
z4	vs2:0c	

Checking the number of array LUNs visible on an FC initiator port

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

Steps

1. Check the visible number on the FC initiator port: `storage array config show -initiator initiator_number`

`storage array config show -initiator 0a`

- If there is more than one array LUN group for a node, add 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.

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 must 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 (DGC_RAID5_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 Group	LUN Count	Array Name	Array Target Port	Initiator
vgv3070f51-01	1	24	HP	50014380025d1508	0a
	2	1	DGC_RAID5_1	200600a0b819e16f	0a
vgv3070f51-02	0	24	HP	50014380025d1508	0a
	2	1	DGC_RAID5_1	200600a0b819e16f	0a

Requirements for connecting to storage arrays

When planning how to connect your ONTAP system to a storage array, your port-to-port connectivity plan must 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 ONTAP 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, so as to avoid 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 ONTAP release.
- If you want to set up a configuration where an FC initiator port is shared with multiple target ports or a target port is shared with multiple FC initiator ports, you must follow the appropriate rules.
- If your storage array supports fewer LUNs per host group per port than the number of LUNs that the ONTAP systems is using, then you need to add additional cables between the ONTAP system and the storage array.

Guidelines for connecting a V-Series system 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.

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 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	NetApp Support
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.
Connecting a disk shelf to a stand-alone V-Series system	The Installation and Setup Instructions for your platform. <ul style="list-style-type: none">• Installation and Setup Instructions 32xx Systems• Installation and Setup Instructions 62xx Systems
Connecting an HA pair to a disk shelf	High-availability configuration
Adding a disk shelf	The appropriate guide for your disk shelf type.
Moving a disk shelf	The appropriate guide for your disk shelf type.
Disk management	Disk and aggregate management

FC initiator ports required for connection to native disks

You must connect a V-Series system to native disk shelves through the FC initiator ports. The number of initiator ports required for the connection depends on whether the V-Series system is a stand-alone system or in an HA pair.

The following table lists the required number of FC initiator ports for connecting a V-Series system to native disk shelves, depending on the system configuration:

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
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 per controller for each loop in the configuration

Example of non-optimized usage of a storage array target port

When the number of I/O requests queued on a given storage array target port exceeds the number of requests that the port can handle, it results in non-optimized usage of the target port.

You can detect such non-optimized usage of a given target port in the output display of the `storage array show` command.

Output example displaying errors on detecting non-optimized usage of a target port

The following example shows the error returned by the `storage array show` command on detecting non-optimized usage of a given target port:

```
vgv3070f50ab::> storage array show -name HP_HSV450_2
```

```
      Name: HP_HSV450_2
      Prefix:
      Vendor: HP
      Model: HSV450
      options:
      Serial Number: 50014380025d1500
      Optimization Policy: iALUA
      Affinity: AAA
```

Errors:

```
Warning: HP_HSV450_2 Detected non optimized usage of a target port. WWPN:
2703750270235, average service time: 215ms, average latency: 30ms
```

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 ONTAP 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 array LUN aggregates

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



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.

How the checksum type is determined for array LUN aggregates

Each 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.

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.

Checksum type considerations for adding spare array LUNs to aggregates

You must account for certain considerations related to checksum types when adding spare array LUNs to aggregates. For example, if you plan to add a spare array LUN to an aggregate by specifying its name, you must ensure that array LUN and the aggregate have the same checksum type.

The following are some checksum type considerations for adding spare array LUNs to aggregates:

- You cannot mix array LUNs of different checksum types in an array LUN aggregate.
- If you specify a number of spare array LUNs to be added to an aggregate, ONTAP by default selects the array LUNs of the same checksum type as the aggregate.
- Array LUNs of the zoned checksum type when added to an existing zoned checksum aggregate continue to be zoned checksum array LUNs.
- Zoned checksum spare array LUNs added to an Advanced Zone Checksum (AZCS) type aggregate use the AZCS checksum scheme.



You can check the checksum type of the spare array LUNs by using the `storage disk show` command. For more information about the command, see the man pages.

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.

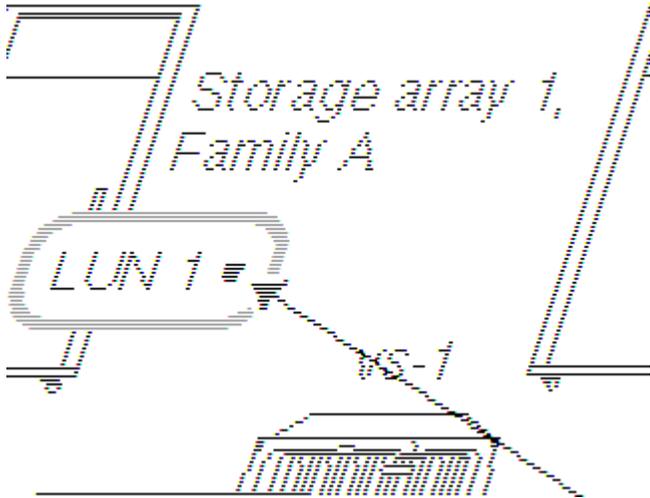
The following examples show some options for laying out array LUNs in aggregates when the storage arrays behind an ONTAP system are in *the same vendor family*.



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

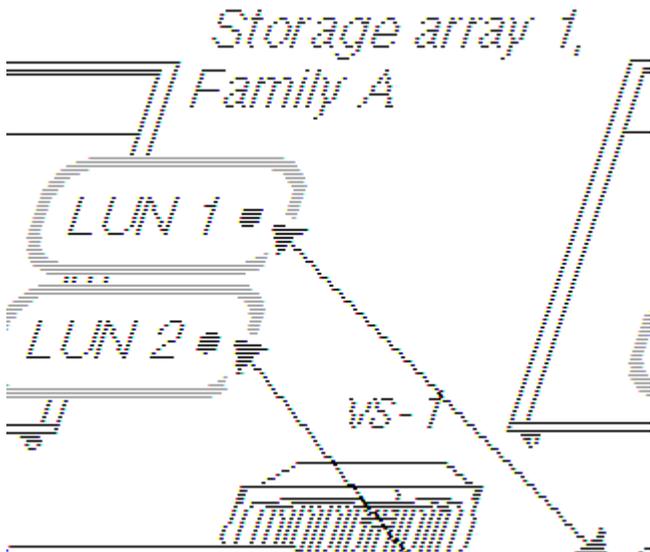
Example 1: Add LUNs from all storage arrays to a single aggregate

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:





This example is not supported if you have storage arrays of the same model and one has Fibre Channel drives and the other storage array has SATA drives. In that case, those storage arrays are not considered to be in the same family.

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.

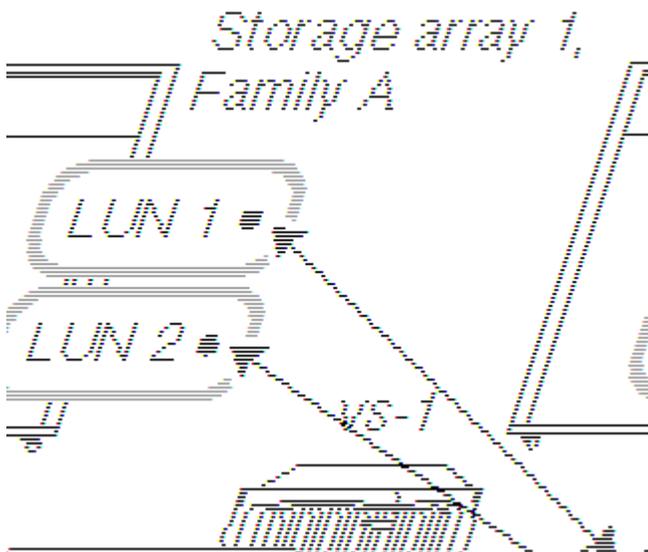


When you create your aggregate, be sure that you explicitly specify the IDs of the array LUNs that you want to add to the aggregate. Do not use the parameters for specifying the number and size of array LUNs to be picked up because the system might automatically pick up LUNs from a different family or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

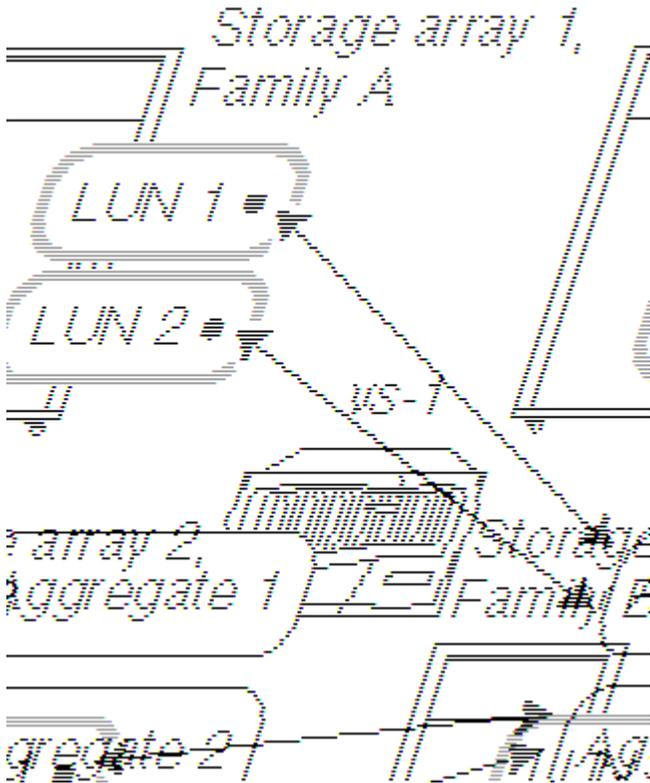
Example 1: LUNs from the two storage arrays are in different aggregates

In this example, some LUNs for ONTAP are from Storage array 1, Family A, while the other LUNs for 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.



Preparing a storage array for use with ONTAP systems

Before you can begin setting up ONTAP systems in a configuration with array LUNs, the storage array administrator must prepare the storage for use with ONTAP.

What you'll need

The storage arrays, firmware, and switches that you plan to use in the configuration must be supported by the specific ONTAP version.

- [NetApp Interoperability](#)

In the IMT, you can use the Storage Solution field to select your MetroCluster solution. You use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

- [NetApp Hardware Universe](#)

You must coordinate with the storage array administrator to perform this task on the storage array.

Steps

1. Create at least four LUNs on the storage array for the ONTAP systems to use.

Each node in the HA pair requires an array LUN for the root volume and an array LUN for core dumps.

2. Configure parameters on the storage array that are required to work with ONTAP.
 - [FlexArray virtualization implementation for third-party storage](#)
 - [FlexArray virtualization implementation for NetApp E-Series storage](#)

Connecting an ONTAP system to a storage array

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

What you'll need

- You must have identified the onboard ports and expansion adapter ports for your ONTAP system for connecting to the storage array.
- You must have located the ports on the storage array that you want to use to connect to the ONTAP system.

This procedure describes how to connect an ONTAP system to a storage array through two paths, with each initiator dedicated to one target port.

Steps

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

2. 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...
<ul style="list-style-type: none"> a. Connect Switch 1 to the storage array controller 1, port 1A. b. Connect Switch 2 to the storage array controller 2, port 2A. 	<ul style="list-style-type: none"> a. Connect Switch 1 to the storage array controller 1, port 1A. b. Connect Switch 2 to the storage array controller 2, port 2A. c. Connect Switch 1 to the storage array controller 1, port 1B. d. Connect Switch 2 to the storage array controller 2, port 2B.

The following illustration shows the connections for an HA pair.

3. **Optional:** Connect the ONTAP system to a tape backup device through a separate FC initiator port or SCSI tape adapter.
4. Verify that the storage array is configured and connected properly, and that it is powered on.

Your configured and connected storage array must be powered on before you power on your ONTAP system. See your storage array documentation for how to power on the storage array.

5. If your deployment includes switches, ensure that all switch IDs are set, and then turn them on at 10-minute intervals.
6. **Optional:** If applicable, turn on any tape backup devices.
7. Power on the ONTAP system and perform the initial network setup and configuration.
8. If the storage array does not automatically discover the ONTAP system WWNs after you connect the ONTAP system to the storage array, you must obtain the WWNs manually.

You must continue with the appropriate setup of ONTAP management features to work with the storage arrays.

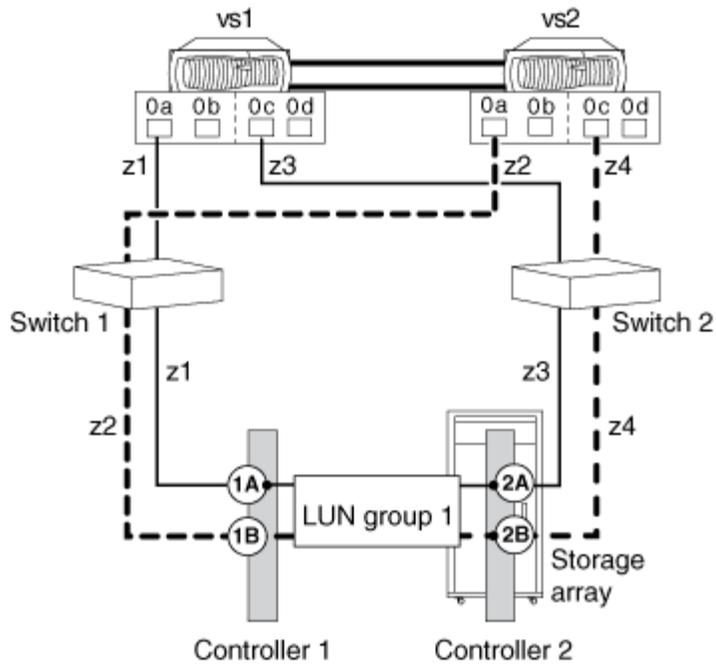
Setting up the switches

Switch configuration is typically done by the storage or SAN administrator. The switches have to be zoned so that the ONTAP systems and the storage arrays can see each other. You must use single-initiator zoning as the zoning strategy.

Steps

1. Log in to the storage array and obtain the WWPNS of the FC adapters of the storage array.
2. Use the Fibre Channel switch commands to zone each switch so that the storage array and the ONTAP system see each other's WWPNS.

Consider the following example of ONTAP systems in an HA pair:



In the example configuration, the zones are as follows:

Zone	ONTAP 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

Configuring LUN security

The storage array administrator must configure the storage array and create host groups so that other hosts cannot access the array LUNs intended for use by ONTAP.

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.

Refer to the storage array documentation for information about setting up the LUN security.

You should also see [FlexArray virtualization implementation for third-party storage](#) for any additional details regarding LUN security for your vendor's storage arrays. Some storage arrays must be dedicated for use by ONTAP systems.

2. Create host groups, or the equivalent, for the ONTAP 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 information

Setting up ONTAP to work with array LUNs

The process of setting up an ONTAP system to work with array LUNs differs on the basis of whether ONTAP systems are ordered with or without disks.

If an ONTAP system is ordered with disk shelves, ONTAP software is installed by the factory. In such a configuration, you do not need to create the root volume and install licenses and ONTAP software.

You can install ONTAP software on systems ordered without disk shelves. On these systems, you must first install ONTAP and then set up a cluster.

Configuring ONTAP on a system that uses only array LUNs

If you want to configure ONTAP for use with array LUNs, you must configure the root aggregate and root volume, reserve space for diagnostics and recovery operations, and set up the cluster.

What you'll need

- The ONTAP system must be connected to the storage array.
- The storage array administrator must have created LUNs and presented them to ONTAP.
- The storage array administrator must have configured the LUN security.

You must configure each node that you want to use with array LUNs. If the node is in an HA pair, then you must complete the configuration process on one node before proceeding with the configuration on the partner node.

Steps

1. Power on the primary node and interrupt the boot process by pressing Ctrl-C when you see the following message on the console: `Press CTRL-C for special boot menu.`
2. Select option 4 (`Clean configuration and initialize all disks`) on the boot menu.

The list of array LUNs made available to ONTAP is displayed. In addition, the array LUN size required for root volume creation is also specified. The size required for root volume creation differs from one ONTAP system to another.

- If no array LUNs were previously assigned, ONTAP detects and displays the available array LUNs, as shown in the following example:

```

mcc8040-ams1::> disk show NET-1.6 -instance
    Disk: NET-1.6
    Container Type: aggregate
    Owner/Home: mcc8040-ams1-01 / mcc8040-ams1-01
    DR Home: -
    Stack ID/Shelf/Bay: - / - / -
    LUN: 0
    Array: NETAPP_INF_1
    Vendor: NETAPP
    Model: INF-01-00
    Serial Number: 60080E50004317B4000003B158E35974
    UID:
60080E50:004317B4:000003B1:58E35974:00000000:00000000:00000000:000000
00:00000000:00000000
    BPS: 512
    Physical Size: 87.50GB
    Position: data
    Checksum Compatibility: block
    Aggregate: eseries
    Plex: plex0

Paths:

          LUN  Initiator Side          Target
Side                               Link
Controller      Initiator    ID  Switch Port          Switch
Port            Acc Use  Target Port          TPGN  Speed
I/O KB/s            IOPS
-----
-----
-----
mcc8040-ams1-01    2c                0  mccb6505-ams1:16    mccb6505-
ams1:18    AO  INU  20330080e54317b4    1  4 Gb/S
0            0
mcc8040-ams1-01    2a                0  mccb6505-ams1:17    mccb6505-
ams1:19    ANO RDY  20320080e54317b4    0  4 Gb/S
0            0

Errors:
-
```

- If array LUNs were previously assigned, for example, through the maintenance mode, they are either marked `local` or `partner` in the list of the available array LUNs, depending on whether the array LUNs were selected from the node on which you are installing ONTAP or its HA partner:

Checksum	Serial Number				
0	vgci9710s202:2-24.0L19	RAID5	DGC	217.3 GB	
Block	6006016083402B0048E576D7				
1	vgbr6510s203:30.126L20	RAID5	DGC	217.3 GB	
Block	6006016083402B0049E576D7				
2	vgci9710s202:2-24.0L21	RAID5	DGC	217.3 GB	
Block	6006016083402B004AE576D7				
3	vgbr6510s203:30.126L22	RAID5	DGC	405.4 GB	local
Block	6006016083402B004BE576D7				
4	vgci9710s202:2-24.0L23	RAID5	DGC	217.3 GB	
Block	6006016083402B004CE576D7				
5	vgbr6510s203:30.126L24	RAID5	DGC	217.3 GB	
Block	6006016083402B004DE576D7				
6	vgbr6510s203:30.126L25	RAID5	DGC	423.5 GB	local
Block	6006016083402B003CF93694				
7	vgci9710s202:2-24.0L26	RAID5	DGC	423.5 GB	
Block	6006016083402B003DF93694				

3. Select the index number corresponding to the array LUN you want to assign as the root volume.

The array LUN must be of sufficient size to create the root volume.

The array LUN selected for root volume creation is marked `local (root)`.

In the following example, the array LUN with index number 3 is marked for root volume creation:

The root volume will be created on switch 0:5.183L33.

ONTAP requires that 11.0 GB of space be reserved for use in diagnostic and recovery operations. Select one array LUN to be used as spare for diagnostic and recovery operations.

Index	Array LUN Name	Model	Vendor	Size	Owner
Checksum	Serial Number				
0	switch0:5.183L1	SYMMETRIX	EMC	266.1 GB	
Block	600604803436313734316631				
1	switch0:5.183L3	SYMMETRIX	EMC	266.1 GB	
Block	600604803436316333353837				
2	switch0:5.183L31	SYMMETRIX	EMC	266.1 GB	
Block	600604803436313237643666				
3	switch0:5.183L33	SYMMETRIX	EMC	658.3 GB	local (root)
Block	600604803436316263613066				
4	switch0:7.183L0	SYMMETRIX	EMC	173.6 GB	
Block	600604803436313261356235				
5	switch0:7.183L2	SYMMETRIX	EMC	173.6 GB	
Block	600604803436313438396431				
6	switch0:7.183L4	SYMMETRIX	EMC	658.3 GB	
Block	600604803436313161663031				
7	switch0:7.183L30	SYMMETRIX	EMC	173.6 GB	
Block	600604803436316538353834				
8	switch0:7.183L32	SYMMETRIX	EMC	266.1 GB	
Block	600604803436313237353738				
9	switch0:7.183L34	SYMMETRIX	EMC	658.3 GB	
Block	600604803436313737333662				

4. Select the index number corresponding to the array LUN you want to assign for use in diagnostic and recovery options.

The array LUN must be of sufficient size for use in diagnostic and recovery options. If required, you can also select multiple array LUNs with a combined size greater than or equal to the specified size. To select multiple entries, you must enter the comma-separated values of all of the index numbers corresponding to the array LUNs you want to select for diagnostic and recovery options.

The following example shows a list of array LUNs selected for root volume creation and for diagnostic and recovery options:

```

Here is a list of the selected array LUNs
Index Array LUN Name      Model      Vendor      Size      Owner
Checksum Serial Number
-----
      2  switch0:5.183L31    SYMMETRIX  EMC        266.1 GB  local
Block      600604803436313237643666
      3  switch0:5.183L33    SYMMETRIX  EMC        658.3 GB  local    (root)
Block      600604803436316263613066
      4  switch0:7.183L0      SYMMETRIX  EMC        173.6 GB  local
Block      600604803436313261356235
      5  switch0:7.183L2      SYMMETRIX  EMC        173.6 GB  local
Block      600604803436313438396431
Do you want to continue (yes|no)?

```



Selecting “no” clears the LUN selection.

- Enter `y` when prompted by the system to continue with the installation process.

The root aggregate and the root volume are created and the rest of the installation process continues.

- Enter the required details to create the node management interface.

The following example shows the node management interface screen with a message confirming the creation of the node management interface:

```
Welcome to node setup.
```

```
You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,  
"back" - if you want to change previously answered questions, and  
"exit" or "quit" - if you want to quit the setup wizard.  
Any changes you made before quitting will be saved.
```

```
To accept a default or omit a question, do not enter a value.
```

```
Enter the node management interface port [e0M]:
```

```
Enter the node management interface IP address: 192.0.2.66
```

```
Enter the node management interface netmask: 255.255.255.192
```

```
Enter the node management interface default gateway: 192.0.2.7
```

```
A node management interface on port e0M with IP address 192.0.2.66 has  
been created.
```

```
This node has its management address assigned and is ready for cluster  
setup.
```

After configuring ONTAP on all of the nodes that you want to use with array LUNs, you should complete the cluster setup process.

Software setup

Installing the license for using array LUNs

The `V_StorageAttach` license must be installed on each ONTAP node that you want to use with array LUNs. It is *not* a single license for the cluster. Array LUNs cannot be used in aggregates until a license is installed.

What you'll need

- The cluster must be installed.
- You must have the license key for the `V_StorageAttach` license.

NetApp Support

You need not perform this procedure if the license key for the `V_StorageAttach` package is already installed. If the ONTAP system is ordered with disks, the factory typically installs the license package for you. Alternatively, many customers install all necessary licenses early in the installation process.

Steps

1. For each ONTAP node in the cluster for use with array LUNs, enter the following command on the node:
`system license add license key`

```

vgv3170f41a> license
Serial Number: nnnnnnnn
Owner: mysystemla
Package          Type      Description          Expiration
-----
V_StorageAttach license Virtual Attached Storage

```

2. Look at the output to confirm that the V_StorageAttach package is shown.

Assigning ownership of array LUNs

On an ONTAP system where the root volume resides on disk shelves, you must assign the ownership of array LUNs to a node before adding them to an aggregate for using them as storage.

What you'll need

- Back-end configuration testing (testing of the connectivity and configuration of devices behind the ONTAP systems) must be completed.
- Array LUNs that you want to assign must be presented to the ONTAP systems.

You can assign ownership of array LUNs that have the following characteristics:

- They are unowned.
- They have no storage array configuration errors, such as the following:
 - The array LUN is smaller than or larger than the size that ONTAP supports.
 - The LDEV is mapped on only one port.
 - The LDEV has inconsistent LUN IDs assigned to it.
 - The LUN is available on only one path.

ONTAP issues an error message if you try to assign ownership of an array LUN with back-end configuration errors that would interfere with the ONTAP system and the storage array operating together. You must fix such errors before you can proceed with array LUN assignment.

ONTAP alerts you if you try to assign an array LUN with a redundancy error: for example, all paths to this array LUN are connected to the same controller or only one path to the array LUN. You can fix a redundancy error before or after assigning ownership of the LUN.

Steps

1. Enter the following command to see the array LUNs that have not yet been assigned to a node:


```
storage disk show -container-type unassigned
```
2. Enter the following command to assign an array LUN to this node:


```
storage disk assign -disk arrayLUNname -owner nodename
```

If you want to fix a redundancy error after disk assignment instead of before, you must use the `-force` parameter with the `storage disk assign` command.

Commands for checking back-end configuration

Several 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 deployment with ONTAP systems using array LUNs. This command shows how storage arrays connect to the cluster. If ONTAP detects an issue that would prevent the ONTAP systems using array LUNs 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 show -name array_name</code>	<p>Displays information about all storage arrays visible to the cluster or about the storage array that you specify. If the number of array LUNs presented exceeds the capacity of the system, the Error text field shows you the number of LUNs that ONTAP could not discover. You should watch for this issue on low memory systems in particular.</p>
<code>storage 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 path resume</code>. Some storage arrays require ceasing I/O for a period of time to remove or move an array LUN.</p> <p> The <code>storage path quiesce</code> command cannot be used with IBM DS storage arrays.</p>

Command	Description
storage path resume	Allows I/O to start flowing again; this is the inverse of quiesce. The <code>storage 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. It is not always necessary to run this command after quiescing a path. For example, ONTAP can discover a newly mapped array LUN.
storage array show	Displays information about the storage arrays that are visible to the cluster, for example, name, vendor, model, and failover type.
storage disk show	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. Entering <code>storage disk show</code> with a disk name or an array LUN name as a parameter 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.
storage disk show -errors **	Entering <code>storage disk show</code> with the <code>-errors</code> parameter provides details of configuration errors at the disk and array LUN levels. Although the command output is similar to that of the <code>storage disk error show</code> command, <code>storage disk show -errors</code> provides additional options to filter the output by parameters such as node, cluster, and bay position.
storage disk error show	Provides details about back-end configuration errors at the disk and array LUN levels. Entering <code>storage disk error show</code> with an array LUN name as a parameter shows details about configuration errors related to the specified array LUN. You must fix these errors before configuring ONTAP to work with storage arrays.

Command	Description
<pre>storage errors show</pre>	<p>Provides details, at the array LUN level, about back-end configuration errors that prevent the ONTAP system and the storage array from operating together. You must fix errors identified by <code>storage errors show</code> before you configure ONTAP to work with storage arrays.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 20px;"> <p> The <code>storage errors show</code> command does not provide information about disk configuration errors. To view information about both disk and array LUN configuration errors, you can use either the <code>storage disk show -errors</code> command or the <code>storage disk error show</code> command.</p> </div>

Back-end configuration errors detected by ONTAP commands

The `storage errors show` command provides details, at the array LUN level, about common back-end configuration errors. You can also use the `storage disk error show` and the `storage disk show -errors` commands to view the errors.

You must fix the back-end configuration errors identified by these commands before you can configure ONTAP to work with array LUNs.



While the `storage errors show` command provides error details pertaining to only array LUNs, the `storage disk error show` and the `storage disk show -errors` commands can provide error details related to array LUNs as well as disks.

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.

List of back-end configuration errors

The `storage errors show`, `storage disk error show`, and `storage disk show -errors` commands can help identify back-end configuration errors such as the following:

- 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 ONTAP maximum LUN size.
- The array LUN does not meet the ONTAP minimum LUN size.
- The block size of an array LUN is invalid.

- An access LUN is presented to ONTAP.

Output examples displaying the back-end configuration errors

The `storage errors show` output is grouped by storage array (if there is more than one storage array behind the ONTAP 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 ONTAP requires two paths to an array LUN.



Four paths to an array LUN are supported for clustered configurations.

The following example shows the errors returned by the `storage errors show` command because of only a single path configured to an array LUN:

```
systemf47ab::*> storage errors show
DGC-1.51          onepath          DGC-1.51
(6006016044d03500ae553e55b09fe311): This array LUN is only available on
one path. Proper configuration requires two paths.
DGC-1.54          onepath          DGC-1.54
(6006016044d03500b4553e55b09fe311): This array LUN is only available on
one path. Proper configuration requires two paths.
DGC-1.55          onepath          DGC-1.55
(6006016044d03500b6553e55b09fe311): This array LUN is only available on
one path. Proper configuration requires two paths.
DGC-1.56          onepath          DGC-1.56
(6006016044d03500b8553e55b09fe311): This array LUN is only available on
one path. Proper configuration requires two paths.
DGC-1.57          onepath          DGC-1.57
(6006016044d03500ba553e55b09fe311): This array LUN is only available on
one path. Proper configuration requires two paths.
DGC-1.58          onepath          DGC-1.58
(6006016044d03500bc553e55b09fe311): This array LUN is only available on
one path. Proper configuration requires two paths.
```

The following example shows similar errors returned by the `storage disk show -errors` command:

```
systemf47ab::*> storage disk show -errors
DGC-1.2          onepath          DGC-1.2
(6006016044d03500e0720e22b59fe311): This array LUN is only available on
one path. Proper configuration requires two paths.
DGC-1.3          onepath          DGC-1.3
(6006016044d03500e2720e22b59fe311): This array LUN is only available on
one path. Proper configuration requires two paths.
DGC-1.4          onepath          DGC-1.4
(6006016044d03500e3720e22b59fe311): This array LUN is only available on
one path. Proper configuration requires two paths.
DGC-1.5          onepath          DGC-1.5
(6006016044d03500e4720e22b59fe311): This array LUN is only available on
one path. Proper configuration requires two paths.
DGC-1.6          onepath          DGC-1.6
(6006016044d03500e5720e22b59fe311): This array LUN is only available on
one path. Proper configuration requires two paths.
DGC-1.7          onepath          DGC-1.7
(6006016044d03500e6720e22b59fe311): This array LUN is only available on
one path. Proper configuration requires two paths.
```

Situations not identified by commands that check back-end configuration

There might be situations that you consider to be an issue but are not errors from the perspective of ONTAP because the situation does not prevent the system from operating. The ONTAP commands that check the back-end configuration do not identify configurations that do not prevent system operation.

Commands such as `storage errors show`, `storage disk error show`, and `storage disk show -errors` do 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 the intended configurations

For example, if you wanted to configure multiple LUN groups and only one was configured, ONTAP does not identify that as an error because a single LUN group is supported.

Verifying installation with storage arrays

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 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

ONTAP requires you to fix common back-end configuration errors that would prevent a storage array and an ONTAP system from operating together normally. Using the `storage array config show` command helps you determine whether there are back-end configuration errors.

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 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 Verifying if the back-end configuration matches the expected output section. (You do not need to continue with the next steps in this procedure.) Verifying if the back-end configuration matches the expected output

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](#) section to learn about the cause of each problem detected by `storage errors show` and how to fix it. [The storage errors show messages and their resolution](#)

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 the expected output.

The storage errors show messages and their resolution

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.

<code>storage errors show</code> 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`.	All paths to an array LUN are on the same storage array controller
NAME (Serial #), port WWPN1: LUN 1 occurs more than once. LUNs cannot be reused on the same array target port.	Duplicated LUN IDs on a target port
NAME (Serial #): This array LUN is an access control LUN. It is not supported and should be masked off or disabled`.	Access control LUN was presented
NAME (Serial #) This array LUN is configured with conflicting failover modes. Each path to this LUN must use the same mode.	Array LUNs are configured with conflicting failover modes
NAME (Serial #): This Array LUN is only available on one path. Proper configuration requires two paths.	Fewer than two paths to an array LUN

storage errors show message	For more information about this message see...
NAME (Serial #): This array LUN is too large and is not usable. The maximum array LUN size supported is xTB.	Array LUN is too small or too large
NAME (Serial #): This array LUN is too small and is not usable. The minimum array LUN size supported is 1GB.	Array LUN is too small or too large
NAME (Serial #): This Array LUN is using multiple LUN IDs. Only one LUN ID per serial number is supported.	LUN IDs for the same LDEV do not match Volume Set Addressing is inconsistent
NAME (Serial #): This array LUN is marked foreign and has a reservation.	Array LUN is marked foreign and has a reservation

Array LUN size is either smaller or larger than the supported values

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

You cannot assign array LUNs with size issues to an ONTAP 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 xGB.
```

Explanation

This message is generated when the array LUN is smaller than the minimum array LUN size supported by ONTAP or exceeds the maximum supported size. Primary reasons for array LUN sizes to be larger or smaller than the supported values might be the following:

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

Some vendors calculate array LUN size limits differently from ONTAP to determine the minimum and maximum array LUN size limits.

The *Hardware Universe* lists the values of the supported maximum and minimum array LUN size limits.

- The array LUNs are for another host whose size limits are different from ONTAP limits.

In an open SAN, ONTAP is exposed to array LUNs intended for other hosts if those array LUNs have not been masked.

ONTAP generates an error message about size issues for any array LUN exposed to it.

Troubleshooting and problem resolution

1. Review the `storage errors show` output so that you find out which array LUN has a size problem.
 - If the array LUN with the size problem is for ONTAP, the storage array administrator must resize the array LUN to meet ONTAP requirements, and then present it again to ONTAP.
 - If the array LUN with the size problem is for another host, the storage array administrator must mask the array LUN so that it is not exposed to ONTAP.
2. After the problem is fixed, run `storage array config show` again to confirm that the error does not persist.

Related information

[NetApp Hardware Universe](#)

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 must be visible to ONTAP systems. The `storage errors show output` identifies the LDEVs whose LUN IDs do not match.

ONTAP does not allow you to assign array LUNs to an ONTAP system if the LUN IDs do not match.

Storage errors show message

```
HIT-1.4 (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 is presented to the same FC initiator port of the ONTAP system from multiple target ports and the LUN IDs are not consistent.
- The LUN IDs of two LDEVs are swapped.

In this case, an error is reported for each array LUN.

- Different LUN IDs for the same LDEV are used when mapping the LDEV to storage array ports that are presenting the LDEV to the ONTAP system.



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.

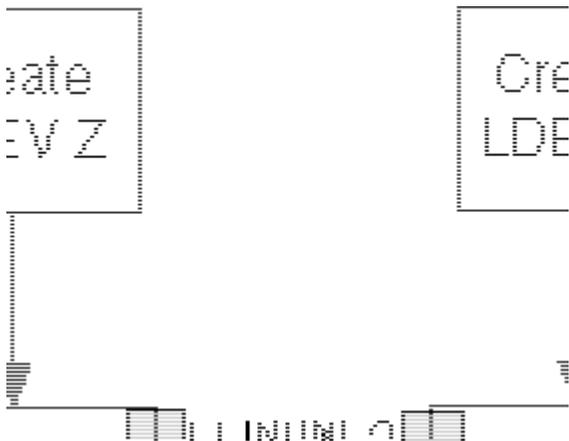
On an EMC Symmetrix storage array, the problem would be that the Volume Set Addressing setting varies on the channel director ports.

Problem scenario

This scenario discusses the case of inconsistent LUN IDs as it applies to most storage arrays. See Volume Set Addressing is inconsistent section for a discussion of this same error message in the context of misconfigured Volume Set Addressing.

Volume Set Addressing is inconsistent

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 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: HIT-1.4
UID: 48495441:43484920:44363030:30353234:30303132:00000000:...
-----

HITACHI_DF600F_1
-----

HIT-1.4 (4849544143484920443630303035323430303132): This Array LUN is
using multiple LUN IDs. Only one LUN ID per serial number is supported.
```



THE UID in this example is
48495441:43484920:44363030:30353234:30303132:00000000:00000000:00000000:00000000:00000000. 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 HIT-1.4
        Disk: HIT-1.4
    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:
          LUN  Initiator Side  Target Side
Controller Initiator ID  Switch Port  Switch Port  Acc Use
Target Port  TPGN...
-----
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:
HIT-1.4 (4849544143484920443630303035323430303132): This Array LUN is
using multiple LUN IDs. Only one LUN ID per serial number is supported.

```



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.

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

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

2. In 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 mysystem1a -initiator 0c -target-wwpn
50060e80004291c1 -lun-number 4
```

The `storage 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, ONTAP can no longer see that LUN.

3. Wait one 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.

4. 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 ONTAP discovery process runs, it discovers the new array LUN. Discovery runs every minute.

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

Volume Set Addressing is inconsistent

ONTAP can detect inconsistent LUN IDs across a path set for storage arrays. 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

```
EMC-1.128 (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 the error message. This explanation is about the display of this message when Volume Set Addressing is set inconsistently.

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



ONTAP does not alert you if Volume Set Addressing is not configured as expected; it alerts you only if the configuration is inconsistent across the channel director ports to which the LUN is mapped.

Troubleshooting and problem resolution

If the `storage errors show` command displays the particular error 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 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 ONTAP.

Related information

[FlexArray virtualization implementation for third-party storage](#)

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 FC initiators of ONTAP systems 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 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.



For some storage arrays, such as HP EVA, 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 are two LUN groups:

Host group	FC initiators in the host group	LDEVs and associated LUN IDs
1	0a0c	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:

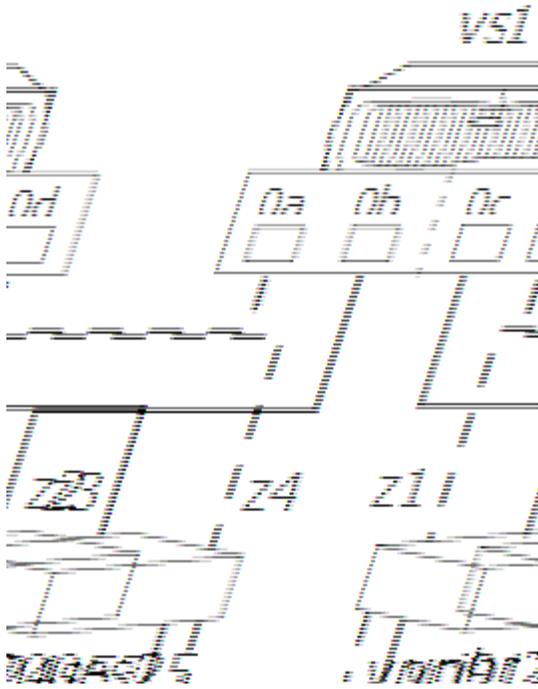
- The initiators in host group 1, 0a and 0c, should be zoned to the target port pair 1A and 2A.
- The initiators in host group 2, 0b and 0d, should be zoned to the target port pair 1B and 2B.

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 are placed in the wrong zone when the zoning is configured, as shown in the following table:

Zone	ONTAP 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:



As you can see from the illustration, two LUN groups are 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.

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 Initiator	LUN Group	LUN Count	Array Name	Array Target Ports	Switch Port
vs1 0a	0	2	DGC_RAID5_1	20:1A:00:a0:b8:0f:ee:04	vnbr200es25:5
0c				20:2A:00:a0:b8:0f:ee:04	vnbr200es25:5
0b	1	2	DGC_RAID5_1	20:1A:00:a0:b8:0f:ee:04	vnbr200es25:5
0d				20:2A:00:a0:b8:0f:ee:04	vnbr200es25:5

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: EMC-1.1
UID: UID-a
-----
EMC-1.1 (UID-a), port WWPN1: LUN 1 occurs more than once.  LUNs cannot be
reused on the same array target port.

Disk: EMC-1.2
UID: UID-b
-----
EMC-1.2 (UID-b), port WWPN1: LUN 2 occurs more than once.  LUNs cannot be
reused on the same array target port.

Disk: EMC-1.3
UID: UID-c
-----
EMC-1.3 (UID-c), port WWPN2: LUN 1 occurs more than once.  LUNs cannot be
reused on the same array target port.

Disk: EMC-1.4
UID: UID-d
-----
EMC-1.4 (UID-d), port WWPN2: 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 IDs, LUN 1 and LUN 2. 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 ONTAP system mapped to it, change the zoning so that the two FC 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 ONTAP and confirm that the error has been fixed. ``

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.

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

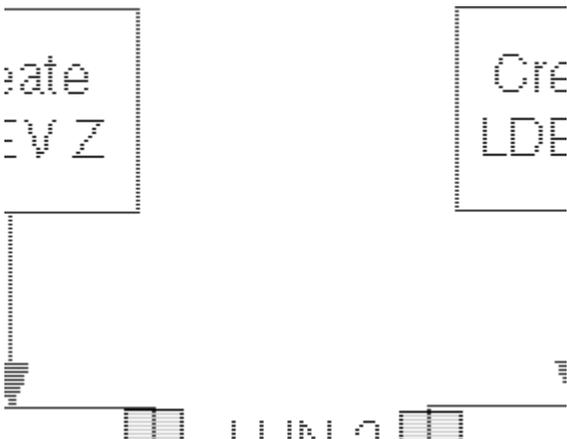
- The LDEV is mapped on only one storage array port.
- The second path to the array LUN is not zoned.
- There is a problem with the host group mapping.
- There is a problem with switch connections.
- The cable dropped out.
- SFPs failed on the adapter.



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
      LUN      LUN
Node      Group  Count  Array Name  Array Target Ports  Switch
Port      Initiator
-----  -
mysystem1a  0      1      DGC_RAID5_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.

```
mysystem1a::> storage errors show
Disk: EMC-1.2
UID: 600508B4:000B6314:00008000:00200000:00000000:00000000:00000000:...
-----
EMC-1.2 (600508b4000b63140000800000200000): This array LUN is only
available on one path. Proper configuration requires two paths.
```

 The UID for this example is 600508B4:000B6314:00008000:00200000:00000000:00000000:00000000:00000000:00000000:00000000. 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 <code>storage array config show</code> 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.

An access control LUN is presented to ONTAP

ONTAP does not support access control array LUNs. The `storage errors show` output alerts you if an 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

ONTAP supports only storage array LUNs. When an array LUN, for example, an E-Series access control LUN is presented to 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 indicates the 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.

```
mssystem1::> storage array config show
```

Node	LUN Group	LUN Count	Array Name	Array Target Port	Initiator
mssystem1	0	1	NETAPP_INF_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 ONTAP, run `storage errors show` again to confirm that the access control LUN is no longer being presented to ONTAP.

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

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.

ONTAP does not allow you to assign array LUNs to an ONTAP 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
```

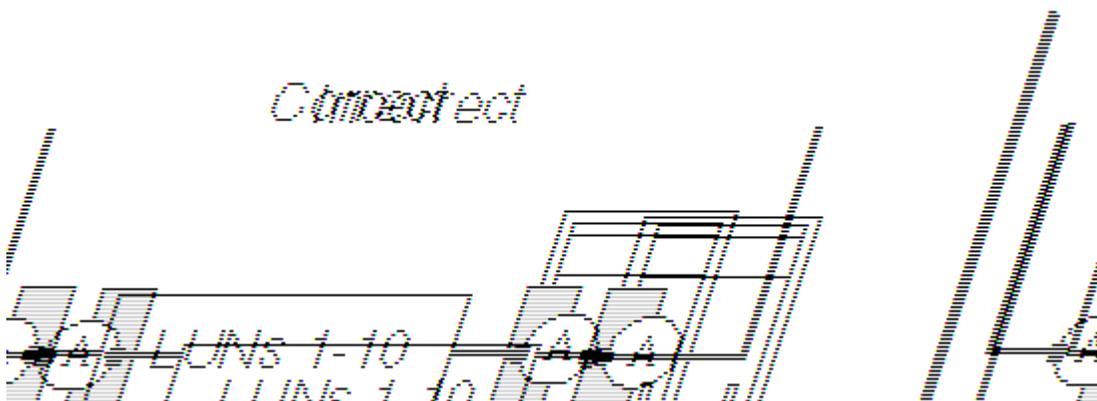
Explanation

This error occurs because the paths to an array LUN are set up to go to the same storage array controller or FRU.



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 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 FEBE board. A FEBE board is considered to be one fault domain because if all paths go through the same FEBE 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: HP-1.15
Container Type: unassigned
Owner/Home: - / -
DR Home: -
Array: HP_HSV450_1
Vendor: HP
Model: HSV450
Serial Number: 600508B4000B63140000800001660000
        UID: 600508B4:000B6314:00008000:01660000:00000000:...
        BPS: 512
Physical Size: -
        Position: present
Checksum Compatibility: block
        Aggregate: -
        Plex: -

```

Paths:

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

Errors:

HP-1.15 (600508b4000b63140000800001660000): All paths to this array LUN are connected to the same fault domain. This is a single point of failure.



The full UID in this example is 600508B4:000B6314:00008000:01660000:00000000:00000000:00000000:00000000:00000000. 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 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

ONTAP requires that array LUNs that are visible to a particular ONTAP 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(UID): 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 EMC CLARiiON 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 ONTAP system. ONTAP does not support inconsistent failover modes for paths to an array LUN from a particular ONTAP 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 ONTAP 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 ONTAP system might not operate properly, failover might not occur, or the system might panic.



ONTAP supports different failover mode settings between nodes running ONTAP. 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 ONTAP discovers during LUN initialization is the failover mode that ONTAP expects for all paths to the LUN from a particular ONTAP system. If the failover mode of subsequent discovered paths does not match the failover mode of the first path, ONTAP issues an error message.

In the following `storage errors show` example, ONTAP tells you that the failover mode for the LUN EMC-1.128, which is visible over `mysystem1` FC initiator `0a`, is *Proprietary*, and that the failover mode is different from the failover mode that ONTAP discovered on the first path for that array LUN.

```
mysystem1::> storage errors show
EMC-1.128 (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: EMC-1.128
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 ONTAP detected on the first path is the failover mode that you want to be used for all paths on that ONTAP system to the array LUN.

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



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 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 ONTAP detected for the first path.	Change, on the storage array, the failover mode for the initiator that 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 ONTAP detected for the first path.	Remove the array LUN from the view of the ONTAP 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 ONTAP, take the second path offline.
 - b. On the storage array, change the failover mode for the initiator that ONTAP identified in the error message.
 - c. In ONTAP, bring the second path back online.
4. If you need to remove the array LUN from the view of the ONTAP 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 ONTAP system does not have to be rebooted.</p> <ol style="list-style-type: none"> In ONTAP, run the following command for each affected spare LUN: <code>disk remove_ownership LUNfullname</code> On the storage array, mask each affected array LUN on all paths to the ONTAP system. Wait for about a minute, then confirm that the array LUNs are no longer visible to the ONTAP system. Set the same failover mode for each of the FC initiators on the ONTAP system. Present all the affected array LUNs to the ONTAP system again. <p>ONTAP should discover the LUNs when it next runs LUN discovery</p> <ol style="list-style-type: none"> Run <code>storage errors show</code> to confirm that there is no longer a failover mode error. 	<p>With this method, the ONTAP system must be rebooted.</p> <ol style="list-style-type: none"> Reboot the ONTAP system and hold it at the LOADER prompt. 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. Reboot the ONTAP system. Run <code>storage errors show</code> to confirm that there is no longer a failover mode error.

Array LUN is marked foreign and has a reservation

You can import data hosted on array LUNs in their native format to ONTAP. The process of importing the contents of such an array LUN that is marked **foreign** might be impacted if the array LUN has external host reservations during import.

Storage errors show message

```
EMC-1.3 (600000e00d1000000010000e00030000): This array LUN is marked
foreign and has a reservation.
```

Explanation

This error condition occurs when the reservations created by external host applications on the LUN are not cleared before the ONTAP administrator initiates the import process. The import operation fails and the message is displayed as the output of the `storage errors show` command.

The reservations must be cleared for the import operation to succeed.

Troubleshooting and problem resolution

To fix the problem, the storage array administrator must remove the persistent reservation from the array LUN.

You can remove the persistent reservation either from the external host that accesses the array LUN data in its

native format or by using ONTAP commands. You can use the following ONTAP commands to remove the reservation:

1. Review the `storage errors show` output to identify the array LUN with reservations, as shown in the following example:

```
systemla::> storage errors show
EMC_SYMMETRIX_1
-----
EMC-1.3 (600000e00d1000000010000e00030000): This array LUN is marked
foreign and has a reservation.
```

2. Run the `set -privilege advanced` command to switch to the *advanced* mode, because you cannot import the array LUN contents in the *admin* mode.
3. Use the `storage disk -remove-reservation` command to remove the persistent reservation from the identified array LUN.

```
systemla::>* storage disk remove-reservation -disk EMC-1.3
```

4. Use the `lun import start` command to initiate the import process of the array LUN contents to ONTAP.

Verifying that the back-end configuration matches the intended configuration

After resolving the back-end configuration errors detected by `storage errors show`, you must check whether the existing back-end configuration matches the configuration you intended, and fix issues arising out of any mismatch.

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:

```

mysystem1::> storage array config show
      LUN  LUN
Node   Group Count  Array Name      Array Target Port  Initiator
-----
mysystem1a   0    10    DGC_RAID5_1    5005076303030124  1a
                                     5005076303088124  1b
                                     5005076303130124  1c
                                     5005076303188124  1d
mysystem1b   0    10    DGC_RAID5_1    5005076303030124  1a
                                     5005076303088124  1b
                                     5005076303130124  1c
                                     5005076303188124  1d

8 entries were displayed.

```

2. Check the `storage array config show` output for each problem statement in the following table, and refer to the appropriate information for solution.

If you find this problem...	See this information...
Empty LUN groups	Reasons for no LUNs in the array LUN group
Missing expected storage arrays	Reasons you might not see all storage arrays you expected
More array LUN groups than expected	Reasons for more array LUN groups than expected
Fewer array LUN groups than expected	Reasons for fewer array LUN groups than expected
Array LUN groups do not show the number of LUNs expected	Reasons for the incorrect number of LUNs in array LUN groups
Fewer paths than expected	Fewer than two paths to an array LUN
More paths than expected	Reasons for more paths to an array LUN than expected

Reasons for no LUNs in the array LUN group

When validating the back-end configuration, you should check the `storage array config show` output to determine whether LUNs are shown in the LUN groups. When the `storage array config show` output shows no LUNs in an array LUN group, ONTAP can see the target port on the fabric but the target port is not presenting array LUNs to ONTAP.

There are a variety of reasons why a target port might not present array LUNs to ONTAP (an *open target port*).

The reasons for an open target port can differ between different storage arrays. In addition, the ways to handle the open target port issues differ with storage arrays. For all storage arrays, troubleshooting the cause of an open target port should include checking the storage array configuration, including the host group configuration.

For example, the cause could be an *empty host group*, which presents itself to ONTAP as an open target port. With an empty host group, the host group defines the FC initiator and target ports, but it does not list any array LUNs (that is, the host group is empty).

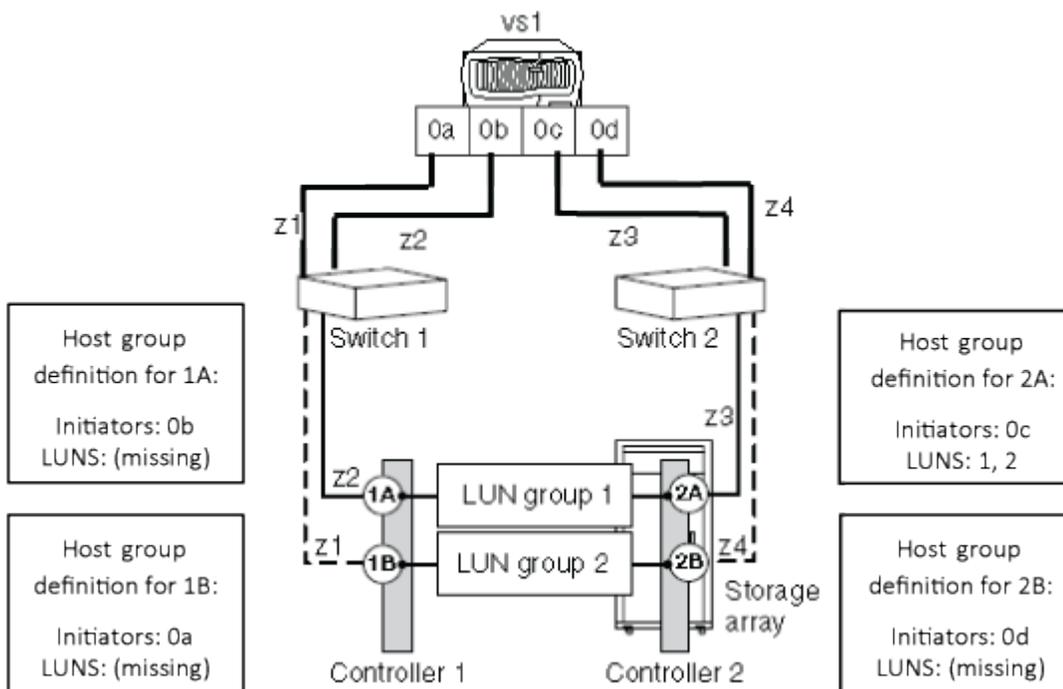


The operation of the ONTAP systems is not impacted by an empty host group.

The following illustration represents an open port target port situation and an empty host group situation. (A stand-alone ONTAP system is used in the illustration for simplicity.)

The open target port is caused by missing LUN IDs in the host group definition for storage controller 1A. In the host group definition for storage controller 1A, the FC initiator 0b is zoned into the target port and there is a host group defined for FC initiator port 0b, but there are no LUN IDs in the host group. As the illustration shows, the host group definition for storage controller 2A includes both initiators and LUNs.

The empty host group is caused by array LUNs not getting listed in the host group definition for storage controllers 1B and 2B. The FC initiators 0a and 0d are zoned into the target ports but no LUNs are shown.



The following example shows the `storage array config show` output for the illustrated situations where the ONTAP system is zoned to the storage array but there are no LUNs in the host group definitions for the storage controller 1A, resulting in an open target port. The controllers 1B and 2B do not have LUNs in their respective host group definitions, resulting in an empty host group.

In the output for the illustrated situations, the FC initiators 0a, 0b, and 0d show no LUNs in the LUN count field. For the FC initiator 0c, two LUNs are shown in LUN group1.

```

cluster-1::> storage array config show
      LUN  LUN
Node   Group Count  Array Name      Array Target Port  Initiator
-----
vs1
      0    2    EMC_SYMMETRIX_1  5006048000001b0    0c
      1    0    EMC_SYMMETRIX_1  5006048000001a0    0b
      2    0    EMC_SYMMETRIX_1  5006048000001b1    0a
      0    0    EMC_SYMMETRIX_1  5006048000001a1    0d

4 entries were displayed.

```



From the output you can determine that the problem is not due to missing FC initiators in the host group. If the FC initiators were missing from the host group, ONTAP would not be able to see the LUN groups having no LUNs.

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 FC initiator port pairs on the ONTAP system. If the LDEV-to-LUN mapping is the same for both FC initiator port pairs, the `storage array config show` output shows one less LUN group than expected.

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

```
mssystem1::> storage array config show
```

Node	LUN Group	LUN Count	Array Name	Array Target Ports	Switch Port
mssystem1 0a	0	2	DGC_RAID5_1	20:1A:00:a0:b8:0f:ee:04	vnbr200es25:5
0c				20:2A:00:a0:b8:0f:ee:04	vnci9124s53:6
0b				20:1A:00:a0:b8:0f:ee:04	vnbr200es25:5
0d				20:2A:00:a0:b8:0f:ee:04	vnci9124s53:6

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

ONTAP does not flag this as an error because an ONTAP system can operate when the LDEV-to-LUN mapping is the same for both FC initiator port pairs. However, multiple FC initiators to the same target port is not supported. You should fix the mapping so that your ONTAP 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 the FC initiator port pair mapping is no longer the same for both FC initiator port pairs on the ONTAP system.
2. On the ONTAP system, run `storage array config show` again and confirm that the number of LUN groups that you expected are shown and that 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 extra 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.



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.

- ONTAP is trying to balance the load over the target ports during migration of array LUNs from one LUN group to another.

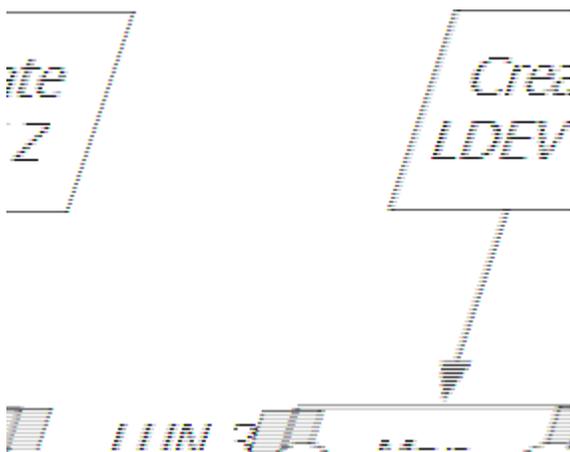
In this case, the LDEV spanning path pairs is a transitional state.

- Too many connections are set up.

ONTAP can run with a single array LUN group or multiple array LUN groups. Therefore, 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 ONTAP releases. During configuration planning, you should check the information in the Interoperability Matrix to verify whether multiple array LUN groups are supported for the storage arrays used in your ONTAP configuration.

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 is created (Group 2), as the following example shows:

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

Node Initiator	LUN Group	LUN Count	Array Name	Array Target Ports	Switch Port
mysystem1 0a	0	3	DGC_RAID5_1	20:1A:00:a0:b8:0f:ee:04	vnbr20es25:5
0c				20:2A:00:a0:b8:0f:ee:04	vnci9124s53:6
0b	1	3	DGC_RAID5_1	20:1B:00:a0:b8:0f:ee:04	vnbr20es25:7
0d				20:2B:00:a0:b8:0f:ee:04	vnci9124s53:8
0a	2	1	DGC_RAID5_1	20:1A:00:a0:b8:0f:ee:04	vnbr20es25:5
0d				20:2B:00:a0:b8:0f:ee:04	vnci9124s53:8

You can infer from the information in the Array Target Ports column for Groups 0 and 1 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, the information in the Array Target Ports column for Group 2 suggests 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 ONTAP system can run with an LDEV spanning path pairs, the message instructing you to run `storage errors show` does not appear after you run the `storage array config show` command. However, this is not a best practice configuration.

Resolving the issue

1. Wait for one 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 ONTAP, run `storage array config show` to confirm that three paths are shown.
5. Remove the incorrect mapping.
6. Wait for one minute while ONTAP discovers the LUN.
7. In ONTAP, run `storage array config show` again to make sure that the extra LUN group is gone.

Related information

[NetApp Interoperability Matrix Tool](#)

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 ONTAP 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.
```



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 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 is not mapped to an ONTAP 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:

```
mssystem1::> storage array config show
```

Node	LUN Group	LUN Count	Array Name	Array Target Port	Initiator
mssystem1	0	50	DGC_RAID5_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 is not mapped to the ONTAP system.
- A mapping error is 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 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 ONTAP system, the storage array administrator has 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 ONTAP 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 is fixed.

Reasons storage arrays are missing from command output

A storage array that is not connected to the ONTAP 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 FC initiators on the ONTAP systems are in the host group.
 - Check that both the storage array and the FC initiator are in the same zone.
2. After the problem is fixed, run `storage array config show` in ONTAP to confirm that the issue was fixed.

Managing array LUNs using ONTAP

If the storage array administrator wants to make configuration changes to an array LUN after it is assigned to a node, for example to resize it, you might need to perform some activities in ONTAP before it is possible to reconfigure the LUN on the storage array.

Modifying assignment of spare array LUNs

You can change the ownership of a *spare* array LUN to another node. You might want to do this for load balancing over the nodes.

Steps

1. At the console of the node that owns the array LUN you want to reassign, enter the following command to see a list of spare array LUNs on the node: `storage disk show -owner local`

The array LUNs owned by the node, both spares and LUNs in aggregates, are listed.

2. Confirm that the LUN you want to reassign to another node is a spare LUN.
3. Enter the following command to assign ownership of the array LUN to another node: `storage disk assign arrayLUNname -owner new_owner_name -force`



The array LUN ownership is not changed if the `-force` option is not used or if the array LUN was already added to an aggregate.

4. Enter the following command to verify that the ownership of the spare array LUN was changed to the other node: `storage disk show -owner local`

The spare array LUN that you changed to the new owner should no longer appear in the list of spares. If the array LUN still appears, repeat the command to change ownership.

5. On the destination node, enter the following command to verify that the spare array LUN whose ownership you changed is listed as a spare owned by the destination node: `storage disk show -owner local`

You must add the array LUN to an aggregate before it can be used for storage.

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.

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 ONTAP selects array LUNs of the same checksum type as the aggregate.



The checksum type of all newly created aggregates using zoned checksum array LUNs is *advanced zoned checksum (AZCS)*. Zoned checksum type 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:

```
storage disk show -fields checksum-compatibility -container-type spare
```

You can add a block checksum array LUN to a block checksum aggregate and a zoned array LUN to an *advanced zoned checksum (AZCS)* aggregate.

Changing the checksum type of an array LUN

You must change the checksum type of an array LUN if you want to add it to an aggregate that has a different checksum type than the checksum type of the LUN.

What you'll need

You must have reviewed the tradeoffs between performance in certain types of workloads and storage capacity utilization of each checksum type.

You can also contact your Sales Engineer for details about using checksums.

- You must assign a **zoned** checksum type to an array LUN that you plan to add to an advanced zoned checksum (AZCS) aggregate. When a zoned checksum array LUN is added to an AZCS aggregate, it becomes an advanced zoned checksum array LUN. Similarly, when a zoned checksum array LUN is added to a zoned aggregate, it is a zoned checksum type.
- You cannot modify the checksum of array LUNs while assigning ownership. You can modify the checksum only on already assigned array LUNs.

Steps

1. Enter the following command to change the checksum type: `storage disk assign -disk disk name -owner owner -c new_checksum_type`

disk name is the array LUN whose checksum type you want to change.

owner is the node to which the array LUN is assigned.

new_checksum_type can be block or zoned.

```
storage disk assign -disk EMC-1.1 -owner system147b -c block
```

The checksum type of the array LUN is changed to the new checksum type you specified.

2. Enter the following command to access the nodeshell: `system node run -node node_name`

node_name is the name of this system.

3. Enter the following command to exit the nodeshell: `exit`

Prerequisites to reconfiguring an array LUN on the storage array

If an array LUN has already been assigned (through ONTAP) to a particular ONTAP system, the information that ONTAP wrote to the array LUN must be removed before the storage administrator attempts to reconfigure the array LUN on the storage array.

When the storage array presents an array LUN to ONTAP, ONTAP collects information about the array LUN (for example, its size) and writes that information to the array LUN. ONTAP cannot dynamically update information that it wrote to an array LUN. Therefore, before the storage array administrator reconfigures an array LUN, you must use ONTAP to change the state of the array LUN to *unused*. The array LUN is unused from the perspective of ONTAP.

While changing the state of the array LUN to unused, ONTAP does the following:

- Terminates I/O operations to the array LUN
- Removes the label for RAID configuration information and the persistent reservations from the array LUN, which makes the array LUN unowned by any ONTAP system

After this process finishes, no ONTAP information remains in the array LUN.

You can do the following after the array LUN's state is changed to unused:

- Remove the mapping of the array LUN to ONTAP and make the array LUN available to other hosts.
- Resize the array LUN or change its composition.

If you want ONTAP to resume using the array LUN after its size or composition is changed, you must present the array LUN to ONTAP again, and assign the array LUN to an ONTAP system again. ONTAP is aware of the new array LUN size or composition.

Changing array LUN size or composition

Reconfiguring the size or composition of an array LUN must be done on the storage array. If an array LUN has already been assigned to an ONTAP system, you must use ONTAP to change the state of the array LUN to unused before the storage array administrator can reconfigure it.

What you'll need

The array LUN must be a spare array LUN before you can change its state to unused.

Steps

1. On the ONTAP system, enter the following command to remove ownership information: `storage disk removeowner -disk arrayLUNname`
2. On the storage array, complete the following steps:
 - a. Unmap (unpresent) the array LUN from the ONTAP systems so that they can no longer see the array LUN.
 - b. Change the size or composition of the array LUN.
 - c. If you want ONTAP to use the array LUN again, present the array LUN to the ONTAP systems again.

At this point, the array LUN is visible to the FC initiator ports to which the array LUN was presented, but it cannot be used by any ONTAP systems yet.

3. Enter the following command on the ONTAP system that you want to be the owner of the array LUN:
`storage disk assign -disk arrayLUNname -owner nodename`

After the ownership information is removed, the array LUN cannot be used by any ONTAP system until the array LUN is assigned again to a system. You can leave the array LUN as a spare or add it to an aggregate. You must add the array LUN to an aggregate before the array LUN can be used for storage.

Removing one array LUN from use by ONTAP

If the storage array administrator no longer wants to use a particular array LUN for ONTAP, you must remove the information that ONTAP wrote to the LUN (for example, size and ownership) before the administrator can reconfigure the LUN for use by another host.

What you'll need

If the LUN that the storage array administrator no longer wants ONTAP to use is in an aggregate, you must take the aggregate offline and destroy the aggregate before starting this procedure. Taking an aggregate offline and destroying it changes the data LUN to a spare LUN.

Step

1. Enter the following command: `storage disk removeowner -disk LUN_name`

LUN_name is the name of the array LUN.

Preparing array LUNs before removing an ONTAP system from service

You must release the persistent reservations on all array LUNs assigned to an ONTAP system before removing the system from service.

When you assign ONTAP ownership of an array LUN, ONTAP places persistent reservations (ownership locks) on that array LUN to identify which ONTAP system owns the LUN. If you want the array LUNs to be available for use by other types of hosts, you must remove the persistent reservations that ONTAP put on those array LUNs; some arrays do not allow you to destroy a reserved LUN if you do not remove the ownership and persistent reservations that ONTAP wrote to that LUN.

For example, the Hitachi USP storage array does not have a user command for removing persistent reservations from LUNs. If you do not remove persistent reservations through ONTAP before removing the ONTAP system from service, you must call Hitachi technical support to remove the reservations.

See the Knowledge Base article [What are SCSI Reservations and SCSI Persistent Reservations](#)

Troubleshooting configurations with storage arrays

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 an ONTAP configuration with array LUNs

While troubleshooting an ONTAP configuration with array LUNs, you should follow a systematic approach to determine the cause of the issue.

This procedure suggests an order for approaching troubleshooting.



As you continue with the troubleshooting steps, you should save any information that you gather about the issue so that you can provide this information to technical support in the case of an escalation.

Steps

1. Determine whether the issue is with the *front end* (an ONTAP issue that affects all the corresponding platforms) or the *back end* (an issue with the switch or storage array configuration).

For example, if you are trying to use an ONTAP feature and it is not working as you expected, the issue is likely to be with the front end.

2. Take appropriate action to address the issue, depending on its nature:

If the ONTAP configuration has a...	Then do this...
Front end issue	Proceed with troubleshooting the ONTAP feature by following instructions in the ONTAP guides. ONTAP 9 Documentation

If the ONTAP configuration has a...	Then do this...
Back end issue	<p>a. Check the Interoperability Matrix to ensure that the following are supported: configuration, storage array, storage array firmware, switch, and switch firmware.</p> <p>NetApp Interoperability Matrix Tool</p> <p>b. Use the <code>storage array config show</code> command to check whether there are any common back-end configuration errors that the system can detect.</p> <p>If ONTAP detects a back-end configuration error, you must run the <code>storage errors show</code> command to obtain details about the error.</p>

3. If the cause of the issue is still not clear, check the following sources to ensure that your system is compliant with the requirements for working with storage arrays:

- [Verifying an installation with storage arrays](#)
- [FlexArray virtualization implementation for third-party storage](#)
- [FlexArray virtualization implementation for NetApp E-Series storage](#)
- [NetApp Interoperability Matrix Tool](#)
- [NetApp Hardware Universe](#)

4. If you still need help resolving the issue, contact technical support.

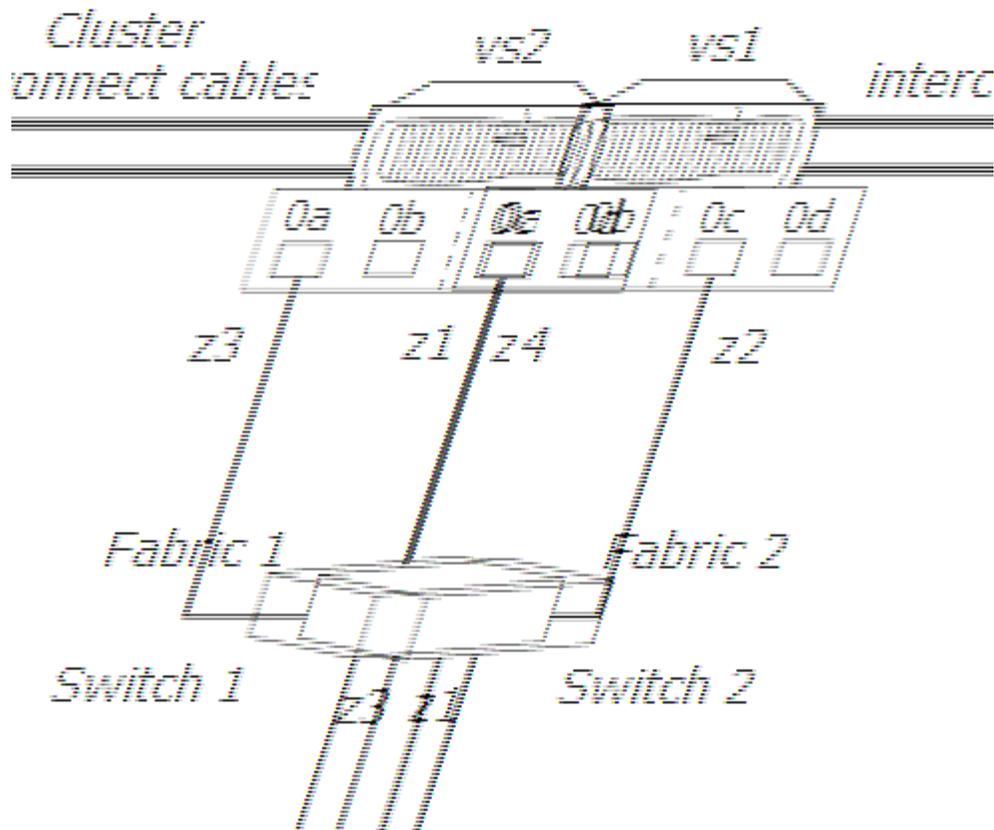
Invalid path setup examples

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 ONTAP requirements.

Invalid path setup: alternate paths are not configured

It is important to set up alternate paths to all array LUNs from both FC initiators on the ONTAP system, thereby avoiding 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 ONTAP systems to each LUN on the storage array. Both FC initiator ports from the same ONTAP 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.

What happens when a link failure occurs

ONTAP monitors a link's usage periodically. The 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...
ONTAP system and the switch	ONTAP receives notification immediately and sends traffic to the other path immediately.
Switch and the storage array	ONTAP is not immediately aware that there is a link failure because the link is still established between the ONTAP system and the switch. ONTAP becomes aware that there is a failure when the I/O times out. 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.

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 FC initiator port on the ONTAP system and the storage array port WWPN or WWNN for that zone. Likewise, when the host group for the ONTAP system is configured on the storage array, the WWPNs of the FC 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 FC initiator port on the ONTAP system 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 ONTAP output, the FC initiator ports on the ONTAP system 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.

How WWPNs are specified	Common error
WWPNs are automatically discovered by the switch	The wrong FC initiator port WWPN is selected from the picklist.



When the ONTAP systems, switches, and storage array are cabled together, the switch automatically discovers the WWPNs of the ONTAP 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 FC 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 FC initiator port labels that are visible on the ONTAP 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 ONTAP output
The FC initiator port in the zone definition is incorrect. This caused the incorrect FC initiator port WWPN to be propagated to the host group configuration.	The WWPN of the FC initiator port shown in the picklist was selected, not the WWPN you intended.	Array LUNs are not visible over the FC initiator port on which the expected LUNs would be visible.
The zone definition includes the correct FC initiator port.	The WWPN in the host group definition is incorrect because of either of the following: <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 FC initiator port. 	

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 FC initiator port 0a of the ONTAP system was put in the zone definition. However, the intention was that the WWPN for FC initiator port 0c was to be put into the zone definition.

2. The host group was created on the storage array.

The WWPN for FC 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 ONTAP, you looked at array LUNs over the FC 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 FC initiator port 0a.



You used the `storage array config show` command to view the array LUN information.

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 ONTAP system.

You realize that you have two zones with the WWPN for FC initiator port 0a in it and no zones with the WWPN for 0c in it.

2. Fix the incorrect zone definitions and activate them.



You would not be able to 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 FC 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 ONTAP system, run `storage array config show` to check the array LUNs over the FC initiator ports to confirm that array LUNs are shown over 0c.

Troubleshooting by checking the host group first

1. From the console of the ONTAP system, run `storage show adapteradapter#``, 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 FC 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.



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 FC 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 FC 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 ONTAP system, run `storage array config show` to check the array LUNs over the FC initiator ports to confirm that array LUNs are shown over 0c.

You should now see access to the LUNs over the FC initiator port.

Additional tasks after installing and testing an ONTAP configuration with array LUNs

After completing the installation and testing of an ONTAP configuration with array LUNs, you can assign additional array LUNs to your ONTAP systems and set up various ONTAP features on your systems.

The following are some of the tasks that you can perform after completing the installation and testing of the ONTAP system with array LUNs:

- Assign additional array LUNs to ONTAP systems as required.
- Create ONTAP aggregates and volumes as required.
- Set up additional ONTAP features on your system as required—for example, features for backup and recovery.

For information about configuring the various ONTAP features, see the appropriate ONTAP guide.

Obtaining WWPNs manually

If the ONTAP system is not connected to the SAN switch, you need to obtain the World Wide Port Names (WWPNs) of the system's FC initiator ports that will be used to connect the system to the switch.

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 system's console connection to a laptop computer.
2. Power on your system.

Interrupt the boot process by pressing Ctrl-c when you see the following message on the console:

```
Press CTRL-C for boot menu
```

3. Select the Maintenance Mode option on the boot options menu.
4. Enter the following command to list the WWPNs of the system's 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 exit the Maintenance mode.

Target queue depth customization

The target queue depth defines the number of ONTAP commands that can be queued (outstanding) on a storage array target port. ONTAP supplies a default value. For most deployments, the default target queue depth is appropriate; however, you can change it to correct performance issues.

The default target queue depth differs with different releases of ONTAP:

- For ONTAP, 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.



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 accessing the storage array port when you are planning the configuration for a specific ONTAP system or a specific host that does not run ONTAP.

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 hosts that do not run ONTAP 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.

Setting the target queue depth (ONTAP prior to 8.2)

The default target queue depth is acceptable for most implementations, but you can change the default value if required.

This setting is per ONTAP system, and it applies to all target ports on all storage arrays. For ONTAP systems running version prior to 8.2, you can use this option.

Step

1. Use the following option to set the target queue depth: `options disk.target_port.cmd_queue_depth value`

Setting the target queue depth

The default target queue depth is acceptable for most implementations, but can be changed if performance issues are encountered.

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>

For more information about these commands, see the man pages.

Displaying target queue depth statistics

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.

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 whether 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` command 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 help you determine whether 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
-----
vgv51-02  0a        21    2    53    0     0      0         0
vgv51-01  0a        21    2    48    1     0      0         0

```

4. You can obtain more detailed information about ports by using the `storage array port show -fields` command with the `average-latency-per-iop`, `average-pending`, `average-waiting`, `max-pending`, or `max-waiting` fields.

Target port utilization policy settings

ONTAP can detect events of resource contention, such as I/O queue being full, commands being timed out or HBA resource getting exhausted, on a target port.

You can set the target port utilization policies by using the `storage array port modify` command if you

detect such events on a given array target port.

The following table describes the two utilization policies associated with a target port:

Policy	Description
normal	When ONTAP detects target port resource contention on a given array target port, it reduces the target port's queue depth and throttles I/O to the target port. In this mode, the reduction of the target port queue depth is lesser than the defer policy for each target port resource contention event. The subsequent rise in the target port queue depth is faster than the defer policy. normal is the default policy.
defer	When ONTAP detects target port resource contention on a given array target port, it reduces the target port's queue depth and throttles I/O to the target port. In this mode, the reduction of the target port queue depth is greater than the normal policy for each target port resource contention event. The subsequent rise in the target port queue depth is slower than the normal utilization policy.

Output examples to view and modify array target port utilization policies

The following command displays the target port utilization policy associated with an array target port:

```

vgv3170_jon::> storage array port show -wwnn 2703750270235
    Array Name: HITACHI_DF600F_1
        WWNN: 2703750270235
        WWPN: 2703750270235
    Connection Type: fabric
        Switch Port: vgbr300s89:9
        Link Speed: 4 GB/s
    Max Queue Depth: 1024
    Utilization Policy: defer

                                LUN
Link
Node                               Initiator  Count  IOPS  KB/s  %busy  %waiting
Errs
-----
-----
0          vgv3170f54a          0a      2     50   1956   85     0
0          vgv3170f54b          0a      2    350  15366  100    40
0

```

By default, the I/O policy for a given array target port is **normal**. You can modify the I/O policy associated with the port by running the following command:

```
vgv3070f50ab::> storage array port modify -wwpn 50014380025d1509
-utilization-policy ?
```

```
normal      This policy aggressively competes for target port resources,
in effect competing with other hosts.
            (normal)
```

```
defer      This policy does not aggressively compete for target port
resources, in effect deferring to other hosts.
```

```
vgv3070f50ab::> storage array port modify -wwpn 50014380025d1509
-utilization-policy defer
1 record updated.
```

Terminology comparison between storage array vendors

Different storage array vendors occasionally use different terms to describe similar concepts. Conversely, the meaning of the same term might differ between array vendors.

The following table provides a mapping between some common vendor terms:

Term	Vendor	Definition
host group	Hitachi	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 ONTAP system 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.
	HP XP	
Storage Group	EMC CX, EMC VNX	
host definition	3PAR	
host	3PAR, HP EVA, HP XP, Hitachi	

Term	Vendor	Definition
parity group	Hitachi, HP XP	The arrangement of disks in the back-end that together form the defined RAID level.
RAID group	ONTAP, EMC CX, EMC VNX	
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	ONTAP	In ONTAP , 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	A hardware component on the storage arrays that contains the ports to which hosts attach.
controller	ONTAP	The component of a storage system that runs the ONTAP operating system and interacts with back-end storage arrays. Controllers are also sometimes called heads or CPU modules.
	Hitachi, HP EVA, HP XP	Hardware on the storage array on which the target ports are located.
node	3-PAR	A hardware component on the storage arrays that contains the ports to which hosts attach.
FEBE Board	EMC Symmetrix	
Storage processor (SP)	EMC CX, EMC VNX	

Term	Vendor	Definition
LUN	Many storage arrays	A grouping of one or more disks or disk partitions into one span of disk storage space. In the ONTAP documentation, this is referred to as <i>array LUN</i> .
LDEV	Hitachi, HP XP	
LUN	ONTAP	The ONTAP system can virtualize the storage attached to it and serve the storage up as LUNs to external applications and clients (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 <i>Vdisk</i> 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	ONTAP documentation, ONTAP storage management tools	The ONTAP documentation uses the term <i>array LUN</i> to distinguish LUNs on the storage arrays from front-end LUNs (ONTAP LUNs).
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	ONTAP	A logical entity that holds user data that is accessible through one or more of the access protocols supported by ONTAP, including Network File System (NFS), Common Internet File System (CIFS), HyperText Transfer Protocol (HTTP), Fibre Channel (FC), and Internet SCSI (iSCSI).

Term	Vendor	Definition
virtual volume	3PAR	A virtual storage unit created by mapping data from one or more logical disks.

Copyright information

Copyright © 2024 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

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