



Analyzing events from dynamic performance thresholds

Active IQ Unified Manager

NetApp
March 26, 2025

Table of Contents

- Analyzing events from dynamic performance thresholds 1
 - Identifying victim workloads involved in a dynamic performance event 1
 - Before you begin 1
 - About this task 1
 - Steps 1
 - Identifying bully workloads involved in a dynamic performance event 1
 - Before you begin 2
 - About this task 2
 - Steps 2
 - Identifying shark workloads involved in a dynamic performance event 2
 - Before you begin 2
 - About this task 2
 - Steps 2
- Performance event analysis for a MetroCluster configuration 2
 - Analyzing a dynamic performance event on a cluster in a MetroCluster configuration 3
 - Analyzing a dynamic performance event for a remote cluster on a MetroCluster configuration 4
- Responding to a dynamic performance event caused by QoS policy group throttling 5
 - Before you begin 5
 - Steps 5
- Responding to a dynamic performance event caused by a disk failure 6
 - Before you begin 6
 - Steps 7
- Responding to a dynamic performance event caused by HA takeover 8
 - Before you begin 8
 - Steps 8

Analyzing events from dynamic performance thresholds

Events generated from dynamic thresholds indicate that the actual response time (latency) for a workload is too high, or too low, compared to the expected response time range. You use the Event details page to analyze the performance event and take corrective action, if necessary, to return performance back to normal.



Dynamic performance thresholds are not enabled on Cloud Volumes ONTAP, ONTAP Edge, or ONTAP Select systems.

Identifying victim workloads involved in a dynamic performance event

In Unified Manager, you can identify which volume workloads have the highest deviation in response time (latency) caused by a storage component in contention. Identifying these workloads helps you understand why the client applications accessing them have been performing slower than usual.

Before you begin

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

About this task

The Event details page displays a list of the user-defined and system-defined workloads, ranked by the highest deviation in activity or usage on the component or most impacted by the event. The values are based on the peaks that Unified Manager identified when it detected and last analyzed the event.

Steps

1. Display the **Event details** page to view information about the event.
2. In the Workload Latency and Workload Activity charts, select **Victim Workloads**.
3. Hover your cursor over the charts to view the top user-defined workloads that are affecting the component, and the name of the victim workload.

Identifying bully workloads involved in a dynamic performance event

In Unified Manager, you can identify which workloads have the highest deviation in usage for a cluster component in contention. Identifying these workloads helps you understand why certain volumes on the cluster have slow response times (latency).

Before you begin

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

About this task

The Event details page displays a list of the user-defined and system-defined workloads ranked by the highest usage of the component or most impacted by the event. The values are based on the peaks that Unified Manager identified when it detected and last analyzed the event.

Steps

1. Display the **Event details** page to view information about the event.
2. In the Workload Latency and Workload Activity charts, select **Bully Workloads**.
3. Hover your cursor over the charts to view the top user-defined bully workloads that are affecting the component.

Identifying shark workloads involved in a dynamic performance event

In Unified Manager, you can identify which workloads have the highest deviation in usage for a storage component in contention. Identifying these workloads helps you determine if these workloads should be moved to a less-utilized cluster.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There are new, acknowledged, or obsolete performance dynamic event.

About this task

The Event details page displays a list of the user-defined and system-defined workloads ranked by the highest usage of the component or most impacted by the event. The values are based on the peaks that Unified Manager identified when it detected and last analyzed the event.

Steps

1. Display the **Event details** page to view information about the event.
2. In the Workload Latency and Workload Activity charts, select **Shark Workloads**.
3. Hover your cursor over the charts to view the top user-defined workloads that are affecting the component, and the name of the shark workload.

Performance event analysis for a MetroCluster configuration

You can use Unified Manager to analyze a performance event for a MetroCluster

configuration. You can identify the workloads involved in the event and review the suggested actions for resolving it.

MetroCluster performance events might be due to *bully* workloads that are over-utilizing the interswitch links (ISLs) between the clusters, or due to link health issues. Unified Manager monitors each cluster in a MetroCluster configuration independently, without consideration of performance events on a partner cluster.

Performance events from both clusters in the MetroCluster configuration are also displayed on the Unified ManagerDashboard page. You can also view the Health pages of Unified Manager to check the health of each cluster and to view their relationship.

Analyzing a dynamic performance event on a cluster in a MetroCluster configuration

You can use Unified Manager to analyze the cluster in a MetroCluster configuration on which a performance event was detected. You can identify the cluster name, event detection time, and the *bully* and *victim* workloads involved.

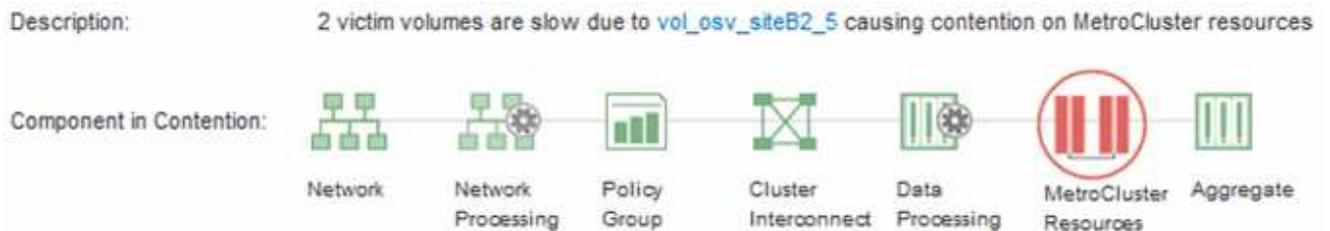
Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events for a MetroCluster configuration.
- Both clusters in the MetroCluster configuration must be monitored by the same instance of Unified Manager.

Steps

1. Display the **Event details** page to view information about the event.
2. Review the event description to see the names of the workloads involved and the number of workloads involved.

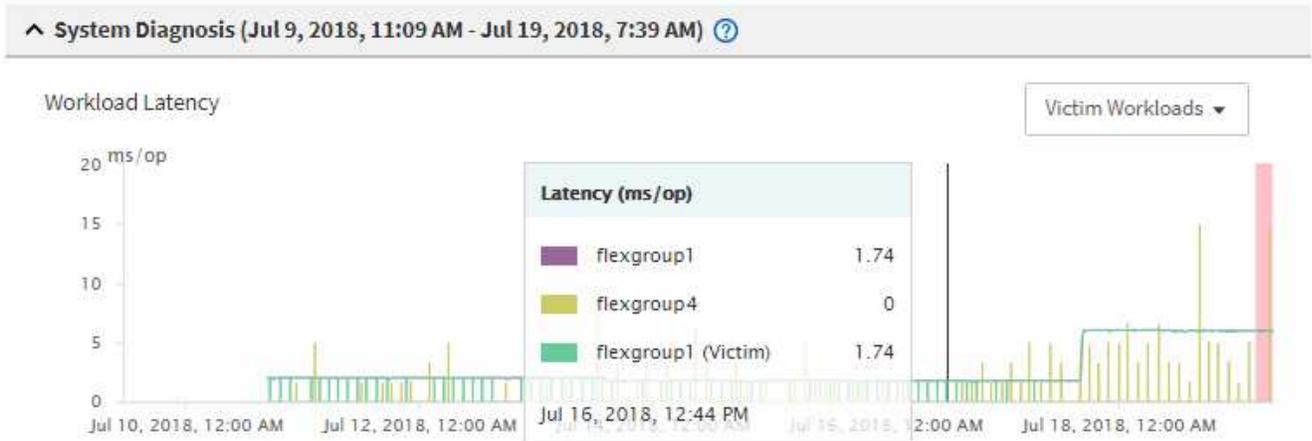
In this example, the MetroCluster Resources icon is red, indicating that the MetroCluster resources are in contention. You position your cursor over the icon to display a description of the icon. At the top of the page in the event ID, the cluster name identifies the name of the cluster on which the event was detected.



3. Make a note of the cluster name and the event detection time, which you can use to analyze performance events on the partner cluster.
4. In the charts, review the *victim* workloads to confirm that their response times are higher than the performance threshold.

In this example, the victim workload is displayed in the hover text. The Latency charts display, at a high-level, a consistent latency pattern for the victim workloads involved. Even though the abnormal latency of the victim workloads triggered the event, a consistent latency pattern might indicate that the workloads are performing within their expected range, but that a spike in I/O increased the latency and triggered the

event.



If you recently installed an application on a client that accesses these volume workloads and that application sends a high amount of I/O to them, you might be anticipating their latencies to increase. If the latency for the workloads returns within the expected range, the event state changes to obsolete, and remains in this state for more than 30 minutes, you can probably ignore the event. If the event is ongoing, and remains in the new state, you can investigate it further to determine whether other issues caused the event.

5. In the Workload Throughput chart, select **Bully Workloads** to display the bully workloads.

The presence of bully workloads indicates that the event might have been caused by one or more workloads on the local cluster overutilizing the MetroCluster resources. The bully workloads have a high deviation in write throughput (MBps).

This chart displays, at a high-level, the write throughput (MBps) pattern for the workloads. You can review the write MBps pattern to identify abnormal throughput, which might indicate that a workload is overutilizing the MetroCluster resources.

If no bully workloads are involved in the event, the event might have been caused by a health issue with the link between the clusters or a performance issue on the partner cluster. You can use Unified Manager to check the health of both clusters in a MetroCluster configuration. You can also use Unified Manager to check for and analyze performance events on the partner cluster.

Analyzing a dynamic performance event for a remote cluster on a MetroCluster configuration

You can use Unified Manager to analyze dynamic performance events on a remote cluster in a MetroCluster configuration. The analysis helps you determine whether an event on the remote cluster caused an event on its partner cluster.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have analyzed a performance event on a local cluster in a MetroCluster configuration and obtained the event detection time.
- You must have checked the health of the local cluster and its partner cluster involved in the performance event and obtained the name of the partner cluster.

Steps

1. Log in to the Unified Manager instance that is monitoring the partner cluster.
2. In the left navigation pane, click **Events** to display the event list.
3. From the **Time Range** selector, select **Last Hour**, and then click **Apply Range**.
4. In the **Filtering** selector, select **Cluster** from the left drop-down menu, type the name of the partner cluster in the text field, and then click **Apply Filter**.

If there are no events for the selected cluster over the last hour, this indicates that the cluster has not experienced any performance issues during the time that the event was detected on its partner.

5. If the selected cluster has events detected over the last hour, compare the event detection time to the event detection time for the event on the local cluster.

If these events involve bully workloads causing contention on the data processing component, one or more of these bullies might have caused the event on the local cluster. You can click the event to analyze it and review the suggested actions for resolving it on the Event details page.

If these events do not involve bully workloads, they did not cause the performance event on the local cluster.

Responding to a dynamic performance event caused by QoS policy group throttling

You can use Unified Manager to investigate a performance event caused by a Quality of Service (QoS) policy group throttling workload throughput (MB/s). The throttling increased the response times (latency) of volume workloads in the policy group. You can use the event information to determine whether new limits on the policy groups are needed to stop the throttling.

Before you begin

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

Steps

1. Display the **Event details** page to view information about the event.
2. Read the **Description**, which displays the name of the workloads impacted by the throttling.



The description can display the same workload for the victim and bully, because the throttling makes the workload a victim of itself.

3. Record the name of the volume, using an application such as a text editor.

You can search on the volume name to locate it later.

4. In the Workload Latency and Workload Utilization charts, select **Bully Workloads**.

5. Hover your cursor over the charts to view the top user-defined workloads that are affecting the policy group.

The workload at the top of the list has the highest deviation and caused the throttling to occur. The activity is the percentage of the policy group limit used by each workload.

6. In the **Suggested Actions** area, click the **Analyze Workload** button for the top workload.
7. In the **Workload Analysis** page, set the Latency chart to view all Cluster Components, and the Throughput chart to view Breakdown.

The breakdown charts are displayed under the Latency chart and the IOPS chart.

8. Compare the QoS Limits in the **Latency** chart to see what amount of throttling impacted the latency at the time of the event.

The QoS policy group has a maximum throughput of 1,000 operations per second (op/sec), which the workloads in it cannot collectively exceed. At the time of the event, the workloads in the policy group had a combined throughput of over 1,200 op/sec, which caused the policy group to throttle its activity back to 1,000 op/sec.

9. Compare the **Reads/writes latency** values to the **Reads/writes/other** values.

Both charts show a high number of read requests with high latency, but the number of requests and amount of latency for write requests is low. These values help you determine whether there is a high amount of throughput or number of operations that increased the latency. You can use these values when deciding to put a policy group limit on the throughput or operations.

10. Use ONTAP System Manager to increase the current limit on the policy group to 1,300 op/sec.
11. After a day, return to Unified Manager and enter the workload that you recorded in Step 3 in the **Workload Analysis** page.
12. Select the Throughput Breakdown chart.

The Reads/writes/other chart is displayed.

13. At the top of the page, point your cursor to the change event icon () for the policy group limit change.
14. Compare the **Reads/writes/other** chart to the **Latency** chart.

The read and write requests are the same, but the throttling has stopped and the latency has decreased.

Responding to a dynamic performance event caused by a disk failure

You can use Unified Manager to investigate a performance event caused by workloads overutilizing an aggregate. You can also use Unified Manager to check the health of the aggregate to see if recent health events detected on the aggregate contributed to the performance event.

Before you begin

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

- There must be new, acknowledged, or obsolete performance events.

Steps

1. Display the **Event details** page to view information about the event.
2. Read the **Description**, which describes the workloads involved in the event and the cluster component in contention.

There are multiple victim volumes whose latency was impacted by the cluster component in contention. The aggregate, which is in the middle of a RAID reconstruct to replace the failed disk with a spare disk, is the cluster component in contention. Under Component in Contention, the Aggregate icon is highlighted red and the name of the aggregate is displayed in parentheses.

3. In the Workload Utilization chart, select **Bully Workloads**.
4. Hover your cursor over the chart to view the top bully workloads that are affecting the component.

The top workloads with the highest peak utilization since the event was detected are displayed at the top of the chart. One of the top workloads is the system-defined workload Disk Health, which indicates a RAID reconstruct. A reconstruct is the internal process involved with rebuilding the aggregate with the spare disk. The Disk Health workload, along with other workloads on the aggregate, likely caused the contention on the aggregate and the associated event.

5. After confirming that the activity from the Disk Health workload caused the event, wait for approximately 30 minutes for the reconstruction to finish and for Unified Manager to analyze the event and detect whether the aggregate is still in contention.
6. Refresh the **Event details**.

After the RAID reconstruction is complete, check that the State is obsolete, indicating that the event is resolved.

7. In the Workload Utilization chart, select **Bully Workloads** to view the workloads on the aggregate by peak utilization.
8. In the **Suggested Actions** area, click the **Analyze Workload** button for the top workload.
9. In the **Workload Analysis** page, set the Time Range to display the last 24 hours (1 day) of data for the selected volume.

In the Event Timeline, a red dot (●) indicates when the disk failure event occurred.

10. In the Node and Aggregate Utilization chart, hide the line for the Node statistics so that just the Aggregate line remains.
11. Compare the data in this chart to the data at the time of the event in the **Latency** chart.

At the time of the event, the Aggregate Utilization shows a high amount of read and write activity, caused by the RAID reconstruction processes, which increased the latency of the selected volume. A few hours after the event occurred, both the reads and writes and the latency have decreased, confirming that the aggregate is no longer in contention.

Responding to a dynamic performance event caused by HA takeover

You can use Unified Manager to investigate a performance event caused by high data processing on a cluster node that is in a high-availability (HA) pair. You can also use Unified Manager to check the health of the nodes to see whether any recent health events detected on the nodes contributed to the performance event.

Before you begin

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

Steps

1. Display the **Event details** page to view information about the event.
2. Read the **Description**, which describes the workloads involved in the event and the cluster component in contention.

There is one victim volume whose latency was impacted by the cluster component in contention. The data processing node, which took over all workloads from its partner node, is the cluster component in contention. Under Component in Contention, the Data Processing icon is highlighted red and the name of the node that was handling data processing at the time of the event is displayed in parentheses.

3. In the **Description**, click the name of the volume.

The Volume Performance Explorer page is displayed. At the top of the page, in the Events time line, a change event icon () indicates the time that Unified Manager detected the start of the HA takeover.

4. Point your cursor to the change event icon for the HA takeover and details about the HA takeover are displayed in hover text.

In the Latency chart, an event indicates that the selected volume crossed the performance threshold due to high latency around the same time as the HA takeover.

5. Click **Zoom View** to display the Latency chart on a new page.
6. In the View menu, select **Cluster Components** to view the total latency by cluster component.
7. Point your mouse cursor to the change event icon for the start of the HA takeover and compare the latency for data processing to the total latency.

At the time of the HA takeover, there was a spike in data processing from the increased workload demand on the data processing node. The increased CPU utilization drove up the latency and triggered the event.

8. After fixing the failed node, use ONTAP System Manager to perform an HA giveback, which moves the workloads from the partner node to the fixed node.
9. After the HA giveback is complete, after the next configuration discovery in Unified Manager (approximately 15 minutes), find the event and workload that triggered by the HA takeover in the **Event Management** inventory page.

The event triggered by the HA takeover now has a state of obsolete, which indicates that the event is

resolved. The latency at the data processing component has decreased, which has decreased the total latency. The node that the selected volume is now using for data processing has resolved the event.

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.