



Understanding performance events and alerts

Active IQ Unified Manager 9.14

NetApp
October 16, 2025

Table of Contents

- Understanding performance events and alerts 1
 - Sources of performance events 1
 - Performance event severity types 2
 - Configuration changes detected by Unified Manager 2
 - Types of system-defined performance threshold policies 3
 - Cluster threshold policies 3
 - Node threshold policies 4
 - Aggregate threshold policies 4
 - Workload latency threshold policies 5
 - QoS threshold policies 5
- Performance event analysis and notification 6
 - Event analysis 6
 - Event state 7
 - Event notification 7
 - Event interaction 7
- How Unified Manager determines the performance impact for an event 7
- Cluster components and why they can be in contention 8
- Roles of workloads involved in a performance event 10

Understanding performance events and alerts

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.

You can configure alerts to send email notification automatically when performance events of certain severity types occur.

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.

Performance event severity types

Each performance event is associated with a severity type to help you prioritize the events that require immediate corrective action.

- **Critical**

A performance event occurred that might lead to service disruption if corrective action is not taken immediately.

Critical events are sent from user-defined thresholds only.

- **Warning**

A performance counter for a cluster object is out of normal range and should be monitored to make sure it does not reach the critical severity. Events of this severity do not cause service disruption, and immediate corrective action might not be required.

Warning events are sent from user-defined, system-defined, or dynamic thresholds.

- **Information**

The event occurs when a new object is discovered, or when a user action is performed. For example, when any storage object is deleted or when there are any configuration changes, the event with severity type Information is generated.

Information events are sent directly from ONTAP when it detects a configuration change.

For more information, see the following links:

- [What happens when an event is received](#)
- [What information is contained in an alert email](#)
- [Adding alerts](#)
- [Adding alerts for performance events](#)

Configuration changes detected by Unified Manager

Unified Manager monitors your clusters for configuration changes to help you determine whether a change might have caused or contributed to a performance event. The Performance Explorer pages display a change event icon (●) to indicate the date and time when the change was detected.

You can review the performance charts in the Performance Explorer pages and in the Workload Analysis page to see whether the change event impacted the performance of the selected cluster object. If the change was detected at or around the same time as a performance event, the change might have contributed to the issue, which caused the event alert to trigger.

Unified Manager can detect the following change events, which are categorized as Informational events:

- A volume moves between aggregates.

Unified Manager can detect when the move is in progress, completed, or failed. If Unified Manager is down during a volume move, when it is back up it detects the volume move and displays a change event for it.

- The throughput (MB/s or IOPS) limit of a QoS policy group that contains one or more monitored workloads changes.

Changing a policy group limit can cause intermittent spikes in the latency (response time), which might also trigger events for the policy group. The latency gradually returns to normal and any events caused by the spikes become obsolete.

- A node in an HA pair takes over or gives back the storage of its partner node.

Unified Manager can detect when the takeover, partial takeover, or giveback operation has been completed. If the takeover is caused by a panicked node, Unified Manager does not detect the event.

- An ONTAP upgrade or revert operation is completed successfully.

The previous version and new version are displayed.

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 load imbalance**

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 a cluster to see if any node has exceeded the 30% threshold value for more than 24 hours. This is a warning event.

- **Cluster capacity imbalance**

Identifies situations in which one aggregate has a much higher used capacity than other aggregates in the cluster, and therefore potentially affecting space required for operations.

It does this by comparing the used capacity value for all aggregates in the cluster to see if there is a difference of 70% between any aggregates. 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:

- **Performance Capacity Used Threshold Breached**

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.

Performance event analysis and notification

Performance events notify you about I/O performance issues on a workload caused by contention on a cluster component. Unified Manager analyzes the event to identify all workloads involved, the component in contention, and whether the event is still an issue that you might need to resolve.

Unified Manager monitors the I/O latency (response time) and IOPS (operations) for volumes on a cluster. When other workloads overuse a cluster component, for example, the component is in contention and cannot perform at an optimal level to meet workload demands. The performance of other workloads that are using the same component might be impacted, causing their latencies to increase. If the latency crosses the dynamic performance threshold, Unified Manager triggers a performance event to notify you.

Event analysis

Unified Manager performs the following analyses, using the previous 15 days of performance statistics, to identify the victim workloads, bully workloads, and the cluster component involved in an event:

- Identifies victim workloads whose latency has crossed the dynamic performance threshold, which is the upper boundary of the latency forecast:
 - For volumes on HDD or Flash Pool hybrid aggregates (local tier), events are triggered only when the latency is greater than 5 milliseconds (ms) and the IOPS are more than 10 operations per second (ops/sec).
 - For volumes on all-SSD aggregates or FabricPool aggregates (cloud tier), events are triggered only when the latency is greater than 1 ms and the IOPS are more than 100 ops/sec.
- Identifies the cluster component in contention.



If the latency of victim workloads at the cluster interconnect is greater than 1 ms, Unified Manager treats this as significant and triggers an event for the cluster interconnect.

- Identifies the bully workloads that are overusing the cluster component and causing it to be in contention.
- Ranks the workloads involved, based on their deviation in utilization or activity of a cluster component, to determine which bullies have the highest change in usage of the cluster component and which victims are the most impacted.

An event might occur for only a brief moment and then correct itself after the component it is using is no longer in contention. A continuous event is one that reoccurs for the same cluster component within a five-minute interval and remains in the active state. For continuous events, Unified Manager triggers an alert after detecting the same event during two consecutive analysis intervals.

When an event is resolved, it remains available in Unified Manager as part of the record of past performance issues for a volume. Each event has a unique ID that identifies the event type and the volumes, cluster, and cluster components involved.



A single volume can be involved in more than one event at the same time.

Event state

Events can be in one of the following states:

- **Active**

Indicates that the performance event is currently active (new or acknowledged). The issue causing the event has not corrected itself or has not been resolved. The performance counter for the storage object remains above the performance threshold.

- **Obsolete**

Indicates that the event is no longer active. The issue causing the event has corrected itself or has been resolved. The performance counter for the storage object is no longer above the performance threshold.

Event notification

The events are displayed on the Dashboard page and on many other pages in the user interface, and alerts for those events are sent to specified email addresses. You can view detailed analysis information about an event and get suggestions for resolving it on the Event details page and on the Workload Analysis page.

Event interaction

On the Event details page and on the Workload Analysis page, you can interact with events in the following ways:

- Moving the mouse over an event displays a message that shows the date and time when the event was detected.

If there are multiple events for the same time period, the message shows the number of events.

- Clicking a single event displays a dialog box that shows more detailed information about the event, including the cluster components that are involved.

The component in contention is circled and highlighted red. You can click **View full analysis** to view the full analysis on the Event details page. If there are multiple events for the same time period, the dialog box shows details about the three most recent events. You can click an event to view the event analysis on the Event details page.

How Unified Manager determines the performance impact for an event

Unified Manager uses the deviation in activity, utilization, write throughput, cluster component usage, or I/O latency (response time) for a workload to determine the level of impact to workload performance. This information determines the role of each workload in the event and how they are ranked on the Event details page.

Unified Manager compares the last analyzed values for a workload to the expected range (latency forecast) of values. The difference between the values last analyzed and the expected range of values identifies the workloads whose performance was most impacted by the event.

For example, suppose a cluster contains two workloads: Workload A and Workload B. The latency forecast for

Workload A is 5-10 milliseconds per operation (ms/op) and its actual latency is usually around 7 ms/op. The latency forecast for Workload B is 10-20 ms/op and its actual latency is usually around 15 ms/op. Both workloads are well within their latency forecast. Due to contention on the cluster, the latency of both workloads increases to 40 ms/op, crossing the dynamic performance threshold, which is the upper bounds of the latency forecast, and triggering events. The deviation in latency, from the expected values to the values above the performance threshold, for Workload A is around 33 ms/op, and the deviation for Workload B is around 25 ms/op. The latency of both workloads spike to 40 ms/op, but Workload A had the bigger performance impact because it had the higher latency deviation at 33 ms/op.

On the Event details page, in the System Diagnosis section, you can sort workloads by their deviation in activity, utilization, or throughput for a cluster component. You can also sort workloads by latency. When you select a sort option, Unified Manager analyzes the deviation in activity, utilization, throughput, or latency since the event was detected from the expected values to determine the workload sort order. For the latency, the red dots (●) indicate a performance threshold crossing by a victim workload, and the subsequent impact to the latency. Each red dot indicates a higher level of deviation in latency, which helps you identify the victim workloads whose latency was impacted the most by an event.

Cluster components and why they can be in contention

You can identify cluster performance issues when a cluster component goes into contention. The performance of workloads that use the component slow down and their response time (latency) for client requests increases, which triggers an event in Unified Manager.

A component that is in contention cannot perform at an optimal level. Its performance has declined, and the performance of other cluster components and workloads, called *victims*, might have increased latency. To bring a component out of contention, you must reduce its workload or increase its ability to handle more work, so that the performance can return to normal levels. Because Unified Manager collects and analyzes workload performance in five-minute intervals, it detects only when a cluster component is consistently overused. Transient spikes of overusage that last for only a short duration within the five-minute interval are not detected.

For example, a storage aggregate might be under contention because one or more workloads on it are competing for their I/O requests to be fulfilled. Other workloads on the aggregate can be impacted, causing their performance to decrease. To reduce the amount of activity on the aggregate, there are different steps you can take, such as moving one or more workloads to a less busy aggregate or node, to lessen the overall workload demand on the current aggregate. For a QoS policy group, you can adjust the throughput limit, or move workloads to a different policy group, so that the workloads are no longer being throttled.

Unified Manager monitors the following cluster components to alert you when they are in contention:

- **Network**

Represents the wait time of I/O requests by the external networking protocols on the cluster. The wait time is time spent waiting for “transfer ready” transactions to finish before the cluster can respond to an I/O request. If the network component is in contention, it means high wait time at the protocol layer is impacting the latency of one or more workloads.

- **Network Processing**

Represents the software component in the cluster involved with I/O processing between the protocol layer and the cluster. The node handling network processing might have changed since the event was detected. If the network processing component is in contention, it means high utilization at the network processing node is impacting the latency of one or more workloads.

When using an All SAN Array cluster in an active-active configuration, the network processing latency value is displayed for both nodes so you can verify the nodes are sharing the load equally.

- **QoS Limit Max**

Represents the throughput maximum (peak) setting of the storage Quality of Service (QoS) policy group assigned to the workload. If the policy group component is in contention, it means all workloads in the policy group are being throttled by the set throughput limit, which is impacting the latency of one or more of those workloads.

- **QoS Limit Min**

Represents the latency to a workload that is being caused by QoS throughput minimum (expected) setting assigned to other workloads. If the QoS minimum set on certain workloads use the majority of the bandwidth to guarantee the promised throughput, other workloads will be throttled and see more latency.

- **Cluster Interconnect**

Represents the cables and adapters with which clustered nodes are physically connected. If the cluster interconnect component is in contention, it means high wait time for I/O requests at the cluster interconnect is impacting the latency of one or more workloads.

- **Data Processing**

Represents the software component in the cluster involved with I/O processing between the cluster and the storage aggregate that contains the workload. The node handling data processing might have changed since the event was detected. If the data processing component is in contention, it means high utilization at the data processing node is impacting the latency of one or more workloads.

- **Volume Activation**

Represents the process that tracks the usage of all active volumes. In large environments where more than 1000 volumes are active, this process tracks how many critical volumes need to access resources through the node at the same time. When the number of concurrent active volumes exceeds the recommended maximum threshold, some of the non-critical volumes will experience latency as identified here.

- **MetroCluster Resources**

Represents the MetroCluster resources, including NVRAM and interswitch links (ISLs), used to mirror data between clusters in a MetroCluster configuration. If the MetroCluster component is in contention, it means high write throughput from workloads on the local cluster or a link health issue is impacting the latency of one or more workloads on the local cluster. If the cluster is not in a MetroCluster configuration, this icon is not displayed.

- **Aggregate or SSD Aggregate Ops**

Represents the storage aggregate on which the workloads are running. If the aggregate component is in contention, it means high utilization on the aggregate is impacting the latency of one or more workloads. An aggregate consists of all HDDs, or a mix of HDDs and SSDs (a Flash Pool aggregate), or a mix of HDDs and a cloud tier (a FabricPool aggregate). An "SSD Aggregate" consists of all SSDs (an all-flash aggregate), or a mix of SSDs and a cloud tier (a FabricPool aggregate).

- **Cloud Latency**

Represents the software component in the cluster involved with I/O processing between the cluster and the cloud tier on which user data is stored. If the cloud latency component is in contention, it means that a

large amount of reads from volumes that are hosted on the cloud tier are impacting the latency of one or more workloads.

- **Sync SnapMirror**

Represents the software component in the cluster involved with replicating user data from the primary volume to the secondary volume in a SnapMirror Synchronous relationship. If the sync SnapMirror component is in contention, it means that the activity from SnapMirror Synchronous operations are impacting the latency of one or more workloads.

Roles of workloads involved in a performance event

Unified Manager uses roles to identify the involvement of a workload in a performance event. The roles include victims, bullies, and sharks. A user-defined workload can be a victim, bully, and shark at the same time.

Role	Description
Victim	A user-defined workload whose performance has decreased due to other workloads, called bullies, that are over-using a cluster component. Only user-defined workloads are identified as victims. Unified Manager identifies victim workloads based on their deviation in latency, where the actual latency, during an event, has greatly increased from its latency forecast (expected range).
Bully	A user-defined or system-defined workload whose over-use of a cluster component has caused the performance of other workloads, called victims, to decrease. Unified Manager identifies bully workloads based on their deviation in usage of a cluster component, where the actual usage, during an event, has greatly increased from its expected range of usage.
Shark	A user-defined workload with the highest usage of a cluster component compared to all workloads involved in an event. Unified Manager identifies shark workloads based on their usage of a cluster component during an event.

Workloads on a cluster can share many of the cluster components, such as aggregates and the CPU for network and data processing. When a workload, such as a volume, increases its usage of a cluster component to the point that the component cannot efficiently meet workload demands, the component is in contention. The workload that is over-using a cluster component is a bully. The other workloads that share those components, and whose performance is impacted by the bully, are the victims. Activity from system-defined workloads, such as deduplication or Snapshot copies, can also escalate into “bullying”.

When Unified Manager detects an event, it identifies all workloads and cluster components involved, including the bully workloads that caused the event, the cluster component that is in contention, and the victim

workloads whose performance has decreased due to the increased activity of bully workloads.



If Unified Manager cannot identify the bully workloads, it only alerts on the victim workloads and the cluster component involved.

Unified Manager can identify workloads that are victims of bully workloads, and also identify when those same workloads become bully workloads. A workload can be a bully to itself. For example, a high-performing workload that is being throttled by a policy group limit causes all workloads in the policy group to be throttled, including itself. A workload that is a bully or a victim in an ongoing performance event might change its role or no longer be a participant in 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.