



Administrative tasks you can perform using Data Warehouse

OnCommand Insight

NetApp
September 13, 2024

Table of Contents

- Administrative tasks you can perform using Data Warehouse 1
 - Managing jobs 1
 - Monitoring Data Warehouse health 2
 - Scheduling daily builds 3
 - Configuring a weekly schedule 4
 - Scheduling daily backups 5
 - Running custom scripts in Data Warehouse 6
 - What you can do using annotations 6
 - Setting email notifications 11
 - Accessing the Reporting Portal 11
 - Viewing the Data Warehouse database schema documentation 12
 - Viewing the Data Warehouse database schema 12
 - Viewing system information 13
 - Advanced options 13

Administrative tasks you can perform using Data Warehouse

OnCommand Insight Data Warehouse is a web-based user interface that enables users to configure and troubleshoot data in OnCommand Insight Data Warehouse and to set up schedules to retrieve data from OnCommand Insight.

Using the Data Warehouse portal, you can perform the following administrative tasks:

- Check the status of current jobs or queries that are running
- Manage annotations
- Configure email notifications
- Access and create custom reports
- Review Data Warehouse documentation and database schema
- Edit the site name
- Identify the Data Warehouse version and upgrade history
- Build the Data Warehouse data from history
- Reset the Data Warehouse database
- Back up and restore the Data Warehouse database
- Troubleshoot Data Warehouse issues and look at OnCommand Insight logs
- Manage user accounts

Managing jobs

You can see a list of current jobs and their status. The first job in a build cycle is in bold type. The build that Data Warehouse performs for each connector and for each data mart is considered a job.

About this task

You can cancel any pending job you have scheduled or begun. You can also clear the history of previously executed jobs. You can clear history of jobs that are not pending, running, or in the process of aborting. You can clear all history or all history except the previous 24 hours to remove all but the last day's entries.

You can see information about the following types of jobs: License, Pre Inventory, Inventory, Post Inventory, Dimensions, Datamarts Preparation, Capacity, Ports, Storage Efficiency, File System Utilization, Performance, Job Status Notification, History build, Dynamic annotations, Connector removal, Skipped build, Phone Home, and Maintenance.

A maintenance job runs weekly and uses MySQL tools to optimize the database.

Steps

1. Log in 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**.

#	Name	Status	Start time	End time
1551	Job Status Notification	COMPLETED	4/9/15 2:12 AM	4/9/15 2:12 AM
1550	Performance	COMPLETED	4/9/15 2:11 AM	4/9/15 2:12 AM
1549	File System Utilization	COMPLETED	4/9/15 2:11 AM	4/9/15 2:11 AM
1548	Storage Efficiency	COMPLETED	4/9/15 2:11 AM	4/9/15 2:11 AM
1547	Ports	COMPLETED	4/9/15 2:11 AM	4/9/15 2:11 AM
1546	Capacity	COMPLETED	4/9/15 2:11 AM	4/9/15 2:11 AM
1545	Datamarts Preparation	COMPLETED	4/9/15 2:10 AM	4/9/15 2:11 AM
1544	Dimensions	COMPLETED	4/9/15 2:10 AM	4/9/15 2:10 AM
1543	Post Inventory	COMPLETED	4/9/15 2:10 AM	4/9/15 2:10 AM
1542	Inventory: connector 'oci-stg01-s08', 'Apr 9, 2015 2:00 AM'	COMPLETED	4/9/15 2:06 AM	4/9/15 2:10 AM
1541	Inventory: connector 'oci-stg06-s08', 'Apr 9, 2015 2:00 AM'	COMPLETED	4/9/15 2:04 AM	4/9/15 2:06 AM
1540	Inventory: connector 'oci-stg03-s08', 'Apr 9, 2015 2:00 AM'	COMPLETED	4/9/15 2:04 AM	4/9/15 2:04 AM

If a Pending status appears, a cancel link appears.

3. To cancel a pending job, click **cancel**.
4. To remove the job history, click **All** or **All but last 24h**.

Monitoring Data Warehouse health

The Data Warehouse (DWH) includes a health monitor that displays information about the state of DWH. Alarm messages are displayed on the **Connectors** and **Jobs** pages of the DWH, as well as sent to the connected Insight server, where they are displayed on the **Admin > Health** page.

DWH collects metrics every ten minutes, and displays an alarm under the following conditions:

- Connection to the Insight server is down
- Disk utilization is greater than 90%
- Reporting (Cognos) service is down
- A query holds a lock on any table for a prolonged time
- A maintenance job is disabled
- Automatic backup is disabled
- Security risk: default encryption keys detected

Health monitor warnings in the Data Warehouse can be suppressed for up to 30 days.

When email notification is enabled, these events are also reported by email. Note that the email does not contain any attachments.

These events are logged in the `dwh_troubleshoot.log` file in the following locations:

- Windows: `<install dir>\SANSscreen\Wildfly\Standalone\Logs`
- Linux: `/var/log/netapp/oci/wildfly/`

Scheduling daily builds

Although you can manually build Data Warehouse by using the Build now control at any time, it is best practice to schedule automatic builds, defining when and how often to build the Data Warehouse database. Data Warehouse performs a build job for each connector and for each data mart. Data Warehouse performs a build job for each connector for licenses and inventory and all other build jobs (for example, capacity) are performed on the consolidated database.

About this task

Whenever the Data Warehouse is built, it performs an inventory job for every connector. After the inventory jobs are complete, Data Warehouse performs jobs for dimensions, capacity, and the remaining data marts.

Steps

1. Log in 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 **Edit Schedule**.

Automatic Schedule

Enabled:	yes	Edit schedule
Schedule:	Daily at 2:00 AM,7:00 PM	
Next run:	4/2/15 7:00 PM	

Build now

3. In the **Build Schedule** dialog box, click **Edit** to add a new schedule.

Type:

Enabled:

Run at:

<input type="checkbox"/> 12:00 AM	<input type="checkbox"/> 1:00 AM	<input checked="" type="checkbox"/> 2:00 AM	<input type="checkbox"/> 3:00 AM	<input type="checkbox"/> 4:00 AM	<input type="checkbox"/> 5:00 AM	<input type="checkbox"/> 6:00 AM	<input type="checkbox"/> 7:00 AM	<input type="checkbox"/> 8:00 AM	<input type="checkbox"/> 9:00 AM	<input type="checkbox"/> 10:00 AM	<input type="checkbox"/> 11:00 AM
<input type="checkbox"/> 12:00 PM	<input type="checkbox"/> 1:00 PM	<input type="checkbox"/> 2:00 PM	<input type="checkbox"/> 3:00 PM	<input type="checkbox"/> 4:00 PM	<input type="checkbox"/> 5:00 PM	<input type="checkbox"/> 6:00 PM	<input checked="" type="checkbox"/> 7:00 PM	<input type="checkbox"/> 8:00 PM	<input type="checkbox"/> 9:00 PM	<input type="checkbox"/> 10:00 PM	<input type="checkbox"/> 11:00 PM

- Choose the frequency - weekly.
- Choose the time of day for each day you want the job to run.
- Chose N/A for days you do not want to run the build.
- To enable the schedule, select **Enabled**.



If you do not check this, the schedule build does not occur.

- Click **Save**.
- To build Data Warehouse outside of the automatic scheduled build, click **Build now**.

Configuring a weekly schedule

Although you can manually build Data Warehouse by using the Build now control at any time, it is best practice to schedule automatic builds, defining when and how often to build the Data Warehouse database. Data Warehouse performs a build job for each connector and for each data mart. Data Warehouse performs a build job for each connector for licenses and inventory and all other build jobs (for example, capacity) are performed on the consolidated database. With a weekly schedule, you can specify the time you want the build to run for each day of the week.

Steps

- Log in to the Data Warehouse Portal at <https://hostname/dwh>, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
- From the navigation pane on the left, click **Edit Schedule**.
- Choose the frequency - weekly.
- Choose the time of day for each day you want the job to run.
- Chose N/A for days you do not want to run the build.
- To enable the schedule, select **Enabled**.



If you do not check this, the schedule build does not occur.

7. Click **Save**.
8. To build Data Warehouse outside of the automatic scheduled build, click **Build now**.

Scheduling daily backups

Although you can manually back up Data Warehouse by using the Backup/Restore control at any time, it is best practice to schedule automatic backups, defining when and how often to back up the Data Warehouse database and Cognos content store. Backups offer protection from data loss, allowing you to restore the Data Warehouse database if needed. You also use a backup when migrating to a new Data Warehouse server or when upgrading to a new Data Warehouse version.

About this task

Scheduling backups during times when the Data Warehouse server is not busy improves backup performance and reduces the impact on users.

Steps

1. Log in 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 **Schedule**.
3. In the **Backup Schedule** dialog box, click **Edit** to add a new schedule.

Backup Enabled:

Backup Location:

Select Backup Configuration:

Run every:

<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday

Run at hour:

Cleanup:

4. To enable the scheduled backups, select **Backup Enabled**.
5. Specify the location where you want to store the backup files.
6. Specify the data you want backed up.
7. Specify the day or days that you want the backup to be performed.
8. Specify what time of day you want the backup started.
9. Specify how many past backup copies you want to retain.

10. Click **Save**.

Running custom scripts in Data Warehouse

Data Warehouse allows customers to create jobs that can run custom scripts that prepare customized data in Data Warehouse.

Before you begin

To prevent the custom script from being deleted during a Data Warehouse upgrade, you must not store the script in the SANscreen directory.

About this task

The job can specify only one script. You can run multiple scripts and commands from one script.

Steps

1. In Data Warehouse, select **DWH Admin > Schedule**.
2. Select the **Script enabled** check box.
3. Enter the absolute path of the script name in the **Script Location** text box.
4. Click **Save**.

Results

The Data Warehouse job engine schedules the task to run a “Custom scripting” job. The job is scheduled to run after an ETL and avoiding other conflicting background processes. The job is not run by a “Build from history” operation.

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.

Edit Annotations

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

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:

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 <code>dwh_inventory.storage</code> .

The dataCenter annotation column appears in the storage_dimensions table.

Setting email notifications

You can have Data Warehouse send email to a specific email address when Data Warehouse jobs do not complete successfully.

Steps

1. Log in 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 **Email Notification**.

3. Enter the following:

- SMTP server address

Specifies the server that is acting as the SMTP server in your organization, identified using either a hostname or an IP address using the `nnn.nnn.nnn.nnn` format. If you specify a host name, ensure that DNS can resolve it.

- SMTP server username and password

Specifies the user name to access the email server and is required only if your SMTP server requires a user to log into the server. This is the same user name you use to log in to the application and access your email.

- Notifications enabled

Yes enables the notifications; **No** disables the notifications.

- Sender's Email

Specifies the email address that is used to send the notifications. This must be a valid email address in your organization.

- Recipient's Email

Specifies the email address or addresses of the person or people who will always receive the email. Separate multiple addresses with commas.

- Email subject

Specifies the subject for the notification.

- Email signature


Specifies the information that displays at the bottom of the email, for example, the department name.

Accessing the Reporting Portal

From the Data Warehouse Portal, you can access the Reporting Portal, where you can create custom reports using report authoring tools such as Workspace Advanced and

Report Studio.


Steps

1. On the Data Warehouse toolbar, click  to open the Insight Reporting Portal.
2. Enter your user name and password and click **Login**.

Viewing the Data Warehouse database schema documentation

You can review Data Warehouse database schema information.


Steps


1. Log in 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. On the Data Warehouse toolbar, click  and select **Schema**.

Viewing the Data Warehouse database schema

You might want to view the database schema to understand how to use the data in another API or to develop SQL queries. The schema option lists all databases, tables, and columns in the schema. You can also review the database schema diagrams showing the table relationships.

Steps

1. Log in 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. Click  on the Data Warehouse toolbar, and select **Documentation**.
3. Select **Database Schema**.
4. For example, in the **Databases** pane, click **DWH_INVENTORY**.
5. In the **All tables** pane, scroll down to **DWH_INVENTORY** section, and click the **annotation_value** table.



The `dwh_inventory.annotation_value` table appears.

Viewing system information

You can view system, module, license, and Data Warehouse upgrade information.

Steps

1. Log in 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 **System Information**.
3. On the **System** tab, review the system information and edit it the site name, if needed, by doing the following:
 - a. Click **Edit Site Name**
 - b. Enter the new site name and click **Save**.
4. To see application information (application name, module, version, and install date), click the **Application Info** tab.
5. To see license information (protocol, code, expiration date, and quantity), click the **Licenses** tab.
6. To see application upgrade information (application name, from date, to date, time, user, and file size), click **Upgrade History**.

Advanced options

Data Warehouse includes various advanced options.

Skipping failed builds

After your first build, sometimes you might encounter an unsuccessful build. To ensure

that all the jobs after an unsuccessful build complete successfully, you can enable the **Skip history build failures** option.

About this task

If a build fails and the **Skip history build failures** option is enabled, Data Warehouse continues building and ignores any failed builds. If this occurs, there will not be a data point in the historical data for any skipped builds.

Use this option only if the build is not successful.

If a build fails in Build from History and the **Skip history build failures** check box is not selected, all subsequent jobs are aborted.

Steps

1. Log in 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 **Build from history**.

Build From History

Target time	Start running	Status
3/13/15 12:00 AM	3/25/15 9:28 AM	COMPLETED
3/14/15 12:00 AM	3/25/15 9:34 AM	COMPLETED
3/15/15 12:00 AM	3/25/15 9:39 AM	COMPLETED
3/16/15 12:00 AM	3/25/15 9:45 AM	COMPLETED
3/17/15 12:00 AM	3/25/15 9:51 AM	COMPLETED
3/18/15 12:00 AM	3/25/15 9:57 AM	COMPLETED
3/19/15 12:00 AM	3/25/15 10:03 AM	COMPLETED
3/20/15 12:00 AM	3/25/15 10:09 AM	COMPLETED
3/21/15 12:00 AM	3/25/15 10:16 AM	COMPLETED
3/22/15 12:00 AM	3/25/15 10:23 AM	COMPLETED
3/23/15 12:00 AM	3/25/15 10:30 AM	COMPLETED
3/24/15 12:00 AM	3/25/15 10:38 AM	COMPLETED
3/25/15 12:00 AM	3/25/15 10:44 AM	COMPLETED

« < 1 2 3 > »

Cancel Pending Jobs **Configure** Run

Skip history build failures:

3. Click **Configure**.
4. Configure the build.
5. Click **Save**.
6. To skip failed builds, check **Skip history build failures**.

You can see this check box only if the **Run** button is enabled.

7. To perform a build outside of the automatic scheduled build, click **Run**.

Resetting the Data Warehouse database or Reporting server

You can delete the contents of the Data Warehouse data marts and delete all configured connectors. You might want to do this if an installation or upgrade did not complete successfully and it left the Data Warehouse database in an intermediate state. You can also delete only the Inventory data model or the Cognos Reporting data model.

Steps

1. Log in 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 **Reset DWH database**.
3. Click one of the following options:

- **Reset DWH Database**

This deletes the contents of all Data Warehouse data marts and all configured connectors and places the Data Warehouse to the default installed state without any custom configurations. You might choose this option, for example, if you changed your connected servers, but restored a different Data Warehouse database accidentally on your server and need to return to a default installed state. This does not delete any reports. (Reports are saved in the Cognos Content Store.)

- **Reset Inventory Only**

This deletes the contents of the Inventory data model only. This does not delete any historical data.

- **Reset Reporting Content**

This resets the content of the reporting server. This deletes any custom reports you may have. Backup your reports before you choose this option.

A warning message displays.

4. To continue, click **Yes**.

Restoring and upgrading reports for versions prior to 6.3

If you are upgrading an Insight version prior to 6.3, you must manually restore your reporting artifacts.

Before you begin

Follow the instructions in the "Upgrading the Data Warehouse (DWH)" and "Backing up custom reports and reporting artifacts" topics.

Steps

1. To restore Reporting artifacts from releases prior to version 6.3, copy the Export Backup.zip file you

created and stored in your <install>\cognos\c10_64\deployment directory.

2. Open a browser and go to <http://<server>:<port>/reporting> for the server and port you used during installation.
3. Enter your user name and password and click **Login**.
4. From the **Launch** menu, select **Insight Reporting Administration**.
5. Click the **Configuration** tab.

Due to changes in the data model, the reports in the old packages may not run and need to be upgraded.

6. Click **Content Administration**.
7. Click the **New Import** button.
8. Make sure that archive you copied to the deployment directory (for example, backup6.0.zip) is selected, and click **Next**.
9. If you entered a password to protect the archive, enter the password and click **OK**.
10. Change the name `Export...` to `Import Backup` and click **Next**.
11. Click on the pencil icon next to each package name and enter a new target name if necessary. For example, add a `_original` suffix to the existing name. Then click **OK**.
12. After you rename the target package names for all packages, select all blue folders and click **Next** to continue.
13. Accept all default values.
14. Click **Finish** and then select **Run**.
15. Check for the details of this import and click **OK**.
16. Click **Refresh** to view the status of the import.
17. Click **Close** after the import is complete.

Results

Two sets of packages appear in the Public Folders tab. For example, one with a `7.0` suffix (for the newer version) and one with a `_original` (or whatever you entered during the backup/restore procedure) suffix which contains your old reports. Due to changes in the data model, the reports in the old packages may not run and need to be upgraded. Your portal tabs now point to the current version of the portal pages.

Accessing MySQL using the command-line interface

In addition to accessing Data Warehouse data elements through the report authoring tools, you can obtain access to Data Warehouse data elements directly by connecting as a MySQL user. You might want to connect as a MySQL user to use the data elements in your own applications.

About this task

There are many ways to connect. The following steps show one way.

When accessing MySQL, connect to the MySQL database on the machine where Data Warehouse is installed. The MySQL port is 3306 by default; however, you can change it during installation. The user name and password is `dwhuser/netapp123`.

Steps

1. On the machine where Data Warehouse is installed, open a command-line window.
2. Access the MySQL directory in the OnCommand Insight directory.
3. Type the following user name and password: `mysql -udwhuser -pnetapp123`

The following is displayed, depending on where Data Warehouse is installed:

```
c:\Program Files\SANscreen\mysql\bin> mysql -udwhuser -pnetapp123
```

```
Welcome to the MySQL monitor. Commands end with ; or \g.  
Your MySQL connection id is 882  
Server version: 5.1.28-rc-community MySQL Community Server (GPL)
```

```
Type 'help;' or '\h' for help. Type '\c' to clear the buffer.
```

4. Show the Data Warehouse databases: `show databases;`

The following is displayed:

```
mysql> show databases;  
+-----+  
| Database |  
+-----+  
| information_schema |  
| dwh_capacity |  
| dwh_capacity_efficiency |  
| dwh_fs_util |  
| dwh_inventory |  
| dwh_performance |  
| dwh_ports |  
+-----+
```

Troubleshooting Data Warehouse

You can do various tasks related to troubleshooting Data Warehouse.

- Use OnCommand Insight ASUP.
- View OnCommand Insight logs.
- Resolve issues related to upgrades and business entities.
- Resolve issues related to the consolidation of multiple OnCommand Insight servers.

You can consolidate multiple OnCommand Insight servers into the same Data Warehouse database. Many configurations may report the same object from multiple connectors (that is, the same switch exists in two OnCommand Insight instances). In such cases, Data Warehouse consolidates the multiple objects into one (a primary connector is chosen and the object's data will be taken from that connector only).

The storage administrator can use the Troubleshooting page to solve problems related to consolidation issues.

Issue resolution with ASUP

You can send ASUP logs to technical support for assistance in troubleshooting. ASUP for Data Warehouse is configured to run automatically. In the Data Warehouse Portal, you can disable the automatic send process, choose to include a backup of the Data Warehouse database, or initiate a transmittal to ASUP.

The information in the logs is forwarded to technical support by using HTTPS protocol. No data is forwarded using ASUP unless you first configure it on the Insight Server.

Data Warehouse sends the logs to the OnCommand Insight Server that is the first connector listed in the Data Warehouse Portal Connectors page. The automatic process sends the following files:

- Data Warehouse logs, which includes the following:
 - boot.log (including backups)
 - dwh.log (including backups such as dwh.log.1)
 - dhw_troubleshoot.log
 - dwh_upgrade.log (including backups)
 - WildFly.log (including backups)
 - ldap.log (including backups)
 - SQL dump of the Data Warehouse management database
 - mysql: my.cnf, .err and slow query logs
 - full innodb status

- Cognos logs, which include the following:

- cognos-logs.zip

Contains the Cognos log files from the <install>\cognos\c10_64\logs directory. It also contains the logs generated by Cognos as well as the OnCommand InsightAP.log file that contains all logging from users logging in to and out of OnCommand Insight reporting.

- DailyBackup.zip

Contains the backup of the reporting artifacts in the Public Folders. The contents of My Folders is not included in this.

- cognos_version_site_name_content_store.zip

Contains a full backup of the Cognos Content Store.

You can generate a troubleshooting report manually. The Troubleshooting Report .zip file contains the following Data Warehouse information:

- boot.log (including backups)
- dwh.log (including backups such as dwh.log.1)
- dwh_upgrade.log (including backups)
- wildfly.log (including backups)
- ldap.log (including backups)
- dump files in c:\Program Files\SANscreen\wildfly\standalone\log\dwh\
- SQL dump of the Data Warehouse management database
- mysql: my.cnf, .err and slow query logs
- full innodb status



ASUP does not automatically send a backup of the OnCommand Insight database to technical support.

Disabling automatic ASUP transmissions

All NetApp products are equipped with automated capabilities to provide the best possible support to troubleshoot issues that occur in your environment. ASUP periodically sends predefined, specific, information to Customer Support. By default, ASUP is enabled for Data Warehouse; however, you can disable it if you no longer want the information sent.

Steps

1. From the navigation pane on the left, click **Troubleshooting**.
2. Click **Disable** to prevent ASUP from sending a daily report.

A message displays saying ASUP is disabled.

Including a backup of the Data Warehouse database

By default, ASUP sends only the Data Warehouse log files to technical support for assistance in troubleshooting; however, you can also choose to include a backup of the Data Warehouse database and select the type of data that is sent.

Steps

1. Log in 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 **Troubleshooting**.
3. To specify that ASUP should include a backup of the Data Warehouse database, click the **Include DWH Database Backup** list and select one of the following options for the type of data that the backup should include:
 - All (including Performance)
 - All Except Performance
 - Inventory Only

4. Click **Update**.

Sending Insight logs to ASUP

You can send ASUP logs to technical support for assistance in troubleshooting. ASUP for Data Warehouse is configured to run automatically. In the Data Warehouse portal, you can disable the automatic send process, choose to include a backup of the Data Warehouse database, or initiate a transmittal to ASUP. When you request an ASUP report, the report request appears as a job in the Data Warehouse portal Jobs page.

About this task

The job is managed by the job queue similar to the processing of other jobs. If an ASUP job is in a Pending or Running state already, an error message appears indicating that the ASUP report request cannot be added to the job request, because the job queue contains pending or running requests.

Steps

1. Log in 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 **Troubleshooting**.
3. In the **OnCommand Insight ASUP** section of the **Troubleshooting** page, click **Download DWH Troubleshooting Report** to retrieve the troubleshooting report.
4. To send the report to the OnCommand Insight Server listed as the first connector in the Data Warehouse Portal **Connectors** page, click **Send Now**.

Viewing OnCommand Insight logs

You can view various Data Warehouse and Cognos logs in OnCommand Insight.

About this task

You can examine troubleshooting and status information in Cognos and Data Warehouse log files.

Steps

1. Log in 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. In the navigation pane on the left, click **Troubleshooting**.
3. In the **Logs** section, click **Log Files**.

The following log files are displayed:

dwh.log
Lists the status of Data Warehouse jobs
wildfly.log

Provides information about the WildFly application server
dwh_upgrade log
Provides information about the upgrade on Data Warehouse
ldap.log
Logs messages related to LDAP authentication
dwh_troubleshoot.log
Logs messages that can help troubleshooting DWH problems
sansscreenap.log
Provides information about connection to the server, authentication and access to the Cognos repository, and information about other processes
cognosserver.log
Cognos log

4. Click on the name of the log file you want to view.

Multiple server chassis consolidation issues

You can view the connectors that report on hosts and adapters and SAN switches and storage arrays. You can also see the various connectors that report on an object and identify the primary connector, which is the connector that was chosen for the object.

Viewing hosts and adapters consolidation issues

The reported data for hosts and their associated adapters is derived from the Inventory data mart.

Steps

1. Log in 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. In the navigation pane to the left, click **Troubleshooting**.
3. In the **Chassis Consolidation** section, click **Hosts and Adapters**.



The configuration shown in this example is not a valid configuration. Principal and Available connectors on localhost suggest that the Insight server and DWH are both installed on the same server. The intent of this example is to familiarize you with the consolidation table.

Hosts and Adapters Consolidation

Host GUID	Host Name	Host IP	Adapter GUID	Adapter WWN	Principal Connector	Available Connectors	Insight ID	Insight Change Time
288	Agassi	192.1.168.71			localhost (1)	localhost (1)	9927	11/18/10 1:36 PM
			576	40:A0:00:00:00:00:84	localhost (1)	localhost (1)	9928	11/18/10 1:36 PM
			577	40:A0:00:00:00:00:85	localhost (1)	localhost (1)	9930	11/18/10 1:36 PM
305	AI_Host1	192.1.168.88			localhost (1)	localhost (1)	12254	11/18/10 1:38 PM
			597	40:A0:00:00:00:00:01:05	localhost (1)	localhost (1)	12255	11/18/10 1:38 PM
306	AI_Host2	192.1.168.89			localhost (1)	localhost (1)	12257	11/18/10 1:38 PM
			598	40:A0:00:00:00:00:01:06	localhost (1)	localhost (1)	12258	11/18/10 1:38 PM
307	AI_Host3	192.1.168.90			localhost (1)	localhost (1)	12260	11/18/10 1:38 PM

For all hosts and adapters there is a row for each connector that reports on them, as well as the Primary Connector from which the host and adapter are taken. For hosts and adapters only, a host that is reported by one connector may have its adapters reported by a different connector.

You can also see the OnCommand Insight change time of a host/adapter for each connector. Using this parameter, you can discover when an update has occurred in OnCommand Insight for the host/adapter and when the same host/adapter has been updated in other OnCommand Insight servers.

- Optionally, filter data in this view by typing a portion of the text and clicking **Filter**. To clear the filter, delete the text in the **Filter** box and click **Filter**. You can filter by host name, host IP, adapter WWN, or OnCommand Insight object ID.

The filter is case sensitive.

- Review the following data:

- **Host GUID**

Global Unique Identifier for this type of consolidated device (hosts)

- **Host Name**

Name of the consolidated host as it appears in the data warehouse

- **Host IP**
IP address of the consolidated host
- **Adapter GUID**
Global Unique identifier for the host adapter
- **Adapter WWN**
WWN of the host adapter
- **Principal Connector**
Name of the OnCommand Insight connector that was the actual source of the data
- **Available Connectors**
All OnCommand Insight connectors where the consolidated host / adapter reside
- **Insight ID**
OnCommand Insight ID of the consolidated host/adapter for the relevant reporting connector
- **Insight Change Time**
When an update has occurred in OnCommand Insight for the host/adapter and when the same host/adapter has been updated in other OnCommand Insight servers

6. To obtain detail about the connector, click on the connector.

You can see the following information for the connector:

- Host name
- The last time a Data Warehouse job was run on that connector
- The last time a change was received from that connector
- The version of the OnCommand Insight server pointed to by that connector

Viewing storage arrays consolidation issues

The reported data for storage arrays is derived from the Inventory data mart. For all storage arrays, there is a row for each connector that reports on them, as well as the Primary Connector from which each array is taken.

Steps

1. Log in 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 **Troubleshooting**.
3. In the **Chassis Consolidation** section, click **SAN Storage Arrays**.
4. Optionally, to filter data in this view, type a portion of the text in the Filter box and click **Filter**. To clear the filter, delete the text in the Filter box and click **Filter**. You can filter by storage name, storage IP, vendor

model, or OnCommand Insight object ID.

The filter is case sensitive.

5. Review the following data:

- **GUID**

Global Unique Identifier for this type of consolidated device (storage array)

- **Name**

Name of the consolidated storage array as it appears in the Data Warehouse

- **IP**

IP address of the consolidated storage array

- **Vendor and Model**

Name of the vendor who sells the consolidated storage array and the manufacturer's model number

- **Principal Connector**

Name of the OnCommand Insight connector that was the actual source of the data

- **Available Connectors**

All OnCommand Insight connectors where the consolidated storage array resides

- **Insight ID**

ID of the consolidated storage array on the OnCommand Insight chassis where the Principal Connector resides

- **Insight Change Time**

When an update has occurred in OnCommand Insight for the storage array and when the same storage array has been updated in other OnCommand Insight servers

Viewing switches consolidation issues

The reported data for switches is derived from the Inventory data mart. For all switches, there is a row for each connector that reports on them, as well as the Primary Connector from which each switch is taken.

Steps

1. Log in 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 **Troubleshooting**.
3. In the **Chassis Consolidation** section, click **SAN Switches**.
4. Optionally, filter data in this view by typing a portion of the text and clicking **Filter**. To clear the filter, clear the Filter box and click **Filter**. You can filter by switch name, switch IP, vendor model, or OnCommand

Insight object ID.

The filter is case sensitive.

5. Review the following data:

- **GUID**

Global Unique Identifier for this type of consolidated device (storage array)

- **Name**

Name of the consolidated storage array as it appears in the data warehouse

- **IP**

IP address of the consolidated storage array

- **Vendor and Model**

Name of the vendor who sells the consolidated storage array and the manufacturer's model number

- **WWN**

WWN for the consolidation switch

- **Principal Connector**

Name of the OnCommand Insight connector that was the actual source of the data

- **Available Connectors**

All OnCommand Insight connectors where the consolidated storage array resides

- **Insight ID**

ID of the consolidated storage array on the OnCommand Insight chassis where the Principal Connector resides

- **Insight Change Time**

When an update has occurred in OnCommand Insight for the storage array and when the same storage array has been updated in other OnCommand Insight servers

Resolving multiple server annotation consolidation issues

The Annotation Consolidation view in the Data Warehouse Troubleshooting view displays a table that contains all the available Annotation Types and the Object Types to which they can be applied.

About this task

The consolidation of annotation values is based on the value of the Annotation Type. A storage array could have two different tier values, each coming from a different connector. Thus, if in one connector there is a tier defined by the name gold and in a second connector a tier is defined with the name goldy, this information

appears in Data Warehouse as two separate tiers.

Because some Annotation Types allow assignment of multiple annotation values to the same object, Data Warehouse allows objects (for example, "host") to have multiple annotation values assigned to them (for example, "data center 1" and "data center 2" could be assigned to the same host).

Tier annotation on volumes functions somewhat differently from the general annotation tables. Potentially, there could be a very large number of volumes in the environment and displaying all of them in the Data Warehouse would affect the usability of the information. Therefore, the Annotations Consolidation view displays only the volumes that have multiple tier values assigned to them, and the storage containing each such volume.

Steps

1. Log in 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 **Troubleshooting**.
3. In the **Annotation Consolidation** section, click **Show** in the row for the object.

The following shows an example of the annotations for `Data_Center`:

Troubleshooting Annotations Consolidation

Annotation Type: `Data_Center`

Object Type: `Host`

Host GUID	Host Name	Host Natural Key	Data_Center Value	Connector
305	AI_Host1	192.1.168.88	New York	localhost (1)
306	AI_Host2	192.1.168.89	New York	localhost (1)
307	AI_Host3	192.1.168.90	New York	localhost (1)

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.