



Understanding and using the Node Failover Planning page

Active IQ Unified Manager

NetApp
March 26, 2025

Table of Contents

- Understanding and using the Node Failover Planning page 1
- Using the Node Failover Planning page to determine corrective actions 1
- Components of the Node Failover Planning page 1
- Performance statistics grid 1
- Comparing pane 2
- Using a threshold policy with the Node Failover Planning page 3
- Using the Performance Capacity Used Breakdown chart for failover planning 3
- About this task 3
- Steps 4

Understanding and using the Node Failover Planning page

The Performance/Node Failover Planning page estimates the performance impact on a node if the node's high-availability (HA) partner node fails. Unified Manager bases the estimates on the historical performance of the nodes in the HA pair.

Estimating the performance impact of a failover helps you to plan in the following scenarios:

- If a failover consistently degrades the takeover node's estimated performance to an unacceptable level, you can consider taking corrective actions to reduce the performance impact due to a failover.
- Before initiating a manual failover to perform hardware maintenance tasks, you can estimate how the failover affects the performance of the takeover node in order to determine the best time to perform the task.

Using the Node Failover Planning page to determine corrective actions

Based on the information that is displayed in the Performance/Node Failover Planning page, you can take actions to ensure that a failover does not cause the performance of an HA pair to drop below an acceptable level.

For example, to reduce the estimated performance impact of a failover, you can move some volumes or LUNs from a node in the HA pair to other nodes in the cluster. Doing so ensures that the primary node can continue to deliver acceptable performance after a failover.

Components of the Node Failover Planning page

The components of the Performance/Node Failover Planning page are displayed in a grid and in the Comparing pane. These sections enable you to assess the impact of a node failover on the performance of the takeover node.

Performance statistics grid

The Performance/Node Failover Planning page displays a grid containing statistics for latency, IOPS, utilization, and performance capacity used.



Latency and IOPS values displayed in this page and in the Performance/Node Performance Explorer page might not match because different performance counters are used to calculate the values to predict node failover.

In the grid, each node is assigned one of the following roles:

- Primary

The node that takes over for the HA partner when the partner fails. The root object is always the Primary node.

- Partner

The node that fails in the failover scenario.

- Estimated Takeover

The same as the Primary node. Performance statistics displayed for this node show the takeover node's performance after it takes over the failed partner.



Although the workload of the takeover node is equivalent to the combined workloads of both nodes after a failover, the statistics for the Estimated Takeover node are not the sum of the statistics of the Primary node and the Partner node. For example, if the latency of the Primary node is 2 ms/op and the latency of the Partner node is 3 ms/op, the Estimated Takeover node might have a latency of 4 ms/op. This value is a calculation that Unified Manager performs.

You can click the name of the Partner node if you want it to become the root object. After the Performance/Node Performance Explorer page is displayed, you can click the **Failover Planning** tab to see how performance changes in this node failure scenario. For example, if Node1 is the Primary node and Node2 is the Partner node, you can click Node2 to make it the Primary node. In this way, you can see how the estimated performance changes depending on which node fails.

Comparing pane

The following list describes the components displayed in the Comparing pane by default:

- **Events charts**

They are displayed in the same format as those in the Performance/Node Performance Explorer page. They pertain to the Primary node only.

- **Counter charts**

They display historical statistics for the performance counter shown in the grid. In each chart, the graph for the Estimated Takeover node shows the estimated performance if a failover had occurred at any given time.

For example, suppose the Utilization chart shows 73% for the Estimated Takeover node at 11 a.m. on February 8. If a failover had occurred at that time, the utilization of the takeover node would have been 73%.

The historical statistics help you find the optimal time for initiating a failover, minimizing the possibility of overloading the takeover node. You can schedule a failover only at times when the predicted performance of the takeover node is acceptable.

By default, statistics for both the root object and the partner node are displayed in the Comparing pane. Unlike in the Performance/Node Performance Explorer page, this page does not display the **Add** button for you to add objects for statistics comparison.

You can customize the Comparing pane in the same way as you do in the Performance/Node Performance Explorer page. The following list shows examples of customizing the charts:

- Click a node name to show or hide the node's statistics in the Counter charts.
- Click **Zoom View** to display a detailed chart for a particular counter in a new window.

Using a threshold policy with the Node Failover Planning page

You can create a node threshold policy so that you can be notified in the Performance/Node Failover Planning page when a potential failover would degrade the performance of the takeover node to an unacceptable level.

The system-defined performance threshold policy named “Node HA pair over-utilized” generates a warning event if the threshold is breached for six consecutive collection periods (30 minutes). The threshold is considered breached if the combined performance capacity used of the nodes in an HA pair exceeds 200%.

The event from the system-defined threshold policy alerts you to the fact that a failover will cause the latency of the takeover node to increase to an unacceptable level. When you see an event that is generated by this policy for a particular node, you can navigate to the Performance/Node Failover Planning page for that node to view the predicted latency value due to a failover.

In addition to using this system-defined threshold policy, you can create threshold policies by using the “Performance Capacity Used - Takeover” counter, and then apply the policy to selected nodes. Specifying a threshold lower than 200% enables you to receive an event before the threshold for the system-defined policy is breached. You can also specify the minimum period of time for which the threshold is exceeded to less than 30 minutes if you want to be notified before the system-defined policy event is generated.

For example, you can define a threshold policy to generate a warning event if the combined performance capacity used of the nodes in an HA pair exceeds 175% for more than 10 minutes. You can apply this policy to Node1 and Node2, which form an HA pair. After receiving a warning event notification for either Node1 or Node2, you can view the Performance/Node Failover Planning page for that node to assess the estimated performance impact on the takeover node. You can take corrective actions to avoid overloading the takeover node if a failover does happen. If you take action when the combined performance capacity used of the nodes is under 200%, the takeover node’s latency does not reach an unacceptable level even if a failover happens during this time.

Using the Performance Capacity Used Breakdown chart for failover planning

The detailed Performance Capacity Used - Breakdown chart shows the performance capacity used for the Primary node and the Partner node. It also shows the amount of free performance capacity on the Estimated Takeover node. This information helps you determine whether you might have a performance issue if the partner node fails.

About this task

In addition to showing the total performance capacity used for the nodes, the Breakdown chart breaks the values for each node into user protocols and background processes.

- User protocols are the I/O operations from user applications to and from the cluster.
- Background processes are the internal system processes involved with storage efficiency, data replication, and system health.

This additional level of detail enables you to determine whether a performance issue is caused by user application activity or background system processes, such as deduplication, RAID reconstruct, disk scrubbing,

and SnapMirror copies.


Steps

1. Go to the **Performance/Node Failover Planning** page for the node that will serve as the Estimated Takeover node.
2. From the **Time Range** selector, choose the period of time for which the historical statistics are displayed in the counter grid and counter charts.

The counter charts with statistics for the Primary node, Partner node, and Estimated Takeover node are displayed.

3. From the **Choose charts** list, select **Perf. Capacity Used**.
4. In the **Perf. Capacity Used** chart, select **Breakdown** and click **Zoom View**.

The detailed chart for Perf. Capacity Used is displayed.



5. Move the cursor over the detailed chart to view the performance capacity used information in the popup window.

The Perf. Capacity Free percentage is the performance capacity available on the Estimated Takeover node. It indicates how much performance capacity is left on the takeover node after a failover. If it is 0%, a failover will cause the latency to increase to an unacceptable level on the takeover node.

6. Consider taking corrective actions to avoid a low performance capacity free percentage.

If you plan to initiate a failover for node maintenance, choose a time to fail the partner node when the

performance capacity free percentage is not at 0.

Copyright information

Copyright © 2025 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.