



Troubleshoot

Upgrade controllers

NetApp
February 22, 2024

Table of Contents

- Troubleshoot..... 1
 - 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

Troubleshoot

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 [Overview of the ARL upgrade](#). Information about the failures that can occur is listed by the phase of the procedure.

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 node2 after completion of the upgrade

At the end of the upgrade procedure, node1 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 correctly, that is, they have node2 as their home node instead of node1, under the following circumstances:

- During Stage 3, when aggregates are relocated from node2 to node1.

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 is left behind on node2.

- After Stage 4, when node2 is replaced with the new system modules.

When node2 is replaced, `aggr_node_1` will come online with node1 as its home node instead of node2.

You can fix the incorrect ownership problem after Stage 6, after you have enabled storage failover by completing the following steps:

Steps

1. Get a list of aggregates:

```
storage aggregate show -nodes node2 -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 the 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 node2:

```
storage aggregate relocation start -node node2 -aggr aggr_node_1 -destination node1
```

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

4. Verify that node1 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 node1 as home owner can be relocated to node1 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 the pre-check phase

Node1 or node2 crashes before the pre-check phase with HA pair still enabled

If either node1 or node2 crashes before the pre-check phase, no aggregates have been relocated yet and the HA pair 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.

Reboots, panics, or power cycles during first resource-release phase

Node1 crashes during the first resource-release phase with HA pair still enabled

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

About this task

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

When node1 enters the `waiting for giveback` state, node2 gives back all of the node1 non- root aggregates.

Steps

1. After node1 is booted up, all the non-root aggregates of node1 have moved back to node1. You must perform a manual aggregate relocation of the aggregates from node1 to node2:

```
storage aggregate relocation start -node node1 -destination node2 -aggregate  
-list * -ndocontroller-upgrade true
```
2. Continue with the node-pair upgrade procedure.

Node1 crashes during the first resource-release phase while HA pair is disabled

Node2 does 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.

Node2 fails during the first resource-release phase with HA pair still enabled

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

About this task

Node1 takes over all of node2's aggregates as well as any of its own aggregates that it had relocated to node2. When node2 boots up, the aggregate relocation is completed automatically.

Steps

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

Node2 crashes during the first resource-release phase and after HA pair is disabled

Node1 does not take over.

Steps

1. Bring up node2.

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

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

Reboots, panics, or power cycles during the first verification phase

Node2 crashes during the first verification phase with HA pair disabled

Node1 does not take over following a node2 crash as the HA pair is already disabled.

Steps

1. Bring up node2.

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

2. Continue with the node-pair upgrade procedure.

Node1 crashes during the first verification phase with HA pair disabled

Node2 does 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.

Reboots, panics, or power cycles during first resource-regain phase

Node2 crashes during the first resource-regain phase during aggregate relocation

Node2 has relocated some or all of its aggregates from node1 to node1. Node1 serves data from aggregates that were relocated. The HA pair is disabled and hence there is no takeover.

About this task

There is client outage for aggregates that were not relocated. On booting up node2, the aggregates of node1 are relocated to node1.

Steps

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

Node1 crashes during the first resource-regain phase during aggregate relocation

If node1 crashes while node2 is relocating aggregates to node1, the task continues after node1 boots up.

About this task

Node2 continues to serve remaining aggregates, but aggregates that were already relocated to node1 encounter client outage while node1 is booting up.

Steps

1. Bring up node1.
2. Continue with the controller upgrade.

Reboots, panics, or power cycles during post-check phase

Node1 or node2 crashes during the post-check phase

The HA pair is disabled hence this is no takeover. There is a client outage for aggregates belonging to the node that rebooted.

Steps

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

Reboots, panics, or power cycles during second resource-release phase

Node1 crashes during the second resource-release phase

If node1 crashes while node2 is relocating aggregates, the task continues after node1 boots up.

About this task

Node2 continues to serve remaining aggregates but aggregates that were already relocated to node1 and node1's own aggregates encounter client outages while node1 is booting.

Steps

1. Bring up node1.
2. Continue with the controller upgrade procedure.

Node2 crashes during the second resource-release phase

If node2 crashes during aggregate relocation, node2 is not taken over.

About this task

Node1 continues to serve the aggregates that have been relocated, but the aggregates owned by node2 encounter client outages.

Steps

1. Bring up node2.
2. Continue with the controller upgrade procedure.

Reboots, panics, or power cycles during the second verification phase

Node1 crashes during the second verification phase

If node1 crashes during this phase, takeover does not happen because the HA pair is already disabled.

About this task

There is a client outage for all aggregates until node1 reboots.

Steps

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

Node2 crashes during the second verification phase

If node2 crashes during this phase, takeover does not happen. Node1 serves data from the aggregates.

About this task

There is an outage for non-root aggregates that were already relocated until node2 reboots.

Steps

1. Bring up node2.
2. 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.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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