



Troubleshoot

AFF and FAS Controller Upgrade

NetApp
October 01, 2021

Table of Contents

- Troubleshoot..... 1
 - Aggregate relocation failures 1
 - Reboots, panics, or power cycles 2
 - Issues that can arise in multiple stages of the procedure 6
 - LIF migration failure 7
 - LIFs are on invalid ports after upgrade 7

Troubleshoot

You might encounter a failure while upgrading the node pair. The node might crash, aggregates might not relocate, or LIFs might not migrate. The cause of the failure and its solution depend on when the failure occurred during the upgrade procedure.

Refer to the table describing the different phases of the procedure in the section [ARL upgrade workflow](#). The information about failures that can occur is listed by the phase of the procedure.

- [Aggregate relocation failures](#)
- [Reboots, panics, or power cycles](#)
- [Issues that can arise in multiple stages of the procedure](#)
- [LIF migration failure](#)
- [LIFs are on invalid ports after upgrade](#)

Aggregate relocation failures

Aggregate relocation (ARL) might fail at different points during the upgrade.

Check for aggregate relocation failure

During the procedure, ARL might fail in Stage 2, Stage 3, or Stage 5.

Steps

1. Enter the following command and examine the output:

```
storage aggregate relocation show
```

The `storage aggregate relocation show` command shows you which aggregates were successfully relocated and which ones were not, along with the causes of failure.

2. Check the console for any EMS messages.
3. Take one of the following actions:
 - Take the appropriate corrective action, depending on the output of the `storage aggregate relocation show` command and the output of the EMS message.
 - Force relocation of the aggregate or aggregates by using the `override-vetoes` option or the `override-destination-checks` option of the `storage aggregate relocation start` command.

For detailed information about the `storage aggregate relocation start`, `override-vetoes`, and `override-destination-checks` options, refer to [References](#) to link to the *ONTAP 9 Commands: Manual Page Reference*.

Aggregates originally on node1 are owned by node4 after completion of the upgrade

At the end of the upgrade procedure, node3 should be the new home node of aggregates that originally had

node1 as the home node. You can relocate them after the upgrade.

About this task

Aggregates might fail to relocate properly, having node1 as their home node instead of node3 under the following circumstances:

- During Stage 3, when aggregates are relocated from node2 to node3.
Some of the aggregates being relocated have node1 as their home node. For example, such an aggregate could be called `aggr_node_1`. If relocation of `aggr_node_1` fails during Stage 3, and relocation cannot be forced, then the aggregate will be left behind on node2.
- After Stage 4, when node2 is replaced with node4.
When node2 is replaced, `aggr_node_1` will come online with node4 as its home node instead of node3.

You can fix the incorrect ownership problem after Stage 6 once storage failover has been enabled by completing the following steps:

Steps

1. Enter the following command to get a list of aggregates:

```
storage aggregate show -nodes <node4> -is-home true
```

To identify aggregates that were not correctly relocated, refer to the list of aggregates with the home owner of node1 that you obtained in the section [Prepare the nodes for upgrade](#) and compare it with output of the above command.

2. Compare the output of [Step 1](#) with the output you captured for node1 in the section [Prepare the nodes for upgrade](#) and note any aggregates that were not correctly relocated.
3. Relocate the aggregates left behind on node4:

```
storage aggregate relocation start -node <node4> -aggr <aggr_node_1>  
-destination <node3>
```

Do not use the `-ndo-controller-upgrade` parameter during this relocation.

4. Enter the following command to verify that node3 is now the home owner of the aggregates:

```
storage aggregate show -aggregate <aggr1,aggr2,aggr3...> -fields home-name
```

`<aggr1,aggr2,aggr3...>` is the list of aggregates that had node1 as the original home owner.

Aggregates that do not have node3 as home owner can be relocated to node3 using the same relocation command in [Step 3](#).

Reboots, panics, or power cycles

The system might crash – reboot, panic or go through a power cycle – during different stages of the upgrade. The solution to these problems depends on when they occur.

Reboots, panics, or power cycles during Stage 2

Crashes can occur before, during, or immediately after Stage 2, during which you relocate aggregates from

node1 to node2, move data LIFs and SAN LIFs owned by node1 to node2, record node1 information, and retire node1.

Node1 or node2 crashes before Stage 2 with HA still enabled

If either node1 or node2 crashes before Stage 2, no aggregates have been relocated yet and the HA configuration is still enabled.

About this task

Takeover and giveback can proceed normally.

Steps

1. Check the console for EMS messages that the system might have issued, and take the recommended corrective action.
2. Continue with the node-pair upgrade procedure.

Node1 crashes during or just after Stage 2 with HA still enabled

Some or all aggregates have been relocated from node1 to node2, and HA is still enabled. Node2 will take over node1's root volume and any non-root aggregates that were not relocated.

About this task

Ownership of aggregates that were relocated looks the same as the ownership of non-root aggregates that were taken over because home owner has not changed.

When node1 enters the `waiting for giveback` state, node2 will give back all the node1 nonroot aggregates.

Steps

1. Complete [Step 1](#) in the section *Relocate non-root aggregates from node1 to node2* again.
2. Continue with the node-pair upgrade procedure.

Node1 crashes after Stage 2 while HA is disabled

Node2 will not take over but it is still serving data from all non-root aggregates.

Steps

1. Bring up node1.
2. Continue with the node-pair upgrade procedure.

You might see some changes in the output of the `storage failover show` command, but that is typical and does not affect the procedure. See the troubleshooting section [Unexpected storage failover show command output](#).

Node2 fails during or after Stage 2 with HA still enabled

Node1 has relocated some or all of its aggregates to node2. HA is enabled.

About this task

Node1 will take over all of node2's aggregates as well any of its own aggregates that it had relocated to node2. When node2 enters the `Waiting for Giveback` state, node1 gives back all of node2's aggregates.

Steps

1. Complete [Step 1](#) in the section *Relocate non-root aggregates from node1 to node2* again.
2. Continue with the node-pair upgrade procedure.

Node2 crashes after Stage 2 and after HA is disabled

Node1 will not take over.

Steps

1. Bring up node2.

A client outage will occur for all aggregates while node2 is booting up.

2. Continue with the rest of the node pair upgrade procedure.

Reboots, panics, or power cycles during Stage 3

Failures can occur during or immediately after Stage 3, during which you install and boot node3, map ports from node1 to node3, move data LIFs and SAN LIFs belonging to node1 and node2 to node3, and relocate all aggregates from node2 to node3.

Node2 crash during Stage 3 with HA disabled and before relocating any aggregates

Node3 will not take over following a node2 crash as HA is already disabled.

Steps

1. Bring up node2.

A client outage will occur for all aggregates while node2 is booting up.

2. Continue with the node-pair upgrade procedure.

Node2 crashes during Stage 3 after relocating some or all aggregates

Node2 has relocated some or all of its aggregates to node3, which will serve data from aggregates that were relocated. HA is disabled.

About this task

There will be client outage for aggregates that were not relocated.

Steps

1. Bring up node2.
2. Relocate the remaining aggregates by completing [Step 1](#) through [Step 3](#) in the section *Relocate non-root aggregates from node2 to node3*.
3. Continue with the node-pair upgrade procedure.

Node3 crashes during Stage 3 and before node2 has relocated any aggregates

Node2 does not take over but it is still serving data from all non-root aggregates.

Steps

1. Bring up node3.

2. Continue with the node-pair upgrade procedure.

Node3 crashes during Stage 3 during aggregate relocation

If node3 crashes while node2 is relocating aggregates to node3, node2 will abort the relocation of any remaining aggregates.

About this task

Node2 continues to serve remaining aggregates, but aggregates that were already relocated to node3 encounter client outage while node3 is booting.

Steps

1. Bring up node3.
2. Complete [Step 3](#) again in the section *Relocate non-root aggregates from node2 to node3*.
3. Continue with the node-pair upgrade procedure.

Node3 fails to boot after crashing in Stage 3

Because of a catastrophic failure, node3 cannot be booted following a crash during Stage 3.

Step

Contact technical support.

Node2 crashes after Stage 3 but before Stage 5

Node3 continues to serve data for all aggregates. The HA pair is disabled.

Steps

1. Bring up node2.
2. Continue with the node-pair upgrade procedure.

Node3 crashes after Stage 3 but before Stage 5

Node3 crashes after Stage 3 but before Stage 5. The HA pair is disabled.

Steps

1. Bring up node3.

There will be a client outage for all aggregates.
2. Continue with the node-pair upgrade procedure.

Reboots, panics, or power cycles during Stage 5

Crashes can occur during Stage 5, the stage in which you install and boot node4, map ports from node2 to node4, move data LIFs and SAN LIFs belonging to node2 from node3 to node4, and relocate all of node2's aggregates from node3 to node4.

Node3 crashes during Stage 5

Node3 has relocated some or all of node2's aggregates to node4. Node4 does not take over but continues to serve non-root aggregates that node3 already relocated. The HA pair is disabled.

About this task

There is be an outage for the rest of the aggregates until node3 boots again.

Steps

1. Bring up node3.
2. Relocate the remaining aggregates that belonged to node2 by repeating [Step 1](#) through [Step 3](#) in the section *Relocate node2's non-root aggregates from node3 to node4*.
3. Continue with the node pair upgrade procedure.

Node4 crashes during Stage 5

Node3 has relocated some or all of node2's aggregates to node4. Node3 does not take over but continues to serve non-root aggregates that node3 owns as well as those that were not relocated. HA is disabled.

About this task

There is an outage for non-root aggregates that were already relocated until node4 boots again.

Steps

1. Bring up node4.
2. Relocate the remaining aggregates that belonged to node2 by again completing [Step 1](#) through [Step 3](#) in *Relocate node2's non-root aggregates from node3 to node4*.
3. Continue with the node-pair upgrade procedure.

Issues that can arise in multiple stages of the procedure

Some issues can occur during different stages of the procedure.

Unexpected "storage failover show" command output

During the procedure, if the node that hosts all data aggregates panics or is rebooted accidentally, you might see unexpected output for the `storage failover show` command before and after the reboot, panic, or power cycle.

About this task

You might see unexpected output from the `storage failover show` command in Stage 2, Stage 3, Stage 4, or Stage 5.

The following example shows the expected output of the `storage failover show` command if there are no reboots or panics on the node that hosts all the data aggregates:


```
cluster::> storage failover show
```

Node	Partner	Takeover	
		Possible	State Description
node1	node2	false	Unknown
node2	node1	false	Node owns partner aggregates as part of the non-disruptive head upgrade procedure. Takeover is not possible: Storage failover is disabled.

The following example shows the output of the `storage failover show` command after a reboot or panic:

```
cluster::> storage failover show
```

Node	Partner	Takeover	
		Possible	State Description
node1	node2	-	Unknown
node2	node1	false	Waiting for node1, Partial giveback, Takeover is not possible: Storage failover is disabled

Although the output says that a node is in partial giveback and that storage failover is disabled, you can disregard this message.

Steps

No action is required; continue with the node-pair upgrade procedure.

LIF migration failure

After you migrate LIFs, they might not come online after migration in Stage 2, Stage 3, or Stage 5.

Steps

1. Verify that the port MTU size is the same as that of the source node.

For example, if the cluster port MTU size is 9000 on the source node, it should be 9000 on the destination node.

2. Check the physical connectivity of the network cable if the physical state of the port is "down".

LIFs are on invalid ports after upgrade

After the upgrade is completed, the FC logical interfaces (LIFs) might be left on incorrect ports if you have a MetroCluster configuration. You can perform a resync operation to reassign the LIFs to the correct ports.

Step

1. Enter the `metrocluster vserver resync` command to reallocate the LIFs to the correct ports.

```
metrocluster vserver resync -vserver <vserver_name> fcp-mc.headupgrade.test.vs
```

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.