



Analyzing events from system-defined performance thresholds

Active IQ Unified Manager 9.14

NetApp
November 11, 2024

Table of Contents

- Analyzing events from system-defined performance thresholds 1
 - Responding to system-defined performance threshold events 1
 - Responding to QoS policy group performance events 2
 - Understanding events from adaptive QoS policies that have a defined block size 3
 - Responding to node resources overutilized performance events 4
 - Responding to cluster imbalance performance events 5

Analyzing events from system-defined performance thresholds

Events generated from system-defined performance thresholds indicate that a performance counter, or set of performance counters, for a certain storage object has crossed the threshold from a system-defined policy. This indicates that the storage object, for example, an aggregate or node, is experiencing a performance issue.

You use the Event details page to analyze the performance event and take corrective action, if necessary, to return performance back to normal.



System-defined threshold policies are not enabled on Cloud Volumes ONTAP, ONTAP Edge, or ONTAP Select systems.

Responding to system-defined performance threshold events

You can use Unified Manager to investigate performance events caused by a performance counter crossing a system-defined warning threshold. You can also use Unified Manager to check the health of the cluster component to see whether recent events detected on the component contributed to the performance event.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There must be new or obsolete performance events.

Steps

1. Display the **Event details** page to view information about the event.
2. Review the **Description**, which describes the threshold breach that caused the event.

For example, the message “Node utilization value of 90 % has triggered a WARNING event based on threshold setting of 85 %” indicates that a node utilization warning event occurred for the cluster object.

3. Make a note of the **Event Trigger Time** so you can investigate whether other events might have occurred at the same time that could have contributed to this event.
4. Under **System Diagnosis**, review the brief description of the type of analysis the system-defined policy is performing on the cluster object.

For some events a green or red icon is displayed next to the diagnosis to indicate whether an issue was found in that particular diagnosis. For other types of system-defined events counter charts display the performance for the object.

5. Under **Suggested Actions**, click the **Help me do this** link to view the suggested actions you can perform to try and resolve the performance event on your own.

Responding to QoS policy group performance events

Unified Manager generates QoS policy warning events when workload throughput (IOPS, IOPS/TB, or MBps) has exceeded the defined ONTAP QoS policy setting and workload latency is becoming affected. These system-defined events provide the opportunity to correct potential performance issues before many workloads are affected by latency.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events.

Unified Manager generates warning events for QoS policy breaches when workload throughput has exceeded the defined QoS policy setting during each performance collection period for the previous hour. Workload throughput may exceed the QoS threshold for only a short period of time during each collection period, but Unified Manager displays only the “average” throughput during the collection period on the chart. For this reason you may receive QoS events while the throughput for a workload might not have crossed the policy threshold shown in the chart.

You can use System Manager or the ONTAP commands to manage policy groups, including the following tasks:

- Creating a new policy group for the workload
- Adding or removing workloads in a policy group
- Moving a workload between policy groups
- Changing the throughput limit of a policy group
- Moving a workload to a different aggregate or node

Steps

1. Display the **Event details** page to view information about the event.
2. Review the **Description**, which describes the threshold breach that caused the event.

For example, the message “IOPS value of 1,352 IOPS on vol1_NFS1 has triggered a WARNING event to identify potential performance problems for the workload” indicates that a QoS Max IOPS event occurred on volume vol1_NFS1.

3. Review the **Event Information** section to see more details about when the event occurred and how long the event has been active.

Additionally, for volumes or LUNs that are sharing the throughput of a QoS policy you can see the names of the top three workloads that are consuming the most IOPS or MBps.

4. Under the **System Diagnosis** section, review the two charts: one for total average IOPS or MBps (depending on the event), and one for latency. When arranged this way you can see which cluster components are most affecting latency when the workload approached the QoS max limit.

For a shared QoS policy event, the top three workloads are shown in the throughput chart. If more than three workloads are sharing the QoS policy, then additional workloads are added together in an “Other workloads” category. Additionally, the Latency chart shows the average latency on all workloads that are part of the QoS policy.

Note that for adaptive QoS policy events that the IOPS and MBps charts show IOPS or MBps values that ONTAP has converted from the assigned IOPS/TB threshold policy based on the size of the volume.

5. Under the **Suggested Actions** section, review the suggestions and determine which actions you should perform to avoid an increase in latency for the workload.

If required, click the **Help** button to view more details about the suggested actions you can perform to try and resolve the performance event.

Understanding events from adaptive QoS policies that have a defined block size

Adaptive QoS policy groups automatically scale a throughput ceiling or floor based on the volume size, maintaining the ratio of IOPS to TBs as the size of the volume changes. Starting with ONTAP 9.5 you can specify the block size in the QoS policy to effectively apply a MB/s threshold at the same time.

Assigning an IOPS threshold in an adaptive QoS policy places a limit only on the number of operations that occur in each workload. Depending on the block size that is set on the client that generates the workloads, some IOPS include much more data and therefore place a much larger burden on the nodes that process the operations.

The MB/s value for a workload is generated using the following formula:

$$\text{MB/s} = (\text{IOPS} * \text{Block Size}) / 1000$$

If a workload is averaging 3,000 IOPS and the block size on the client is set to 32 KB, then the effective MB/s for this workload is 96. If this same workload is averaging 3,000 IOPS and the block size on the client is set to 48 KB, then the effective MB/s for this workload is 144. You can see that the node is processing 50% more data when the block size is larger.

Let's look at the following adaptive QoS policy that has a defined block size and how events are triggered based on the block size that is set on the client.

Create a policy and set the peak throughput to 2,500 IOPS/TB with a block size of 32KB. This effectively sets the MB/s threshold to 80 MB/s ((2500 IOPS * 32KB) / 1000) for a volume with 1 TB used capacity. Note that Unified Manager generates a Warning event when the throughput value is 10% less than the defined threshold. Events are generated under the following situations:

Used Capacity	Event is generated when throughput exceeds this number of ...	
	IOPS	MB/s
1 TB	2,250 IOPS	72 MB/s
2 TB	4,500 IOPS	144 MB/s
5 TB	11,250 IOPS	360 MB/s

If the volume is using 2TB of the available space, and the IOPS is 4,000, and the QoS block size is set to 32KB on the client, then the MB/ps throughput is 128 MB/s $((4,000 \text{ IOPS} * 32 \text{ KB}) / 1000)$. No event is generated in this scenario because both 4,000 IOPS and 128 MB/s are below the threshold for a volume that is using 2 TB of space.

If the volume is using 2TB of the available space, and the IOPS is 4,000, and the QoS block size is set to 64KB on the client, then the MB/s throughput is 256 MB/s $((4,000 \text{ IOPS} * 64 \text{ KB}) / 1000)$. In this case the 4,000 IOPS does not generate an event, but the MB/s value of 256 MB/s is above the threshold of 144 MB/s and an event is generated.

For this reason, when an event is triggered based on a MB/s breach for an adaptive QoS policy that includes the block size, a MB/s chart is displayed in the System Diagnosis section of the Event details page. If the event is triggered based on an IOPS breach for the adaptive QoS policy, an IOPS chart is displayed in the System Diagnosis section. If a breach occurs for both IOPS and MB/s you will receive two events.

For more information on adjusting QoS settings, see [Performance management overview](#).

Responding to node resources overutilized performance events

Unified Manager generates node resources overutilized warning events when a single node is operating above the bounds of its operational efficiency, and therefore potentially affecting workload latencies. These system-defined events provide the opportunity to correct potential performance issues before many workloads are affected by latency.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There must be new or obsolete performance events.

Unified Manager generates warning events for node resources overutilized policy breaches by looking for nodes that are using more than 100% of their performance capacity for more than 30 minutes.

You can use System Manager or the ONTAP commands to correct this type of performance issue, including the following tasks:

- Creating and applying a QoS policy to any volumes or LUNs that are overusing system resources
- Reducing the QoS maximum throughput limit of a policy group to which workloads have been applied
- Moving a workload to a different aggregate or node
- Increasing capacity by adding disks to the node, or by upgrading to a node with a faster CPU and more RAM

Steps

1. Display the **Event details** page to view information about the event.
2. Review the **Description**, which describes the threshold breach that caused the event.

For example, the message “Perf. Capacity Used value of 139% on simplicity-02 has triggered a WARNING event to identify potential performance problems in the data processing unit.” indicates that performance capacity on node simplicity-02 is overused and affecting node performance.

3. Under the **System Diagnosis** section, review the three charts: one for performance capacity used on the node, one for average storage IOPS being used by the top workloads, and one for latency on the top workloads. When arranged in this way you can see which workloads are the cause of the latency on the node.

You can view which workloads have QoS policies applied, and which do not, by moving your cursor over the IOPS chart.

4. Under the **Suggested Actions** section, review the suggestions and determine which actions you should perform to avoid an increase in latency for the workload.

If required, click the **Help** button to view more details about the suggested actions you can perform to try and resolve the performance event.

Responding to cluster imbalance performance events

Unified Manager generates cluster imbalance warning events when one node in a cluster is operating at a much higher load than other nodes, and therefore potentially affecting workload latencies. These system-defined events provide the opportunity to correct potential performance issues before many workloads are affected by latency.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

Unified Manager generates warning events for cluster imbalance threshold policy breaches by comparing the performance capacity used value for all nodes in the cluster to see if there is a load difference of 30% between any nodes.

These steps help you identify the following resources so that you can move high-performing workloads to a lower utilized node:

- The nodes on the same cluster that are less utilized
- The aggregates on the new node that are the least utilized
- The highest-performing volumes on the current node

Steps

1. Display the **Event** details page to view information about the event.
2. Review the **Description**, which describes the threshold breach that caused the event.

For example, the message “The performance capacity used counter indicates a load difference of 62% between the nodes on cluster Dallas-1-8 and has triggered a WARNING event based on the system threshold of 30%” indicates that performance capacity on one of the nodes is overused and affecting node performance.

3. Review the text in the **Suggested Actions** to move a high-performing volume from the node with the high performance capacity used value to a node with the lowest performance capacity used value.
4. Identify the nodes with the highest and lowest performance capacity used value:
 - a. In the **Event Information** section, click the name of the source cluster.
 - b. In the **Cluster / Performance Summary** page, click **Nodes** in the **Managed Objects** area.

- c. In the **Nodes** inventory page, sort the nodes by the **Performance Capacity Used** column.
 - d. Identify the nodes with the highest and lowest performance capacity used value and write down those names.
5. Identify the volume using the most IOPS on the node that has the highest performance capacity used value:
- a. Click the node with the highest performance capacity used value.
 - b. In the **Node / Performance Explorer** page, select **Aggregates on this Node** from the **View and Compare** menu.
 - c. Click the aggregate with the highest performance capacity used value.
 - d. In the **Aggregate / Performance Explorer** page, select **Volumes on this Aggregate** from the **View and Compare** menu.
 - e. Sort the volumes by the **IOPS** column, and write down the name of the volume using the most IOPS, and the name of the aggregate where the volume resides.
6. Identify the aggregate with the lowest utilization on the node that has the lowest performance capacity used value:
- a. Click **Storage > Aggregates** to display the **Aggregates** inventory page.
 - b. Select the **Performance: All Aggregates** view.
 - c. Click the **Filter** button and add a filter where “Node” equals the name of the node with the lowest performance capacity used value that you wrote down in step 4.
 - d. Write down the name of the aggregate that has the lowest performance capacity used value.
7. Move the volume from the overloaded node to the aggregate you identified as having low utilization on the new node.

You can perform the move operation by using ONTAP System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After a few days, check to see whether you are receiving the same cluster imbalance event from this cluster.

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.