



Understanding performance events and alerts

OnCommand Unified Manager 9.5

NetApp
October 23, 2024

Table of Contents

Understanding performance events and alerts	1
Sources of performance events	1
Performance event severity types	1
Configuration changes detected by Unified Manager	2
What happens when an event is received	3
What information is contained in an alert email	4
Adding alerts	5
Adding alerts for performance events	7
Types of system-defined performance threshold policies	8

Understanding performance events and alerts

Performance events are notifications that Unified Manager generates automatically when a predefined condition occurs, or when a performance counter value crosses a threshold. Events help you identify performance issues in the clusters that are monitored.

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 volume 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.

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 Performance/Volume Details 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 (MBps 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 back 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.

What happens when an event is received


When Unified Manager receives an event, it is displayed in the Dashboards/Overview page, in the Summary and Explorer tabs of the Performance/Cluster page, in the Events inventory page, and in the object-specific inventory page (for example, the Health/Volumes inventory page).

When Unified Manager detects multiple continuous occurrences of the same event condition for the same cluster component, it treats all occurrences as a single event, not as separate events. The duration of the event is incremented to indicate that the event is still active.

Depending on how you configure settings in the Configuration/Alerting page, you can notify other users about these events. The alert causes the following actions to be initiated:

- An email about the event can be sent to all Unified Manager Administrator users.
- The event can be sent to additional email recipients.
- An SNMP trap can be sent to the trap receiver.
- A custom script can be executed to perform an action.

This workflow is shown in the following diagram.



What information is contained in an alert email

Unified Manager alert emails provide the type of event, the severity of the event, the name of the policy that was breached to cause the event, and a description of the event. The email message also provides a hyperlink for each event that enables you to view the details page for the event in the UI.

Alert emails are sent to all users who have subscribed to receive alerts.

If a performance counter or capacity value has a large change during a collection period, it could cause both a critical and a warning event to be triggered at the same time for the same threshold policy. In this case, you may receive one email for the warning event and one for the critical event. This is because Unified Manager enables you to subscribe separately to receive alerts for warning and critical threshold breaches.



After upgrading to Unified Manager 7.2, or greater, links to events and alerts from emails that were send from older versions of Unified Manager will no longer work because of a change in the event and alert URLs.

A sample alert email is shown below:

From: 10.11.12.13@company.com
Sent: Tuesday, May 1, 2018 7:45 PM
To: sclaus@company.com; user1@company.com
Subject: Alert from OnCommand Unified Manager: Thin-Provisioned Volume Space At Risk (State: New)

A risk was generated by 10.11.12.13 that requires your attention.

Risk - Thin-Provisioned Volume Space At Risk

Impact Area - Capacity

Severity - Warning

State - New

Source - svm_n1:/sm_vol_23

Cluster Name - fas3250-39-33-37

Cluster FQDN - fas3250-39-33-37-cm.company.com

Trigger Condition - The thinly provisioned capacity of the volume is 45.73% of the available space on the host aggregate. The capacity of the volume is at risk because of aggregate capacity issues.

Event details:

<https://10.11.12.13:443/events/94>

Source details:

<https://10.11.12.13:443/health/volumes/106>

Alert details:

<https://10.11.12.13:443/alerting/1>

Adding alerts

You can configure alerts to notify you when a particular event is generated. You can configure alerts for a single resource, for a group of resources, or for events of a particular severity type. You can specify the frequency with which you want to be notified and associate a script to the alert.

Before you begin

- You must have configured notification settings such as the user email address, SMTP server, and SNMP trap host to enable the Unified Manager server to use these settings to send notifications to users when an event is generated.
- You must know the resources and events for which you want to trigger the alert, and the user names or email addresses of the users that you want to notify.
- If you want to have a script execute based on the event, you must have added the script to Unified Manager by using the Management/Scripts page.

- You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can create an alert directly from the Event details page after receiving an event in addition to creating an alert from the Configuration/Alerting page, as described here.

Steps

1. In the left navigation pane, click **Configuration > Alerting**.
2. In the **Configuration/Alerting** page, click **Add**.
3. In the **Add Alert** dialog box, click **Name**, and enter a name and description for the alert.
4. Click **Resources**, and select the resources to be included in or excluded from the alert.

You can set a filter by specifying a text string in the **Name contains** field to select a group of resources. Based on the text string that you specify, the list of available resources displays only those resources that match the filter rule. The text string that you specify is case-sensitive.

If a resource conforms to both the include and exclude rules that you have specified, the exclude rule takes precedence over the include rule, and the alert is not generated for events related to the excluded resource.

5. Click **Events**, and select the events based on the event name or event severity type for which you want to trigger an alert.



To select more than one event, press the Ctrl key while you make your selections.

6. Click **Actions**, and select the users that you want to notify, choose the notification frequency, choose whether an SNMP trap will be sent to the trap receiver, and assign a script to be executed when an alert is generated.



If you modify the email address that is specified for the user and reopen the alert for editing, the Name field appears blank because the modified email address is no longer mapped to the user that was previously selected. Also, if you modified the email address of the selected user from the Management/Users page, the modified email address is not updated for the selected user.

You can also choose to notify users through SNMP traps.

7. Click **Save**.

Example of adding an alert

This example shows how to create an alert that meets the following requirements:

- Alert name: HealthTest
- Resources: includes all volumes whose name contains “abc” and excludes all volumes whose name contains “xyz”
- Events: includes all critical health events
- Actions: includes “sample@domain.com”, a “Test” script, and the user has to be notified every 15 minutes

Perform the following steps in the Add Alert dialog box:

1. Click **Name**, and enter **HealthTest** in the **Alert Name** field.
2. Click **Resources**, and in the Include tab, select **Volumes** from the drop-down list.
 - a. Enter **abc** in the **Name contains** field to display the volumes whose name contains "abc".
 - b. Select **<<All Volumes whose name contains 'abc'>>** from the Available Resources area, and move it to the Selected Resources area.
 - c. Click **Exclude**, and enter **xyz** in the **Name contains** field, and then click **Add**.
3. Click **Events**, and select **Critical** from the Event Severity field.
4. Select **All Critical Events** from the Matching Events area, and move it to the Selected Events area.
5. Click **Actions**, and enter **sample@domain.com** in the Alert these users field.
6. Select **Remind every 15 minutes** to notify the user every 15 minutes.

You can configure an alert to repeatedly send notifications to the recipients for a specified time. You should determine the time from which the event notification is active for the alert.

7. In the Select Script to Execute menu, select **Test** script .
8. Click **Save**.

Adding alerts for performance events

You can configure alerts for individual performance events just like any other events received by Unified Manager. Additionally, if you want to treat all performance events alike and have email sent to the same person, you can create a single alert to notify you when any critical or warning performance events are triggered.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The example below shows how to create an event for all critical latency, IOPS, and MBps events. You can use this same methodology to select events from all performance counters, and for all warning events.

Steps

1. In the left navigation pane, click **Configuration > Alerting**.
2. In the **Configuration/Alerting** page, click **Add**.
3. In the **Add Alert** dialog box, click **Name**, and enter a name and description for the alert.
4. Do not select any resources on the **Resources** page.

Because no resources are selected, the alert is applied to all clusters, aggregates, volumes, and so on, for which these events are received.

5. Click **Events** and perform the following actions:

- a. In the Event Severity list, select **Critical**.
 - b. In the Event Name Contains field, enter `latency` and then click the arrow to select all the matching events.
 - c. In the Event Name Contains field, enter `iops` and then click the arrow to select all the matching events.
 - d. In the Event Name Contains field, enter `mbps` and then click the arrow to select all the matching events.
6. Click **Actions** and then select the name of the user who will receive the alert email in the **Alert these users** field.
7. Configure any other options on this page for issuing SNMP traps and executing a script.
8. Click **Save**.

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 individual policies from the Configuration/Manage Events page.

Node threshold policies

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

- **Node resources over-utilized**

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

For nodes installed with ONTAP 8.3.x and earlier software, it does this by looking for nodes that are using more than 85% of their CPU and RAM resources (node utilization) for more than 30 minutes.

For nodes installed with ONTAP 9.0 and later software, it does this by looking for nodes that are using more than 100% of their performance capacity for more than 30 minutes.

- **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. This is an informational event.

For nodes installed with ONTAP 8.3.x and earlier software, it does this by looking at the CPU and RAM usage for the two nodes in the HA pair. If the combined node utilization of the two nodes exceeds 140% for more than one hour, then a controller failover will impact workload latencies.

For nodes installed with ONTAP 9.0 and later software, 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 one hour, then a controller failover will impact workload latencies.

- **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.

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 MBps). Unified Manager triggers an event when the workload throughput value is 15% less than the configured QoS value.

- **QoS Max IOPS or MBps threshold**

Identifies volumes and LUNs that have exceeded their QoS maximum IOPS or MBps 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 MBps 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 MBps value based on the size of each volume. Then it looks for volumes that have exceeded the QoS max MBps during each performance collection period for the previous hour.



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

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—with 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.