



Analyzing performance events

Active IQ Unified Manager 9.11

NetApp
January 31, 2025

Table of Contents

- Analyzing performance events 1
 - Displaying information about performance events 1
 - Analyzing events from user-defined performance thresholds 2
 - Analyzing events from system-defined performance thresholds 3
 - Analyzing events from dynamic performance thresholds 8

Analyzing performance events

You can analyze performance events to identify when they were detected, whether they are active (new or acknowledged) or obsolete, the workloads and cluster components involved, and the options for resolving the events on your own.

Displaying information about performance events

You can use the Event Management inventory page to view a list of all the performance events on the clusters being monitored by Unified Manager. By viewing this information you can determine the most critical events and then drill down to detailed information to determine the cause of the event.

What you'll need

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

The list of events is sorted by detected time, with the most recent events listed first. You can click a column header to sort the events based on that column. For example, you can sort by the Status column to view events by severity. If you are looking for a specific event, or for a specific type of event, you can use the filter and search mechanisms to refine the list of events that appear in the list.

Events from all sources are displayed on this page:

- User-defined performance threshold policy
- System-defined performance threshold policy
- Dynamic performance threshold

The Event Type column lists the source of the event. You can select an event to view details about the event in the Event details page.

Steps

1. In the left navigation pane, click **Event Management**.
2. From the View menu, select **Active performance events**.

The page displays all New and Acknowledged Performance events that have been generated in the past 7 days.

3. Locate an event that you want to analyze and click the event name.

The details page for the event displays.



You can also display the details page for an event by clicking the event name link from the Performance Explorer page and from an alert email.

Analyzing events from user-defined performance thresholds

Events generated from user-defined thresholds indicate that a performance counter for a certain storage object, for example, an aggregate or volume, has crossed the threshold you defined in the policy. This indicates that the cluster object 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.

Responding to user-defined performance threshold events

You can use Unified Manager to investigate performance events caused by a performance counter crossing a user-defined warning or critical threshold. You can also use Unified Manager to check the health of the cluster component to see whether recent health 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 “Latency value of 456 ms/op has triggered a WARNING event based on threshold setting of 400 ms/op” indicates that a latency warning event occurred for the object.

3. Hover your cursor over the policy name to display details about the threshold policy that triggered the event.

This includes the policy name, the performance counter being evaluated, the counter value that must be breached to be considered a critical or warning event, and the duration by which the counter must exceed the value.

4. 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.
5. Follow one of the options below to further investigate the event, to determine whether you need to perform any actions to resolve the performance problem:

Option	Possible investigation actions
Click the Source object name to display the Explorer page for that object.	This page enables you to view the object details and compare this object with other similar storage objects to see whether other storage objects have a performance issue around the same time. For example, to see whether other volumes on the same aggregate are also having a performance issue.
Click the cluster name to display the Cluster Summary page.	This page enables you to view the details for the cluster on which this object resides to see whether other performance issues have occurred around the same time.

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.

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.

What you'll need

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

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).

What you'll need

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

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.

What you'll need

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

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 Manager Dashboard 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.

What you'll need

- 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.

Description:

2 victim volumes are slow due to `vol_osv_siteB2_5` causing contention on MetroCluster resources

Component in Contention:




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.

^ System Diagnosis (Jul 9, 2018, 11:09 AM - Jul 19, 2018, 7:39 AM) ?

Workload Latency



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 (MB/s).

This chart displays, at a high-level, the write throughput (MB/s) pattern for the workloads. You can review the write MB/s pattern to identify abnormal throughput, which might indicate that a workload is over-utilizing 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.

What you'll need

- 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.

What you'll need

- 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.

What you'll need

- 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.

What you'll need

- 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.