



Collecting Host and VM file system utilization data

OnCommand Insight

NetApp
October 24, 2024

Table of Contents

Collecting Host and VM file system utilization data	1
Configure Insight for file system collection	1
File system chargeback and reporting	3

Collecting Host and VM file system utilization data

The Host and VM File Systems data source, combined with the Host Utilization license, enables reporting and chargeback at the file system level for known Hosts and VMs.

OnCommand Insight collects data from storage devices, most of which report their volumes as block devices. This allows Insight to report on utilization at the storage level, but not at the file system level. Storage arrays typically know which blocks have been written to, but not which blocks have been freed.

Client hosts and VMs implement file systems (ntfs, ext*...) on top of these block devices. Most file systems keep a table of contents containing directory and file metadata. When files are deleted, their entries are simply removed from the table of contents. Blocks consumed by those files are now eligible for re-use by the file system, but the storage array doesn't know this. In order for Insight to report on filesystem usage, it must be collected from the client host or VM point of view for accurate chargeback.

Insight allows this level of file system utilization data collection through the **NetApp Host and VM File System** data source, in combination with the **Host Utilization** license. VM's must be annotated with the appropriate **Compute Resource Group** name, and associated storage arrays must be annotated with appropriate **Tier** annotations with proper costs for accurate cost reporting.



The Host Utilization License is resource-based, as opposed to capacity-based as other Insight licenses.

Configure Insight for file system collection

To configure Insight for collection of file system utilization data, you must install the Host Utilization Pack license and configure the NetApp Host and VM File Systems data source.

Before you begin

If you haven't already, install the Host Utilization Pack license. You can check for the license in the **Admin > Setup** page, on the **Licenses** tab.

The Host and VM File Systems data source only reports file system utilization and file system metadata for known **Compute Resources** (hosts and VMs) currently being collected or discovered in Insight:

- Virtual Machines are collected by hypervisor data sources such as Hyper-V and VMware.
- Hosts are discovered via device resolution.

The proper Tier annotations must be present on the appropriate storage resources.

The following connected block storage devices are supported:

- NetApp Clustered Data OnTap (cDOT)
- NetApp 7-Mode
- Clariion
- Windows: VMWare virtual disks (VMDKs) for FC, iSCSI

- Linux: VMWare VMDKs (iSCSI and FC not supported)

A **Compute Resource Group** is an annotation that allows grouping of hosts and/or virtual machines that share a common administrative credential.

Steps

1. First, annotate the hosts and/or virtual machines to be included in your **Compute Resource Group**. Go to **Queries > +New query** and search for *Virtual Machine* assets.
You will need to repeat these steps for *Host* assets.
2. Click on the column selector on the right of the table and select the **Compute Resource Group** column to display it in the query results table.
3. Select the virtual machines you wish to add to the desired compute resource group. You can use a filter to search for specific assets.
4. Click on the **Actions** button and choose **Edit annotation**.
5. Select the *Compute Resource Group* annotation, then choose the desired resource group name in the *Value* field.

The resource group annotation is added to the selected VMs. The resource group name must match the name you will configure in the Host and VM File Systems data source later.

6. To configure the Host and VM File Systems data source for a compute resource group, click on **Admin > Data sources** and **Add** the *NetApp Host and VM File Systems* data source.

The screenshot shows the 'Add Data Source' configuration dialog. The 'Settings' tab is active, displaying fields for Name (empty), Vendor (NetApp), Model (Host and VM File Systems), Where to run (Clustered Data ONTAP 8.1.1+), and What to collect (Host and VM File Systems). The 'Model' dropdown is expanded, showing a list of options including Host and VM File Systems, Clustered Data ONTAP 8.1.1+, Clustered Data ONTAP 8.1.1+ (Unified Manager 6.0+), Data ONTAP 7-Mode, E-Series (Firmware 6.x), E-Series (Firmware 7.x+), Host and VM File Systems (selected), SolidFire 8.1+, and StorageGrid. Below the settings are tabs for Configuration, Advanced configuration, and Test. At the bottom are 'Cancel' and 'Save' buttons.

7. In the **Configuration** section, enter a **User Name** and **Password** for an operating system user with appropriate rights to retrieve file system data. For Windows operating system users, this must include the domain prefix if your Windows environment uses it.

Note that an Insight Acquisition Unit (AU) installed on Linux can report on Linux compute resources, while an AU installed on Windows can talk to either Linux or Windows compute resources.

8. Enter the name of the **Compute Resource Group** for the assets from which you will want to collect file system utilization data. This name must match the resource group name you used to annotate the assets above.

If you leave the Compute Resource Group field empty, the data source will collect data for hosts or VMs that have no Compute Resource Group annotation.

9. In the **Advanced Configuration** section, enter the desired polling interval for this data source. The default of 6 hours is usually adequate.
10. It is recommended to **Test** the data source connection before saving it. A successful connection result will also show you how many compute resource targets are contained in the group.
11. Click **Save**. The Host and VM File Systems data source will begin collecting data on its next poll.
12. Once file system data is being collected, you can view it on the host's or VM's asset page, in the File System widget:

File Systems				
Name	Capacity (Used / Total GB)	Type	Storage Resource	
/	9.15% (11.0 / 120.0)	xfs	 vifasname:...:vm_ocl_	
/boot	23.79% (0.1 / 0.5)	xfs	 vifasname:...:vm_ocl_	
/dev/dm-1	7.8	swap	 vifasname:...:vm_ocl_	

Showing 1 to 3 of 3 entries

13. Repeat these steps for each Compute Resource Group you will have. Each compute resource group must be associated with its own Host and VM File Systems data source.

Note that file system information will be collected for hosts and VM's that are already being acquired by any traditional VMware or Hyper-V data sources in your environment.

File system chargeback and reporting

Chargeback for file systems is always performed from the storage perspective. Storage arrays associated with virtual machines annotated for a particular compute resource group will be included in chargeback reports for that resource group.

Before you begin

Any virtual machines which you wish to include in file system utilization chargeback must be annotated with the appropriate compute resource group name. Storage arrays associated with those virtual machines must be annotated with the appropriate Tier annotations. ETL to data warehouse must have occurred after these annotations are in place.

Steps

1. Open a browser to your Reporting server, usually <https://<host> or IP>:9300/p2pd> or <http://<host> or IP>:9300/bi> (7.3.3 or later) and log in.

2. Choose the **File System Utilization** package and create a new report.

List of all packages:

Cognos > Public Folders > **Packages**

Name
Application Volume Hourly Performance
Chargeback
File System Utilization
Host Volume Hourly Performance
Internal Volume Capacity

3. Drag and drop items from your data mart(s) to build your report.

The example below is a very simple report. You can create complex reports built around your specific business needs.

Name	Type	Allocated Capacity GB	Used Capacity GB	Tier Name	Cost	Storage Name
/	xfs	119.96	9.96	N/A		vifasnane05,vifasnane06
/	xfs	5,492.53	799.63	Tier 1	100	vifasnane
/boot	xfs	0.48	0.17	N/A		vifasnane05,vifasnane06
/boot	xfs	8.72	2.41	Tier 1	100	vifasnane
/dev/dm-1	swap	7.81	0.00	N/A		vifasnane05,vifasnane06
/dev/dm-1	swap	140.61	0.78	Tier 1	100	vifasnane
C:\	NTFS	948.27	331.98	Tier 1	100	vifasnane
PHYSICALDRIVE0: System Reserved	NTFS	1.70	1.41	Tier 1	100	vifasnane

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.