



Insights

Cloud Insights

NetApp
July 26, 2024

Table of Contents

- Insights 1
 - Insights 1
 - Insights: Shared Resources Under Stress 1
 - Insights: Kubernetes Namespaces Running out of Space 4
 - Insights: Reclaim ONTAP Cold Storage 5

Insights

Insights

Insights allow you to look into things like resource usage and how it affects other resources, or time-to-full analyses.

A number of Insights are available. Navigate to **Dashboards > Insights** to start diving in. You can view active Insights (Insights that are currently occurring) on the main tab, or inactive Insights on the *Inactive Insights* tab. Inactive Insights are those that were previously active but are no longer occurring.

Insight Types

Shared Resources Under Stress

High-impact workloads can reduce the performance of other workloads in a shared resource. This puts the shared resource under stress. Cloud Insights provides tools to help you investigate resource saturation and impact in your environment. [Learn More](#)

Kubernetes Namespaces Running Out of Space

The Kubernetes Namespaces Running Out of Space Insight gives you a view into workloads on your Kubernetes namespaces that are at risk of running out of space, with an estimate for the number of days remaining before each space becomes full. [Learn More](#)

Reclaim ONTAP Cold Storage

The *Reclaim ONTAP Cold Storage* Insight provides data about cold capacity, potential cost/power savings and recommended action items for volumes on ONTAP systems. [Learn More](#)



This is a *Preview* feature and may change over time as improvements are made. [Learn more](#) about Cloud Insights Preview features.

Insights: Shared Resources Under Stress

High-impact workloads can reduce the performance of other workloads in a shared resource. This puts the shared resource under stress. Cloud Insights provides tools to help you investigate resource saturation and impact in your environment.

Terminology

When talking about workload or resource impact, the following definitions are useful.

A **Demanding Workload** is a workload that is currently identified as impacting other resources in the shared storage pool. These workloads drive higher IOPS (for example), reducing IOPS in the Impacted Workloads. Demanding workloads are sometimes called *high-consuming workloads*.

An **Impacted Workload** is a workload that is affected by a high-consuming workload in the shared Storage Pool. These workloads are experiencing reduced IOPS and/or higher latency, caused by the Demanding Workloads.

Note that if Cloud Insights has not discovered the leading compute workload, the volume or internal volume itself will be recognized as the workload. This applies to both demanding and impacted workloads.

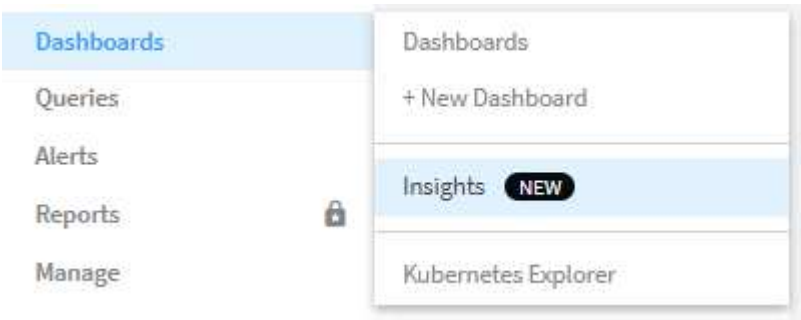
Shared Resource Saturation is the ratio of impacting IOPS to *baseline*.

Baseline is defined as the maximum reported data point for each workload in the hour immediately preceding the detected saturation.

A **Contention** or **Saturation** occurs when IOPS are determined to be affecting other resources or workloads in the shared storage pool.

Demanding Workloads

To start looking into Demanding and impacted workloads in your shared resources, click on **Dashboards > Insights** and select the **Shared Resources Under Stress** Insight.



Cloud Insights displays a list of any workloads where a saturation has been detected. Note that Cloud Insights will show workloads where at least one *demanding resource* **or** *impacted resource* has been detected.

Click on a workload to view the details page for it. The top chart shows the activity on the shared resource (for example, a storage pool) on which the contention/saturation is occurring.



Below that are two charts showing the *demanding* workloads and the workloads that are *impacted* by those demanding workloads.

Demanding Workloads (1) ⓘ

Potentially impacted the shared resource and other related workloads

Contributing IOPS ▾

iops.total (IOPS)



| Workload | Current Contributing IOPS (IOPS) ↓ | Change Since Detection (IOPS) |
|---------------------|------------------------------------|-------------------------------|
| internal-volume-331 | 500.00 | +190.00 |

Impacted Workloads (1) ⓘ

Impacted by changed workloads on the shared resource

Latency ▾

latency.total (ms)



| Workload | Current Latency (ms) ↓ | Change Since Detection (ms) |
|---------------------|------------------------|-----------------------------|
| internal-volume-332 | 200.00 | +110.00 |

Below each table is a list of workloads and/or resources affecting or affected by the contention. Clicking on a resource (for example, a VM) opens a detail page for that resource. Clicking on a workload opens a query page showing the pods involved. Note that if the link opens an empty query, it may be because the affected pod is no longer part of the active contention. You can modify the query's time range to view the pod list in greater or more focused time range.

What do I do to resolve saturation?

There are a number of steps you can take to reduce or eliminate the chance of saturation in your environment. These are shown by expanding the **+Show Recommendations** link on the page. Here are a few things you can try.

- Move high-IOPS consumers

Move the "greedy" workloads to less-saturated Storage Pools. It is recommended to assess the tier and capacity of these pools before moving the workloads, to avoid unnecessary costs or additional contentions.

- Implement a quality of service (QoS) policy

Implementing a QoS policy per workload to ensure enough free resources available will alleviate saturation on the Storage Pool. This is a long-term solution.

- Add additional resources

If the shared resource (for example, Storage Pool) has reached the IOPS saturation point, adding more or faster disks to the pool will ensure enough free resources available to alleviate saturation.

Finally, you can click the **Copy Insight Link** to copy the page url to the clipboard, to more easily share with colleagues.

Insights: Kubernetes Namespaces Running out of Space

Running out of space in your environment is never a good situation. Cloud Insights helps you predict the time you have before Kubernetes persistent volumes become full.

The *Kubernetes Namespaces Running Out of Space* Insight gives you a view into workloads on your Kubernetes namespaces that are at risk of running out of space, with an estimate for the number of days remaining before each persistent volume becomes full.

You can view this Insight by navigating to **Dashboards > Insights**.




Kubernetes Namespaces Running Out of Space (3)

| Description | Estimated Days to Full | Workloads at Risk | Detected ↓ |
|--|------------------------|-------------------|------------|
| 1 workload at risk on es | 35 | 1 | 2 days ago |
| 1 workload at risk on manager | 24 | 1 | 2 days ago |
| 2 workloads at risk on cloudinsights | 1 | 2 | 2 days ago |

Click on a workload to open a detail page for the Insight. On this page you will see a graph showing the workload capacity trends as well as a table showing the following:

- Workload Name
- Persistent Volume affected
- Predicted Time-to-Full in days
- Persistent Volume capacity
- Backend Storage Resource affected, with current capacity used out of total capacity. Clicking this link will open the detailed landing page for the backend volume.

Workloads at risk (2)

|  Workloads | Persistent Volume (pvClaim) | Time to Full (Days) ↓ | Persistent Volume Capacity (GiB) | Backend Storage Resource (Capacity Used) |
|---|-----------------------------|-----------------------|----------------------------------|--|
|  multi (1) | pv1 (pvc1) | 1 | 4.00 | internal-volume-601 60.00% (3.00/5.00 GiB) |
|  taskmanager (1) | pv1 (pvc1) | 1 | 4.00 | internal-volume-601 60.00% (3.00/5.00 GiB) |

What can I do if I'm running out of space?


On the Insight page, click the **+Show Recommendations** to view possible solutions. The easiest option when running out of space is always to add more capacity, and Cloud Insights shows you the optimal capacity to add to increase time-to-full to a target 60-day prediction. Other recommendations are also shown.

Show Recommendations

- 1 Get time to full back up to 60 days by adding more capacity to backend resources
Add to the following resources to bring time-to-full up to ideal capacity.

| Backend Resource ↓ | Current Capacity (time to full) | Recommended Capacity to Add | | Ideal Capacity (time to full) |
|-------------------------------------|---------------------------------|-----------------------------|------------|-------------------------------|
| internal-volume-601 | 2.00 GiB 1 Days | + | 518.79 GiB | = 520.79 GiB 60 Days |

- 2 Use NetApp Astra Trident with your K8s to automatically grow capacity
Astra Trident can keep your capacity lean without risk of running out of space.

[Learn more about](#)  **Astra Trident**

 [Copy Insight Link](#)

It is here also that you can copy a convenient link to this Insight, to bookmark the page or to easily share with your team.

Insights: Reclaim ONTAP Cold Storage

The *Reclaim ONTAP Cold Storage* Insight provides data about cold capacity, potential cost/power savings and recommended action items for volumes on ONTAP systems.

To view these Insights, navigate to **Dashboards > Insights** and take a look at the *Reclaim ONTAP Cold Storage* Insight. Note that this Insight will only list affected storages if Cloud Insights has detected cold storage, otherwise you will see an "all clear" message.

Keep in mind that cold data less than 30 days old is not shown.

Reclaim ONTAP Cold Storage (3)

| Description | Cold data storage(TiB) | Workloads with cold data | Detected ↓ |
|---|------------------------|--------------------------|-------------|
| 0.30 TiB of cold data on storage rtp-sa-cl04 | 0.30 | 45 | an hour ago |
| 1.22 TiB of cold data on storage umeng-aff300-01-02 | 1.22 | 84 | 16 days ago |
| 11.62 TiB of cold data on storage rtp-sa-cl01 | 11.62 | 171 | 16 days ago |

The Insight description gives a quick indication of the amount of data detected as "cold" and which storage that data resides on. The table also provides a count of workloads with cold data.


Selecting an Insight from the list opens a page showing more details, including recommendations to move data to the Cloud or cycle down unused disks, as well as estimated cost and power savings you could potentially realize from implementing those recommendations. The page even provides a handy link to [NetApp's TCO Calculator](#) so you can experiment with the numbers.



150 Workloads on storage **rtp-sa-cl01** contains a total of 9.5 TiB of cold data.

You could lower costs 9.3% a year and reduce your carbon footprint by moving cold storage to the cloud.

Detected: 2 months ago, 9:21 AM
(ACTIVE)
May 19, 2023 10:05AM



Recommendations

On the Insight page, expand the **Recommendations** to explore the following options:

- Move unused workloads (zombies) to a lower cost storage tier (HDD)

Utilizing the zombie flag, cold storage and number of days, find the coldest and largest amount of data and move the workload to a lower cost storage tier (such as a storage pool using hard disk storage). A workload is considered a "zombie" when it has not received any significant IO requests for 30 days or more.

- Delete unused workloads

Verify which workloads are not in use and consider archiving them or remove them from the storage system.

- Consider NetApp's Fabric Pool Solution


NetApp's [Fabric Pool Solution](#) automatically tiers cold data to low cost cloud storage, thus increasing the efficiency of your performance tier as well as providing remote data protection.

Visualize and Explore


The graphs and table provide additional trending information as well as allow you to drill into the individual workloads.

Cluster Cold Storage Trend [Show Details](#)

Cold Data (TiB)



Cold Storage by Days Cold (TiB)



Workloads with cold data (150) [View all workloads](#)

Filter...

| Workloads | # Days cold | ↑ | Total Size (GiB) | Cold Data Size (GiB) | Percent Cold (%) | Is Zombie | i Disk Type |
|---------------------------|-------------|---|------------------|----------------------|------------------|-----------|-----------------------------|
| SelectPool | 31 | | 8,192.00 | 1,714.21 | 20.93 | N A | SAS |
| nj_UCS_VMw_Infrastructure | 31 | | 5,120.00 | 934.74 | 18.26 | N A | SAS |
| Oracle_SAP_DS_220 | 31 | | 2,048.00 | 861.97 | 42.09 | N A | SSD |
| rtp_sa_workspace | 31 | | 13,000.00 | 741.32 | 5.70 | N A | SAS |
| vc220_migrate | 31 | | 4,311.58 | 685.30 | 15.89 | N A | SAS |
| H01_shared | 31 | | 998.25 | 646.55 | 64.77 | N A | SSD |
| ProdSelectPool | 31 | | 8,192.00 | 555.30 | 6.78 | N A | SAS |
| vcenter_migrate | 31 | | 6,144.00 | 475.99 | 7.75 | N A | SAS |
| rtp_sa_mgmt_apps | 31 | | 4,096.00 | 449.26 | 10.97 | N A | SAS |
| SOFTWARE | 31 | | 600.00 | 365.54 | 60.92 | N A | SAS |
| DP_Migrate | 31 | | 7,168.00 | 347.20 | 4.84 | N A | SAS |

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.