



What performance events are

Active IQ Unified Manager 9.7

NetApp
August 30, 2024

Table of Contents

- What performance events are 1
- Sources of performance events 1
- Types of system-defined performance threshold policies 1

What performance events are

Performance events are incidents related to workload performance on a cluster. They help you identify workloads with slow response times. Together with health events that occurred at the same time, you can determine the issues that might have caused, or contributed to, the slow response times.

When Unified Manager detects multiple occurrences of the same event condition for the same cluster component, it treats all occurrences as a single event, not as separate events.

Sources of performance events

Performance events are issues related to workload performance on a cluster. They help you identify storage objects with slow response times, also known as high latency. Together with other health events that occurred at the same time, you can determine the issues that might have caused, or contributed to, the slow response times.

Unified Manager receives performance events from the following sources:

- **User-defined performance threshold policy events**

Performance issues based on custom threshold values that you have set. You configure performance threshold policies for storage objects; for example, aggregates and volumes, so that events are generated when a threshold value for a performance counter has been breached.

You must define a performance threshold policy and assign it to a storage object to receive these events.

- **System-defined performance threshold policy events**

Performance issues based on threshold values that are system-defined. These threshold policies are included with the installation of Unified Manager to cover common performance problems.

These threshold policies are enabled by default, and you might see events shortly after adding a cluster.

- **Dynamic performance threshold events**

Performance issues that are the result of failures or errors in an IT infrastructure, or from workloads overutilizing cluster resources. The cause of these events might be a simple issue that corrects itself over a period of time or that can be addressed with a repair or configuration change. A dynamic threshold event indicates that the workloads on an ONTAP system are slow due to other workloads with high usage of shared cluster components.

These thresholds are enabled by default, and you might see events after three days of collecting data from a new cluster.

Types of system-defined performance threshold policies

Unified Manager provides some standard threshold policies that monitor cluster performance and generate events automatically. These policies are enabled by default, and they generate warning or information events when the monitored performance

thresholds are breached.



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

If you are receiving unnecessary events from any system-defined performance threshold policies, you can disable the events for individual policies from the Event Setup page.

Cluster threshold policies

The system-defined cluster performance threshold policies are assigned, by default, to every cluster being monitored by Unified Manager:

- **Cluster imbalance threshold**

Identifies situations in which one node is operating at a much higher load than other nodes in the cluster, and therefore potentially affecting workload latencies.

It does this 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. This is a warning event.

Node threshold policies

The system-defined node performance threshold policies are assigned, by default, to every node in the clusters being monitored by Unified Manager:

- **Node resources over-utilized**

Identifies situations in which a single node is operating above the bounds of its operational efficiency, and therefore potentially affecting workload latencies.

It does this by looking for nodes that are using more than 100% of their performance capacity for more than 12 hours. This is a warning event.

- **Node HA pair over-utilized**

Identifies situations in which nodes in an HA pair are operating above the bounds of the HA pair operational efficiency.

It does this by looking at the performance capacity used value for the two nodes in the HA pair. If the combined performance capacity used of the two nodes exceeds 200% for more than 12 hours, then a controller failover will impact workload latencies. This is an informational event.

- **Node disk fragmentation**

Identifies situations in which a disk or disks in an aggregate are fragmented, slowing key system services and potentially affecting workload latencies on a node.

It does this by looking at certain read and write operation ratios across all aggregates on a node. This policy might also be triggered during SyncMirror resynchronization or when errors are found during disk scrub operations. This is a warning event.



The “Node disk fragmentation” policy analyzes HDD-only aggregates; Flash Pool, SSD, and FabricPool aggregates are not analyzed.

Aggregate threshold policies

The system-defined aggregate performance threshold policy is assigned by default to every aggregate in the clusters being monitored by Unified Manager:

- **Aggregate disks over-utilized**

Identifies situations in which an aggregate is operating above the limits of its operational efficiency, thereby potentially affecting workload latencies. It identifies these situations by looking for aggregates where the disks in the aggregate are more than 95% utilized for more than 30 minutes. This multicondition policy then performs the following analysis to help determine the cause of the issue:

- Is a disk in the aggregate currently undergoing background maintenance activity?

Some of the background maintenance activities a disk could be undergoing are disk reconstruction, disk scrub, SyncMirror resynchronization, and reparity.

- Is there a communications bottleneck in the disk shelf Fibre Channel interconnect?
- Is there too little free space in the aggregate? A warning event is issued for this policy only if one (or more) of the three subordinate policies are also considered breached. A performance event is not triggered if only the disks in the aggregate are more than 95% utilized.



The “Aggregate disks over-utilized” policy analyzes HDD-only aggregates and Flash Pool (hybrid) aggregates; SSD and FabricPool aggregates are not analyzed.

Workload latency threshold policies

The system-defined workload latency threshold policies are assigned to any workload that has a configured Performance Service Level policy that has a defined “expected latency” value:

- **Workload Volume/LUN Latency Threshold Breached as defined by Performance Service Level**

Identifies volumes (file shares) and LUNs that have exceeded their “expected latency” limit, and that are affecting workload performance. This is a warning event.

It does this by looking for workloads that have exceeded the expected latency value for 30% of the time during the previous hour.

QoS threshold policies

The system-defined QoS performance threshold policies are assigned to any workload that has a configured ONTAP QoS maximum throughput policy (IOPS, IOPS/TB, or MB/s). Unified Manager triggers an event when the workload throughput value is 15% less than the configured QoS value:

- **QoS Max IOPS or MB/s threshold**

Identifies volumes and LUNs that have exceeded their QoS maximum IOPS or MB/s throughput limit, and that are affecting workload latency. This is a warning event.

When a single workload is assigned to a policy group, it does this by looking for workloads that have exceeded the maximum throughput threshold defined in the assigned QoS policy group during each collection period for the previous hour.

When multiple workloads share a single QoS policy, it does this by adding the IOPS or MB/s of all workloads in the policy and checking that total against the threshold.

- **QoS Peak IOPS/TB or IOPS/TB with Block Size threshold**

Identifies volumes that have exceeded their adaptive QoS peak IOPS/TB throughput limit (or IOPS/TB with Block Size limit), and that are affecting workload latency. This is a warning event.

It does this by converting the peak IOPS/TB threshold defined in the adaptive QoS policy into a QoS maximum IOPS value based on the size of each volume, and then it looks for volumes that have exceeded the QoS max IOPS during each performance collection period for the previous hour.



This policy is applied to volumes only when the cluster is installed with ONTAP 9.3 and later software.

When the “block size” element has been defined in the adaptive QoS policy, the threshold is converted into a QoS maximum MB/s value based on the size of each volume. Then it looks for volumes that have exceeded the QoS max MB/s during each performance collection period for the previous hour.



This policy is applied to volumes only when the cluster is installed with ONTAP 9.5 and later software.

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.