



Planning for paths to array LUNs

ONTAP FlexArray

NetApp
September 20, 2021

This PDF was generated from https://docs.netapp.com/us-en/ontap-flexarray/install/concept_when_to_check_for_redundant_paths_to_array_luns.html on September 20, 2021. Always check docs.netapp.com for the latest.

Table of Contents

- Planning for paths to array LUNs 1
 - Requirements for redundant setup of components in a path 1
 - Required number of paths to an array LUN 2
 - Benefits of using multiple LUN groups 3
 - Array LUN name format 6
 - Pre-cluster array LUN name format 6
 - How the array LUN name is displayed in ONTAP 7
 - Valid pathing: stand-alone system with a single 2-port array LUN group 8
 - Valid pathing: a single 4-port array LUN group in a fabric-attached configuration 9
 - Valid pathing: eight-port array LUN group configuration 10

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

Zone	FC initiator port on the ONTAP system	Storage array
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. <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>
Fabric-attached	<code>node-name:switch-name:port.idlun-id</code>	<p><code>node-name</code> 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. <code>switch-name</code> is the name of the switch.</p> <p><code>port</code> is the switch port that is connected to the target port (the end point).</p> <p><code>id</code> is the device ID.</p> <p><code>lun-id</code> is the array LUN number that the storage array presents to hosts.</p> <p>Example: <code>node1:brocade3:6.126L1</code></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, it is helpful to check your command output against command output in this guide to verify that the number of LUN groups is what you intended and that there are redundant paths.

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

Zone	FC initiator port on the ONTAP system	Storage array
z3	vs1, Port 0b	Controller 1, Port 1C
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.

Copyright Information

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system- without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

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

Trademark Information

NETAPP, the NETAPP logo, 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.