



What you can do using annotations

OnCommand Insight

NetApp
April 01, 2024

Table of Contents

- What you can do using annotations 1
 - Preparing annotations in OnCommand Insight 1
 - Importing user-defined annotations into Data Warehouse 2
 - Viewing the Annotations job in the Jobs list 3
 - Displaying annotation changes in the database schema 4

What you can do using annotations

Annotations provide a method for defining information that relates to objects in your environment and then allows you to track objects based on the annotation. For example, you could add building or floor number annotations to devices in your environment and then create a query that returns all of the devices on the first floor of a data center.

Additionally, you might want to look at all devices in a specific data center or business entity and determine which business entity is using the most tier 1 storage. To do this, you assign a data center, business entity, or tier annotation to the device using the OnCommand Insight web UI. Then, you can bring selected user-defined annotations from OnCommand Insight into Data Warehouse. You want to do this to see the annotation values assigned to objects appear in your custom reports.

You can specify which user-defined annotations propagate to Data Warehouse. Annotations are added as additional columns to the object table in the inventory, and to the relevant dimension table in the data marts. When you update the annotations on resources using the OnCommand Insight user interface and initiate or wait for the next Data Warehouse build, you see the results in the following tables:

- `dwh_inventory.annotation_value`
- `dwh_inventory.object_to_annotation`

Ensuring annotations entered in OnCommand Insight are included in Data Warehouse requires the following major processes:

- Before you import annotations into Data Warehouse, you must ensure that they are prepared in OnCommand Insight.

To do this, you can manually run the **Troubleshooting > Force Update of Annotations for Data Warehouse** option or wait until the next scheduled transient data run process. When you force the update of annotations, you force the OnCommand Insight server to calculate and place the transient data (such as annotation values) into database tables so that the Data Warehouse ETL process can read the data. The update of annotations data occurs automatically every fifteen minutes; however, you can force it to happen more frequently.

- You then import annotations into Data Warehouse by using the Data Warehouse **Annotations** option.
- If you want to include annotations in reports that you create by using the OnCommand Insight Reporting Portal report authoring tools, you must update the OnCommand Insight reporting metadata model.

When you upgrade Data Warehouse, the annotations job runs automatically during the database restore process. The annotations job runs automatically also when WildFly starts up.



WildFly is an application server where the OnCommand Insight Java code runs and is needed for both for the OnCommand Insight server and for Data Warehouse.

Preparing annotations in OnCommand Insight

Annotations must be prepared in OnCommand Insight before they can be imported into Data Warehouse.

Steps

1. Log in to the OnCommand Insight Portal as administrator `https://hostname`, where `hostname` is the name of the system where OnCommand Insight is installed.
2. Click on **Admin > Troubleshooting**. At the bottom of the page, click on **Advanced Troubleshooting**.
3. In the **Actions** tab, click **Update DWH Annotations (include deleted)**.

Importing user-defined annotations into Data Warehouse

After forcing an annotation update in OnCommand Insight, you need to select the annotations you want in Data Warehouse and initiate a Data Warehouse build. You can wait until the next scheduled build or initiate a build now.

Steps

1. Log in as an administrator to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Annotations**.

Annotations

Annotation	Column Name	Target Object	Published
Compute_Resource_Group	Compute_Resource_Group	Virtual Machine	
Data_Center	dataCenter	Host	✓
Data_Center	dataCenter	Storage	✓
Data_Center	dataCenter	Switch	✓
Note	Note	Switch	
Switch_Level	switchLevel	Switch	✓
Tier	Tier	Internal Volume	
Tier	Tier	Qtree	
Tier	Tier	Storage	
Tier	Tier	Storage Pool	
Tier	Tier	Volume	

Edit

The list displays a row for every annotation type and a target object to which the annotation can be assigned. A check mark in the Published column indicates that the annotation was already selected for the particular target object and is already available through the Data Warehouse data marts.

3. Click **Edit** to edit how annotations will be imported from OnCommand Insight.

Annotation	Column Name	Target Object	Published All / None	Init With Current All / None
Compute_Resource_Group	Compute_Resource_Group	Virtual Machine	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Data_Center	dataCenter	Host	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Data_Center	dataCenter	Storage	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Data_Center	dataCenter	Switch	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Note	Note	Switch	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Switch_Level	switchLevel	Switch	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Tier	Tier	Internal Volume	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Tier	Tier	Qtree	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Tier	Tier	Storage	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Tier	Tier	Storage Pool	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Tier	Tier	Volume	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Save Cancel

4. To edit the annotation process, do the following:

- Select **Published** to add annotations retrieved from OnCommand Insight into the Data Warehouse database. Click **All** to select all annotations on all objects. Click **None** to ensure that all options are not selected.



Uncheck this option to remove the annotation column from the specific object's inventory table and associated data marts. If any custom-designed reports use annotation data, the reports do not run successfully.

- Check **Init with Current** to initialize historical data in Data Warehouse dimension tables with the current annotation value. Click **All** to select all annotations on all objects. Click **None** to ensure that all options are not selected. This check box is disabled after an annotation is published; the check box is enabled for annotations that are not published. For example, if a host is annotated with annotation type "floor" and gets the value "1", and there are 3 rows for that host in the host_dimension table, then selecting **Init with Current** associates the value "1" in the "floor" column for all 3 rows in the host_dimension table. If **Init with Current** is not selected, then only the latest row for that host will have the value "1" in the floor column.

5. Click **Save**.

A warning message appears indicating that this will cause changes to the structure of the data or data loss, if you are removing annotations.

6. To continue, click **Yes**.

Data Warehouse initiates an asynchronous annotations job that applies the requested changes. You can see the job in the Jobs page. You can also see the changes in the Data Warehouse database schema.

Viewing the Annotations job in the Jobs list

You can view the Annotations job in the Jobs list and apply the annotation changes to Data Warehouse data marts.

Steps

1. Log in as an administrator to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Jobs**.

Displaying annotation changes in the database schema

The database schema reflects the changes in the specific table.


About this task

For example, if you add annotations to a storage array, they appear in the storage or switch table in the inventory or other data marts.

If you update the annotations on resources using the OnCommand Insight user interface and initiate or wait for the next Data Warehouse build, you see a new column added or removed in the corresponding object in inventory (`dwh_inventory`) and in the corresponding dimension table as well (in the appropriate data mart). You see the results in the following tables:

- `dwh_inventory.annotation_value`
- `dwh_inventory.object_to_annotation`

Steps

1. Click  on the Data Warehouse tool bar and select **Documentation**.
2. Select **Database Schema**.
3. In the **Database Schema** pane on the left, scroll to the **DWH_INVENTORY** section and click **switch**.

<div>Database Schema</div> <div>Databases</div> <div> storage_port storage_to_applica switch switch_port switch_port_to_app switch_to_applicati tape tape_controller tape_port tier violation virtual_switch virtual_to_backend vm_to_application volume volume_in_storage </div>	dwh_inventory.switch			
	Column	Type	Nullable	Description
	id	int(11)	false	GUID of the switch.
	fabricId	int(11)	true	GUID of the fabric on which this switch is configured to operate. References: <ul style="list-style-type: none"> id in dwh_inventory.fabric
	identifier	varchar (255)	false	Identifier of the device.
	wwn	varchar (255)	false	WWN of the switch.
	ip	varchar (255)	false	IP address of the switch.
	Name	varchar (255)	false	Name of the switch.
	Manufacturer	varchar (255)	true	Manufacturer of the switch
	Model	varchar (255)	true	Manufacturer's model of the switch.
	Firmware	varchar (255)	true	Firmware version running on the switch.

4. The **dwh_inventory.switch** table reflects the changes:

<div>Database Schema</div> <div>Databases</div> <div> host_group_dimen internal_volume_co internal_volume_di qtree_capacity_fac qtree_dimension service_level_dime storage_dimension storage_pool_dime tier_dimension vm_capacity_fact vm_dimension volume_fact_curre </div>	dwh_capacity.storage_dimension			
	Column	Type	Nullable	Description
	tk	int(11)	false	TK of this storage array row.
	name	varchar (255)	false	Name of the storage array.
	identifier	varchar (255)	false	Identifier of the device.
	ip	varchar (255)	false	IP address of the storage array.
	model	varchar (255)	true	Manufacturer's model of the storage array.
	manufacturer	varchar (255)	true	Manufacturer of the storage array.
	serialNumber	varchar (255)	true	Serial number for the storage array.
	microcodeVersion	varchar (255)	true	Version of the firmware running on the storage array.
	family	varchar (255)	true	Family name of the storage array (e.g. Clariion, Symmetrix etc).
	id	int(11)	true	GUID of the storage array in dwh_inventory.storage .

The dataCenter annotation column appears in the storage_dimensions table.

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.