



Monitoring performance using the Performance Inventory pages

OnCommand Unified Manager 9.5

NetApp
February 12, 2024

This PDF was generated from <https://docs.netapp.com/us-en/oncommand-unified-manager-95/performance-checker/concept-object-monitoring-using-the-object-inventory-performance-pages.html> on February 12, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Monitoring performance using the Performance Inventory pages 1
 - Object monitoring using the Performance object inventory pages..... 1
 - Refining Performance inventory page contents 2
 - Understanding the Unified Manager recommendations to tier data to the cloud 4

Monitoring performance using the Performance Inventory pages

The object inventory performance pages display performance information, performance events, and object health for all objects within an object type category. This provides you with an at-a-glance overview of the performance status of each object within a cluster, for example, for all nodes or all volumes.

Object inventory performance pages provide a high-level overview of object status, enabling you to assess the overall performance of all objects and compare object performance data. You can refine the content of object inventory pages by searching, sorting, and filtering. This is beneficial when monitoring and managing object performance, because it enables you to quickly locate objects with performance issues and to begin the troubleshooting process.

Performance / Nodes

Last updated: 07:43 AM, 03 Nov

 Refresh

Latency, IOPS, MBps, Utilization are based on hourly samples averaged over the previous 72 hours

Search Node data

Filtering

Export

Assign Performance Threshold Policy

Clear Performance Threshold Policy

<input type="checkbox"/>	Status	Node	Latency	IOPS	MBps	Flash Cache F	Perf. Capacity	Utilization	Free Capacity	Total Capacity	Cluster	Policy	
<input type="checkbox"/>	<div><div></div><div></div></div>	opm-mobility-02	0.704 ms/op	5,011 IOPS	49.2 MBps	N/A	23%	21%	93,708 GB	103,748 GB	opm-m...lity		
<input type="checkbox"/>	<div><div></div><div></div></div>	opm-vitality-02	0.357 ms/op	< 1 IOPS	46.8 MBps	0%	N/A	20%	972 GB	3,563 GB	opm-vitality		
<input type="checkbox"/>	<div><div></div><div></div></div>	opm-longevity-01	0.523 ms/op	456 IOPS	20.9 MBps	N/A	N/A	6%	2,162 GB	2,953 GB	opm-lo...vity		
<input type="checkbox"/>	<div><div></div><div></div></div>	opm-mobility-01	61.3 ms/op	2,750 IOPS	25.7 MBps	N/A	9%	8%	80,175 GB	90,361 GB	opm-m...lity	headroom	
<input checked="" type="checkbox"/>	<div><div></div><div></div></div>	opm-vitality-01	15.2 ms/op	3,575 IOPS	146 MBps	0%	N/A	25%	2,835 GB	4,800 GB	opm-vitality		
<input type="checkbox"/>	<div><div></div><div></div></div>	opm-longevity-02	0.106 ms/op	< 1 IOPS	7.93 MBps	N/A	N/A	8%	5,743 GB	6,762 GB	opm-lo...vity		

By default, objects on the performance inventory pages are sorted based on object performance criticality. Objects with new critical performance events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed. All performance data is based on a 72-hour average.

You can easily navigate from the object inventory performance page to an object details page by clicking the object name in the object name column. For example, on the Performance/Nodes inventory page, you would click a node object in the **Nodes** column. The object details page provides in-depth information and detail about the selected object, including side-by-side comparison of active events.

Object monitoring using the Performance object inventory pages

The Performance object inventory pages enable you to monitor object performance based on the values of specific performance counters or based on performance events. This is beneficial because identifying objects with performance events enables you to investigate the cause of cluster performance issues.

The Performance object inventory pages display the associated counters, associated objects, and performance threshold policies for all objects in all clusters. These pages also enable you to apply

performance threshold policies to objects. You can sort the page based on any column, and you can search across all object names or data.

You can export data from these pages to a comma-separated values (`.csv`) file by using the **Export** button, and then use the exported data to build reports.

Refining Performance inventory page contents

The inventory pages for performance objects contain tools to help you refine object inventory data content, enabling you to locate specific data quickly and easily.

Information contained within the Performance object inventory pages can be extensive, often spanning multiple pages. This kind of comprehensive data is excellent for monitoring, tracking, and improving performance; however, locating specific data requires tools to enable you to quickly locate the data for which you are looking. Therefore, the Performance object inventory pages contain functionality for searching, sorting, and filtering. Additionally, searching and filtering can work together to further narrow your results.

Searching on Object Inventory Performance pages

You can search strings on Object Inventory Performance pages. Use the **Search** field located at the top right of the page to quickly locate data based on either object name or policy name. This enables you to quickly locate specific objects and their associated data, or to quickly locate policies and view associated policy object data.

Steps

1. Perform one of the following options, based on your search requirements:

To locate this...	Type this...
A specific object	The object name into the Search field, and click Search . The object for which you searched and its related data is displayed.
A user-defined performance threshold policy	All or part of the policy name into the Search field, and click Search . The objects assigned to the policy for which you searched are displayed.

Sorting on the Object Inventory Performance pages

You can sort all data on Object Inventory Performance pages by any column in ascending or descending order. This enables you to quickly locate object inventory data, which is helpful when examining performance or beginning a troubleshooting process.

About this task

The selected column for sorting is indicated by a highlighted column heading name and an arrow icon indicating the sorting direction at the right of the name. An up arrow indicates ascending order; a down arrow indicates descending order. The default sort order is by **Status** (event criticality) in descending order, with the most critical performance events listed first.

Steps

1. You can click a column name to toggle the sort order of the column in ascending or descending order.

The Object Inventory Performance page contents are sorted in ascending or descending order, based on the selected column.

Filtering data in the Object Inventory Performance pages

You can filter data in the Object Inventory Performance pages to quickly locate data based on specific criteria. You can use filtering to narrow the contents of the Object Inventory Performance pages to show only the results you have specified. This provides a very efficient method of displaying only the performance data in which you are interested.

About this task

You can use the Filtering panel to customize the grid view based on your preferences. Available filter options are based on the correlated object type being viewed in the grid. If filters are currently applied, an asterisk (*) displays at the left of the Filtering control.

Four types of filter parameters are supported.

Parameter	Validation
String (text)	The operators are contains and starts with .
Number	The operators are greater than and less than .
Resource	The operators are name contains and name starts with .
Status	The operators are is and is not .

All three fields are required for each filter; the available filters reflect the filterable columns on the current page. The maximum number of filters you can apply is four. Filtered results are based on combined filter parameters. Filtered results apply to all pages in your filtered search, not just the page currently displayed.

You can add filters using the Filtering panel.

1. At the top of the page, click **Filtering**. The Filtering panel displays.
2. In the Filtering panel, click the left drop-down list, and select an object name: for example, *Cluster*, or a performance counter.
3. Click the center drop-down list, and select the boolean operator **name contains** or **name starts with** if the first selection was an object name. If the first selection was a performance counter, select **greater than** or **less than**. If the first selection was **Status**, select **is** or **is not**.
4. If your search criteria requires a numeric value, up and down arrow buttons display in the field at the right. You can click the up and down arrow buttons to display your desired numeric value.
5. If required, type your non-numeric search criteria in the text field at the right.

6. To add filters, click **Add Filter**. An additional filter field displays. Complete this filter using the process described in the preceding steps. Note that upon adding your fourth filter, the **Add Filter** button no longer displays.
7. Click **Apply Filter**. The filter options are applied to the grid and an asterisk (*) is displayed in the Filtering button.
8. Use the Filtering panel to remove individual filters by clicking the trash icon at the right of the filter to be removed.
9. To remove all filters, click **Reset** at the bottom of the filtering panel.

Filtering example

The illustration shows the Filtering panel with three filters. The **Add Filter** button displays when you have fewer than the maximum of four filters.

The screenshot shows a 'Filtering' panel with three filter rows. Each row has three columns: a field type dropdown, a comparison operator dropdown, and a value input field. The first row is 'MBps' (field type), 'greater than' (operator), and '5' (value). The second row is 'Node' (field type), 'name starts with' (operator), and 'test' (value). The third row is 'Type' (field type), 'is' (operator), and 'FCP Port' (value). To the right of each value field is a trash icon for removal. Below the filter rows is a '+ Add Filter' button. At the bottom right are 'Cancel' and 'Apply Filter' buttons.

After clicking **Apply Filter**, the Filtering panel closes and applies your filters.



Understanding the Unified Manager recommendations to tier data to the cloud

The Performance/Volumes inventory page displays information related to the size of the user data stored on the volume that is inactive (cold). In some cases, Unified Manager identifies certain volumes that would benefit by tiering the inactive data to the cloud tier (cloud provider or StorageGRID) of a FabricPool-enabled aggregate.



FabricPool was introduced in ONTAP 9.2, so if you are using a version of ONTAP software prior to 9.2, the Unified Manager recommendation to tier data requires upgrading your ONTAP software. Additionally, the `auto` tiering policy was introduced in ONTAP 9.4, so if the recommendation is to use the `auto` tiering policy, you must upgrade to ONTAP 9.4 or greater.

The following three fields on Performance/Volumes inventory page provide information about whether you can improve your storage system's disk utilization and save space on the performance tier by moving inactive data to the cloud tier.

- **Tiering Policy**

The tiering policy determines whether the data on the volume remains on the performance tier or whether some of the data is moved from the performance tier to the cloud tier.

The value in this field indicates the tiering policy set on the volume, even if the volume does not currently reside on a FabricPool aggregate. The tiering policy takes effect only when the volume is on a FabricPool aggregate.

- **Cold Data**

The cold data displays the size of the user data stored on the volume that is inactive (cold).

A value is displayed here only when using ONTAP 9.4 or greater software because it requires that the aggregate on which the volume is deployed has the `inactive data reporting` parameter set to `enabled`, and that the minimum number of cooling days threshold has been met (for volumes that use the `snapshot-only` or `auto` tiering policy). Otherwise the value is listed as “N/A”.

- **Cloud Recommendation**

After enough information has been captured about the data activity on the volume, Unified Manager may determine there is no action required, or that you could save space on the performance tier by tiering inactive data to the cloud tier.



The Cold Data field is updated every 15 minutes, but the Cloud Recommendation field is updated every 7 days when the cold data analysis is performed on the volume. Therefore, the exact amount of cold data may differ between the fields. The Cloud Recommendation field displays the date when the analysis was run.

When Inactive Data Reporting is enabled, the Cold Data field displays the exact amount of inactive data. Without the inactive data reporting capability Unified Manager uses performance statistics to determine if data is inactive on a volume. The amount of inactive data is not displayed in the Cold Data field in this case, but it is displayed when you hover your cursor over the word **Tier** to view the cloud recommendation.

The cloud recommendations you will see are:

- **Learning.** Not enough data has been collected to make a recommendation.
- **Tier.** Analysis has determined that the volume contains inactive (cold) data and that you should configure the volume to move that data to the cloud tier. In some cases this may require that you move the volume to a FabricPool-enabled aggregate first. In other cases where the volume is already on a FabricPool aggregate, you just have to change the tiering policy.
- **No Action.** Either the volume has very little inactive data, the volume is already set to the “auto” tiering policy on a FabricPool aggregate, or the volume is a data protection volume. This value is also displayed when the volume is offline or when it is being used in a MetroCluster configuration.

To move a volume, or to change the volume tiering policy or the aggregate inactive data reporting settings, use OnCommand System Manager, the ONTAP CLI commands, or a combination of these tools.

If you are logged in to Unified Manager with the OnCommand Administrator or Storage Administrator role, the **Configure Volume** link is available in the cloud recommendation when you hover your cursor over the word **Tier**. Click this button to open the Volumes page in System Manager to make the recommended change.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.