



Managing performance using performance capacity and available IOPS information

OnCommand Unified Manager 9.5

NetApp
October 23, 2024

Table of Contents

- Managing performance using performance capacity and available IOPS information 1
 - What performance capacity used is 1
 - What the performance capacity used value means 2
 - What available IOPS is 3
 - Viewing node and aggregate performance capacity used values 4
 - Viewing node and aggregate available IOPS values 5
 - Viewing performance capacity counter charts to identify issues 6
 - Performance capacity used performance threshold conditions 8
 - Using the performance capacity used counter to manage performance 8

Managing performance using performance capacity and available IOPS information

Performance capacity indicates how much throughput you can get out of a resource without surpassing the useful performance of that resource. When viewed using existing performance counters, performance capacity is the point at which you get the maximum utilization from a node or aggregate before latency becomes an issue.

Unified Manager collects performance capacity statistics from nodes and aggregates in each cluster. *Performance capacity used* is the percentage of performance capacity that is currently being used, and *performance capacity free* is the percentage of performance capacity that is still available.

While performance capacity free provides a percentage of the resource that is still available, *available IOPS* tells you the number of IOPS that can be added to the resource before reaching the maximum performance capacity. By using this metric, you can be sure that you can add workloads of a predetermined number of IOPS to a resource.

Monitoring the performance capacity information has the following benefits:

- Assists with workflow provisioning and balancing.
- Helps you prevent overloading a node or pushing its resources beyond the optimal point, thus reducing the need to troubleshoot.
- Helps you determine with greater precision where additional storage equipment might be needed.

What performance capacity used is

The performance capacity used counter helps you to identify whether the performance of a node or an aggregate is reaching a point where the performance might degrade if the workloads increase. It can also show you if a node or aggregate is currently being overused during specific periods of time. Performance capacity used is similar to utilization, but the former provides more insight about the available performance capabilities in a physical resource for a specific workload.



Performance capacity data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.

The optimal used performance capacity is the point at which a node or an aggregate has optimal utilization and latency (response time), and is being used efficiently. A sample latency versus utilization curve is shown for an aggregate in the following figure.



In this example, the *operational point* identifies that the aggregate is currently operating at 50% utilization with latency of 1.0 ms/op. Based on the statistics captured from the aggregate, Unified Manager determines that additional performance capacity is available for this aggregate. In this example, the *optimal point* is identified as the point when the aggregate is at 80% utilization with latency of 2.0 ms/op. Therefore, you can add more volumes and LUNs to this aggregate so that your systems are used more efficiently.

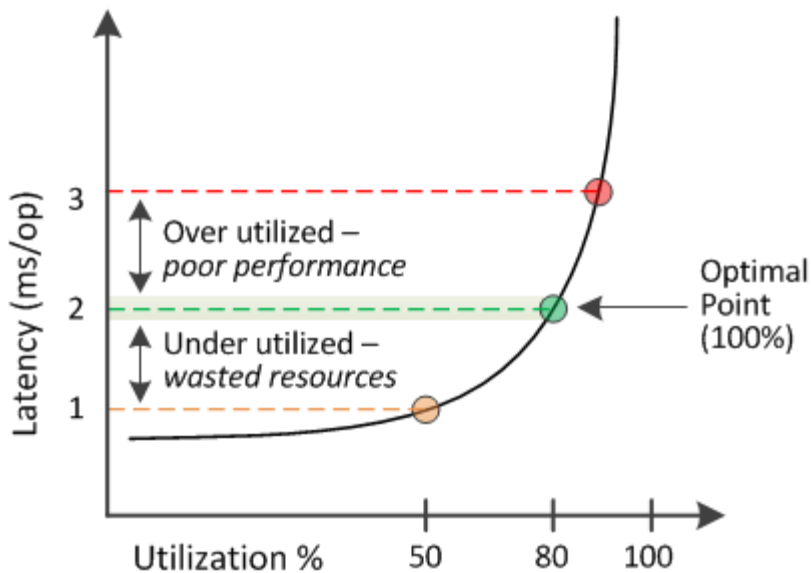
The performance capacity used counter is expected to be a larger number than the “utilization” counter because performance capacity adds in the impact on latency. For example, if a node or aggregate is 70% used, the performance capacity value may be in the 80% to 100% range, depending on the latency value.

In some cases, however, the utilization counter may be higher on the Dashboards/Performance page. This is normal because the dashboard refreshes the current counter values at each collection period; it does not display averages over a period of time like the other pages in the Unified Manager user interface. The performance capacity used counter is best used as an indicator of performance averaged over a period of time, whereas the utilization counter is best used for determining the instantaneous usage of a resource.

What the performance capacity used value means

The performance capacity used value helps you identify the nodes and aggregates that are currently being overutilized or underutilized. This enables you to redistribute workloads in order to make your storage resources more efficient.

The following figure shows the latency versus utilization curve for a resource and identifies, with colored dots, three areas where the current operational point could be located.



- A performance capacity used percentage equal to 100 is at the optimal point.

Resources are being used efficiently at this point.

- A performance capacity used percentage above 100 indicates that the node or aggregate is overutilized, and that workloads are receiving sub-optimal performance.

No new workloads should be added to the resource, and the existing workloads may need to be redistributed.

- A performance capacity used percentage below 100 indicates that the node or aggregate is underutilized, and that resources are not being used effectively.

More workloads can be added to the resource.



Unlike utilization, the performance capacity used percentage can be above 100%. There is no maximum percentage, but resources will typically be in the 110% to 140% range when they are being overutilized. Higher percentages would indicate a resource with serious issues.

What available IOPS is

The available IOPS counter identifies the remaining number of IOPS that can be added to a node or an aggregate before the resource reaches its limit. The total IOPS that a node can provide is based on the physical characteristics of the node—for example, the number of CPUs, the CPU speed, and the amount of RAM. The total IOPS that an aggregate can provide is based on the physical properties of the disks—for example, a SATA, SAS, or SSD disk.

While the performance capacity free counter provides the percentage of a resource that is still available, the available IOPS counter tells you the exact number of IOPS (workloads) can be added to a resource before reaching the maximum performance capacity.


For example, if you are using a pair of FAS2520 and FAS8060 storage systems, a performance capacity free value of 30% means that you have some free performance capacity. However, that value does not provide

visibility into how many more workloads you can deploy to those nodes. The available IOPS counter may show that you have 500 available IOPS on the FAS8060, but only 100 available IOPS on the FAS2520.



Available IOPS data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.

A sample latency versus IOPS curve for a node is shown in the following figure.



The maximum number of IOPS that a resource can provide is the number of IOPS when the performance capacity used counter is at 100% (the optimal point). The operational point identifies that the node is currently operating at 100K IOPS with latency of 1.0 ms/op. Based on the statistics captured from the node, Unified Manager determines that the maximum IOPS for the node is 160K, which means that there are 60K free or available IOPS. Therefore, you can add more workloads to this node so that your systems are used more efficiently.



When there is minimal user activity in the resource, the available IOPS value is calculated assuming a generic workload based on approximately 4,500 IOPS per CPU core. This is because Unified Manager lacks the data to accurately estimate the characteristics of the workload being served.

Viewing node and aggregate performance capacity used values

You can monitor the performance capacity used values for all nodes or for all aggregates in a cluster, or you can view details for a single node or aggregate.

Performance capacity used values appear in the Performance Dashboard, Performance Inventory pages, Top Performers page, Create Threshold Policy page, Performance Explorer pages, and in detail charts. For example, the Performance/Aggregate Inventory page provides a column Perf. Capacity Used to view the performance capacity used value for all aggregates.

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

Filtering No filter applied

Search Aggregates Data Search

Assign Threshold Policy
Clear Threshold Policy

Status	Aggregate	Latency	IOPS	MBps	Perf. Capacity Used ↓↑	Utilization	Free Capacity	Total Capacity	Cluster	Node	Policy
✓	opm_mo..._agg0	16.3 ms/op	124 IOPS	< 1 MBps	45%	9%	154 GB	3,179 GB	opm-mobility	opm-m...-02	
✓	rt_aggr2	19.8 ms/op	290 IOPS	< 1 MBps	45%	15%	6,692 GB	6,693 GB	opm-mobility	opm-m...-02	
✓	aggr_snap_mirror	13.9 ms/op	267 IOPS	< 1 MBps	38%	12%	6,692 GB	6,693 GB	opm-mobility	opm-m...-02	
✓	sdot_aggr	17.3 ms/op	745 IOPS	< 1 MBps	24%	11%	26,621 GB	26,774 GB	opm-mobility	opm-m...-02	
✓	aggr1	15.5 ms/op	434 IOPS	< 1 MBps	16%	6%	4,390 GB	20,080 GB	opm-mobility	opm-m...-01	
✓	rt_aggr1	22.3 ms/op	267 IOPS	< 1 MBps	11%	6%	6,691 GB	6,693 GB	opm-mobility	opm-m...-01	
✓	aggr2	15.6 ms/op	259 IOPS	1.03 MBps	11%	5%	18,472 GB	20,080 GB	opm-mobility	opm-m...-02	
✓	aggr2	9.52 ms/op	87 IOPS	20.8 MBps	Not Supported	5%	847 GB	984 GB	opm-lo...vity	opm-lo...ty-01	aggr_IOPS
⚠	RTaggr	7.62 ms/op	199 IOPS	34.7 MBps	Not Supported	6%	1,292 GB	1,477 GB	opm-lo...vity	opm-lo...ty-01	aggr_IOPS

The status “N/A” is displayed when nodes are not installed with ONTAP 9.0 or later software.

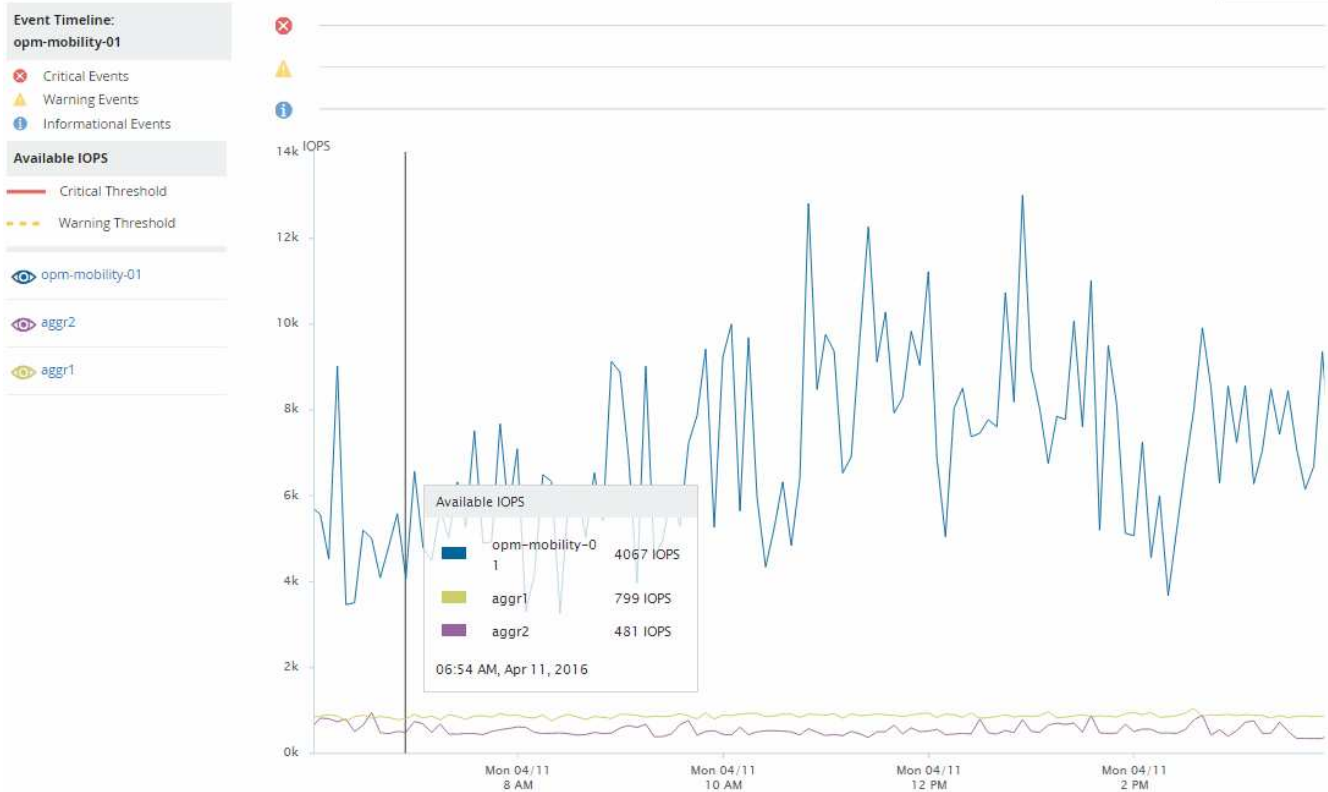
Monitoring the performance capacity used counter enables you to identify the following:

- Whether any nodes or aggregates on any clusters have a high performance capacity used value
- Whether any nodes or aggregates on any clusters have active performance capacity used events
- The nodes and aggregates that have the highest and lowest performance capacity used value in a cluster
- Latency and utilization counter values in conjunction with nodes or aggregates that have high performance capacity used values
- How the performance capacity used values for nodes in an HA pair will be affected if one of the nodes fails
- The busiest volumes and LUNs on an aggregate that has a high performance capacity used value

Viewing node and aggregate available IOPS values

You can monitor the available IOPS values for all nodes or for all aggregates in a cluster, or you can view details for a single node or aggregate.

Available IOPS values appear in the Performance Explorer page charts. For example, when viewing a node in the Performance/Node Explorer page, you can select the “Available IOPS” counter chart from the list so you can compare the available IOPS values for multiple aggregates on that node.



Monitoring the available IOPS counter enables you to identify:

- The nodes or aggregates that have the greatest available IOPS values to help determine where future workloads can be deployed.
- The nodes or aggregates that have the smallest available IOPS values to identify the resources you should monitor for potential future performance issues.
- The busiest volumes and LUNs on an aggregate that has a small available IOPS value.

Viewing performance capacity counter charts to identify issues

You can view performance capacity used charts for nodes and aggregates on the Performance Explorer page. This enables you to view detailed performance capacity data for the selected nodes and aggregates for a specific timeframe.

About this task

The standard counter chart displays the performance capacity used values for the selected nodes or aggregates. The Breakdown counter chart displays the total performance capacity values for the root object separated into usage based on user protocols versus background system processes. Additionally, the amount of free performance capacity is also shown.




Because some background activities associated with system and data management are identified as user workloads and categorized as user protocols, the user protocols percentage may appear artificially high when those processes run. These processes typically run around midnight when cluster usage is low. If you see a spike in user protocol activity around midnight, verify if cluster backup jobs or other background activities are configured to run at that time.

Steps

1. Select the **Explorer** tab from a node or aggregate **Landing** page.
2. In the **Counter Charts** pane, click **Choose charts**, and then select the **Perf. Capacity Used** chart.
3. Scroll down until you can view the chart.

The colors of the standard chart show when the object is in the optimal range (yellow), when the object is underutilized (green), and when the object is overutilized (red). The Breakdown chart shows detailed performance capacity details for the root object only.



4. If you want to view either chart in a full size format, click **Zoom View**.

In this manner you can open multiple counter charts in a separate windows to compare performance capacity used values with IOPS or MBps values over the same timeframe.

Performance capacity used performance threshold conditions

You can create user-defined performance threshold policies so that events are triggered when the performance capacity used value for a node or aggregate exceeds the defined performance capacity used threshold setting.

Additionally, nodes can be configured with a “Performance capacity used takeover” threshold policy. This threshold policy totals the performance capacity used statistics for both nodes in an HA pair to determine whether either node would lack sufficient capacity if the other node fails. Because the workload during failover is the combination of the two partner nodes’ workloads, the same performance capacity used takeover policy can be applied to both nodes.



This performance capacity used equivalency is generally true between nodes. However, if there is significantly more cross-node traffic destined for one of the nodes through its failover partner, the total performance capacity used when running all workloads on one partner node versus the other partner node could be slightly different depending on which node has failed.

The performance capacity used conditions can also be used as secondary performance threshold settings to create a combination threshold policy when defining thresholds for LUNs and volumes. The performance capacity used condition is applied to the aggregate or node on which the volume or LUN resides. For example, you can create a combination threshold policy using the following criteria:

Storage object	Performance counter	Warning threshold	Critical threshold	Duration
Volume	Latency	15 ms/op	25 ms/op	20 minutes

Combination threshold policies cause an event to be generated only when both conditions are breached for the entire duration.

Using the performance capacity used counter to manage performance

Typically, organizations want to operate with a performance capacity used percentage below 100 so that resources are being efficiently used while reserving some additional performance capacity to support peak period demands. You can use threshold policies to customize when alerts are sent for high performance capacity used values.

You can establish specific goals based on your performance requirements. For example, financial services firms might reserve more performance capacity to guarantee the timely execution of trades. These companies might want to set performance capacity used thresholds in the 70-80 percent range. Manufacturing companies with smaller margins might choose to reserve less performance capacity if they are willing to risk performance to better manage IT costs. These companies might set performance capacity used thresholds in the 85-95

percent range.

When the performance capacity used value exceeds the percentage set in a user-defined threshold policy, Unified Manager sends an alert email and adds the event to the Event Inventory page. This enables you to manage potential problems before they impact performance. These events can also be used as indicators that you need to make workload moves and changes within your nodes and aggregates.

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.