NetApp

Implementing third-party storage

ONTAP FlexArray

NetApp
February 11, 2024

Table of Contents

FlexArray® Virtualization Implementation for Third-Party Storage

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

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…
Tasks for ONTAP implementation with storage arrays, including the following: • Planning the implementation • Connecting the ONTAP systems and the array • Verifying the installation	FlexArray virtualization installation requirements and reference
ONTAP working with devices, including the following: • 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	NetApp Interoperability Matrix Tool

For information about…	You should look here…
ONTAP limits for releases and platforms, including the following: • 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	[NetApp Hardware Universe](#)

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](#)

Support for advanced features of storage arrays

ONTAP does not support advanced features of storage arrays unless otherwise specified in the Interoperability Matrix.

Related information

[NetApp Interoperability Matrix Tool](#)

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

General configuration guidelines for all storage arrays

There are certain configuration guidelines that you must follow for any of the storage arrays supported by ONTAP.

Requirement for 8-Gb array port initialization

In a typical Fibre Channel (FC) fabric, storage array ports initialize as N-Ports (node ports) and their corresponding switch ports initialize as F-Ports (fabric ports). In some cases, these F-ports might go offline and resume operation as L-Ports (loop ports), resulting in fabric errors.

This behavior is caused by topology auto-negotiation issues in the Brocade switch.

To ensure that switch ports always initialize as F-ports, you must run the `portcfggport `<portnumber> 1` command on Brocade switch ports that are connected to storage arrays. You must run this command when you start configuring the system to avoid this problem, and you must fix the problem if it occurs.

To promote efficiency on the fabric as a whole, you can also use this workaround on any Brocade ports attached to ONTAP initiator ports.

Required fill word settings for 8 Gb Brocade switch ports

If you are using 8 Gb Brocade switches running firmware version FOS 6.3.1 or later in your storage environment, you must verify that appropriate fill word settings are configured on the switch ports connected to ONTAP systems and that the switch ports are connected to storage arrays. Providing incorrect fill words might lead to link initialization failures in the switch ports.

For ONTAP systems, you must set the fill word on the corresponding switch ports to 3.

For storage arrays, you must set the fill word according to the following recommended values, using the Brocade switch documentation for instructions:

Storage array	Fill word value on the Brocade switch ports
Hitachi	2
HP XP	
All other storage arrays supported with Data ONTAP systems	3

Related information

EMC VNX storage arrays

You must meet specific requirements when configuring EMC VNX storage arrays to work with ONTAP systems that use the array LUNs. These requirements include setting configuration parameters on your storage arrays and deploying only supported configurations.

Required parameters for EMC VNX storage arrays with ONTAP systems

Certain parameter settings are required on the storage array for the storage array to work successfully with ONTAP systems.

Host configuration parameters that must be set on the storage array are listed in the following table.

Parameter	Setting
Initiator type	CLARiiON/VNX
Array Com Path	Enabled
Failover mode	4
Unit serial number	LUN
Host name	User-supplied host name and port number
IP address	Unique fake IP address You must ensure that you have not entered this IP address anywhere else in the storage array configuration, and that it is not an IP address that is present on the network.

How EMC VNX storage array control access to data

EMC VNX array use Storage Groups to control access to data. A Storage Group is one or more LUNs within the storage array that can be accessed only by the host or hosts that you associate with the array LUNs. A host cannot access or modify data in any array LUNs that are not part of its Storage Group.

ONTAP supports multiple Storage Groups if these rules are followed:

- Switch zoning must define which target ports the FC initiator ports on the ONTAP system use to access each array LUN group.
- LUN masking must be used to restrict host access to array LUNs.
- Storage Groups must define which array LUN groups are presented to each FC initiator port.

• One FC initiator port pair on each ONTAP system is required for each array LUN group.

If array LUN neighborhoods are used, the V-Series systems in the same neighborhood must be in the same Storage Group.

Limitation on array LUN numbering with EMC VNX array

EMC VNX storage array only support array LUN numbering from 0 through 255. Array LUNs that are numbered beyond this range are not visible to ONTAP.

Guidelines for enabling ALUA on EMC VNX storage arrays

Support has been added for Asymmetric Logical Unit Access (ALUA) with CLARiiON and VNX arrays.

By default, support for ALUA is enabled in ONTAP. But to use it, ALUA must be enabled on the storage array.

You should enable ALUA only on new configurations. You should not enable ALUA in an existing configuration.

If you are setting up ALUA on the storage array, you must ensure that all hosts in a Storage Group are in the same failover mode, that is failover mode 4 (ALUA).

ALUA behavior of VNX2 storage arrays

Although all EMC CLARiiON and VNX arrays use ALUA (asymmetric active-active) failover mode, VNX2 storage arrays might differ in their ALUA behavior.

Like all other EMC array LUNs, VNX2 array LUNs presented to the ONTAP system from dynamic disk pools use ALUA failover mode. However, VNX2 array LUNs presented to the ONTAP system from traditional RAID groups use active-active failover mode, with all the paths being reported as `AO` (active optimized). This behavior of the array LUNs changes during certain operations, such as Snapshot creation in the backend array.

Therefore, to maintain consistent behavior of VNX2 array LUNs, ONTAP treats these LUNs as ALUA, irrespective of whether these array LUNs are presented from traditional RAID groups or from dynamic disk pools.

In such a situation, the I/O requests for a given VNX2 LUN are distributed only over the paths reported to be in `INU` (in-use) state and not over all the paths reported as `AO` (active-optimized).

For example, if you have four paths to a VNX2 array LUN presented to the ONTAP system from a traditional RAID group, all of the paths are reported as `AO`; however, only two of the paths are in `INU` state, while the two other paths are not used but are in `RDY` (ready) state.

EMC VNX storage array families

ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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

The following EMC VNX storage array families are supported with ONTAP systems:

- Family 1: VNX1
- Family 2: VNX2

These arrays use the asymmetric active-active (ALUA) failover mode.

The Interoperability Matrix is the final authority for information about the storage arrays supported for each vendor.

Related information

NetApp Interoperability Matrix Tool

EMC Symmetrix storage arrays

You must meet specific requirements when configuring storage arrays to work with ONTAP systems that use array LUNs. These requirements include setting configuration parameters on your storage arrays and deploying only supported configurations.

Required Parameter settings on the storage array for ONTAP systems

Certain parameter settings are required on the storage array for the storage array to work successfully with ONTAP systems.

Required host channel director port configuration parameters

The host channel director port configuration parameters that must be set on the storage array are shown in the following table:

Parameter (names might differ between GUI and CLI)	Setting
`Common SN` (Common Serial Number or C-bit parameter)	Enable
`PP` (Point-to-Point parameter)	Enable
`SC3` (SCSI-3)	Enable
`SPC-2` (SCS2_Protocol_version, SCSI Primary Command 2 parameter, or `Allow inquiry data to be compiled to the standard`)	Enable

Parameter (names might differ between GUI and CLI)	Setting
UWN (Unique Worldwide Name)	Enable
Volume Set Addressing	Disable

The `Volume Set Addressing` parameter must be set the same way on all channel director ports to which the LUN is mapped. If the settings are different, ONTAP reports this as a LUN ID mismatch in `storage errors show` output and in an EMS message.

Related information

[FlexArray virtualization installation requirements and reference](#)

Requirements for implementing LUN security on EMC Symmetrix storage arrays

You use LUN security to eliminate the possibility of a host writing data to a LUN that is not owned by that host.

To eliminate the possibility of a non-ONTAP host overwriting EMC Symmetrix array LUNs owned by an ONTAP system or vice versa, you must present the Symmetrix logical devices through the host (channel) director ports in one of the following ways:

- Present only the Symmetrix logical devices for ONTAP on specific Symmetrix host (channel) director ports that are dedicated to ONTAP use.

  If ports cannot be dedicated to ONTAP, you should confirm that all other hosts using those ports are compatible with ONTAP requirements. This is because each host connected to the Symmetrix array has requirements for different port attribute settings. Sharing ports between multiple hosts that are connected to the Symmetrix storage array might result in an impossible-to-implement configuration.

- For VMAX storage arrays, create masking views for required mapping and masking by creating port groups, Storage Groups, and initiator groups.

  To achieve this, you must first enable the ACLX port attribute on the VMAX storage array ports.

ⓘ Do not present the VCMDB LUN to all hosts by default. Configure the global setting to restrict visibility to the VCMDB unless it has been specifically made visible to a particular host.

Caution about using the VCMDB LUN

For the VCMDB (Volume Configuration Management Database) to be enabled, the VCMDB LUN must exist. The VCMDB LUN is a `command` type LUN, not a storage LUN. The VCMDB is typically mapped to LUN 0, but can be mapped to an array LUN other than LUN 0.

If the VCMDB LUN is mapped to an ONTAP system, ONTAP periodically logs a message that the VCMDB LUN is less than the minimum size required, and it marks the VCMDB LUN as failed. The ONTAP system continues to function normally after logging this error message, but it cannot use the LUN.

A VCMDB LUN should be unmapped from an ONTAP system.

Caution about using the ACLX LUN

On VMAX arrays, the ACLX (Access Control Logix) LUN is created during initialization if the customer requests that ACLX be used. The ACLX LUN is not a storage LUN, so it should not be mapped to ONTAP.

If the ACLX LUN is mapped to an ONTAP system, ONTAP logs a message that the ACLX LUN is less than the minimum size required, and it marks the LUN as failed. The ONTAP system continues to function normally after logging this error message, but it cannot use the LUN.

The ACLX LUN should be unmapped from the front-end director ports on the ONTAP system.

Restriction on using gatekeeper LUNs

If a gatekeeper logical device (LUN) is presented, you must not map it to the ONTAP system. ONTAP systems cannot use gatekeeper LUNs. A gatekeeper LUN is a Symmetrix logical device through which SYMAPI or the ControlCenter agent communicates with the storage array.

EMC Symmetrix storage array families

ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

The following EMC Symmetrix storage array families are supported with ONTAP systems:

- Family 1: VMAX
- Family 2: VMAX3

The Interoperability Matrix is the final authority for information about the storage arrays supported for each vendor.

Related information

NetApp Interoperability Matrix Tool

Hitachi storage arrays

You must meet certain requirements when configuring Hitachi storage arrays to work with ONTAP systems that use array LUNs. These requirements include setting configuration parameters on storage arrays and deploying only supported configurations.

Required parameters for Hitachi storage arrays with ONTAP systems

Certain parameter settings are required on the storage array for the storage array to work successfully with ONTAP systems.

Required system parameters for AMS and HUS storage arrays

The parameters that must be set are listed in the following table:

Parameter	Setting
`Mapping mode`	Enabled
`Host group security`	Enabled
`Mode Setting > Common Setting`	Standard Mode
`Platform`	Not specified
`Middleware`	Not specified

Required system parameters for VSP and HUS VM storage arrays

The parameters that must be set are listed in the following table:

Parameter	Setting
`Host mode`	Standard or 00
`Port security`	Enabled

> ⓘ A host group that is separate from the default host groups should be created for each initiator-target port pair.

The storage array parameters should be set and LUNs should be mapped to the host groups before the switch is zoned to connect the ONTAP system to the Hitachi storage array.

If ONTAP FC initiator ports are zoned with Hitachi array target ports before the storage array parameters are set and the LUNs are mapped to the host groups, you might not be able to see any LUNs presented to the ONTAP interface.

In this situation, ONTAP generates one of the following EMS messages:

```
[fci.device.login.reject:info] The PLOGI transmitted by port XX was
rejected by port_id 0xYYYY with reason code 0x9 'Invalid R_CTL Field',
explanation code 0x29 'Insufficient Resources to Support Login'
```

```
[fci.device.login.reject:info]: The PLOGI transmitted by port XX was
rejected by port_id 0xYYYY with reason code 0x3 'Nx_Port Not Available,
Temporary', explanation code 0x29 'Insufficient Resources to Support
Login'
```

Workaround

If zoning was configured prior to setting the parameters and mapping the LUNs to the host groups, you can resolve the problem of LUNs not being visible by performing a path re-initialization to force a rediscovery of the array LUNs that are mapped from the Hitachi storage array. See Workaround when LUNs are not visible with Hitachi storage arrays section for the workaround to this problem.

Workaround when LUNs are not visible with Hitachi storage arrays

Workaround when LUNs are not visible with Hitachi storage arrays

The storage array parameters specified in the Required parameters for Hitachi storage arrays with systems running ONTAP section should be set and LUNs should be mapped to the host groups before the switch is zoned to connect the ONTAP system that uses array LUNs to a Hitachi storage array.

Required parameters for Hitachi storage arrays with systems running ONTAP

If zoning was configured prior to setting the parameters and mapping the LUNs to the host groups, you can resolve the problem of LUNs not being visible by performing a path reinitialization. The path reinitialization forces a rediscovery of the array LUNs that are mapped from a Hitachi storage array.

The procedure you use depends on whether aggregates with Hitachi array LUNs have already been configured on the ONTAP system. (In the following procedures, assume that Node-1 and Node-2 are the names of the ONTAP nodes.)

If aggregates with Hitachi array LUNs are already created on an ONTAP system that uses array LUNs

1. Run the following command to ensure that HA configuration is enabled on your ONTAP systems: `cluster ha show`

   The following output is displayed:

   ```
   High Availability Configured: true
   ```

2. Run the following command to ensure that storage failover is configured on your ONTAP system: `storage failover show`

   The following output is displayed:

```
           Takeover
 Node             Partner         Possible State Description
 -------------- -------------- --------
 --------------------------------------
 Node-1    Node-2    true     Connected to Node-2
 Node-2    Node-1    true     Connected to Node-1
 2 entries were displayed.
```

3. Run the following command to perform a takeover operation on Node-1: **storage failover takeover -ofnode Node-1**

4. Run the following command again to verify that Node-2 is ready for a giveback operation: **storage failover show**

   The following output is displayed:

```
           Takeover
 Node             Partner         Possible State Description
 -------------- -------------- --------
 --------------------------------------
 Node-2    Node-1    false    In takeover, Auto giveback will be
                                        initiated in 348 seconds
 Node-1    Node-2    -        Waiting for giveback (HA mailboxes)
 2 entries were displayed.
```

5. Run the following command to perform a giveback operation on Node-1: **storage failover giveback -ofnode Node-1 -require-partner-waiting true**

6. From other nodes in the cluster, repeat Step 3 through Step 5.

If aggregates with Hitachi array LUNs have not been created on an ONTAP system that uses array LUNs

1. Run the following command to list all the ONTAP FC initiator ports that are connected to the Hitachi storage array: **system node run -node <node name> -command "sysconfig -v"**

   For example: **system node run -node Node-1 -command "sysconfig -v"**

2. Run the following command to take all the listed ports offline: **system node run -node <node name> -command "priv set advanced;fcadmin offline <adapter>"**

   For example: **system node run -node Node-1 -command "priv set advanced;fcadmin offline 0a"**

3. Run the following command to bring all the listed ports online: **system node run -node <node name> -command "priv set advanced;fcadmin online <adapter>"**

   For example: **system node run -node Node-1 -command "priv set advanced;fcadmin**

```
online 0a"
```

Requirements for creating LUNs in dynamic pools

When creating LUNs in the dynamic pools, do not over-provision the pool capacity.

Requirements for formatting array LUNs

You can perform different types of array LUN formatting. Each type has an impact on the availability of LUNs or the entire storage array.

If you must present a new LUN to the ONTAP system immediately, you should use quick formatting. If you plan to present LUNs to the ONTAP system in the future, quick formatting is recommended but Online formatting is acceptable.

Quick formatting is recommended because both Online and Offline formatting delay disk discovery until after the LUN formatting is complete, as described in the following list:

- Online formatting takes the array LUN that is being formatted completely offline, and the array LUN does not respond to any SCSI commands.

  Formatting is performed on one array LUN at a time until all array LUNs are formatted.

- Offline formatting takes the entire storage array offline and the storage array does not respond to any SCSI commands.

  The array LUNs are formatted six at a time per controller until the formatting is complete.

  The storage array does not respond until all formatting is complete.

- If you are creating array LUNs (for any host) after initial setup of the ONTAP system and you are using Offline formatting, the system panics.

  Any array LUNs that are mapped to the ONTAP system, including array LUNs in the root volume, become unavailable until formatting is complete.

- Online and quick formatting of LUNs for other hosts do not impact the ONTAP system.

Requirements for using disks external to the storage array

ONTAP supports using the same disks behind the storage arrays that the vendor supports (that is, disks external to the storage array). When you are deploying storage arrays that use external disks, ONTAP accepts the configuration between the storage array and the external disks.

The following are restrictions and best practice recommendations for deploying storage arrays that use external disks:

- Root volumes and spare core LUNs are supported on external disks.
- Array LUNs from external and internal disks should not be in the same ONTAP aggregate.
- Array LUNs from SATA drives and array LUNs from FC drives should not be in the same aggregate.

- Rules for the number of paths to array LUNs on external disks are the same as for array LUNs on disks on the storage array.

  In the case of external disks, the paths go from the ONTAP system through the storage array to the external disks.

- Only block checksum is supported for array LUNs on external disks.

Related information

[FlexArray virtualization installation requirements and reference](#)

Guidelines for configuring ports on Hitachi storage arrays

You should follow certain guidelines when configuring ports on Hitachi storage arrays.

The guidelines for configuring ports on Hitachi storage arrays are as follows:

- You should add only one host group per port.
- You should ensure that a host group includes all the FC initiator ports of the V-Series systems in an array LUN neighborhood.

Hitachi storage array families

ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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

The following list shows the Hitachi storage array families supported with ONTAP systems:

- Family 1: VSP
- Family 2: VSP Gx00
- Family 3: AMS 2x00
- Family 4: HUS
- Family 5: HUS VM

The Interoperability Matrix is the final authority for information about the storage arrays supported for each vendor.

Related information

[NetApp Interoperability Matrix Tool](#)

HP EVA storage arrays

You must meet certain requirements when configuring storage arrays to work with ONTAP systems. These requirements include setting configuration parameters on storage arrays and deploying only supported configurations.

Required parameters for HP EVA arrays with ONTAP systems

Certain parameter settings are required on the storage array for the storage array to work successfully with ONTAP systems.

System parameter settings are provided in the following table:

Parameter	Setting
Redundancy	Any RAID level except VRAID0
Preferred path/mode	No preference (default)
Host type	SUN Solaris

Requirements for configuring the HP EVA Storage Management Server

Sharing FC initiator ports present on an ONTAP system and initiator ports present on an HP EVA Storage Management Server with common EVA target ports causes compatibility issues.

The compatibility issues result from the different host settings for the FC initiator ports present on an ONTAP system and the EVA Storage Management Server initiator ports.

HP recommends that any HP EVA Storage Management Server that resides in a fabric be configured in a zone that is separate from all operating systems.

Preparing HP EVA storage arrays for NDU

HP EVA storage arrays have a tendency to bring target ports offline during a firmware upgrade. ONTAP provides commands that increase ONTAP resiliency so that the ONTAP system is not disrupted if target ports go offline during the firmware upgrade. These commands can be used only with storage systems running ONTAP.

You must follow the guidelines provided by HP for performing the firmware upgrade, including guidelines about the load recommendation (HP EVA). You can upgrade only to firmware that HP supports. This procedure pertains to the HP EVA firmware upgrade type *online upgrade*.

ⓘ   During the firmware upgrade, both HP EVA controllers are rebooted.

This procedure instructs you to use ONTAP commands that increase ONTAP resiliency throughout the firmware upgrade. After the firmware upgrade is complete, you use the ONTAP commands again to return the

ports to their normal operating mode.

Steps

1. On the ONTAP system, set the privilege level of the command session to advanced: `set -privilege advanced`

2. Set the `is-upgrade-pending` parameter to **true** for the HP EVA storage array that is undergoing a firmware upgrade: `storage array modify -name array_name -is-upgrade-pending true`

3. On the storage array, start the firmware upgrade.

4. After the firmware upgrade is complete, set the privilege level to advanced again, if necessary, and then on the ONTAP system, set the `is-upgrade-pending` parameter to **false** to return the storage array ports to normal operation: `storage array modify -name array_name -is-upgrade-pending false`

   If you did not exit advanced mode in Step 2 you do not need to access it again.

   The `is-upgrade-pending` parameter contains the following features to ensure proper use of the command:

   - If you attempt to set `is-upgrade-pending` to **false** while the upgrade is in progress on the storage array, the command fails and an EMS message is returned.

   - If the `is-upgrade-pending` status is not returned to **false** within 60 minutes from the time it is set to **true**, EMS messages are logged hourly until `is-upgrade-pending` is returned to **false**.

ALUA support for HP EVA storage arrays

Support has been added for asymmetric logical unit access (ALUA) with HP EVA arrays.

By default, support for ALUA is enabled in ONTAP as well as on all the HP EVA storage arrays.

HP EVA storage array families

ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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

The following list shows the HP EVA storage arrays organized by family:

- Family 1: P6xxxx
- Family 2: HP EVA x100
- Family 3 : HP EVA x400

The Interoperability Matrix is the final authority for information about the storage arrays supported for each vendor.

Related information

HP XP storage arrays

You must meet certain requirements when configuring storage arrays to work with ONTAP systems that use array LUNs. These requirements include setting configuration parameters on storage arrays and deploying only supported configurations.

Required parameters for HP XP storage arrays with ONTAP

Certain parameter settings are required on the storage array for this storage array to work successfully with ONTAP systems.

Required system parameters for HP XP storage arrays

The required parameter settings are shown in the following table:

Parameter	Setting
Host mode	Standard or 00
Port security	Enabled

> ⓘ A host group that is separate from the default host groups should be created for each initiator-target port pair.

The storage array parameters should be set and LUNs should be mapped to the host groups before the switch is zoned to connect the ONTAP system to an HP XP storage array.

If ONTAP FC initiator ports are zoned with HP XP array target ports before the storage array parameters are set and the LUNs are mapped to the host groups, you might not be able to see any LUNs presented to the ONTAP interface.

In this situation, ONTAP generates one of the following EMS messages:

```
[fci.device.login.reject:info] The PLOGI transmitted by port XX was
rejected by port_id 0xYYYY with reason code 0x9 'Invalid R_CTL Field',
explanation code 0x29 'Insufficient Resources to Support Login'
```

```
[fci.device.login.reject:info]: The PLOGI transmitted by port XX was
rejected by port_id 0xYYYY with reason code 0x3 'Nx_Port Not Available,
Temporary', explanation code 0x29 'Insufficient Resources to Support
Login'
```

Workaround

If zoning was configured prior to setting the parameters and mapping the LUNs to the host groups, you can resolve the problem of LUNs not being visible by performing a path re-initialization to force a rediscovery of the array LUNs that are mapped from the HP XP storage array.

Workaround when LUNs are not visible with HP XP storage arrays

Workaround when LUNs are not visible with HP XP storage arrays

The storage array parameters specified in the Required parameters for HP XP storage arrays with ONTAP systems section should be set and LUNs should be mapped to the host groups before the switch is zoned to connect the ONTAP system that uses array LUNs to an HP XP storage array.

Required parameters for HP XP storage arrays with ONTAP systems

If zoning was configured prior to setting the parameters and mapping the LUNs to the host groups, you can resolve the problem of LUNs not being visible by -performing a path reinitialization. The path reinitialization forces a rediscovery of the array LUNs that are mapped from a HP XP storage array.

The procedure you use depends on whether aggregates with HP XP array LUNs have already been configured on the ONTAP system. (In the following procedures, assume that Node-1 and Node-2 are the names of the ONTAP nodes.)

If aggregates with HP XP array LUNs are already created on an ONTAP system that uses array LUNs

1. Run the following command to ensure that HA configuration is enabled on your ONTAP systems: `cluster ha show`

    The following output is displayed:

    ```
    High Availability Configured: true
    ```

2. Run the following command to ensure that storage failover is configured on your ONTAP systems: `storage failover show`

    The following output is displayed:

    ```
            Takeover
     Node           Partner        Possible State Description
     -------------- -------------- --------
     ------------------------------------
     Node-1     Node-2     true     Connected to Node-2
     Node-2     Node-1     true     Connected to Node-1
     2 entries were displayed.
    ```

3. Perform a takeover operation on Node-1 by running the following command: `storage failover takeover -ofnode Node-1`

4. Run the following command again to verify that Node-2 is ready for a giveback operation: **storage failover show**

   The following output is displayed:

```
         Takeover
Node            Partner        Possible State Description
-------------- -------------- --------
--------------------------------------
Node-2    Node-1    false    In takeover, Auto giveback will be
                                      initiated in 348 seconds
Node-1    Node-2    -         Waiting for giveback (HA mailboxes)
2 entries were displayed.
```

5. Perform a giveback operation on Node-1 by running the following command: **storage failover giveback -ofnode Node-1 -require-partner-waiting true**

6. From other nodes in the cluster, repeat Step 3 through Step 5.

If aggregates with HP XP array LUNs have not been created on an ONTAP systems that uses array LUNs

1. Run the following command to list all the ONTAP FC initiator ports that are connected to the HP XP storage array: **system node run -node <node name> -command "sysconfig -v"**

   For example: **system node run -node Node-1 -command "sysconfig -v"**

2. Run the following command to take all the listed ports offline: **system node run -node <node name> -command "priv set advanced;fcadmin offline <adapter>"**

   For example: **system node run -node Node-1 -command "priv set advanced;fcadmin offline 0a"**

3. Run the following command to bring all the listed ports online: **system node run -node <node name> -command "priv set advanced;fcadmin online <adapter>"**

   For example: **system node run -node Node-1 -command "priv set advanced;fcadmin online 0a"**

Requirements for creating LUNs in dynamic pools

When creating LUNs in the dynamic pools, do not over-provision the pool capacity.

Requirements for formatting array LUNs

You can perform different types of array LUN formatting. Each type has an impact on the availability of LUNs or the entire storage array.

If you must present a new LUN to the ONTAP system immediately, you should use quick formatting. If you plan

to present LUNs to the ONTAP system in the future, quick formatting is recommended but Online formatting is acceptable.

Quick formatting is recommended because both Online and Offline formatting delay disk discovery until after the LUN formatting is complete, as described in the following list:

- Online formatting takes the array LUN that is being formatted completely offline, and the array LUN does not respond to any SCSI commands.

  Formatting is performed on one array LUN at a time until all array LUNs are formatted.

- Offline formatting takes the entire storage array offline and the storage array does not respond to any SCSI commands.

  The array LUNs are formatted six at a time per controller until the formatting is complete.

  The storage array does not respond until all formatting is complete.

- If you are creating array LUNs (for any host) after initial setup of the ONTAP system and you are using Offline formatting, the system panics.

  Any array LUNs that are mapped to the ONTAP system, including array LUNs in the root volume, become unavailable until formatting is complete.

- Online and quick formatting of LUNs for other hosts do not impact the ONTAP system.

Requirements for using disks external to the storage array

ONTAP supports using the same disks behind the storage arrays that the vendor supports (that is, disks external to the storage array). When you are deploying storage arrays that use external disks, ONTAP accepts the configuration between the storage array and the external disks.

The following are restrictions and best practice recommendations for deploying storage arrays that use external disks:

- Root volumes and spare core LUNs are supported on external disks.
- Array LUNs from external and internal disks should not be in the same ONTAP aggregate.
- Array LUNs from SATA drives and array LUNs from FC drives should not be in the same aggregate.
- Rules for the number of paths to array LUNs on external disks are the same as for array LUNs on disks on the storage array.

  In the case of external disks, the paths go from the ONTAP system through the storage array to the external disks.

- Only block checksum is supported for array LUNs on external disks.

Related information

FlexArray virtualization installation requirements and reference

Guidelines for configuring ports on HP XP storage arrays

You should follow certain guidelines when configuring ports on HP XP storage arrays.

The guidelines for configuring ports on HP XP storage arrays are as follows:

- You should add only one host group per port.
- You should ensure that a host group includes all the FC initiator ports of the V-Series systems in an array LUN neighborhood.

HP XP storage array families

ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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

The following list shows the storage arrays organized by family:

The following list shows the HP XP storage array family supported with ONTAP systems:

- Family 1
    - P9500
    - XP7

The Interoperability Matrix is the final authority for information about the storage arrays supported for each vendor.

Related information

NetApp Interoperability Matrix Tool

IBM DS storage arrays

You must meet certain requirements when configuring storage arrays to work with ONTAP systems that use array LUNs. These requirements include setting configuration parameters on storage arrays and deploying only supported configurations.

Required settings for IBM DS storage arrays with ONTAP systems

Certain parameter settings are required on the storage array for the storage array to work successfully with ONTAP systems.

Required host type for DS8xxx storage arrays

The required host type setting is shown in the following table:

Parameter	Setting
Host type	N series Gateway

Requirements for configuring DS8300 9A2 LPAR models

When setting up the DS8300 9A2 LPAR (system logical partition) model to interact with ONTAP systems, you must set up access to each array LUN so that the redundant paths are both accessing the same LPAR.

Requirements for volume groups on DS8xxx storage arrays

You must be aware of limitations on the number of volume groups that you can use with a LUN group when configuring DS8xxx arrays.

You must use a single volume group for each LUN group on a DS8xxx storage array to guarantee that the DS8xxx array LUNs are consistently presented to all the FC initiators (that are on the ONTAP system) that access them.

ⓘ   If array LUNs are not presented consistently, there is a potential for data corruption.

IBM DS storage array families

ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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

The following list shows arrays organized by family:

- Family 1
  - DS8300
  - DS8100
- Family 2
  - DS8800
  - DS8700
  - DS8870

The Interoperability Matrix is the final authority for information about the storage arrays supported for each

vendor.

IBM XIV storage arrays

You must meet certain requirements when configuring storage arrays to work with ONTAP systems that use array LUNs. These requirements include setting configuration parameters on storage arrays and deploying only supported configurations.

Requirements for configuring IBM XIV arrays with ONTAP systems

Certain parameter settings are required on the storage array for the storage array to work successfully with ONTAP systems.

System parameter settings

The following table lists the system parameter settings for IBM XIV Gen3 arrays:

Parameter	Setting
`Type`	Default
`Host`	User-supplied host name
`Cluster`	User-supplied cluster name

IBM XIV storage array families

ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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.

IBM XIV Gen3 storage arrays are supported with ONTAP systems.

The Interoperability Matrix is the final authority for information about the storage arrays supported for each vendor.

3PAR storage arrays

You must meet certain requirements when configuring storage arrays to work with ONTAP systems that use array LUNs. These requirements include setting configuration parameters on storage arrays and deploying only supported configurations.

Required parameters for 3PAR arrays with ONTAP systems

Certain parameter settings are required on the storage array for the storage array to work successfully with ONTAP systems.

Host persona settings

For 3PAR arrays with firmware 2.3.1 and later, the required host persona values must be set as shown in the following table:

Connection type	Host persona
Direct connect	10
Fabric attached	10

The host definition, with the required host persona, should be created on a 3PAR storage array before the switch is zoned to connect the ONTAP system to a 3PAR storage array. If zoning is configured before the host definition, ONTAP discovers a 3PAR LUN as LUN 254 and product type SES, which ONTAP cannot use. The result is that newly mapped 3PAR LUNs cannot be assigned to ONTAP system and some 3PAR LUNs that might have already been added to an aggregate cannot be used because ONTAP sees a LUN 254 instead of a 3PAR LUN. If this problem occurs, ONTAP generates the following EMS message:

```
mlm.array.unknown:warning - Array LUN XXXXXXXX.126L254 [S/N 'XXXXXXX'
vendor '3PARdata' product 'SES'] is unknown and is not supported in this
version of Data ONTAP.
```

If zoning was configured prior to the host definition, you can resolve the problem by doing path re-initialization to force a rediscovery of the array LUNs that are mapped from the 3PAR storage array. See the Workaround for the unknown device (L254/SES) problem with 3PAR storage arrays section for the workaround for this problem.

Workaround for the unknown device (L254/SES) problem with 3PAR storage arrays

Port persona settings

For 3PAR arrays with firmware 2.2.x, the required port persona values must be set as shown in the following table:

Connection type	Port persona
Direct connect	18
Fabric attached	19

Workaround for the unknown device (L254/SES) problem with 3PAR storage arrays

The host definition, with the host persona specified in the required parameters for 3PAR arrays with ONTAP systems section, should be created on a 3PAR storage array before

the switch is zoned to connect the ONTAP system to a 3PAR storage array. If zoning is configured before the host definition, ONTAP discovers a 3PAR array LUN as LUN 254 and product type SES, which ONTAP cannot use.

Required parameters for 3PAR arrays with ONTAP systems

When ONTAP discovers 3PAR array LUNs as LUN 254 and product type SES, newly mapped 3PAR LUNs cannot be assigned to ONTAP systems and some 3PAR LUNs that might have already been added to an aggregate cannot be used because ONTAP sees a LUN 254. ONTAP generates EMS messages for the LUN it discovered as LUN 254 and product type SES, reporting some LUNs as unknown devices.

If zoning was configured before the host definition, you can resolve the LUN 254/SES unknown device problem by forcing a rediscovery of the array LUNs that are mapped from the 3PAR storage array. The procedure you use depends on whether aggregates with 3PAR array LUNs have already been configured on the ONTAP system. (In the following procedures, assume that Node-1 and Node-2 are the names of the ONTAP nodes.)

If aggregates with 3PAR array LUNs were already created on an ONTAP system

1. Run the following command to ensure that an HA configuration is enabled on your ONTAP systems:
   cluster ha show

   The following output is displayed:

   ```
   High Availability Configured: true
   ```

2. Run the following command to ensure that storage failover is configured on your ONTAP systems:
   storage failover show

   The following output is displayed:

   ```
             Takeover
   Node            Partner        Possible State Description
   -------------- -------------- --------
   --------------------------------------
   Node-1    Node-2    true     Connected to Node-2
   Node-2    Node-1    true     Connected to Node-1
   2 entries were displayed.
   ```

3. Perform a takeover operation on Node-1 by running the following command: **storage failover takeover -ofnode Node-1**

4. Run the following command again to verify that Node-2 is ready for a giveback operation: **storage failover show**

   The following output is displayed:

```
          Takeover
Node           Partner        Possible State Description
-------------- -------------- --------
--------------------------------------
Node-2    Node-1    false     In takeover, Auto giveback will be
                                      initiated in 348 seconds
Node-1    Node-2    -         Waiting for giveback (HA mailboxes)
2 entries were displayed.
```

5. Perform a giveback operation on Node-1 by running the following command: **`storage failover giveback -ofnode Node-1 -require-partner-waiting true`**

6. From other nodes in the cluster, repeat Step 3 through Step 5.

If aggregates with 3PAR array LUNs have not been created on an ONTAP system

1. Run the following command to list all of the ONTAP FC initiator ports that are connected to the 3PAR storage array: **`system node run -node <node name> -command "sysconfig -v"`**

   The following example shows the command with a specific node name included: **`system node run -node Node-1 -command "sysconfig -v"`**

2. Run the following command to take all of the listed ports offline: **`system node run -node <node name> -command "priv set advanced;fcadmin offline <adapter>"`**

   The following example shows the command with a specific node name and adapter included: **`system node run -node Node-1 -command "priv set advanced;fcadmin offline 0a"`**

3. Run the following command to bring all of the listed ports online: **`system node run -node <node name> -command "priv set advanced;fcadmin online <adapter>"`**

   The following example shows the command with a specific node name and adapter included: **`system node run -node Node-1 -command "priv set advanced;fcadmin online 0a"`**

3PAR storage array families

ONTAP does not support mixing some types of storage in aggregates. To help you determine the array LUNs that can be mixed in an aggregate, the storage arrays from each vendor are grouped into families. When you create aggregates, you cannot mix array LUNs from different vendors and from different storage array families in the same aggregate.

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

The following list shows the supported 3PAR storage arrays organized by different families:

- Family 1: 3PAR InServ Storage
  - Tx00
- Family 2: 3PAR InServ Storage
  - Fx00
- Family 3: 3PAR StoreServ Storage
  - 20000
  - 208xx

The Interoperability Matrix is the final authority for information about the storage arrays supported for each vendor.

Related information

[NetApp Interoperability Matrix Tool](#)

Supported fabric-attached configurations

Various fabric-attached configurations are supported for all the storage array vendors unless mentioned otherwise.

The Interoperability Matrix contains additional information about specific array models.

Related information

[NetApp Interoperability Matrix Tool](#)

Stand-alone basic configuration

The stand-alone basic configuration of an ONTAP system that uses array LUNs is a simple, fabric-attached configuration with a single FC initiator port pair accessing a single LUN group.

This configuration is supported for use with all storage arrays listed in the Interoperability Matrix as supported for the release of ONTAP running on your system.

The following illustration shows this configuration:

Related information

[NetApp Interoperability Matrix Tool](#)

Stand-alone system with two 2-port array LUN groups

In a stand-alone ONTAP system that is in a fabric-attached simple configuration, each FC initiator port pair present on the ONTAP system accesses a separate array LUN group.

This configuration is supported for use with all storage arrays listed in the Interoperability Matrix as supported for the release of ONTAP running on your system.

The following illustration shows fabric-attached simple configuration:
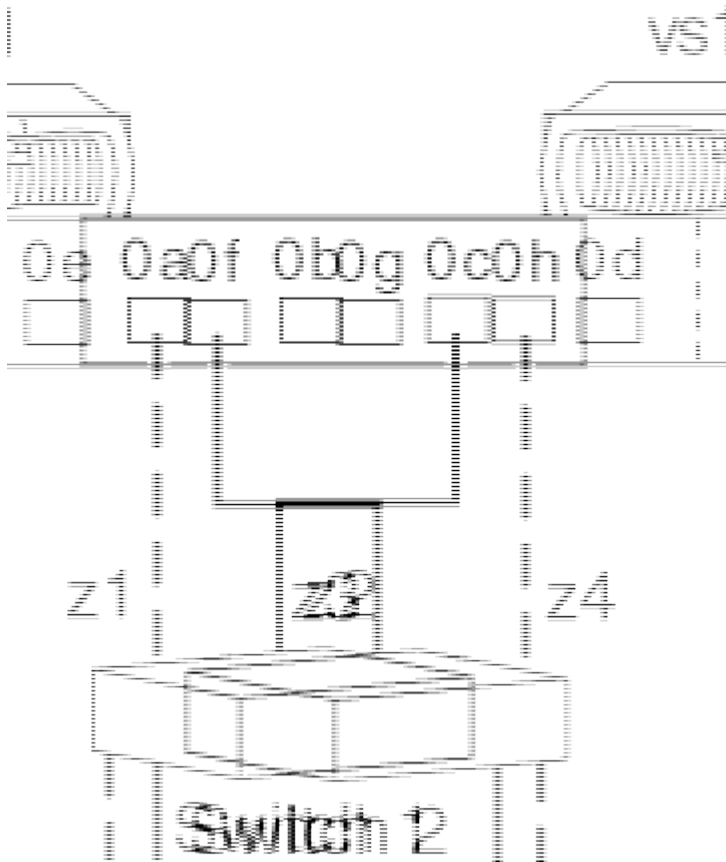
Related information

NetApp Interoperability Matrix Tool

Single 4-port array LUN group configuration

This configuration contains a single, 4-port LUN group with each target port accessed by a single ONTAP FC initiator port from the HA pair. Due to zoning, only two paths are allowed to a specific array LUN from each of the ONTAP system.

This configuration is supported for use with all storage arrays listed in the Interoperability Matrix as supported for the release of ONTAP running on your system.

The following illustration shows this configuration:
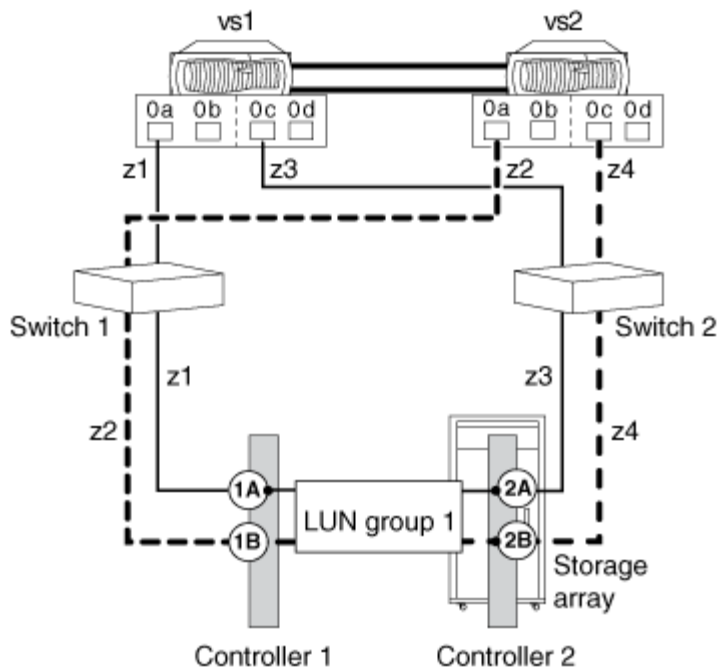
Related information

Two 4-port array LUN group configuration

In this configuration, each ONTAP FC initiator port pair accesses a separate array LUN group. The zoning is a single ONTAP FC initiator to a single array target port.

This configuration is supported for use with all storage arrays listed in the Interoperability Matrix as supported for the release of ONTAP running on your system.

The following illustration shows a block diagram of this configuration:
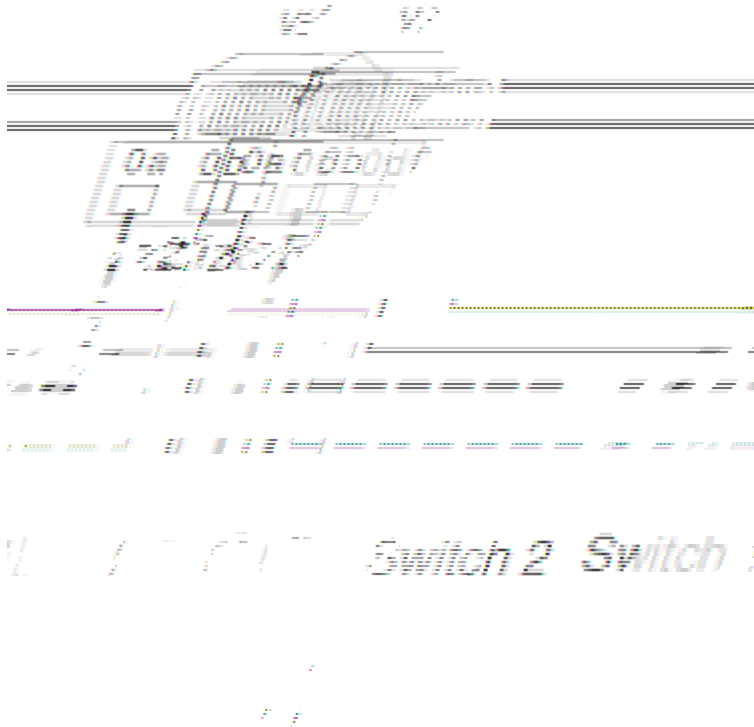
Related information

Eight-port array LUN group configuration

An eight-port LUN group configuration is supported on clustered V-Series systems and on ONTAP systems that can use array LUNs.

This configuration is supported for use with all storage arrays listed in the Interoperability Matrix as supported for the release of ONTAP running on your system.

There are two ways of deploying this configuration: crossed and uncrossed back-end connections.

Crossed back-end connections

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

In this illustration of crossed back-end connections, note how the nodes 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. This optimizes the use of switch ports and array ports, which reduces the impact of a switch or storage array controller failure.
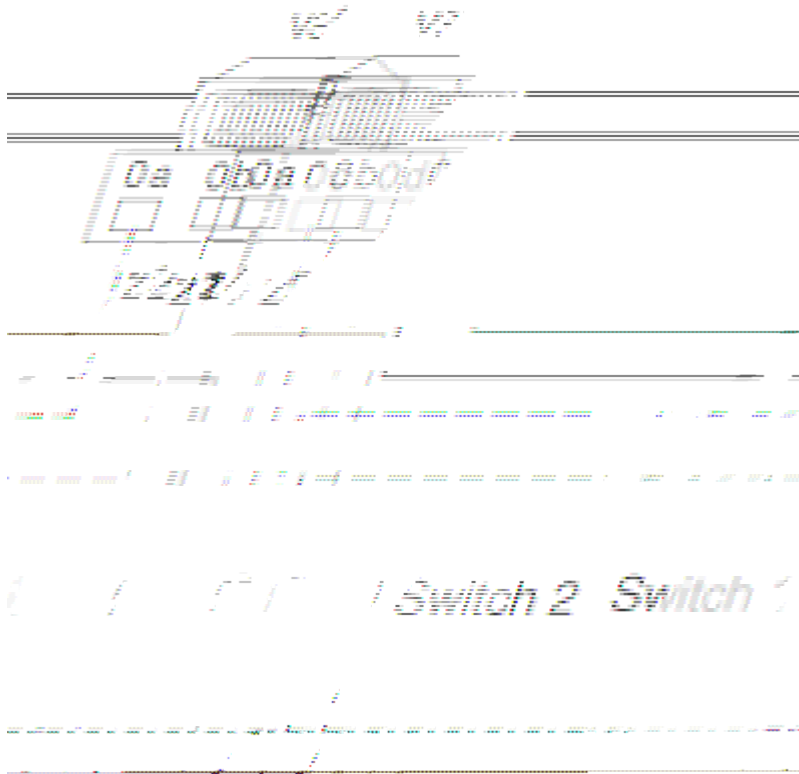
Back-end connections are not crossed

In this 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 this configuration when the back-end connections are not crossed.
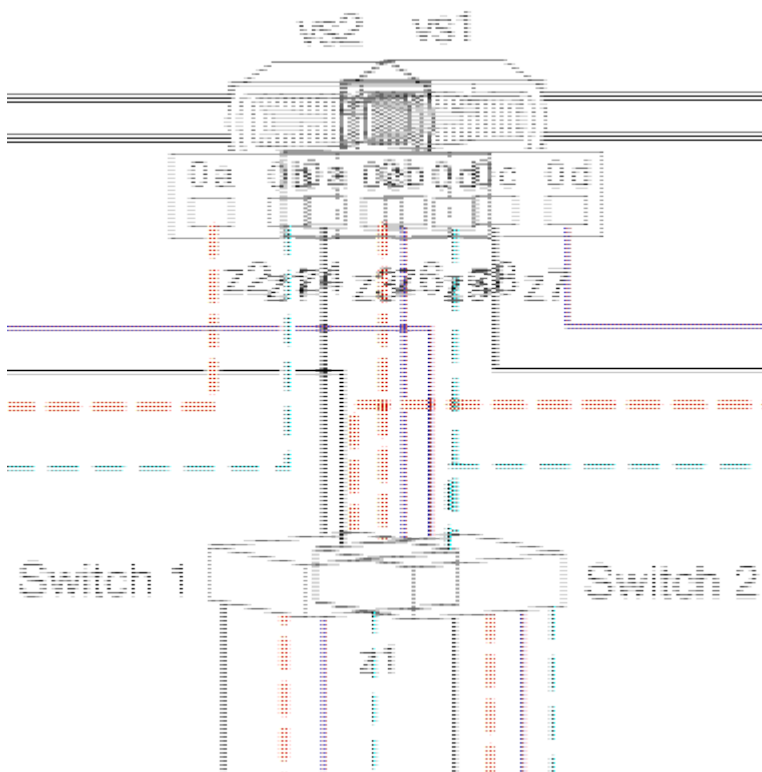
Related information

NetApp Interoperability Matrix Tool

Fabric-attached MetroCluster installation and configuration

Attaching FC initiator port to multiple targets ports

You can connect an FC initiator port on an ONTAP system to multiple target ports on separate storage arrays within the same family. This configuration is supported with MetroCluster configurations in ONTAP.

ONTAP supports sharing an FC initiator port with multiple target ports in both HA pairs and in stand-alone systems. This configuration is supported for use with all storage arrays listed in the Interoperability Matrix as supported for the release of ONTAP running on your system.

> (i)   For details about sharing an FC initiator port with multiple target ports and zoning, see the information in the *FlexArray Virtualization Installation Requirements and Reference*

Single ONTAP FC initiator port connecting to target ports on separate storage arrays

The following example shows an HA pair in which a single ONTAP FC initiator port connects to multiple target ports on different storage arrays:

Related information

Sharing a target port with two FC initiator ports

A maximum of two ONTAP FC initiator ports across nodes can be attached to a single target port on the storage array. This configuration is supported with MetroCluster configurations in ONTAP.

This configuration is supported for use with all storage arrays listed in the Interoperability Matrix as supported for the release of ONTAP running on your system.

The following illustration shows a shared target port configuration with an HA pair. Sharing a target port with two FC initiator ports is supported with both stand-alone systems and HA pairs.

> ⓘ For details about sharing an FC initiator port with multiple target ports and zoning, see the *FlexArray Virtualization Installation Requirements and Reference*

Shared target port attached to FC initiator ports

The following example shows an HA pair in which a single target port connects to two FC initiator ports:
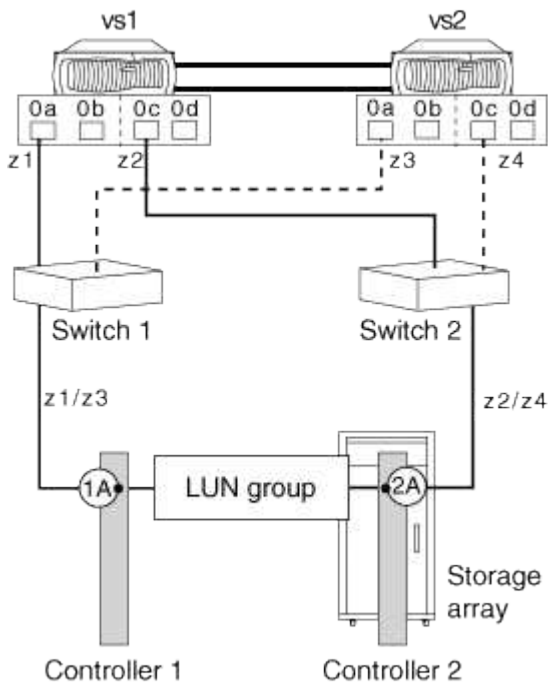
The initiator ports 0a of controllers vs1 and vs2 are connected to the storage array port 1A and ports 0c of the controllers are connected to the storage array port 2A.

Related information