



Data Warehouse administration

OnCommand Insight

NetApp
March 08, 2024

Table of Contents

- Data Warehouse administration 1
 - Welcome to OnCommand Insight Data Warehouse 1
 - Getting started with Data Warehouse 7
 - Administrative tasks you can perform using Data Warehouse 28

Data Warehouse administration

Welcome to OnCommand Insight Data Warehouse

The OnCommand Insight Data Warehouse is a centralized repository that stores data from multiple OnCommand Insight servers and transforms data into a common, multidimensional data model for querying and analysis.

The OnCommand Insight Data Warehouse enables access to an open database consisting of several data marts that let you generate custom capacity and performance reports such as chargeback reports, trending reports with historical data, consumption analyses, and forecasting reports.

Data Warehouse features

The OnCommand Insight Data Warehouse is an independent database made up of several data marts.

Data Warehouse includes the following features:

- Current and historical configuration and inventory data that enables you to create trending reports useful for forecasting and planning
- Several multidimensional historical data marts and an additional current-only inventory data mart
- An optimized database for predefined queries or user-defined queries
- A platform for integration with third-party reporting and business intelligence engines, including:
 - Configuration management databases
 - Financial accounting systems
 - Asset management systems

Data Warehouse components

Data Warehouse contains several components.

- Data Warehouse Portal
- OnCommand Insight Reporting Portal
- Report authoring tools

What you can do using the Data Warehouse Portal

The Data Warehouse Portal is a web-based user interface that you use to configure options and set up fixed schedules to retrieve data. From the Data Warehouse Portal, you can also access the OnCommand Insight reporting portal.

Using Data Warehouse portal, you can do the following:

- Access the OnCommand Insight reporting portal to view predesigned reports or to create custom reports using report authoring tools.

- Consolidate multiple OnCommand Insight databases.
- Manage connections to OnCommand Insight servers.
- Check the status of current jobs or queries that are running.
- Schedule Data Warehouse builds.
- Edit the site name.
- View Data Warehouse version and upgrade history, including specific information such as module versions, sites, and licenses.
- Import annotations.
- Configure a build from history.
- View Data Warehouse documentation and the database schema.
- Reset the Data Warehouse database.
- Back up and restore the Data Warehouse database.
- Troubleshoot Data Warehouse issues.
- Manage user accounts.

Data Warehouse software components

OnCommand Insight Data Warehouse includes several software components.

- **MySQL database**
The back-end repository for data mart tables
- **IBM Cognos**
The reporting engine for OnCommand Insight
- **Apache Derby Database**
Used for storing Cognos configuration and content
- **WildFly**
The Java Enterprise application server that hosts OnCommand Insight components

Data Warehouse processes

Data Warehouse performs many types of processes.

- **ETL process**
The Extract Transform and Load (ETL) process retrieves data from multiple OnCommand Insight databases, transforms the data, and saves it into the data mart. The Data Warehouse build process is an ETL process.
- **Jobs**
Data Warehouse performs and reports on jobs such as these: inventory, dimensions, capacity, port

capacity, VM capacity, file system utilization, performance, capacity efficiency, licenses, history build, dynamic annotations, connector removal, skipped build, ASUP option, and maintenance jobs.

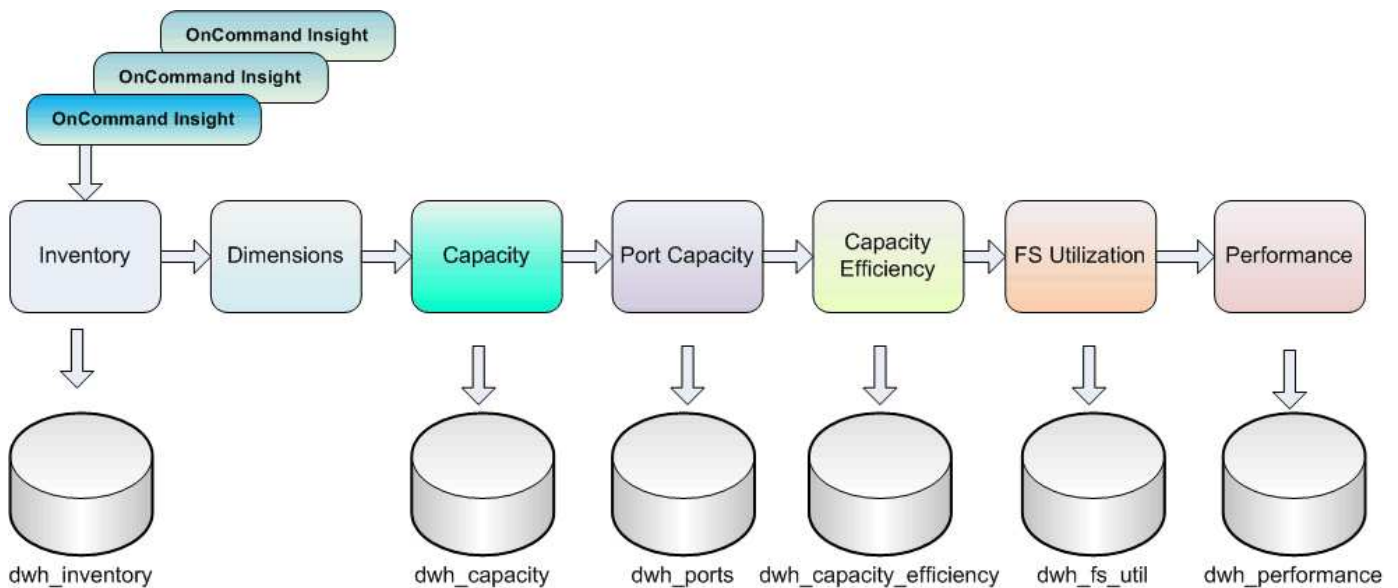
- **Consolidation process**

Data Warehouse supports the consolidation of multiple OnCommand Insight servers into the same Data Warehouse database. In many configurations it might happen that the same object is reported from multiple connectors (that is, the same switch exists in two OnCommand Insight instances). In that case, Data Warehouse consolidates the multiple objects into one (a primary connector is chosen and the object's data is taken from that connector only).

How Data Warehouse extracts data

The Extract, Transform, and Load (ETL) process retrieves data from multiple OnCommand Insight databases, transforms the data, and saves it into the data marts.

OnCommand Insight connectors invoke a series of batch jobs to extract data from multiple OnCommand Insight mySQL databases and publish the data in various data marts, as shown in the following diagram.



The ETL process includes these individual processes:

- **Extract**

This process takes data from multiple OnCommand Insight databases, transforms the data, and saves it into the data mart. The process is performed against each OnCommand Insight instance at the same time. To ensure that data cleansing and deduplication is performed, it is not possible to split the ETL process into multiple scheduled ETL operations.

- **Transform**

This process applies business logic rules or functions to extract the data from the OnCommand Insight database.

- **Load**

This process loads the transformed data into public data marts.

ETL frequency and date data

You should run the Extract, Transform, and Load (ETL) process at least once per day; however, you choose to run ETL numerous times if needed.

By default, the Cognos reporting engine treats all capacity and performance facts as additive. As a result, there is a risk of double counting capacity data if the ETL process is run multiple times per day without the proper time filters.

Two date data elements in the Date dimension are related to the daily ETL process. The Date dimension, which is used in several data models, includes the following data elements that are affected by the ETL:

- **Is Day Representative**

The "Is Day Representative" data element is set to a value of 1 (true) during the first ETL process run during any given day. If the first ETL process is run at 1:00 a.m., Is Day Representative is set to 1 for all of the data loaded during the 1:00 a.m. ETL process. If a second ETL is scheduled later (for example, 1:00 p.m.), Is Day Representative is set to 0 (false) for the data loaded during that ETL process.

- **Is Latest**

The "Is Latest" member is set to a value of 1 (true) after each ETL process completes. If the first ETL process is run at 1:00 a.m., Is Latest is set to 1 for all of the data loaded during the 1:00 a.m. ETL process. If another ETL process is scheduled later (for example, 1:00 p.m.), Is Latest is set to 1 for data loaded during the 1 p.m. ETL process. The ETL process also sets the 1:00 a.m. ETL load's Is Latest entry to 0 (false).

How historical data is retained in Data Warehouse

Data is maintained in Data Warehouse according to a schedule. As data gets older, the data record retention is reduced.

Data Warehouse retains historical data based on the data marts and granularity of the data, as shown in the following table.

Data mart	Measured object	Granularity	Retention period
Performance marts	Volumes and internal volumes	Hourly	14 days
Performance marts	Volumes and internal volumes	Daily	13 months
Performance marts	Application	Hourly	13 months
Performance marts	Host	Hourly	13 months
Performance marts	Switch performance for port	Hourly	5 weeks

Performance marts	Switch performance for host, storage, and tape	Hourly	13 months
Performance marts	Storage node	Hourly	14 days
Performance marts	Storage node	Daily	13 months
Performance marts	VM performance	Hourly	14 days
Performance marts	VM performance	Daily	13 months
Performance marts	Hypervisor performance	Hourly	14 days
Performance marts	Hypervisor performance	Daily	13 months
Performance marts	VMDK performance	Hourly	14 days
Performance marts	VMDK performance	Daily	13 months
Performance marts	Disk performance	Hourly	14 days
Performance marts	Disk performance	Daily	13 months
Capacity marts	All (except individual volumes)	Daily	13 months
Capacity marts	All (except individual volumes)	Monthly representative	14 months and beyond
Inventory marts	Individual volumes	Current state	1 day (or until next ETL)

After 13 months (which is configurable), Data Warehouse retains only one record per month instead of one record per day for capacity, performance, and resource data in the following fact tables:

- Chargeback fact table (dwh_capacity.chargeback_fact)
- File System Utilization fact table (dwh_fs_util.fs_util_fact)
- Host fact table (dwh_sa.sa_host_fact)
- Internal Volume Capacity fact table (dwh_capacity.internal_volume_capacity_fact)
- Ports fact table (dwh_ports.ports_fact)
- Qtree Capacity fact table (dwh_capacity.qtree_capacity_fact)
- Storage and Storage Pool Capacity fact table (dwh_capacity.storage_and_storage_pool_capacity_fact)
- Volume Capacity fact table (dwh_capacity.vm_capacity_fact)
- Storage Node Hourly Performance (storage_node_hourly_performance_fact) and Storage Node Daily Performance (storage_node_daily_performance_fact) fact tables

Data retention, ETL, and time periods

OnCommand Insight Data Warehouse retains data obtained from the Extract, Transform, and Load (ETL) process for different time periods based on the different data marts and time granularity of the data.

Performance Marts and hourly granularity for volumes and internal volumes

The OnCommand Insight Data Warehouse records the hourly averages, hourly maximums, and access bit for each hour of the day (24 data points) for 14 days. The access bit is a Boolean value that is true if the volume is accessed or false if the volume is not accessed during the hourly interval. All 24 data points for the preceding day are obtained during the first ETL process of the day.

You do not need to run one ETL process per hour to gather the hourly data points. Running additional ETL processes during the day does not obtain any performance information from the OnCommand Insight Servers.

Performance Marts and daily granularity for volumes and internal volumes

Each day when the ETL is processed, the daily averages for the preceding day are calculated and populated within Data Warehouse. The daily average is a summary of the 24 data points for the previous day. The performance data marts retain daily summaries for volumes and internal volumes for 13 months.

Capacity marts and daily granularity

The Capacity marts provide daily measurements for various capacity facts on a daily basis for a period of 13 months. The capacity facts in Data Warehouse are current as of the last data source acquisition for the device prior to the ETL.

Capacity marts and monthly granularity

Data Warehouse retains daily capacity data for 13 months. After the 13-month threshold is reached, the capacity data is summarized on a monthly basis. The monthly data is based on the values reflected by the date that is the month representative date.

The following table shows which monthly data is included in the monthly summary:

Date	Is Month Representative value	Allocated capacity
Jan 1	1 (True)	50 TB
Jan 2	0 (False)	52 TB
...
Jan 31	0 (False)	65 TB
Feb 1	1 (True)	65 TB

Based on the table, a monthly report would show 50 TB allocated for January and 65 TB allocated for February. All of the other capacity values for January would not be included in the monthly summary.

Inventory mart

The Inventory data mart is not historical. Each time an ETL process is run, the Inventory mart is erased and rebuilt. Therefore, any reports generated out of the Inventory mart do not reflect historical inventory configuration.

Getting started with Data Warehouse

OnCommand Insight Data Warehouse enables you to configure options needed before generating reports that include your data. Data Warehouse contains many features; however, you need to use only a few of them to get started. To set up Data Warehouse, you use options in the Data Warehouse Portal.

About this task

To set up OnCommand Insight Data Warehouse, a storage administrator should complete the following procedures:

- Accessing the Data Warehouse portal
- Connecting Data Warehouse to OnCommand Insight servers
- Building the database from history
- Setting up backup and restore processes

Additionally, a storage administrator might want to complete the following procedures.

- Accessing MySQL using the command line interface
- Scheduling daily builds
- Setting up multiple tenancy in reporting
- Troubleshooting setup issues
 - Why can't I see my annotations?
 - What should I do with failing historical build points?

If this is the first time you are using the Data Warehouse Portal, you must set up Data Warehouse before any information can appear on the Jobs page. You also need to repeat this setup process after resetting the Data Warehouse database.

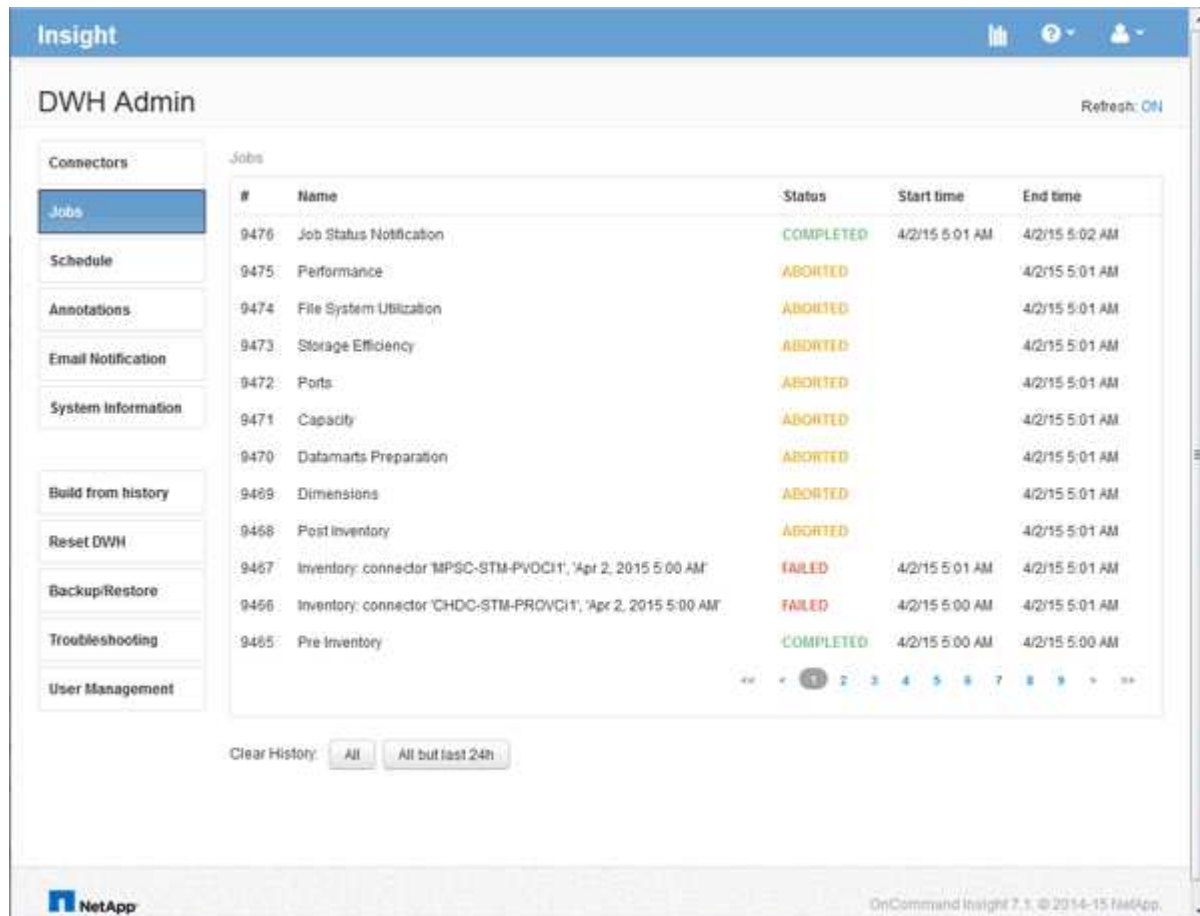
Accessing the Data Warehouse portal

The OnCommand Insight Data Warehouse portal is a web-based user interface that you can use to update connector information, view job queues, schedule daily builds, select annotations, set up email notifications, view system information, build the database, reset Data Warehouse, back up and restore the database, troubleshoot issues, manage Data Warehouse and Reporting portal user accounts, and access documentation and schema diagrams.

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. Enter your user name and password.
3. Click **Login**.

The Data Warehouse portal opens:



Managing Data Warehouse and Reporting user accounts

User accounts, user authentication, and user authorization for the OnCommand Insight reporting tools are defined and managed from the Data Warehouse (DWH). Based on these configurations, users and administrators gain access to some or all of the available OnCommand Insight reports.

Access to the User Management in the Data Warehouse requires an account with System Administrator privileges. This includes:


- Full administrative capabilities for the Data Warehouse
- Configuration and maintenance of all user accounts
- Read access to the database
- Capability to set up connectors in the ETL, schedule Data Warehouse jobs, reset the database, assign or

change roles, and add and remove user accounts

Accessing the Data Warehouse and Reporting portal

The Data Warehouse portal provides access to administration options. From the Data Warehouse portal, you can also access the Reporting portal.

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. On the Data Warehouse toolbar, click  to open the Reporting portal.

Reporting user roles

Each user account is assigned a role with a set of permissions. The number of users is limited by the number of Reporting licenses attached to each role.

Each role can perform the following actions:

- **Recipient**

Views OnCommand Insight Reporting portal reports and sets personal preferences such as those for languages and time zones.



Recipients cannot create reports, run reports, schedule reports, export reports, nor perform administrative tasks.

- **Business Consumer**

Runs reports and performs all Recipient options.

- **Business Author**

Views scheduled reports, runs reports interactively, creates stories, in addition to performing all Business Consumer options.

- **Pro Author**

Creates reports, creates packages and data modules, in addition to performing all Business Author options.

- **Administrator**

Performs reporting administrative tasks such as the import and export of report definitions, configuration of reports, configuration of data sources, and the shutdown and restart of reporting tasks.

The following table shows the privileges and the maximum number of users allowed for each role:

Feature	Recipient	Business Consumer	Business Author	Pro Author	Admin
---------	-----------	-------------------	-----------------	------------	-------

View reports in the Team Content tab	Yes	Yes	Yes	Yes	Yes
Run reports	No	Yes	Yes	Yes	Yes
Schedule reports	No	Yes	Yes	Yes	Yes
Upload external files	No	No	Yes	Yes	No
Create stories	No	No	Yes	Yes	No
Create reports	No	No	Yes	Yes	No
Create Packages and Data Modules	No	No	No	Yes	No
Perform administrative tasks	No	No	No	No	Yes
Number of users	Number of OnCommand Insight users	20	2	1	1

When you add a new Data Warehouse and Reporting user, if you exceed the limit in a role, the user is added as “deactivated,” and you need to deactivate or remove another user with that role to give a new user membership.



Report authoring capabilities require Insight Plan license. You can add additional Business Author and Pro Author users by purchasing the ARAP (Additional Report Authoring Package). Contact your OnCommand Insight representative for assistance.

These reporting user roles do not affect direct database access. These reporting user roles do not impact your ability to create SQL queries using the data marts.

Adding a Reporting user

You must add a new user account for each person who requires access to the Reporting portal. Having a different user account for each person provides a way of controlling access rights, individual preferences, and accountability.

Before you begin

Before adding a Reporting user, you must have allocated a unique user name, determined what password to use, and verified the correct user role or roles. These roles are specialized in the Reporting portal.

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 **User Management**.
3. In the **User Management** window, click **Add New User**.
4. Enter the following information for the new Reporting user:

- **User name**

User name (alphanumeric, including a-z, A-Z, and 0-9) for the account

- **E-mail Address**

Email address associated with the user account and required if the user subscribes to any reports

- **Password**

Password to log in to OnCommand Insight with this user account, which is typically selected by the user and confirmed in the interface

- **Insight role**

Roles available to the user with appropriate permissions



The options for the OnCommand Insight role are shown only if OnCommand Insight is installed on the same machine as the reporting facilities, which is not typical.

- **Reporting roles**

Reporting role for this user account (for example, Pro Author)



The Administrator role is unique. You can add this role to any user.

5. Click **Add**.

Managing user accounts

You can configure user accounts, user authentication, and user authorization from the Data Warehouse portal. Each user account is assigned a role with one of the following permission levels. The number of users is limited by the number of Reporting licenses attached to each role.

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 **User Management**.

Name	OnCommand Insight roles			Reporting roles					E-mail				
	Guest	User	Administrator	Recipient	Business Consumer	Business Author	Pro Author	Administrator					
guest	X									Edit	Delete	Change password	Deactivate
user	X	X								Edit	Delete	Change password	Deactivate
admin	X	X	X				X	X		Edit		Change password	
oadmin	X	X	X							Edit		Change password	Deactivate

LDAP Configuration

Add New User

Change DWH User password

The following table shows the privileges for each reporting role:

Feature	Recipient	Business Consumer	Business Author	Pro Author	Administrator
View reports (in Public Folder tab, My Folders)	Yes	Yes	Yes	Yes	Yes
Run reports	No	Yes	Yes	Yes	Yes
Schedule Reports	No	Yes	Yes	Yes	Yes
Create reports in Query Studio	No	No	Yes	Yes	No
Create reports in Workspace (Standard)	No	Yes	Yes	Yes	No
Create reports in Workspace (Advanced)	No	No	Yes	Yes	No
Create reports in Report Studio	No	No	No	Yes	No
Perform administrative tasks	No	No	No	No	Yes

3. Do one of the following:

- To edit an existing user, select the row for the user and click **Edit**.
- To change a user's password, select the row for the user and click **Change password**.
- To delete a user, select the row for the user and click **Delete**

4. To activate or deactivate a user, select the row for the user and click **Activate** or **Deactivate**.

Configuring LDAP for Reporting

From the Data Warehouse portal, the Administrator can configure LDAP usage for Data Warehouse and Reporting.

Before you begin

You must log in to Insight as an Administrator to perform this task.

For all Secure Active Directory (i.e. LDAPS) users, you must use the AD server name exactly as it is defined in the certificate. You can not use IP address for secure AD login.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system on which OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **User Management**.
3. Click **LDAP Configuration**.
4. Select **Enable LDAP** to start the LDAP user authentication and authorization process.
5. Make whatever changes are necessary to configure LDAP.

The majority of the fields contain default values. The default settings are valid for the Active Directory.

- **User principal name attribute**

Attribute that identifies each user in the LDAP server. Default is `userPrincipalName`, which is globally unique. OnCommand Insight attempts to match the contents of this attribute with the username

that has been supplied above.

- **Role attribute**

LDAP attribute that identifies the user's fit within the specified group. Default is `memberOf`.

- **Mail attribute**

LDAP attribute that identifies the user's email address. Default is `mail`. This is useful if you want to subscribe to reports available from OnCommand Insight. Insight picks up the user's email address the first time each user logs in and does not look for it after that.



If the user's email address changes on the LDAP server, be sure to update it in Insight.

- **Distinguished name attribute**

LDAP attribute that identifies the user's distinguished name. default is `distinguishedName`.

- **Referral**

Indicates whether to follow the path to other domains if there are multiple domains in the enterprise. You must always use the default `follow` setting.

- **Timeout**

Length of time to wait for a response from the LDAP server before timing out, in milliseconds. default is 2,000, which is adequate in all cases and should not be modified.

- **LDAP servers**

This is the IP address or DNS name to identify the LDAP server. To identify a specific port, where `ldap-server-address` is the name of the LDAP server, you can use the following format:

```
ldap://ldap-server-address:port
```

To use the default port, you can use the following format:

```
ldap://ldap-server-address
```



When entering multiple LDAP servers in this field, separate entries with a comma, and ensure that the correct port number is used in each entry.

+ To import the LDAP certificates, click **Import Certificates** and automatically import or manually locate the certificate files.

- **Domain**

LDAP node where OnCommand Insight should start looking for the LDAP user. Typically this is the top-level domain for the organization. For example:

```
DC=<enterprise>,DC=com
```

- **Insight server admins group**

LDAP group for users with Insight Server Administrator privileges. Default is `insight.server.admins`.

- **Insight administrators group**

LDAP group for users with Insight Administrator privileges. Default is `insight.admins`.

- **Insight users group**

LDAP group for users with Insight User privileges. Default is `insight.users`.

- **Insight guests group**

LDAP group for users with Insight Guest privileges. Default is `insight.guests`.

- **Reporting administrators group**

LDAP group for users with Insight Reporting administrator privileges. Default is `insight.report.admins`.

- **Reporting pro authors group**

LDAP group for users with Insight Reporting pro authors privileges. Default is `insight.report.proauthors`.

- **Reporting business authors group**

LDAP group for users with Insight Reporting business authors privileges. Default is `insight.report.business.authors`.

- **Reporting business consumers group**

LDAP group for users with Insight Reporting business consumers privileges. Default is `insight.report.business.consumers`.

- **Reporting recipients group**

LDAP group for users with Insight Reporting recipient privileges. Default is `insight.report.recipients`.

6. Enter values in the **Directory lookup user** and **Directory lookup user password** fields if you made any changes.

If you do not enter the revised values in these fields, your changes are not saved.

7. Retype the directory lookup user password in the **Confirm directory lookup user password** field, and

click **Validate Password** to validate the password on the server.

8. Click **Update** to save the changes. Click **Cancel** to remove changes.

Connecting Data Warehouse to OnCommand Insight servers

Connectors establish connections from the OnCommand Insight Data Warehouse to the OnCommand Insight servers. You can connect Data Warehouse with one or more OnCommand Insight servers. You can add or remove connections to or from OnCommand Insight databases.

About this task

Data Warehouse assigns a global unique ID to the connector that is used in conjunction with the connector name. After adding a connector, Data Warehouse queries the OnCommand Insight database for the OnCommand Insight site name and version.

You can choose to connect to a data source with or without SSL. Choosing the secure data source forces the connection to use SSL when communicating with the OnCommand Insight remote database.

Data Warehouse can provide a consolidated view of data from multiple OnCommand Insight installations. This consolidated database provides the following information:

- Globally Unique Identifiers

Each object is assigned a globally unique ID that is independent of the IDs used by individual sites, to avoid conflicting IDs and enable duplicate detection. These IDs are shared between all the data marts. This ID is the Globally Unique ID (GUID) in the Comment column of the Inventory data mart tables.

- No duplication

Entities that exist in multiple OnCommand Insight databases are registered only once in the consolidated database.

- Current record

The data in the consolidated database (Inventory data mart) is always the most up-to-date possible.

When you add or edit a connection, you can also test the connection. The test does the following:

- Verifies the host IP address, user name, and password and ensures that a connection can be established.

Invalid connections appear in red.

- Compares the OnCommand Insight version to the Data Warehouse version.

If the versions are not compatible, an error message appears.

- Verifies that the OnCommand Insight database has not been changed or restored to a different database as seen by the last Data Warehouse processing. If there has been a change, an error message appears.

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 **Connectors**.

The Connectors table appears blank at first and shows connector information after you add a connector.

3. Click **New** to add a new connector.

4. Enter the following:

- **Encryption**

To enable Data Warehouse requests to be made using SSL encryption, select `Enabled`.

- **Name**

A connector name that will identify the connector on the Connectors view.

- **Host**

Host IP address

- **User name**

"inventory"



Using this user name and password, you can log in to the remote OnCommand Insight database and perform queries on the data.

- **Password**

"sanscreen"

5. To specify the port to use for TCP connections to the host, click **Advanced** and enter the TCP port number.
6. To specify the port (other than the default port) to use for HTTPS connections to the host, click **Advanced** and enter the port number.
7. Click **Test**.

Data Warehouse tests the connection.

8. Click **Save**.

If you enter multiple connections for multiple installations, Data Warehouse invokes independent build processes, one for each database from which data should be extracted. Each such build process extracts data from an OnCommand Insight database and loads it into the consolidated database.

Data Warehouse database build from history overview

You can build the Data Warehouse database using historical data in your OnCommand Insight server. Data Warehouse extracts data from the OnCommand Insight servers and builds the Data Warehouse data marts according to the build from history schedule.

This option does not require a special license and inventory data is included in the build. However, to build capacity information, the OnCommand Insight Plan and OnCommand Insight Perform licenses are required.

If any build (from history or current) has already been performed, the build cannot be done on dates before the last job. This means if you perform a current build, you cannot build from history. More specifically, if you performed builds from history that ended on Jan 1, 2012, you cannot perform any build on the year 2011.

If the history build does not include a day or two of any unsuccessful ETL processes, do not try building history for just these few days. Historical data is for longer periods and a day or two is not going to change trending significantly. If you do want to rebuild from history, rebuild the entire history.

The Build from History view displays all build jobs from all connectors. For example, the view might display an inventory job for every connector, a port capacity job for every build run, and an annotations job.

Before you configure the Build from History, the following must occur:

- Connectors must be configured.
- Annotations should be entered in OnCommand Insight and can be manually updated using the **Force Update of Annotations for DWH option** in the old OnCommand Insight Portal or will be automatically updated 15 minutes after they are set.

Adding a job that builds a Data Warehouse database from history

You can build the Data Warehouse database using historical data that is kept in your OnCommand Insight server, which enables you to run projection reports.

Before you begin

You must have updated annotations in the OnCommand Insight server and forced an update of annotation information for Data Warehouse.

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

Cancel Pending Jobs

Configure

Run

Skip history build failures: ☒

3. Click **Configure**.

Configure Build From History

Start time:	11	February
	2015	...
End time:	02	April
	2015	...
Interval:	<input checked="" type="radio"/> Daily <input type="radio"/> Weekly <input type="radio"/> Monthly <input type="radio"/> Quarterly	
Hour:	12:00 AM	
<input type="button" value="Save"/> <input type="button" value="Reset"/> <input type="button" value="Cancel"/>		

4. Enter the start and end times.

To display a calendar from which you can select these dates, click the down arrow near the month name.

The time format depends upon the locale of the Data Warehouse server.

The start and end times must be within the range of history contained in all the OnCommand Insight servers to which Data Warehouse is connected, as set in the Data Warehouse portal Connectors option. The default start and end times reflect the maximum valid period. The Data Warehouse build job runs automatically at the time you specify.



Configuring a non-realistic schedule such as “Daily for 4 years” results in 1460 build cycles, which could take 10 days to complete.

5. Choose the interval.

If you select a monthly or weekly interval, the Day field appears. If you selected monthly, then Day is a date. If you selected weekly, Day is Sunday through Saturday.

6. Choose the hour when the build will take place.

7. Optionally, to return the options to default settings, click **Reset**.

8. Click **Save**.

9. From the **Build from History** page, to perform a build outside of the automatic schedule build, click **Run**.

The Target Time column displays the time that this entry was built. The Status column displays whether the build was completed or failed.

Canceling a build from history job

You can cancel all planned jobs. The job status becomes “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**.
3. Click **Cancel**.

Backing up the Data Warehouse database

You can back up the Data Warehouse database, which also includes a Cognos backup, to a file and later restore it using the Data Warehouse portal. Such a backup enables you to migrate to a different Data Warehouse server or upgrade to a new Data Warehouse version.

Steps

1. Log in to the Data Warehouse Portal at `https://fqdn/dwh`.
2. From the navigation pane on the left, select **Backup/Restore**.
3. Click **Backup** and select your backup configuration:

- a. All Datamarts except Performance Datamart
- b. All Datamarts

This operation can take 30 minutes or more.

+ Data Warehouse creates a backup file and displays its name.

4. Right-click the backup file and save it to a location you want.

You might not want to change the file name; however, you should store the file outside the Data Warehouse installation path.

The Data Warehouse backup file includes the DWH instance's MySQL; custom schemas (MySQL DBs) and tables; LDAP configuration; the data sources that connect Cognos to the MySQL database (not the data sources that connect the Insight server to devices to acquire data); import and export tasks that imported or exported reports; reporting security roles, groups, and namespaces; user accounts; any modified Reporting Portal reports; and any custom reports, regardless of where they are stored, even in the My Folders directory. Cognos system configuration parameters, such as SMTP server setting, and Cognos custom memory settings are not backed up.

The default schemas where custom tables are backed up include the following:

dwh_capacity
dwh_capacity_staging
dwh_dimensions
dwh_fs_util
dwh_inventory
dwh_inventory_staging
dwh_inventory_transient
dwh_management
dwh_performance
dwh_performance_staging
dwh_ports
dwh_reports
dwh_sa_staging

Schemas where custom tables are excluded from backup include the following:

information_schema
acquisition

cloud_model
host_data
innodb
inventory
inventory_private
inventory_time
logs
management
mysql
nas
performance
performance_schema
performance_views
sansscreen
scrub
serviceassurance
test
tmp
workbench

In any backup initiated manually, a .zip file is created that contains these files:

- A daily backup .zip file, which contains Cognos report definitions
- A reports backup .zip file, which contains all the reports in Cognos, including those in the My Folders directory

- A Data Warehouse database backup file In addition to manual backups, which you can perform at any time, Cognos creates a daily backup (automatically generated each day to a file called `DailyBackup.zip`) that includes the report definitions. The daily backup includes the top folders and packages shipped with the product. The My Folders directory and any directories that you create outside the product's top folders are not included in the Cognos backup.



Due to the way Insight names the files in the `.zip` file, some unzip programs show that the file is empty when opened. As long as the `.zip` file has a size greater than 0 and does not end with a `.bad` extension, the `.zip` file is valid. You can open the file with another unzip program like 7-Zip or WinZip®.

Backing up custom reports and report artifacts

If you created custom reports in a version of Insight earlier than 7.0, and you want to upgrade to the newest Insight version, you should back up your reports and report artifacts before the upgrade installation and restore them after the upgrade installation. You should also pay attention to the folders that you are using to store report artifacts.

About this task

If you made changes to the predesigned reports, create your own copies of those reports in a separate folder. That way, when you update the predesigned artifacts, you do not overwrite your changes.

If you have reports in the My Folders area, you should copy them to the Custom Reports folders so that they are not lost.

Restoring the Data Warehouse database

You can restore a Data Warehouse database by using the `.zip` file that was created when you backed up that Data Warehouse database.

About this task

When you restore a Data Warehouse database, you have the option to restore user account information from the backup as well. User management tables are used by the Data Warehouse report engine in a Data Warehouse only installation.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system on which OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Backup/Restore**.
3. In the **Restore Database and Reports** section, click **Browse**, and locate the `.zip` file that contains the Data Warehouse backup.
4. If you want to restore reports or user account data, select one or both of the following check boxes:
 - **Restore database**

Includes Data Warehouse settings, data marts, connections, and user account information.

◦ **Restore reports**

Includes custom reports, predesigned reports, changes that you made to predesigned reports, and reporting settings that you created in the Reporting Portal.



If your database backup contains a custom report that has a slash (/) or an open bracket ([) in its name (for example, US IT Center Switch Port Boston/July), the restore operation renames the report, replacing the slash or open bracket with an underscore (for example, US IT Center Switch Port Boston_July).

5. Click **Restore**.

After the restore process is completed, a message is displayed below the Restore button. If the restore process is successful, the message indicates success. If the restore process fails, the message reports the specific exception that caused the failure. If an exception occurs and the restore process fails, the original database is automatically reset.

Setting up multiple tenancy in reporting

OnCommand Insight Data Warehouse accommodates multiple tenancy (often shortened to “multi-tenancy” or “multitenancy”) in reporting by allowing you to associate users with one or more business entities. With this feature, administrators can separate data or reports according to user attributes or user affiliation.

Business entities use a hierarchy for the purposes of capacity chargeback using the following values:

- Tenant: Primarily used by service providers to associate resources with a customer, for example, NetApp.
- Line of Business (LOB): A line of business within a company, for example "Hardware" or "Software."
- Business Unit: A traditional business unit such as "Sales" or "Marketing."
- Project: A project to which you might want to assign capacity chargeback.

The process of configuring multiple tenancy involves the following major steps:

- Configure a Data Warehouse user account.
- Create a group in Reporting Portal.
- Assign users to one or more groups, which represent business entities.
- Assign users to one or more business entities. For example, users associated with "NetApp" obtain access to all business entities that have “NetApp” as a tenant.
- Test that users can see only those reports that they should see.

The following points summarize how users access reporting data:

- A user, not assigned to any group, gets access to all the data.
- A user, assigned to any group, will not be able to get access to records without business entity.

For example, you might have the following departments and need to separate reports for users within these departments.

User	Engineering	Support	Finance	Legal
------	-------------	---------	---------	-------

User1	X	X		
User2			X	X
User3		X		

Configuring user accounts

You must complete several steps to configure user accounts.


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 **User Management**.
3. Configure each user account.

Assigning users to business entities

You must complete a series of steps to assign users to business entities. Data Warehouse accommodates multiple tenancy (often shortened to “multi-tenancy” or “multitenancy”) in reporting by allowing you to associate users with one or more business entities. This enables administrators to separate data or reports according to user attributes or user affiliation.

Steps

1. Log in to the Data Warehouse Portal as administrator 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  to open the Reporting Portal.
3. Enter your user name and password and click **Login**.
4. From the Launch menu, select **IBM Cognos Administration**.
5. Click the **Security** tab:
6. In the Directory, select **Cognos**.
7. Create a new subfolder in the Cognos folder called “BEs”, for business entities.
8. Open the BEs folder.
9. Click the **New Group** icon to add groups that correspond to different permission levels.

These permission levels can be either the full name of the business entity (for example, NetApp.N/A) or a prefix (for example, NetApp.N/A.Finance). Either of these formats enables access to all projects within the business entity (NetApp.N/A.Finance).

The New Group wizard displays.

10. Complete the pages of the wizard.
11. Select a business entity and click **More**.

12. Click **Set members**.
13. Click **Add**.
14. Select the SANscreen directory.
15. From the list of users, select each user that you want to include in the Business Entity and add the user to the Selected Entries box.
16. Click **OK**.
17. Repeat the process of adding members to each of the business entity groups.

Troubleshooting setup issues

There are several common issues with annotations, builds, and reports that you may face during setup. You can troubleshoot these issues by following the steps outlined.

Why I cannot see my annotations

If you cannot see annotations in Data Warehouse, you might need to force an update of annotations and then initiate a Data Warehouse build.

Missing annotations affect the way data is imported into Data Warehouse and is displayed in the reports. For example, if the annotation “Tier” is not available, you will not be able to group storage systems by tier in Data Warehouse reports.

Forcing an update of annotations for Data Warehouse

You can initiate an update of annotations from OnCommand Insight to Data Warehouse.

About this task

You can update annotations using one of two options:

- Including deleted objects: This includes data about devices that no longer exist such as hosts, storage arrays, or switches that were removed. This is needed if you want to build Data Warehouse data with historical data points.
- Not including deleted objects: Choose this option if you want to exclude deleted objects.

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)**.

Generating a manual Data Warehouse build

After forcing an annotations update (running transient data) in OnCommand Insight, you need to 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 **Schedule**.
3. Click **Build now**.

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.

What to do with failing historical build points

You can build from history, omitting any failed builds by enabling the **Skip history build failures** option.

If you do this, the build from history continues. If a build fails and this option is enabled, Data Warehouse continues building and ignores any failed builds. In such cases, there is no data point in the historical data for any skipped builds. If you do not enable this option and the build fails, all subsequent jobs are aborted.

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

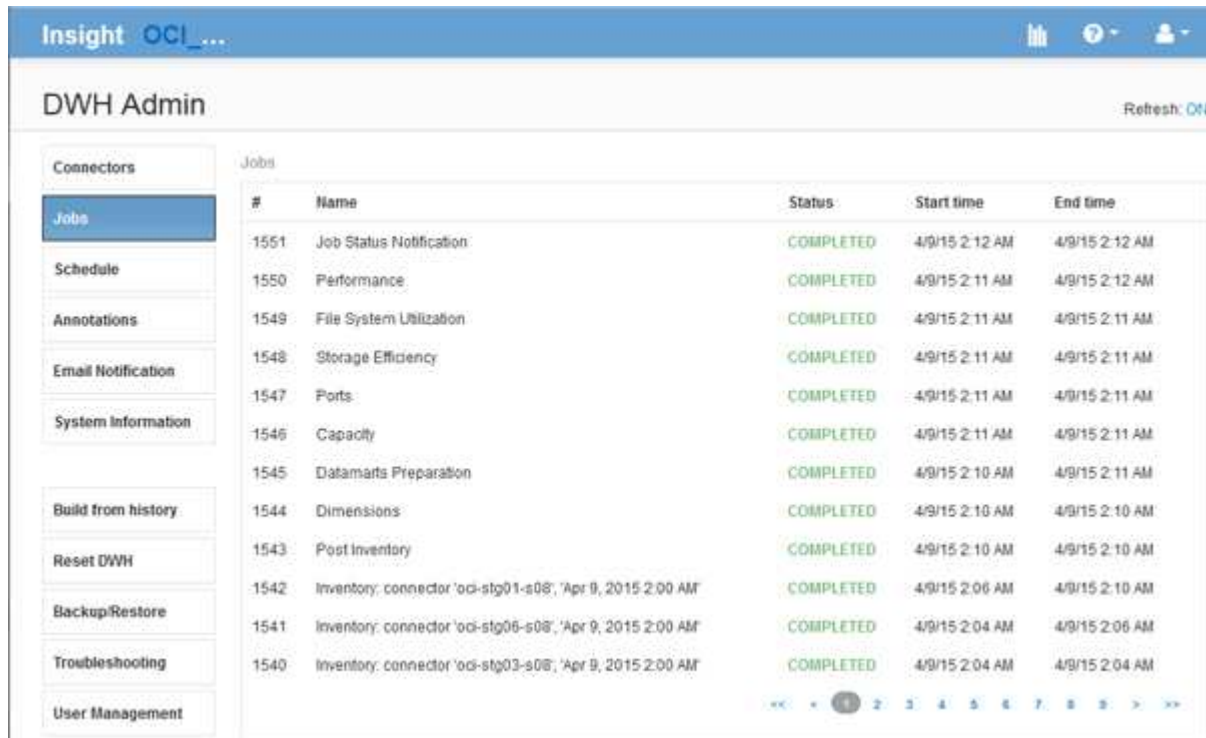
You can cancel any pending job you have scheduled or begun. You can also clear the history of previously executed jobs. You can clear history of jobs that are not pending, running, or in the process of aborting. You can clear all history or all history except the previous 24 hours to remove all but the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



The screenshot shows the 'DWH Admin' interface with the 'Jobs' tab selected in the left navigation pane. The main area displays a table of jobs, all with a 'COMPLETED' status. The table has columns for Job ID, Name, Status, Start time, and End time. A 'Refresh: ON' button is in the top right. Below the table is a pagination control showing page 1 of 1.

#	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:	<input type="text" value="yes"/>	<input type="button" value="Edit schedule"/>
Schedule:	<input type="text" value="Daily at 2:00 AM,7:00 PM"/>	
Next run:	<input type="text" value="4/2/15 7:00 PM"/>	

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.

- 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

- 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**.

Database Schema Databases 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	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:

Database Schema Databases 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	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.

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

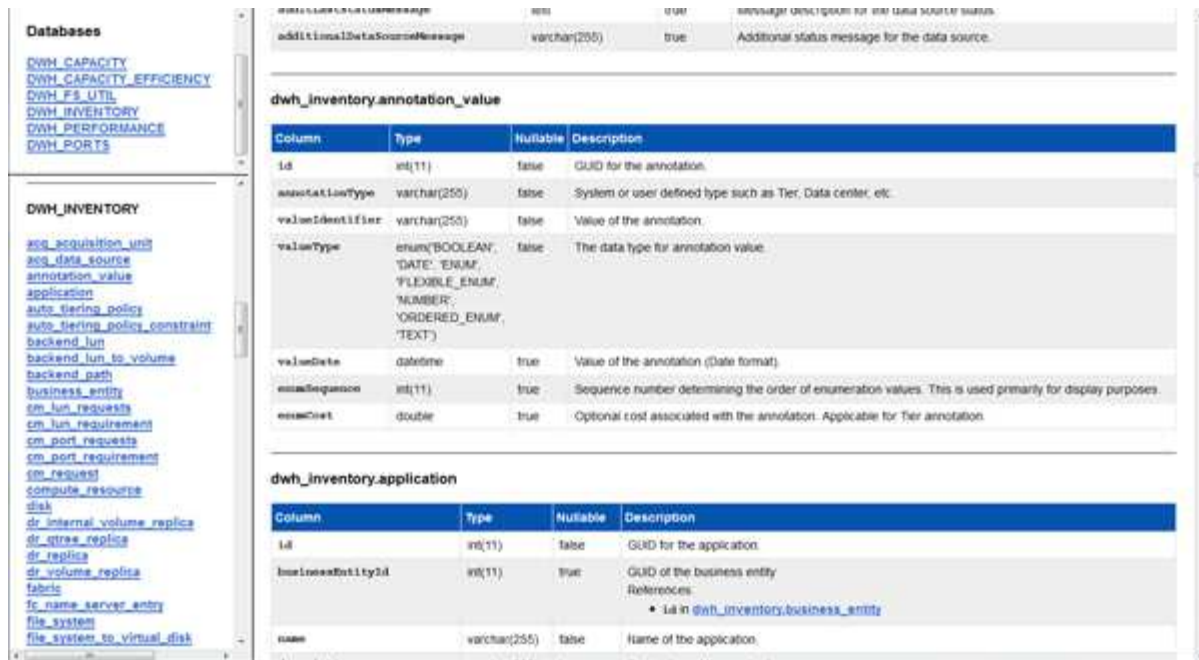
1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. On the Data Warehouse toolbar, click  and select **Schema**.

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 dw_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

Filter

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.