OnCommand® Insight 7.2

Configuration and Administration Guide

April 2016 | 215-10390_B0
doccomments@netapp.com

NetApp®

Contents

OnCommand Insight overview 7
Insight architecture 7
Insight Data Warehouse architecture 8
Firewall-friendly architecture 8
How Insight is used by administrators, managers, and planners 9
Finding answers to your Insight questions 10
Accessing Help directly 10
Locating product information on the NetApp Support Site 10
Setting up Insight 12
Accessing the web UI 12
Installing your Insight licenses 13
Updating expired licenses 14
OnCommand Insight licenses 14
Setting up and managing user accounts 17
Insight user roles 18
Configuring user definitions using LDAP 18
Changing user passwords 20
Editing a user definition 21
Deleting a user account 21
Smart Card and certificate login support 21
Configuring hosts for Smart Card and certificate login 22
Configuring a client to support Smart Card and certificate login 23
Configuring Data Warehouse for Smart Card and certificate login 24
Configuring Data Warehouse for Smart Card and certificate login for Cognos 25
Setting up weekly backups for your Insight database 26
Configuring your email 27
Email configuration prerequisites 27
Configuring your email for Insight 28
Configuring SNMP notifications 28
Enabling the syslog facility 29
Insight syslog contents 30
Configuring performance and assure violation notifications 30
Configuring your ASUP processing 30
Creating custom dashboards 32
Deleting a custom dashboard 34
How OnCommand Insight anomaly detection works 34
Installing the anomaly detection software 35
Using the Anomalies section 36
Defining applications 38
Monitoring OnCommand Insight system health 39

Your business entities hierarchy 39
Designing your business entities hierarchy 40
Creating business entities 41
Defining applications 42
Assigning business entities to assets 44
Defining annotations 45
Using annotations to monitor your environment 45
Querying assets 57
Creating a query 57
Viewing queries 58
Exporting query results to a .CSV file 58
Modifying queries 59
Deleting queries 59
Assigning multiple applications to or removing multiple applications from assets 59
Editing or removing multiple annotations from assets 60
Managing performance policies 61
Creating performance policies 62
Performance policy evaluation precedence 64
Editing performance policies 64
Deleting performance policies 65
Insight data source management 66
Setting up your data sources in Insight 66
Data source support information 67
Adding data sources 67
Cloning a data source 69
Testing the data source configuration 70
Vendor-specific data source reference 71
3PAR InServ Firmware 2.2.2+ (SSH) data source 71
Brocade EFCM (9.5, 9.6, 9.7) data sources 72
Brocade FC Switch Firmware (SSH) data source 74
Brocade Sphereon/Intrepid Switch (SNMP) data source 76
Cisco FC Switch Firmware (SNMP) data source 78
Dell Compellent Storage Center data source 81
EMC Atmos (HTTP) data source 82
EMC Celerra data source 83
EMC CLARiiON (NaviCLI) data source 85
EMC Data Domain data source 87
EMC ECC StorageScope data sources 88
EMC XtremIO (HTTP) data source 88
EMC Isilon data source 89
EMC Solutions Enabler (CLI) with SMI-S Performance data source 90
EMC VNX (SSH) data source 93
EMC VPLEX data source 94
Fujitsu Eternus data source 95

HDS HiCommand Device Manager data source 97
HDS HiCommand Device Manager Lite data source 99
HP CommandView AE data source 99
HP CommandView AE Lite data source 101
HP EVA Storage (SSSU) data source 102
IBM DS (CLI) data source 104
IBM PowerVM data source configuration and commands 105
IBM SVC (SSH) data source 107
IBM Tivoli Monitoring data source 109
IBM TotalStorage DS4000 data source 109
IBM XIV (CLI) data source 110
Microsoft Hyper-V (WMI) data source 111
NetApp Clustered Data ONTAP data source 114
NetApp Clustered Data ONTAP 8.1.x for Unified Manager data source ... 114
NetApp Data ONTAP operating in 7-Mode data source 115
NetApp E-Series data source 118
NetApp Host and VM File Systems data source 119
QLogic FC Switch (SNMP) data source 122
Violin Memory 6000 Series Flash Memory Array (HTTP) data source 122
VMware vSphere (Web Services) data source 124
Changing data source credentials 125
Changes causing data collection problems 126
Examining one data source in detail 126
Researching a failed data source 127
Controlling data source polling 129
Editing data source information 129
Editing information for multiple data sources 130
Deleting a data source 131
What data source patches are 132
Applying a data source patch 132
Installing a patch on one type of data source 133
Managing patches 133
Committing a data source patch 134
Rolling back a data source patch 135
Maintaining Insight 136
Managing Insight 136
Monitoring OnCommand Insight system health 137
Auditing system and user activities 137
Monitoring the violations in your network 139
Fibre Channel environment monitoring 140
Port metrics 141
Acquisition unit status 143
Restoring the Insight database 144
Updating expired licenses 145
Licenses no longer compliant 145

Replacing licenses for older Insight versions 145
Applying a service pack 145
Preparing a special troubleshooting report 146
Manually backing up the OnCommand Insight database 146
Bundling logs for Support 146
Sending information to NetApp Support 147
Advanced configuration tools 147
Restarting the OnCommand Insight Server 148
Configuring the number of hours to ignore dynamic data 149
Generating logs for Customer Support 150
Displaying system information 151
Listing installed OnCommand Insight components 151
Calculating the number of database objects 151
Moving MySQL data using the migrate option 152
Forcing annotation updates 153
Configuring Client settings 153
Checking a tier upgrade 154
Checking the status of server resources 154
Scrubbing data for transfer to support 154
Finding ghost data sources 157
Adding a missing disk model 157
Specialized administration toolkit 159
Copyright information 162
Trademark information 163
How to send comments about documentation and receive update notifications 164
Index 165

OnCommand Insight overview

OnCommand Insight enables you to simplify operational management of complex private and hybrid cloud and virtual IT environments. Insight is a single solution to enable cross-domain, multi-vendor resource management and analysis across networks, storage, and servers in physical and virtual environments.

Insight provides a “single pane of glass” for reporting on storage costs and provides the transparency needed to make decisions about performance and efficiency.

Insight architecture

OnCommand Insight enables you to administer your product easily, using a streamlined system architecture that includes the Insight Server, a collection engine, web-based and Java UIs, and data warehousing.

The major components of the Insight architecture are shown in this diagram and described after it:

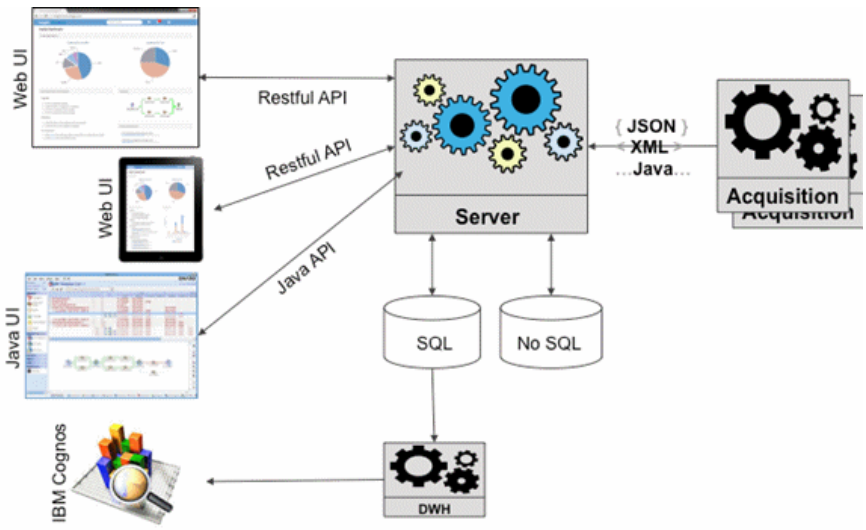

OnCommand Insight Server

The OnCommand Insight Server is the “brain” of the application. It includes main data repository and analysis components. The server is continuously building an end-to-end topology of the environment, analyzing it, and generating alerts when an incident or violation is detected.

Acquisition units

The Insight collection engine is built of one or more acquisition units. Each acquisition unit is a service running in the network that accesses (through modules called *data sources*) and collects data from different devices in the data center. Information collected by the acquisition units is then sent to the server (in an XML, JSON, or native Java format) for analysis.

The collection engine is designed to be highly modular and easily patched.

Web UI

The HTML5 web-based user interface (UI) for Insight enables you to set up your monitoring environment and data sources. You then use the web UI Asset Dashboard and asset pages to identify and research potential problems.

Java UI

This is the OnCommand Insight user interface (UI) or Client. You can use the Java UI to research issues like Fibre Channel mappings in your environment.

Data Warehouse (DWH)

Consolidates and prepares data for reporting for one or multiple installations of Insight. This includes history, trending, inventory, chargeback, show back and presenting the data in different ways to enable long-term planning of the data center's infrastructure.

IBM Cognos

This software is a reporting engine that provides a user interface for creating enterprise-level reports.

Insight Data Warehouse architecture

In a large environment, the OnCommand Insight Data Warehouse (DWH) consolidates data across different installations and hence different Insight data centers.

As shown in this diagram, the architecture enables users to view their entire environment and generate meaningful reports through a “single pane of glass” interface:

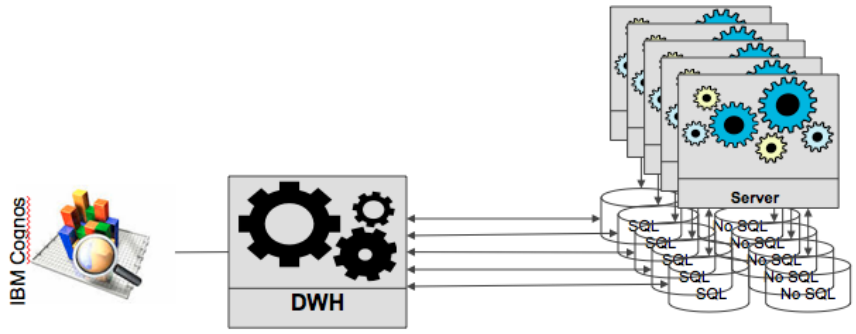

Firewall-friendly architecture

The OnCommand Insight architecture enables you to easily create firewalls around the product so that your assets are more secure.

As shown in the diagram, the architecture has these firewall features:

- All OCI clients use HTTPS to communicate with the server.
- The acquisition units enable only *outgoing* HTTPS connections; no ports are opened on the acquisition unit processes.

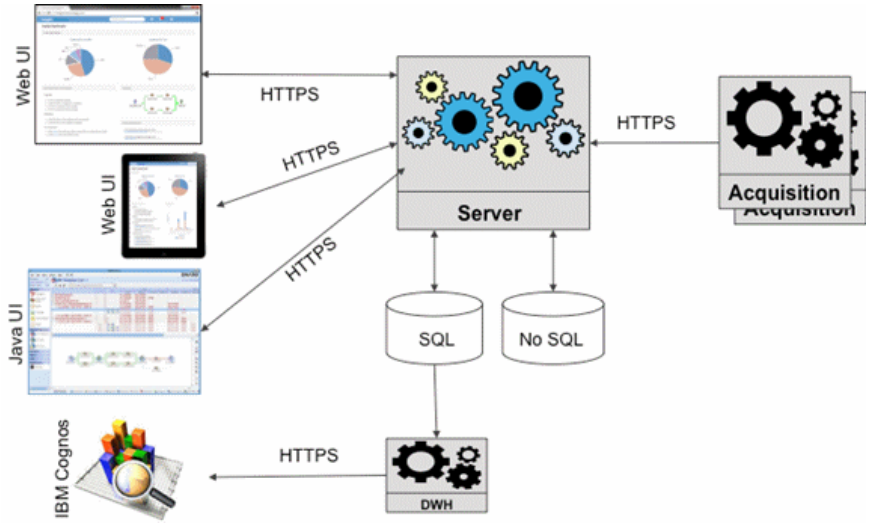

How Insight is used by administrators, managers, and planners

OnCommand Insight supplies information that is vital for storage administrators, managers, and storage architects to perform troubleshooting and analysis.

Experienced storage administrators use OnCommand Insight along with their network storage knowledge to accomplish these typical tasks:

- Manage the SAN and NAS environment.
- Work with SAN engineers on network concerns.
- Evaluate, test, and integrate new storage technologies into the environment.
- Troubleshoot performance issues, alerts, policy breaches, violations, and vulnerabilities.

Managers and network planners use OnCommand Insight to perform these business tasks:

- Capacity planning
- Develop project budgets and timelines.
- Evaluate and revise project plans to meet changing project demands.
- Manage project planning and expenses.
- Purchase hardware and software.
- Provide business reports for capacity management, charge back billing, right sizing, and service level agreements.

Finding answers to your Insight questions

OnCommand Insight provides you with product documentation as online Help and PDF guides. The web UI Help includes all of the Insight installation and upgrade instructions and information about the new web UI.

About this task

You can access the Insight documentation from the Help menu.

Steps

1. To display the documentation within the web UI, open Insight in your browser.
2. On the toolbar, click the question mark and select either of these options from the menu:
   - **Help topics**
     Insight documentation including installation and configuration
   - **Java UI Help**
     Instructions for using the Insight Client

Accessing Help directly

If you want to search the OnCommand Insight Help without opening the software, you can run both WebHelp versions in your browser from their server locations.

Before you begin

OnCommand Insight 7.x must be installed.

Steps

1. Open your browser.
2. Enter the URL for your Insight server with the Help directory using these formats for the corresponding Insight versions:
   - Web UI Help: `https://<oci_server_name>/documentation/oci_webhelp/index.html`
   - Java UI Help: `https://<oci_server_name>/legacy/help/webhelp/index.html`
3. Use the Search function or Index to locate the information you want.

Locating product information on the NetApp Support Site

The NetApp Support Site maintains full documentation sets in different formats. Videos showing you how to perform tasks in Insight are available on the Communities site.

About this task

You must have a NetApp Customer Support login name and password to access the Support and Communities sites.

Steps

1. To view all Insight product documentation, navigate to the product libraries on the NetApp Support Site.

a. Select OnCommand Insight.

b. Locate the Release number that you installed in the list of Insight releases.

c. Select the documents and formats that you want.

2. To run demonstration videos, open the NetApp Communities site.

a. Navigate to the Insight Communities site.

b. Select the **Content** tab.

c. Select the **Videos** tab.

Setting up Insight

To set up Insight, you must activate Insight licenses, set up your data sources, define users and notifications, enable backups, and perform any required advanced configuration steps.

After the OnCommand Insight system is installed, you must perform these setup tasks:

- Install your Insight licenses.
- Set up your data sources in Insight.
- Set up user accounts.
- Configure your email.
- Define your SNMP, email, or syslog notifications if needed.
- Enable automatic weekly backups of your Insight database.
- Perform any advanced configuration steps required, including defining annotations and thresholds.

Related concepts

Licenses no longer compliant on page 145

Related tasks

Updating expired licenses on page 14

Accessing the web UI

After you install OnCommand Insight, you must install your licenses and then set up Insight to monitor your environment. To do this, you use a web browser to access the Insight client (web UI).

Steps

1. Do one of the following:
   - Open Insight on the Insight server:

     `https://`*`fqdn`*
   - Open Insight from any other location:

     `https://`*`fqdn`*`:`*`port`*

     The port number is either 443 or another port configured when the Insight server was installed. The port number defaults to 443 if you do not specify it in the URL.

   The OnCommand Insight dialog box displays:

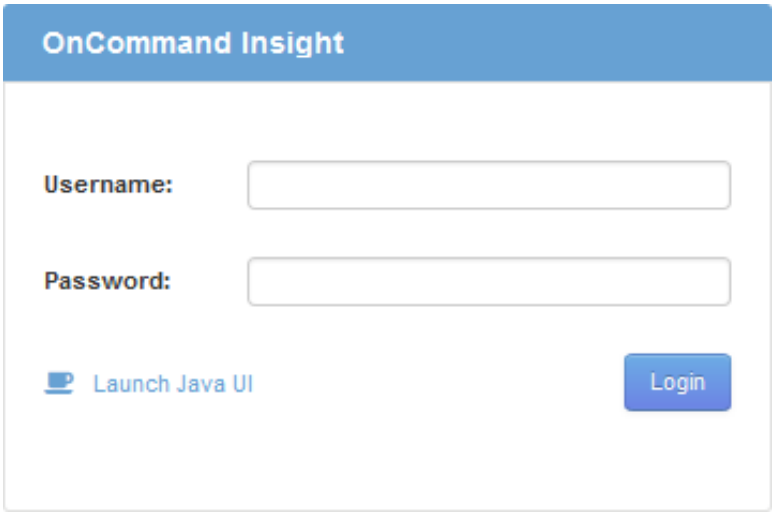

2. Enter your user name and password and click **Login**.

   The following table shows the default user name and password. Change these defaults as soon as possible after installation:

Data	Value
Default user name	`admin`
Default password	`admin123`

   If the licenses have been installed, the data source setup page displays.

   Note: An Insight browser session that is inactive for 30 minutes is timed out and you are logged out of the system.

Related tasks

Updating expired licenses on page 14

Installing your Insight licenses

After you receive the license file containing the Insight license keys from NetApp, you can use the setup features to install all of your licenses at the same time.

About this task

Insight license keys are stored in a `.txt` or `.lcn` file.

Steps

1. Open the license file in a text editor and copy the text.
2. Open Insight in your browser.
3. Click the **Admin** icon on the toolbar.
4. Click **Setup**.
5. Click the **Licenses** tab.
6. Click **Update License**.

7. Copy the license key text into the **License** text box.
8. Select the **Update (most common)** operation.
9. Click **Save**.

After you finish

After installing the licenses, you can perform these configuration tasks:

- Configure data sources.
- Create OnCommand Insight user accounts.

See the Insight Help or the *OnCommand Insight Configuration and Administration Guide* for instructions.

Updating expired licenses

If one or more of your Insight licenses expired, you can update the licenses quickly using the same procedure as you did to install the licenses originally.

Steps

1. In a text editor, such as Notepad, open the new license file you received from NetApp Support and copy the license key text to your Windows Clipboard.
2. Open OnCommand Insight in your browser.
3. Click the **Admin** icon on the toolbar.
4. Click **Setup**.
5. Click the **Licenses** tab.
6. Click **Update License**.
7. Copy the license key text into the **License** text box.
8. Select the **Update (most common)** operation.

   This operation adds your new licenses to any currently active Insight licenses.
9. Click **Save**.

OnCommand Insight licenses

OnCommand Insight operates with licenses that enable specific features on the Insight Server.

Discover

Discover is the basic Insight license that supports inventory. You must have a Discover license to use OnCommand Insight, and the Discover license must be paired with at least one of the Assure, Perform, or Plan licenses.

Assure

An Assure license provides support for assurance functionality, including global and SAN path policy, and violation management. An Assure license also enables you to view and manage vulnerabilities.

Perform

A Perform license supports performance monitoring on asset pages, dashboard widgets, queries, and so on, as well as managing performance policies and violations.

Plan

A Plan license supports planning functions, including resource usage and allocation.

Host Utilization pack

A Host Utilization license supports file system utilization on hosts and virtual machines.

Report Authoring

A Report Authoring license supports additional authors for reporting. This license requires the Plan license.

OnCommand Insight modules are licensed for annual term or perpetual:

- By terabyte of monitored capacity for Discover, Assure, Plan, Perform modules
- By number of hosts for Host Utilization pack
- By number of additional units of Cognos pro-authors required for Report Authoring

License keys are a set of unique strings that are generated for each customer. You can obtain license keys from your OnCommand Insight representative.

Your installed licenses control the following options that are available in the software:

Discover

Acquire and manage inventory (Foundation)

Monitor changes and manage inventory policies

Assure

View and manage SAN path policies and violations

View and manage vulnerabilities

View and manage tasks and migrations

Plan

View and manage requests

View and manage pending tasks

View and manage reservation violations

View and manage port balance violations

Perform

Monitor performance data, including data in dashboard widgets, asset pages, and queries

View and manage performance policies and violations

The following tables provide details of the features that are available with and without the Perform license for admin users and non-admin users.

Feature (admin)	With Perform license	Without Perform license
Application	Yes	No performance data or charts; no anomaly detection-related widgets
Virtual machine	Yes	No performance data or charts
Hypervisor	Yes	No performance data or charts
Host	Yes	No performance data or charts
Datastore	Yes	No performance data or charts
VMDK	Yes	No performance data or charts
Internal volume	Yes	No performance data or charts

Feature (admin)	With Perform license	Without Perform license
Volume	Yes	No performance data or charts
Storage pool	Yes	No performance data or charts
Disk	Yes	No performance data or charts
Storage	Yes	No performance data or charts
Storage node	Yes	No performance data or charts
Fabric	Yes	No performance data or charts
Switch port	Yes	No performance data or charts; “Port Errors” shows “N/A”
Storage port	Yes	Yes
NPV port	Yes	No performance data or charts
Switch	Yes	No performance data or charts
NPV switch	Yes	No performance data or charts
Search	Yes	Yes
Admin	Yes	Yes
Dashboard	Yes	Yes
Widgets	Yes	Partially available (only asset, query, and admin widgets are available)
Violations dashboard	Yes	Hidden
Assets dashboard	Yes	Partially available (storage IOPS and VM IOPS widgets are hidden)
Manage performance policies	Yes	Hidden
Manage annotations	Yes	Yes
Manage annotation rules	Yes	Yes
Manage applications	Yes	Yes
Queries	Yes	Yes
Manage business entities	Yes	Yes

Feature	User - with Perform license	Guest - with Perform license	User - without Perform license	Guest - without Perform license
Assets dashboard	Yes	Yes	Partially available (storage IOPs and VM IOPs widgets are hidden)	Partially available (storage IOPs and VM IOPs widgets are hidden)
Custom dashboard	View only (no create, edit, or save options)	View only (no create, edit, or save options)	View only (no create, edit, or save options)	View only (no create, edit, or save options)

Feature	User - with Perform license	Guest - with Perform license	User - without Perform license	Guest - without Perform license
Manage performance policies	Yes	Hidden	Hidden	Hidden
Manage annotations	Yes	Hidden	Yes	Hidden
Manage applications	Yes	Hidden	Yes	Hidden
Manage business entities	Yes	Hidden	Yes	Hidden
Queries	Yes	View and edit only (no save option)	Yes	View and edit only (no save option)

In some cases, you can also install a Replication Assurance license, which enables you to collect disaster recovery (DR) information. DR information can be monitored in the Java UI in the following views:

- Inventory > DR Paths
- Inventory > Replications Dashboard
- Assurance > DR Rules
- Assurance > DR Violations

Setting up and managing user accounts

User accounts, user authentication, and user authorization can be defined and managed in either of two ways: in Microsoft Active Directory (Version 2 or 3) LDAP (Lightweight Directory Access Protocol) server, or in an internal OnCommand Insight user database. Having a different user account for each person provides a way of controlling the access rights, individual preferences, and accountability. Use an account that has Administrator privileges for this operation.

Before you begin

You must have completed the following tasks:

- Install your OnCommand Insight licenses.
- Allocate a unique user name for each user.
- Determine what passwords to use.
- Assign the correct user roles.

Steps

1. Open Insight in your browser.
2. Click the **Admin** icon on the toolbar.
3. Click **Setup**.
4. Select the **Users** tab.

5. To create a new user, click the **Actions** button and select **Add user**.

   You enter the **Name**, **Password**, **Email** address, and select one of the user **Roles** as Administrator, User, or Guest.

6. To change a user's information, select the user from the list and click the **Edit user account** symbol to the right of the user description.

7. To remove a user from the OnCommand Insight system, select the user from the list and click **Delete user account** to the right of the user description.

Result

When a user logs in to OnCommand Insight, the server first attempts to authenticate through LDAP, if LDAP is enabled. If OnCommand Insight cannot locate the user on the LDAP server, it searches in local Insight database.

Insight user roles

Each user account is assigned one of the three possible permission levels.

- Guest permits you to log into Insight and to view the various pages. If your user account is defined to the OnCommand Insight local user database (and not through LDAP), you can also modify your own password. This account type does not allow you to perform actions such as identifying generic devices and defining the policies in the Java UI.
- User permits all guest-level privileges, as well as access to Insight operations such as defining policy and identifying generic devices. The User account type does not allow you to perform data source operations, nor to add or edit any user accounts other than your own.
- Administrator permissions vary depending on whether you use LDAP:
  - If you authenticate and authorize users through LDAP, this level of permission allows you manage data sources.
  - If you use only the local database to manage users, this level of permission allows you perform any operation, including adding new users and managing data sources.

Best Practice: Limit the number of users with Administrator permissions by creating most accounts for users or guests.

Configuring user definitions using LDAP

To configure OnCommand Insight (OCI) for user authentication and authorization from an LDAP server, you must be defined in the LDAP server as the OnCommand Insight server administrator.

About this task

This procedure assumes that you are using Microsoft Active Directory Version 2 or 3 LDAP (Lightweight Directory Access Protocol).

LDAP users display along with the locally defined users in the Users list.

Steps

1. On the toolbar, click the **Admin** icon.
2. Click **Setup**.
3. Click the **Users** tab.
4. Scroll to the bottom of the page to display the LDAP setup, as shown here.

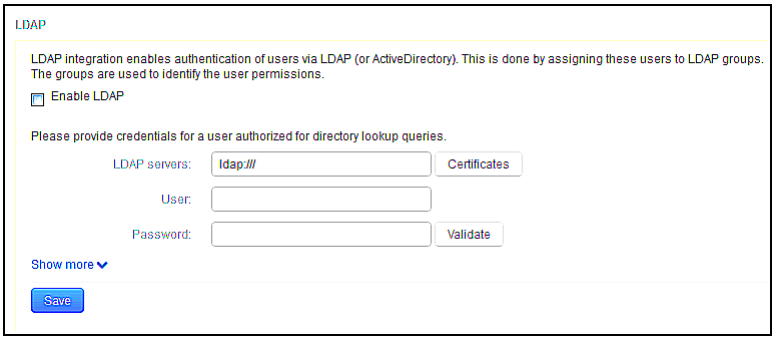

5. Click **Enable LDAP** to allow the LDAP user authentication and authorization.
6. Enter the credentials for a user authorized for directory lookup queries in these fields:
   - LDAP servers (To import the Insight certificate for this field, click **Certificates**.)
   - User name
   - Password (To confirm this password on the LDAP server, click **Validate**.)

   Note: The IP address or DNS name to identify the LDAP server, typically is entered in this format:

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

   or use the default port:

   ```
   ldap://<ldap-server-address>
   ```

7. If you want to define this LDAP user more precisely, click **Show more** and enter information for any of these attributes.

   The default settings are valid for the Active Directory.

   Admins group

   Value of the Role Attribute that identifies the user as having Administrator privileges (defaults to `sanscreen.admin`).

   Users group

   Value of the Role Attribute that identifies the user as having User privileges (defaults to `sanscreen.users`).

   Guests group

   Value of the Role Attribute that identifies the user as having Guest privileges (defaults to `sanscreen.guests`).

   Server admins group

   Value of the Role Attribute that identifies the user as having Server Administrator privileges (defaults to `sanscreen.server.admin`).

   Timeout

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

   Domain

LDAP node where OnCommand Insight should start looking for the LDAP user (typically the top-level domain for the organization). For example:

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

User principal name attribute

Attribute that identifies each user in the LDAP server (defaults to userPrincipalName, which is globally unique). OnCommand Insight attempts to match the contents of this attribute with the username that has been supplied.

Role attribute

LDAP attribute that identifies the user role (defaults to `memberOf`).

Mail attribute

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

Note: If the email address changes on the LDAP server, make sure to change it manually in OnCommand Insight.

Distinguished name attribute

LDAP attribute that identifies the user distinguished name (defaults to `distinguishedName`).

8. Click **Save**.

Changing user passwords

A user with administrator privileges can change the password for any OnCommand Insight user account defined on the local server.

Before you begin

The following items must have been completed:

- Notifications to anyone who logs into the user account you are modifying.
- New password to be used after this change.

About this task

When using this method, you cannot change the password for a user who is validated through LDAP.

Steps

1. Log in with administrator privileges.
2. Click the **Admin** icon on the toolbar.
3. Click **Setup**.
4. Click the **Users** tab.
5. Locate the row that displays the user account you want to modify.
6. To the right of the user information, click **Edit user account**.
7. Enter the new **Password** and then enter it again in the verification field.
8. Click **Save**.

Editing a user definition

A user with administrator privileges can edit a user account to change the email address or roles for OnCommand Insight or DWH and reporting functions.

Before you begin

Determine the type of user account (OnCommand Insight, DWH or a combination) that needs to be changed.

About this task

For LDAP users, you can only modify the email address using this method.

Steps

1. Log in with administrator privileges.
2. Click the **Admin** icon on the toolbar.
3. Click **Setup**.
4. Click the **Users** tab.
5. Locate the row that displays the user account you want to modify.
6. To the right of the user information, click the **Edit user account** icon.
7. Make the necessary changes.
8. Click **Save**.

Deleting a user account

Any user with Administrator privileges can delete a user account, either when it is no longer used (for a local user definition) or to force OnCommand Insight to rediscover the user information the next time the user logs in (for an LDAP user).

Steps

1. Log into OnCommand Insight with Administrator privileges.
2. Click the **Admin** icon on the toolbar.
3. Click **Setup**.
4. Click the **Users** tab.
5. Locate the row that displays the user account you want to delete.
6. To the right of the user information, click the **Delete user account** "**x**" icon.
7. Click **Save**.

Smart Card and certificate login support

OnCommand Insight supports use of Smart Cards (CAC) and certificates to authenticate users logging in to the Insight servers. You must configure hosts and clients to enable these features.

After configuring the hosts and clients to support CAC and certificates, navigating to a new session of OnCommand Insight results in the browser displaying a native dialog providing the user with a list

of personal certificates to choose from. These certificates are filtered based on the set of personal certificates that have been issued by CAs trusted by the OnCommand Insight server. Most often, there is a single choice. By default, Internet Explorer skips this dialog if there is only one choice.

For the Java client, a custom dialog appears with all personal certificates, regardless of issuer.

Note: For CAC users, smart cards contain multiple certificates, only one of which can match the trusted CA. The CAC certificate for `identification` should be used.

Configuring hosts for Smart Card and certificate login

You must make modifications to the OnCommand Insight host configuration to support Smart Card (CAC) and certificate logins.

Before you begin

- LDAP must be enabled on the system.
- The LDAP `User principal account name` attribute must match the LDAP field that contains a user's government ID number.

About this task

The common name (CN) stored on government-issued CACs is normally in the following format: `first.last.id`. For some LDAP fields, such as `sAMAccountName`, this format is too long. For these fields, OnCommand Insight only extracts the ID number from the CNs.

Steps

1. Use the `regedit` utility to modify registry values in `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\SANscreen Server`:
   a. Change the JVM_Option `DclientAuth=false` to `DclientAuth=want.`
   b. Change the JVM_Option `DauthClass=com.onaro.sanscreen.authenticate.CustomBasicAuthenticator` to `DauthClass=com.onaro.sanscreen.authenticate.CustomSSLAuthenticator.`
2. Restart the server.
3. Open a command prompt specifying
   `Run as administrator`
4. Change directory to `C:\Program Files\SANscreen\jboss\server\onaro\cert`.
5. Delete the self-generated certificate:
   ```
   keytool -delete -alias "ssl certificate" -keystore "C:\Program Files
   \SANscreen\jboss\server\onaro\cert\server.keystore
   ```
6. Generate a new certificate:
   ```
   keytool -genkey -alias "alias_name" -keyalg RSA -sigalg SHA1withRSA -
   keysize 2048 -keystore "C:\Program Files\SANscreen\jboss\server\onaro
   \cert\server.keystore" -dname
   "CN=commonName,OU=orgUnit,O=orgName,L=localityNameI,S=stateName,C=countr
   yName"
   ```
7. Generate a certificate signing request:
   ```
   keytool -certreq -sigalg SHA1withRSA -alias "alias_name" -keystore "C:
   \Program Files\SANscreen\jboss\server\onaro\cert\server.keystore" -file
   C:\temp\server.csr"
   ```

8. After the CSR is returned, import the certificate and then export the certificate in Base-64 format and place it in `"C:\temp" named servername.cer`.

9. Extract the certificate from the keystore:

   keytool -v -importkeystore -srckeystore ""C:\Program Files\SANscreen \jboss\server\onaro\cert\server.keystore"" -srcalias "alias_name" -destkeystore "C:\temp\file.p12" -deststoretype PKCS1

10. Extract a private key from the p12 file:

    openssl pkcs12 -in "C:\temp\file.p12" -out "C:\temp \servername.private.pem"

11. Merge the DoD-issued certificate with private key:

    openssl pkcs12 -export -in "C:\temp\DoD-issued.cer" -inkey "C:\temp \servername.private.pem -out "C:\temp\servername.new.p12" -name "servername.abc.123.yyy.zzz"

12. Import the merged certificate into the keystore:

    keytool -importkeystore -destkeystore "C:\Program Files\SANscreen\jboss \server\onaro\cert\server.keystore" -srckeystore "C:\temp \servername.new.p12" -srcstoretype PKCS12 -alias "alias_name"

13. Import the DoD root certificate:

    keytool -importkeysotre -keystore "C:\Program Files\SANscreen\java64\lib \security\cacerts" -file "C:\DoD_Root_CA_2__05__DoD_Root_CA_2.cer -trustcacerts -alias dod_root"

14. Import the DoD intermediate email certificate:

    keytool -importkeysotre -keystore "C:\Program Files\SANscreen\java64\lib \security\cacerts" -file "C:\DoD_Root_CA_2__45__DOD_EMAIL_CA-24.cer -trustcacerts -alias dod_email_ca24

    Repeat this step for all intermediate email certificates.

15. Import the DoD intermediate certificate:

    keytool -importkeysotre -keystore "C:\Program Files\SANscreen\java64\lib \security\cacerts" -file "C:\DoD_Root_CA_2__47__DOD_CA-24.cer -trustcacerts -alias dod_ca24

    Repeat this step for all intermediate certificates.

16. Specify the domain in LDAP to match this example.

Configuring a client to support Smart Card and certificate login

Client machines require middleware and modifications to browsers to enable the use of Smart Cards and for certificate login. Customers who are already using Smart Cards should not require additional modifications to their client machines.

About this task

The following are the common client configuration requirements:

- Installing Smart Card middleware, such as ActivClient (see http://militarycac.com/activclient.htm)
- Modifying the IE browser (see http://militarycac.com/files/Making_AKO_work_with_Internet_Explorer_color.pdf)
- Modifying the Firefox browser (see https://militarycac.com/firefox2.htm)

Additional configuration steps include the following:

Steps

1. Disable unnecessary dialog pop-ups when launching the Java client:
   a. Open the “Java” control panel.
   b. Select the “Advanced” tab.
   c. Select “Advanced Security Settings”.
   d. Clear the “Use certificates and keys in browser keystore” option.
2. Import soft certificates into a user's personal certificates using the Chrome browser:
   a. Navigate to “Advanced Settings”.
   b. Select “Manage Certificates”
   c. Select the “Personal" certificates tab”.
   d. Select “Import...”.
   e. Choose the .p12 file.
   f. Complete the wizard, making sure to fill in the password field with the password that was used to create the certificate.

   Similar steps are required to configure the Internet Explorer and Firefox browsers.

Configuring Data Warehouse for Smart Card and certificate login

You must modify the OnCommand Insight Data Warehouse configuration to support Smart Card (CAC) and certificate logins.

Before you begin

- LDAP must be enabled on the system.
- The LDAP `User principal account name` attribute must match the LDAP field that contains a user's government ID number.
  The common name (CN) stored on government-issued CACs is normally in the following format: `first.last.ID`. For some LDAP fields, such as `sAMAccountName`, this format is too long. For these fields, OnCommand Insight extracts only the ID number from the CNs.

Steps

1. Use regedit to modify registry values in `HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\services\SANscreen Server.`
   a. Change the JVM_Option `-DclientAuth=false` to `-DclientAuth=want`.
   b. Change the JVM_Option `DdwhauthClass=org.apache.catalina.authenticator.FormAuthenticator` to `-DdwhAuthClass=com.onaro.sanscreen.authenticate.CustomSSLAuthenticator`
2. Add certificate authorities (CAs) to the Data Warehouse trust store:
   a. In a command window, go to `..\SANscreen\jboss\server\onaro\cert\`.
   b. Use the `keytool` utility to list the trusted CAs:

```
..\..\..\..\java64\bin\keytool.exe -list -keystore server.truststore -
storepass changeit
```

The first word in each line indicates the CA alias.

c. If necessary, supply a CA certificate file, usually a `.pem` file. To include customer's CAs with Data Warehouse trusted CAs go to `..\SANscreen\jboss\server\onaro\cert` and use the `keytool` import command:

```
..\..\..\..\java64\bin\keytool.exe -importcert -keystore
server.trustore -alias my_alias -file 'path/to/my.pem' -v -
trustcacerts
```

my_alias is usually an alias that would easily identify the CA in the `keytool -list` operation.

Configuring Data Warehouse for Smart Card and certificate login for Cognos

You must modify the OnCommand Insight Data Warehouse configuration to support Smart Card (CAC) and certificate logins for the Cognos server.

Before you begin

- LDAP must be enabled on the system.
- The LDAP `User principal account name` attribute must match the LDAP field that contains a user's government ID number.

Steps

1. Add certificate authorities (CAs) to the Cognos trust store.
   a. In a command window, go to `..\SANscreen\cognos\c10_64\configuration\certs\`
   b. Use the `keytool` utility to list the trusted CAs:

```
..\..\bin64\jre\7.0\bin\keytool.exe -list -keystore
CognosTruststore.jks -storepass changeit
```

The first word in each line indicates the CA alias.

c. If no suitable files exist, supply a CA certificate file, usually a `.pem` file. To include customer's CAs with OnCommand Insight trusted CAs, go to `..\SANscreen\cognos\c10_64\configuration\certs\`.

d. Optional: To include customer's CAs with OnCommand Insight trusted CAs, go to `..\SANscreen\cognos\c10_64\configuration\certs\`.

e. Use the `keytool` utility to import the `.pem` file

```
..\..\bin64\jre\7.0\bin\keytool.exe -importcert -keystore
CognosTruststore.jks -alias my_alias -file 'path/to/my.pem' -v -
trustcacerts
```

.

my_alias is usually an alias that would easily identify the CA in the `keytool -list` operation.

f. When prompted for a password, enter `changeit`, the default password.

g. Answer `yes` when prompted to trust the certificate.

2. Modify the reporting portal registry:

a. Use `regedit` to modify `HKEY_LOCAL_MACHINE\SOFTWARE\JavaSoft\Prefs\com\netapp\sanscreen\reporting`.

b. Change the portal_url value to `https:\\DWHServer.name:8080/ibmcognos`.

3. Redirect the reporting portal:

a. Open the `..\SANscreen\jboss\server\onaro\deploy\dwh-redirect.war\redirect.html` file in edit mode.

b. Change the URL value from`https:\\DWHServer.name:9300/p2pd/servlet/dispatch` to `https:\\DWHServer.name:8080/ibmcognos`.

4. Enable CAC mode:

a. Open the `..\SANscreen\cognos\c10_6\configuration\SANscreenAP.properties` file in edit mode.

b. Change `authentication.mode=form` to `authentication.mode=cac`.

c. Save the file.

d. Restart the server.

5. Start the ServletGateway:

a. Go to `..\SANscreen\cognos\c10_64\wlp\bin`.

b. Set the Java Home path:

`set java_home=..\SANscreen\cognos\c10_64\bin64\jre\7.0`

c. Start the ServletGateway: **`server start servletgateway`**

6. Changing the ServletGateway port from the default port (8080) to a custom port requires repeating the following steps:

a. Update the reporting portal in registry entry.

b. Update the reporting portal redirect.

c. Restart the ServletGateway port.

After you finish

- The ServletGateway must be stopped and started manually every time the server is rebooted or a command prompt session is closed because no ServletGateway service is available.
- ETL must be built to reflect the new reporting portal URL in the OnCommand Insight server.
- All of the manual steps might need to be performed after a Data Warehouse Cognos upgrade.

Setting up weekly backups for your Insight database

You might want to set up automatic weekly backups for your Insight database to protect your data. These automatic backups overwrite the files in the specified backup directory.

About this task

Best practice: When you are setting up the weekly backup of the OCI database, you need to store the backups on a different server than Insight is using, in case that server fails. Do not store any manual backups in the weekly backup directory because each weekly backup overwrites the files in the directory.

Steps

1. On the Insight toolbar, click the **Admin** icon.
2. Click **Setup**.
3. Click the **Backup/ASUP** tab.
4. Select **Enable weekly backup** as shown in this example.

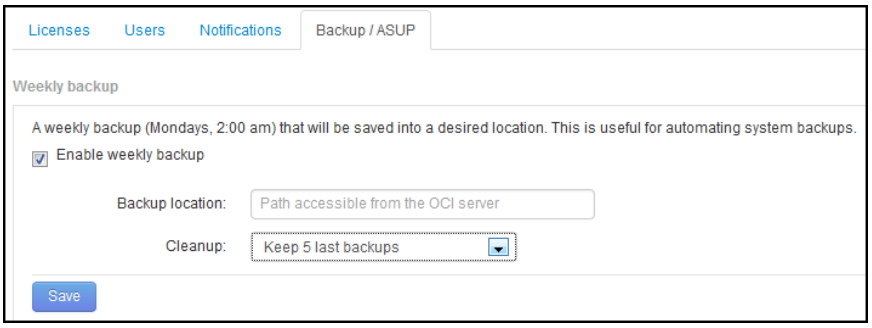

5. Enter the path to the **Backup location** on a different server that is accessible from the Insight server.
6. Select the **Cleanup** option to keep either the last two or the last five backups.
7. Click **Save**.

Related tasks

Restoring the Insight database on page 144

Configuring your email

You must configure OnCommand Insight to access your email system so that theOnCommand Insight Server can use your email to deliver reports, to which you subscribe, and transport support information for troubleshooting to NetApp technical support.

Email configuration prerequisites

Before you can configure OnCommand Insight to access your email system, you need to discover the host name or IP address to identify the (SMTP or Exchange) mail server and allocate an email account for OnCommand Insight reports.

Ask your email administrator to create an email account for OnCommand Insight. You will need the following information:

- The host name or IP address to identify the (SMTP or Exchange) mail server used by your organization. You can find this information through the application you use to read your email. In Microsoft Outlook, for example, you can find the name of the server by viewing your account configuration: Tools - E-mail accounts - View or change existing email account.
- Name of email account through which OnCommand Insight will send regular reports. The account must be a valid email address in your organization. (Most mail systems will not send messages unless they are sent from a valid user.) If the email server requires a user name and password in order to send mail, obtain this information from your system administrator.

Configuring your email for Insight

If your users want to receive Insight reports in their email accounts, you need to configure your email server to enable this feature.

Steps

1. Click **Admin**, and select **Notifications**.
2. Scroll down to the **Email** section of the page.
3. In the **Server** box, enter the name of your SMTP server in your organization, which is identified using either a hostname or an IP address (*nnn.nnn.nnn.nnn* format).

   If you specify a hostname, ensure that the name can be resolved through DNS.
4. In the **User name** box, enter your user name.
5. In the **Password** box, enter the password for accessing the email server, which is required only if your SMTP server is protected by a password.

   This is the same password that you use to log into the application that lets you read your email. If a password is required, you must enter it a second time for verification.
6. In the **Sender email** box, enter the sender email account that will be identified as the sender on all OnCommand Insight reports.

   This account must be a valid email account within your organization.
7. In the **Email signature** box, enter the text that you want to be inserted in every email that is sent.
8. In the Recipients box, click , enter an email address, and click **OK**.

   To edit an email address, select the address, and click . To delete an email address, select the address, and click .
9. To email a test message to specified recipients, click .
10. Click **Save**.

Configuring SNMP notifications

OnCommand Insight supports SNMP notifications for changes and violations. For example, SNMP notifications are sent when data source thresholds are exceeded.

Before you begin

The following must have been completed:

- Identifying the IP address of the server that consolidates traps for each type of event.
  You might have to consult with your system administrator to obtain this information.
- Identifying the port number through which the designated machine obtains SNMP traps, for each type of event.
  The default port for SNMP traps is 162.
- Compiling the MIB at your site.
  The proprietary MIB comes with the installation software to support OnCommand Insight traps. The NetApp MIB is compatible with all standard SNMP management software.

Steps

1. Click **Admin** and select **Notifications**.
2. Scroll down to the **SNMP** section of the page.
3. Click **Actions** and select **Add trap source**.
4. In the **Add SNMP trap recipients** dialog box, enter these values:

   IP

   The IP address to which OnCommand Insight sends SNMP trap messages.

   Port

   The port number to which OnCommand Insight sends SNMP trap messages.

   Community String

   Use “public” for SNMP trap messages.
5. Click **Save**.

Enabling the syslog facility

You can identify a location for the log of the OnCommand Insight violations and performance alerts and activate the logging process.

Before you begin

- You must have the IP address of the server to store the system log.
- You must know the facility level that corresponds to the type of program that is logging the message, such as LOCAL1 or USER.

About this task

The syslog includes the following types of information:

- Violation changes
- Performance alerts

The following units are used in the syslog:

- Utilization metrics: percentage
- Traffic metrics: MB
- Traffic rate: MB/s

Steps

1. Click **Admin**, and select **Notifications**.
2. Scroll down to the **Syslog** section of the page.
3. Select the **Syslog Enabled** check box.
4. In the **Server** field, enter the IP address of the log server.
5. In the **Facility** field, select the facility level that corresponds to the type of program that is logging the message.
6. Click **Save**.

Insight syslog contents

You can enable a syslog on a server to collect Insight violation and performance alert messages that include utilization and traffic data.

Message types

The Insight syslog lists three types of messages:

- SAN path violations
- General violations
- Performance alerts

Data provided

Violation descriptions include the elements involved, time of the event, and relative severity or priority of the violation.

Performance alerts include these data:

- Utilization percentages
- Traffic types
- Traffic rate measured in MB

Configuring performance and assure violation notifications

OnCommand Insight supports notifications for performance and assure violations. By default, Insight does not send notifications for these violations; you must configure Insight to send email, to send syslog messages to the syslog server, or to send SNMP notifications when a violation occurs.

Before you begin

You must have set up the email, syslog, and SNMP sending methods.

Steps

1. Click **Admin** and select **Notifications**.
2. Click the **Events** tab.
3. In the **Performance Violations events** or **Assure Violations events** section, click the list for the notification method (**Email**, **Syslog**, or **SNMP**0 you want, and select the severity level (**Warning and above** or **Critical**) for the violation.

   For example, if you want to receive an email notification when a violation is critical, click **Email** and select **Critical**.
4. Click **Save**.

Configuring your ASUP processing

All NetApp products are equipped with automated capabilities to provide the best possible support for the customers. This automated support (ASUP) periodically sends predefined and specific

information to the Customer Support. You can control what information is forwarded to NetApp, and how often it is sent.

Before you begin

You must configure OnCommand Insight to forward data before any data is sent.

About this task

The information is forwarded to NetApp by using FTP, HTTPS, or email depending on your configuration selections.

The information in the logs is forwarded to technical support by using HTTPS, FTP, or email, according to your choice. Although the AutoSupport tool supports FTP and email for the delivery of AutoSupport messages, should you use HTTPS because it provides the best security and supports the latest AutoSupport features.

Steps

1. Click the **Admin** icon on the toolbar.
2. Click **Setup**.
3. Click the **Backup/ASUP** tab.
4. In the ASUP section, select **Enable ASUP** to activate facility.
5. You might want to change your corporate information in these fields:
   - **Company name**
   - **Site name**
6. Select your preferences for these options:
   - **What to send**: Logs, Inventory data, Performance data
   - **How to send**: HTTP, HTTPS, FTP, email
7. At this point, if you want to test the ASUP configuration, click **Test Connection** and click **Save** if your connection test was successful.

   If you require more detailed information for your ASUP, enter any information that you require in the shaded fields on this page.
8. Enter the size and split file limits:
   - **Size limit**: for email ASUP data transmission, enter the maximum data size to be sent.
   - **Split Files if bigger than**: for email ASUP data transmission, indicate the cutoff point for file size. The default split file size is 1 MB.
9. If you have selected the HTTP or HTTPS protocol for ASUP data transmission to connect to the Internet, enable the proxy and define it.

   If you have selected FTP or email transmission, skip to **Test Connection**:
   - **Enable Proxy**: select for the HTTP or HTTPS ASUP data transmissions.
   - **Proxy Host**: identify the proxy server used to connect to the Internet.
   - **Proxy Port**: identify the proxy port for the ASUP data transmission.
   - **Proxy User**: specify the credentials to use for ASUP data transmission.

- **Proxy Password**: specify the credentials to use for ASUP data transmission.

10. Click **Test Connection** to be certain the connection you specified works.
11. Click **Save**.

Creating custom dashboards

In addition to using the default Assets Dashboard and Violations Dashboard, you can create custom dashboards that contain only the specific data that you want to see for your network. After you create a custom dashboard, you can view it by selecting it from the **Dashboards** menu.

Before you begin

You must use an account that has Administrator privileges.

About this task

It is possible to have the following scenarios when two users with administrator privileges (Administrator A and Administrator B) are creating and modifying dashboards simultaneously:

- If Administrator A creates a new dashboard, it immediately appears as an option on Administrator A's **Dashboard** menu; however, the new dashboard does not appear as an option on Administrator B's Dashboard menu until an action occurs that causes a refresh of the **Dashboard** menu.
- When Administrator A and Administrator B both try to save changes to the same dashboard at the same time, one receives a message saying that the changes were not made because of the conflict. If you receive this message, you should reload the page and make your changes again.

Steps

1. Open Insight in your browser.
2. From the **Dashboards** menu, select **Create Dashboard**.

   The New Dashboard page displays.

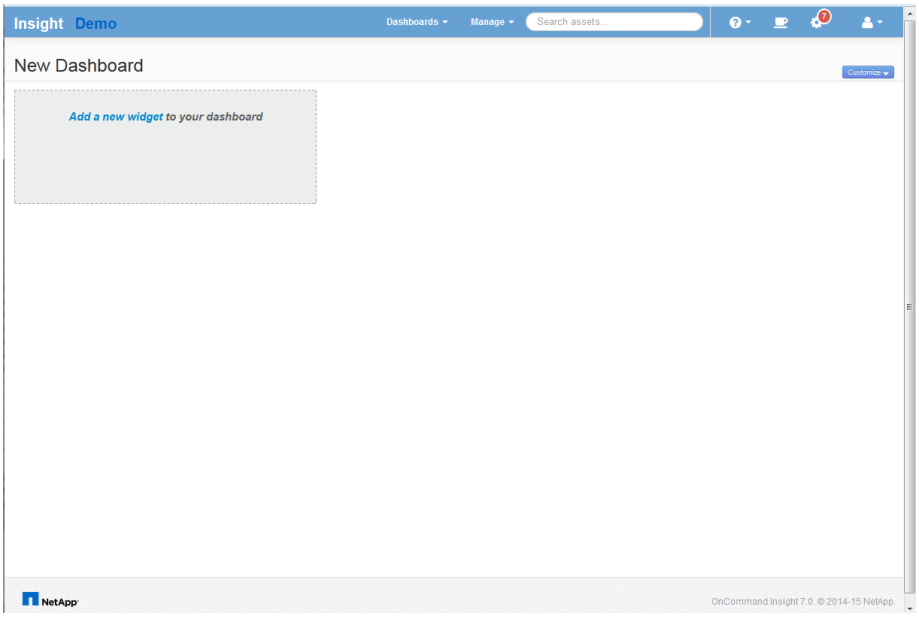

3. From the **Customize** menu, select **Rename Dashboard**, type a name, and then click **OK**.
4. Use one of these methods to begin to add widgets to the dashboard:
   - Click the **Add a new widget** link.
   - From the **Customize** menu, select **Add Widget**.

   Either method opens the Widget Library dialog box, which lists all the available widgets. By default, it displays all the widgets; you can click a category to display only the widgets of a certain type.
5. To add a widget to the dashboard, click **Add Widget** under any widget.
6. To add additional widgets to the dashboard, select **Add Widgets** from the **Customize** menu, click **Add Widget**, and repeat this action until you are satisfied with the contents of the dashboard.

   You can add the same widget multiple times. After you add three widgets, a new row is created in the dashboard. There is no limit to the number of widgets that the dashboard can contain.
7. To modify the settings of a widget, perform the following actions:
   a. Position your pointer in the over the widget title, and click .

      The Widget Settings dialog box lists the settings that you can change for the widget, as shown in this example for the Asset Performance Table:

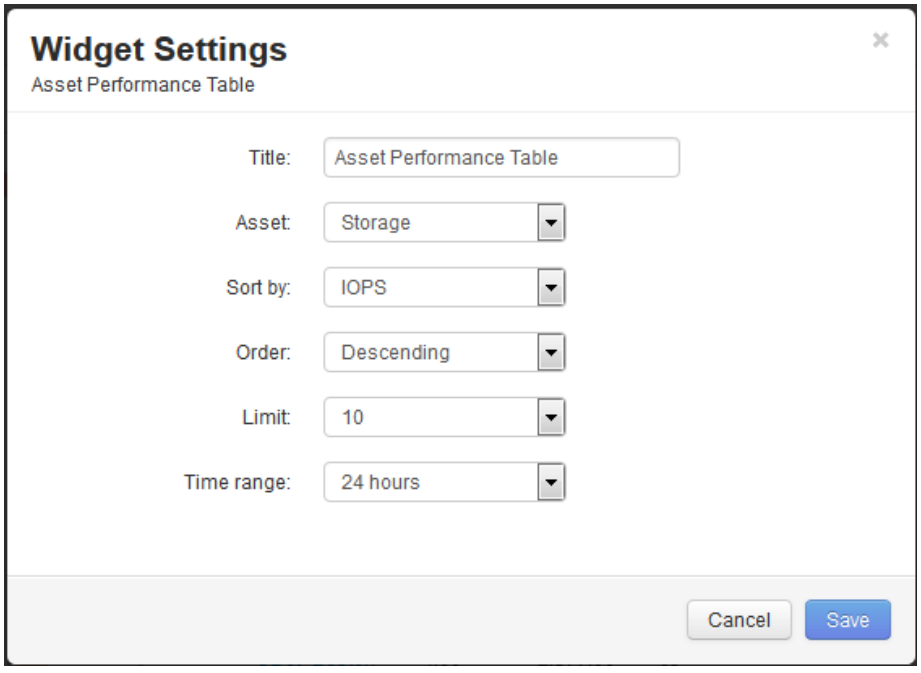

   b. Make the changes you want, and then click **Save**.
8. To delete a widget from the dashboard, perform the following actions:
   a. Click in the upper-right corner of the widget.
   b. In the confirmation message, click **OK**.
9. To change the order in which the widgets are placed on the dashboard, perform the following actions:

a. From the **Customize** menu, select **Arrange Widgets**.

b. In the **Arrange Widgets** dialog box, click the up arrow or down arrow next to a particular widget to change its order in the dashboard, and then click **Save**.

Deleting a custom dashboard

You can delete a custom dashboard when it is no longer needed.

Before you begin

You must use an account that has Administrator privileges.

Steps

1. Log in to OnCommand Insight.
2. From the **Dashboards** menu, select the custom dashboard that you want to delete.
3. From the **Customize** menu, select **Delete Dashboard**.
4. In the confirmation message, click **OK**.

   The dashboard is deleted, and the default dashboard, the Assets Dashboard, displays.

How OnCommand Insight anomaly detection works

OnCommand Insight contains machine-learning anomaly detection analytics to identify the normal operating workload range for an application and identify when changes in performance are outside of expected levels. The application anomaly detection engine ingests performance metrics collected by Insight and identifies anomalies in the application infrastructure.

Anomaly analysis provides proactive visibility into the infrastructure supporting business critical applications and informs you of performance anomalies before they become service disruptions.

Anomalies are items, events, or observations that do not conform to an expected pattern or other items in a dataset. You use anomaly detection to perform these tasks:

- Monitor selected critical applications
- Identify significant changes in the application infrastructure behavior
- Track an application through the entire stack and across multiple counters
- Diagnose the cause of periodically recurring events in an application workload
- Provide a more comprehensive set of information compared to data from performance counters using static threshold values
- Monitor in an improved reactive mode, with an end-to-end view of the topology and the ability to focus on the most anomalous behaviors.

The anomaly detection engine uses Insight data for application analysis. When monitoring is first started, up to 14 days of historical performance data can be ingested by the analysis engine. Data is collected for weeks or even months, providing more accurate data about a given resource. The data includes totals for the following counters:

Objects	Counter
VM	Latency, IOPS
Hypervisor	CPU utilization, IOPS

Objects	Counter
Edge port	BB credit zero
Storage Node	Latency, utilization, IOPS
Volume	Latency, IOPS
Internal volume	Latency, IOPS
Storage pool	IOPS, utilization

Installing the anomaly detection software

OnCommand Insight contains software that applies machine-learning anomaly detection techniques to your Insight data. You install this software separately from other OnCommand Insight products.

Before you begin

- Your environment must include a dedicated virtual machine (VM) with the following specifications:

Operating system	Red Hat Enterprise Linux 7
Memory	32 GB
CPU	8 cores
Disk space	200 GB

- You must be using OnCommand Insight 7.2 with a valid Perform license.
- You must know the IP address of the Insight server on which you want to perform anomaly detection.
- The Insight server on which you want to perform anomaly detection must reside on the same network, or at least in the same site or data center as that of the server running the anomaly detection engine.
  The anomaly detection software does not support configuration in a WAN.
- An alternate port number must be available on the Insight server if you do not accept the default port.
- TCP ports 8080 and 9200 must be open on the VM.
- You must have the user name and password for an account with Administrator privileges on the VM.
- A user name containing a single backslash (company\user) must be entered with two backslashes (company\\user).
- The user name cannot contain a “t” following a backslash (company\\tom).

Steps

1. Log in as an Administrator to the Linux server on which you want to run the anomaly detection software.
2. Copy the `.zip` file containing the anomaly detection software to the Linux server.
3. Extract the files to the `oci-prelert-7.2.0.0.494-linux-x86_64` directory.
4. Navigate to the directory where the installer is located:

```
cd oci-prelert-7.2.0.0.494-linux-x86_64
```

5. Install the anomaly detection software:

   ```
   sudo ./oci-prelert-install.sh
   ```

   During the installation, you are prompted to enter the server name or IP address of the OnCommand Insight server, and the user name and password for an account with Administrator privileges.

   You can remove the anomaly detection software using the following command:

   ```
   sudo /usr/bin/oci-prelert-uninstall.sh
   ```

Result

The anomaly detection software is automatically registered with the instance of OnCommand Insight that is specified in the installation. The software can communicate only with the OnCommand Insight instance that it is registered with, and only one instance of the software can be registered with an OnCommand Insight server.

Restarting the server that is running the anomaly detection software or the Insight server automatically restarts the anomaly detection process.

Using the Anomalies section

The Anomalies section, which displays only on an application asset page, enables you to turn on anomaly detection. Anomaly detection provides insight into the application infrastructure and enables you to identify abnormal behaviors that could result in service disruptions.

Before you begin

- You must have the Perform license configured to install the anomaly detection engine.
- You must have installed and configured the anomaly detection engine.
  For more information about installing and configuring the anomaly detection engine, see the *OnCommand Insight Configuration and Administration Guide*.
- You must have assigned an application to the compute resource (virtual machine or host) for any assets for which you want anomaly detection information.
  Insight discovers the rest of the resources automatically.
  For more information about assigning applications to assets, see the *OnCommand Insight Configuration and Administration Guide*.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage** and select **Applications**.
3. Locate the application assigned to the computer resource.
4. In the **Anomalies** section, ensure that the **Monitor** control box in the upper right of the section shows **ON** rather than **OFF**; if not, click anywhere in the **Monitor** control box to switch the control to **ON**.

   You are able to concurrently monitor up to 48 applications for anomalies.

   The status in the section displays “Collecting data” for at least an hour; the anomaly detection engine preloads up to 14 days of any existing performance data so the engine can learn the behavior of your environment. After the Insight server performs the first poll of your environment's assets, the anomaly detection engine calculates and returns a score for the application within that time.

5. While the anomaly detection engine is running, you can view the following information in the **Anomalies** section of the application asset page:

   - Anomaly detection score
     The score (for example, 50) that displays is the highest anomaly score in the bucket of time (for example, 1:30 am - 1:59 am) in the time period (for example, either 3 hours, 24 hours, 3 days, 7 days, or a custom time period) that is selected on the application's asset page.
     A score of 0-50 indicates mild anomalous behavior raising awareness to new factors affecting the infrastructure performance, a score of 51-75 indicates some anomalous behavior worth further investigation, and a score of 76-100 could indicate a significant change in the application infrastructure behavior.

   - Application assets view
     Located below the anomaly detection score, this shows how all the assets contribute to the current anomaly score. If an asset is not contributing significantly to the anomaly score, it is represented by . If the asset is contributing to the significance of the anomaly score, either 1, 2, or three blue bars appear, for example, .
     Both assets that are assigned to an application (hosts, virtual machines, volumes, internal volumes, and hypervisors) and also assets in the infrastructure that are somehow connected to the assigned assets (for example, nodes, ports, and so on). Both assets that contribute and those that do not contribute to the specific anomaly are displayed, because the data from the assigned application resources may flow through them.
     If there are more than 60 assets of one type, the **More** button displays, which you can click to display the remaining assets. That status of those remaining assets, for example,

     , appears next to the **More** button. This example indicates that of a total of 129 assets not shown, 124 of those assets did not contribute to the anomaly score; however, five of those assets did contribute some level of significance to the anomaly score.

   - Anomaly graph
     The anomaly graph displays the highest anomaly score, indicated by the orange circle, for the time range (3 hours, 24 hours, 3 days, 7 days, or a custom time range) selected for the application asset page. You can drag your pointer over the chart to see the anomaly scores, indicated by a blue circle, for other points during the time range and then click the blue circle, which will update the application assets view to display the assets that contributed to the score at that point in time.

6. To aid you in troubleshooting asset performance, you can add an asset from the application assets view to the expert view chart in the application asset page by doing the following:

   a. Click an asset icon in the applications assets view.

      The following displays: The name of the asset, which is a hyperlink to its asset page, the metric that contributed to the anomaly, the significance of the asset's contribution, and a button that enables you to add the asset to the expert view chart.

   b. Click .

   The asset is added to the expert view chart.

Defining applications

If you want to track data associated with specific applications running in your environment, you need to define those applications.

Before you begin

If you want to associate the application with a business entity, you must have already created the business entity.

About this task

You can associate applications with the following assets: host, virtual machines, volumes, internal volumes, and hypervisors. If you want to associate applications with qtrees or shares, you must use the Java UI. For more information, refer to the *OnCommand Insight Inventory User Guide for the Java UI*.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage** and select **Applications**.

   After you define an application, the Applications page displays the application's name, its priority, and, if applicable, the business entity associated with the application.

   Note: If the anomaly detection engine is installed, the Applications Dashboard page also shows an Anomaly Score for the application. For more information about installing and configuring the anomaly detection engine, see the *OnCommand Insight Configuration and Administration Guide*.

3. Click **Add**.

   The Add Application dialog box displays.

4. Enter a unique name for the application in the **Name** box.
5. Click **Priority** and select the priority (critical, high, medium, or low) for the application in your environment.
6. If you plan to use this application with a business entity, click **Business Entity** and select the entity from the list.
7. **Optional**: If you do not use volume sharing, click to clear the **Validate volume sharing** box.

   This requires the Assure license. Set this when you want to ensure each host has access to the same volumes in a cluster. For example, hosts in high-availability clusters often need to be masked to the same volumes to allow for failover; however, hosts in unrelated applications usually have no need to access the same physical volumes. Additionally, regulatory policies might require you to explicitly disallow unrelated applications from accessing the same physical volumes for security reasons.

8. Click **Save**.

   The application appears in the Applications page. If you click the application's name, Insight displays the asset page for the application.

After you finish

After defining an application, you can go to an asset page for host, virtual machine, volume, internal volume, or hypervisor to assign an application to an asset.

Related tasks

Assigning applications to assets on page 43
Assigning multiple applications to or removing multiple applications from assets on page 59

Monitoring OnCommand Insight system health

You can check the current status of your Insight system components by viewing the health page, which shows a component status and alerts you when there is an issue.

Steps

1. Log in to the Insight web UI.
2. Click **Admin** and select **Health**.

   The Health page displays.
3. View the summary of the current status of the components, paying particular attention to any attention status in the **Details** column that is preceded by a red circle, which indicates an issue that requires your immediate attention.

   The Health page displays information about the following Insight components:

   - Insight servers
   - Acquisition units
   - Data Warehouse servers (ETL)
   - Anomaly detection engine
     If the anomaly detection engine is installed and configured for monitored objects, see the `prelert.log` file in the following location: *`disk drive`*`:\`*`install directory`*`\SANscreen\jboss\server\onaro\log`.

Your business entities hierarchy

You can define business entities to track and report on your environment data at a more granular level.

In OnCommand Insight, the business entities hierarchy contains these levels:

- **Tenant** is primarily used by service providers to associate resources with a customer, for example, NetApp.
- **Line of Business (LOB)** is a line of business or product line within a company, for example, Data Storage.
- **Business Unit** represents a traditional business unit such as Legal or Marketing.
- **Project** is often used to identify a specific project within a business unit for which you want capacity chargeback. For example, "Patents" might be a project name for the Legal business unit and "Sales Events" might be a project name for the Marketing business unit. Note that level names may include spaces.

You are not required to use all of the levels in the design of your corporate hierarchy.

Designing your business entities hierarchy

You need to understand the elements of your corporate structure and what needs to be represented in the business entities because they become a fixed structure in your OnCommand Insight database. You can use the following information to set up your business entities. Remember you do not need to use all of the hierarchy levels to gather data in these categories.

Steps

1. Examine each level of the business entities hierarchy to determine if that level should be included in your business entity hierarchy for your company:
   - **Tenant** level is needed if your company is an ISP and you want to track customer usage of resources.
   - **Line of Business (LOB)** is needed in the hierarchy if the data for different product lines needs to be tracked.
   - **Business Unit** is required if you need to track data for different departments. This level of the hierarchy is often valuable in separating a resource that one department uses that other departments do not.
   - **Project** level can be used for specialized work within a department. This data might be useful to pinpoint, define, and monitor a separate project's technology needs compared to other projects in a company or department.
2. Create a chart showing each business entity with the names of all of the levels within the entity.
3. Check the names in the hierarchy to be certain they will be self-explanatory in OnCommand Insight views and reports.
4. Identify all applications that are associated with each business entity.

Defining applications

If you want to track data associated with specific applications running in your environment, you need to define those applications.

Before you begin

If you want to associate the application with a business entity, you must have already created the business entity.

About this task

You can associate applications with the following assets: host, virtual machines, volumes, internal volumes, and hypervisors. If you want to associate applications with qtrees or shares, you must use the Java UI. For more information, refer to the *OnCommand Insight Inventory User Guide for the Java UI.*

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage** and select **Applications**.

   After you define an application, the Applications page displays the application's name, its priority, and, if applicable, the business entity associated with the application.

Note: If the anomaly detection engine is installed, the Applications Dashboard page also shows an Anomaly Score for the application. For more information about installing and configuring the anomaly detection engine, see the *OnCommand Insight Configuration and Administration Guide*.

3. Click **Add**.

   The Add Application dialog box displays.

4. Enter a unique name for the application in the **Name** box.

5. Click **Priority** and select the priority (critical, high, medium, or low) for the application in your environment.

6. If you plan to use this application with a business entity, click **Business Entity** and select the entity from the list.

7. **Optional**: If you do not use volume sharing, click to clear the **Validate volume sharing** box.

   This requires the Assure license. Set this when you want to ensure each host has access to the same volumes in a cluster. For example, hosts in high-availability clusters often need to be masked to the same volumes to allow for failover; however, hosts in unrelated applications usually have no need to access the same physical volumes. Additionally, regulatory policies might require you to explicitly disallow unrelated applications from accessing the same physical volumes for security reasons.

8. Click **Save**.

   The application appears in the Applications page. If you click the application's name, Insight displays the asset page for the application.

After you finish

After defining an application, you can go to an asset page for host, virtual machine, volume, internal volume, or hypervisor to assign an application to an asset.

Related tasks

Assigning applications to assets on page 43
Assigning multiple applications to or removing multiple applications from assets on page 59

Creating business entities

After designing the business entities hierarchy for your company, you can set up applications and then associate the business entities with the applications. This process creates the business entities structure in your OnCommand Insight database.

About this task

Associating applications with business entities is optional; however, it is a best practice.

Steps

1. Log in to the Insight web UI.

2. Click **Manage** and select **Business entities**.

   The Business Entities page displays.

3. Click

to begin building a new entity.

   The **Add Business Entity** dialog box displays.

4. For each entity level (Tenant, Line of Business, Business Unit, and Project), you can do any of the following:
   - Click the entity level list and select a value.
   - Type a new value and press Enter.
   - Leave the entity level value as N/A if you do not want to use the entity level for the business entity.
5. Click **Save**.

Defining applications

If you want to track data associated with specific applications running in your environment, you need to define those applications.

Before you begin

If you want to associate the application with a business entity, you must have already created the business entity.

About this task

You can associate applications with the following assets: host, virtual machines, volumes, internal volumes, and hypervisors. If you want to associate applications with qtrees or shares, you must use the Java UI. For more information, refer to the *OnCommand Insight Inventory User Guide for the Java UI*.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage** and select **Applications**.

   After you define an application, the Applications page displays the application's name, its priority, and, if applicable, the business entity associated with the application.

   Note: If the anomaly detection engine is installed, the Applications Dashboard page also shows an Anomaly Score for the application. For more information about installing and configuring the anomaly detection engine, see the *OnCommand Insight Configuration and Administration Guide*.
3. Click **Add**.

   The Add Application dialog box displays.
4. Enter a unique name for the application in the **Name** box.
5. Click **Priority** and select the priority (critical, high, medium, or low) for the application in your environment.
6. If you plan to use this application with a business entity, click **Business Entity** and select the entity from the list.
7. **Optional**: If you do not use volume sharing, click to clear the **Validate volume sharing** box.

   This requires the Assure license. Set this when you want to ensure each host has access to the same volumes in a cluster. For example, hosts in high-availability clusters often need to be masked to the same volumes to allow for failover; however, hosts in unrelated applications usually have no need to access the same physical volumes. Additionally, regulatory policies might require you to explicitly disallow unrelated applications from accessing the same physical volumes for security reasons.

8. Click **Save**.

   The application appears in the Applications page. If you click the application's name, Insight displays the asset page for the application.

After you finish

After defining an application, you can go to an asset page for host, virtual machine, volume, internal volume, or hypervisor to assign an application to an asset.

Related tasks

Assigning applications to assets on page 43
Assigning multiple applications to or removing multiple applications from assets on page 59

Assigning applications to assets

After defining applications with or without business entities, you can associate the applications with assets.

About this task

If you want to assign an application to qtrees or shares, you must use the Java UI. For more information, see the *OnCommand Insight Inventory User Guide for the Java UI*.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate the asset (host, virtual machine, volume, or internal volume) to which you want to apply the application by doing either of the following:
   - Click **Dashboard**, select **Assets Dashboard**, and click the asset.
   - Click on the toolbar to display the **Search assets** box, type the name of the asset, and then select the asset from the list.
3. In the **User Data** section of the asset page, position your cursor over the name of the application currently assigned to the asset (if there is no application assigned, **None** displays instead) and then click (Edit application).

   The list of available applications for the selected asset display. The applications that are currently associated with the asset are preceded by a check mark.
4. You can type in the Search box to filter the application names, or you can scroll down the list.
5. Select the applications you want to associate with the asset.

   You can assign multiple applications to host, virtual machine, and internal volume; however, you can only assign one application to volume.
6. Click to assign the selected application or applications to the asset.

   The application names appear in the User Data section; if the application is associated with a business entity, the name of the business entity appears in this section also.

Related tasks

Assigning multiple applications to or removing multiple applications from assets on page 59

Editing applications

You might want to change an application's priority, the business entity associated with an application, or the status of volume sharing.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage** and select **Applications**.
3. Position your cursor over the application you want to edit and click .

   The Edit Application dialog box displays.
4. Do any of the following:
   - Click **Priority** and select a different priority.

     Note: You cannot change the application's name.
   - Click **Business Entity** and select a different business entity to associate the application with or select **None** to remove the association of the application with the business entity.
   - Click to clear or select **Validate volume sharing**.

     Note: This option is only available if you have the Assure license.
5. Click **Save**.

Deleting applications

You might want to delete an application when it no longer fulfills a need in your environment.

Steps

1. Log in to the Insight web UI.
2. Click **Manage** and select **Applications**.
3. Position your cursor over the application you want to delete and click .

   A confirmation dialog box is displayed, asking if you want to delete the application.
4. Click **OK**.

Assigning business entities to assets

You can assign a business entity to an asset (host, virtual machine, volume, or internal volume) without having associated the business entity to an application; however, business entities are assigned automatically to an asset if that asset is associated with an application related to a business entity. You must use the Java UI to assign business entities to qtrees or shares.

Before you begin

You must have already created a business entity.

About this task

While you can assign business entities directly to assets, it is recommended that you assign applications to assets and then assign business entities to assets.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate the asset (host, virtual machine, volume, or internal volume) to which you want to apply the business entity by doing either of the following:
   - Click on the asset in the Assets Dashboard.
   - Click on the toolbar to display the **Search assets** box, type the name of the asset, and then select the asset from the list.
3. In the **User Data** section of the asset page, position your cursor over **None** next to **Business Entities** and then click .

   The list of available business entities display.
4. Type in the **Search** box to filter the list for a specific entity or scroll down the list; select a business entity from the list.

   If the business entity you choose is associated with an application, the application name is displayed. In this case, the word “derived” appears next to the business entity name. If you want to maintain the entity for only the asset and not the associated application, you can manually override the assignment of the application.
5. To override an application derived from a business entity, place your cursor over the application name and click , select another business entity, and select another application from the list.

Defining annotations

When customizing OnCommand Insight to track data for your corporate requirements, you can define any specialized annotations needed to provide a complete picture of your data: for example, asset end of life, data center, building location, storage tier, or volume, and internal volume service level.

Steps

1. List any industry terminology to which environment data must be associated.
2. List corporate terminology to which environment data must be associated, which is not already being tracked using the business entities.
3. Identify any default annotation types that you might be able to use.
4. Identify which custom annotations you need to create.

Using annotations to monitor your environment

When customizing OnCommand Insight to track data for your corporate requirements, you can define specialized notes, called *annotations*, and assign them to your assets. For example, you can annotate assets with information such as asset end of life, data center, building location, storage tier, or volume service level.

Using annotations to help monitor your environment includes the following high-level tasks:

- Creating or edit definitions for all annotation types.
- Displaying asset pages and associating each asset with one or more annotations.

For example, if an asset is being leased and the lease expires within two months, you might want to apply an end-of-life annotation to the asset. This helps prevent others from using that asset for an extended time.

- Creating rules to automatically apply annotations to multiple assets of the same type.
- Using the annotation import utility to import annotations.
- Filter assets by their annotations.
  You can sort assets by annotations; however, you can only do this in the Java UI.
  See the *OnCommand Insight Inventory User Guide for the Java UI* for more information about grouping assets by annotation.
- Grouping data in reports based on annotations and generate those reports.
  See the *OnCommand Insight Reporting Guide* for more information about reports.

Managing annotation types

OnCommand Insight provides some default annotation types, such as asset life cycle (birthday or end of life), building or data center location, and tier, that you can customize to show in your reports. You can define values for default annotation types or create your own custom annotation types. You can later edit those values.

Default annotation types

OnCommandInsight provides some default annotation types. These annotations can be used to filter or group data and to filter data reporting.

You can associate assets with default annotation types such as the following:

- Asset life cycle, such as birthday, sunset, or end of life
- Location information about a device, such as data center, building, or floor
- Classification of assets, such as by quality (tiers), by connected devices (switch level), or by service level
- Status, such as hot (high utilization)

The following table lists the default annotation types. You can edit any of these annotation names to suit your needs.

Annotation types	Description	Type
Alias	User-friendly name for a resource.	Text
Birthday	Date when the device was or will be brought online.	Date
Building	Physical location of host, storage, switch, and tape resources.	List
City	Municipality location of host, storage, switch, and tape resources.	List
Compute Resource Group	Group assignment used by the Host and VM Filesystems data source.	List
Continent	Geographic location of host, storage, switch, and tape resources.	List
Country	National location of host, storage, switch, and tape resources.	List

Annotation types	**Description**	**Type**
Data Center	Physical location of the resource and is available for hosts, storage arrays, switches, and tapes.	List
Direct Attached	Indicates (Yes or No) if a storage resource is connected directly to hosts.	Boolean
End of Life	Date when a device will be taken offline, for example, if the lease expired or the hardware is being retired.	Date
Fabric Alias	User-friendly name for a fabric.	Text
Floor	Location of a device on a floor of a building. Can be set for hosts, storage arrays, switches, and tapes.	List
Hot	Devices already in heavy use on a regular basis or at the threshold of capacity.	Boolean
Note	Comments that you want associated with a resource.	Text
Rack	Rack in which the resource resides.	Text
Room	Room within a building or other location of host, storage, switch, and tape resources.	List
SAN	Logical partition of the network. Available on hosts, storage arrays, tapes, switches, and applications.	List
Service Level	A set of supported service levels that you can assign to resources. Provides an ordered options list for internal volumes, qtree, and volumes. Edit service levels to set performance policies for different levels.	List
State/Province	State or province in which the resource is located.	List
Sunset	Threshold set after which no new allocations can be made to that device. Useful for planned migrations and other pending network changes.	Date
Switch Level	Includes predefined options for setting up categories for switches. Typically, these designations remain for the life of the device, although you can edit them, if needed. Available only for switches.	List

Annotation types	Description	Type
Tier	Can be used to define different levels of service within your environment. Tiers can define the type of level, such as speed needed (for example, gold or silver). This feature is available only on internal volumes, qtrees, storage arrays, storage pools, and volumes.	List
Violation Severity	Rank (for example, major) of a violation (for example, missing host ports or missing redundancy), in a hierarchy of highest to lowest importance.	List

Note: Alias, Data Center, Hot, Service Level, Sunset, Switch Level, Service Level, Tier, and Violation Severity are system-level annotations, which you cannot delete or rename; you can change only their assigned values.

How annotations are assigned

You can assign annotations manually or automatically using annotation rules. OnCommand Insight also automatically assigns some annotations on acquisition of assets and by inheritance. Any annotations that you assign to an asset appear in the User Data section of the asset page.

Annotations are assigned in the following ways:

- You can assign an annotation manually to an asset.
  If an annotation is assigned directly to an asset, the annotation appears as normal text on an asset page. Annotations that are assigned manually always take precedence over annotations that are inherited or assigned by annotation rules.
- You can create an annotation rule to automatically assign annotations to assets of the same type.
  If the annotation is assigned by rule, Insight displays the rule name next to the annotation name on an asset page.
- Insight automatically associates a tier level with a storage tier model to expedite the assignment of storage annotations to your resources on acquisition of assets.
  Certain storage resources are automatically associated with a predefined tier (Tier 1 and Tier 2). For example, the Symmetrix storage tier is based on the Symmetrix and VMAX family and is associated with Tier 1. You can change the default values to match your tier requirements. If the annotation is assigned by Insight (for example, Tier), you see “System-defined” when you position your cursor over the annotation's name on an asset page.
- A few resources (children of an asset) can derive the predefined Tier annotation from their asset (parent).
  For example, if you assign an annotation to a storage, the Tier annotation is derived by all the storage pools, internal volumes, volumes, qtrees, and shares belonging to the storage. If a different annotation is applied to an internal volume of the storage, the annotation is subsequently derived by all the volumes, qtrees, and shares. “Derived” appears next to the annotation name on an asset page.

Associating costs with annotations

Prior to running cost-related reports, you should associate costs with the Service Level, Switch Level, and Tier system-level annotations, which enables chargeback to the storage users based on their

actual usage of production and replicated capacity. For example, for the Tier level, you might have gold and silver tier values and assign a higher cost to the gold tier than to the silver tier.

Steps

1. Log in to the Insight web UI.
2. Click Manage and select **Annotations**.

   The Annotation page displays.
3. Position your cursor over the Service Level, Switch Level, or Tier annotation, and click .

   The Edit Annotation dialog box displays.
4. Enter the values for any existing levels in the **Cost** field.

   The Tier and Service Level annotations have Auto Tier and Atmos values, respectively, which you cannot remove.
5. Click

to add additional levels.
6. Click **Save** when you finish.

Creating custom annotations

Using annotations, you can add custom business-specific data that matches your business needs to assets. While OnCommand Insight provides a set of default annotations, you might find that you want to view data in other ways. The data in custom annotations supplements device data already collected, such as switch manufacturer, number of ports, and performance statistics. The data you add using annotations is not discovered by Insight.

Steps

1. Log in to the Insight web UI.
2. Click **Manage** and select **Annotations**.

   The Annotations page displays the list of annotations.
3. Click

.

   The **Add Annotation** dialog box displays.
4. Enter a name and a description in the **Name** and **Description** fields.

   You can enter up to 255 characters in these fields.
5. Click **Type** and then select one of the following options that represents the type of data allowed in this annotation:

   - Boolean
     This creates a drop-down list with the choices of yes and no. For example, the “Direct Attached” annotation is Boolean.
   - Date
     This creates a field that holds a date. For example, if the annotation will be a date, select this.
   - List
     This can create either of the following:
     - A drop-down fixed list
       When others are assigning this annotation type on a device, they cannot add more values to the list.

    - A drop-down flexible list
      If you select the **Add new values on the fly** option when you create this list, when others are assigning this annotation type on a device, they can add more values to the list.

- Number
  This creates a field where the user assigning the annotation can enter a number. For example, if the annotation type is “Floor”, the user could select the Value Type of “number” and enter the floor number.
- Text
  This creates a field that allows free-form text. For example, you might enter “Language” as the annotation type, select “Text” as the value type, and enter a language as a value.

Note: After you set the type and save your changes, you cannot change the type of the annotation. If you need to change the type, you have to delete the annotation and create a new one.

6. If you select **List** as the annotation type, do the following:
   a. Select **Add new values on the fly** if you want the ability to add more values to the annotation when on an asset page, which creates a flexible list.

      For example, suppose you are on an asset page and the asset has the City annotation with the values Detroit, Tampa, and Boston. If you selected the **Add new values on the fly** option, you can add additional values to City like San Francisco and Chicago directly on the asset page instead of having to go to the Annotations page to add them. If you do not choose this option, you cannot add new annotation values when applying the annotation; this creates a fixed list.

   b. Enter a value and a name in **Value** and **Description** fields.

   a. Click

 to add additional values.

   b. Click to remove a value.
7. Click **Save**.
   Your annotations appear in the list on the Annotations page.

After you finish

In the web UI, the annotation is available immediately for use. To use the annotation in the Java client, you must restart any running Java client before you can apply the annotation. If you are adding or changing annotation values, or deleting an annotation type, you do not need to restart the Java client.

Manually assigning annotations to assets

Assigning annotations to assets helps you sort, group, and report on assets in ways that are relevant to your business. Although you can assign annotations to assets of a particular type automatically, using annotation rules, you can assign annotations to an individual asset by using its asset page.

Before you begin

You must have created the annotation you want to assign.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate the asset to which you want to apply the annotation by doing either of the following:

- Click the asset in the Assets Dashboard.
- Click on the toolbar to display the **Search assets** box, type the type of or name of the asset, and then select the asset from the list that displays.
  The asset page displays.

3. In the **User Data** section of the asset page, click .

   The Add Annotation dialog box displays.

4. Click **Annotation** and select an annotation from the list.
5. Click **Value** and do either of the following, depending on type of annotation you selected:
   - If the annotation type is list, date, or Boolean, select a value from the list.
   - If the annotation type is text, type a value.
6. Click **Save**.

   The annotation is attached to the asset. You can later filter assets by annotation using a query. If you want to group assets by annotations, you must use the Java UI. Refer to the *OnCommand Insight Inventory User Guide for the Java UI* for more information.

7. If you want to change the value of the annotation after you assign it, click and select a different value.

   If the annotation is of list type for which the **Add values dynamically upon annotation assignment** option is selected, you can type to add a new value in addition to selecting an existing value.

Related tasks

Editing or removing multiple annotations from assets on page 60
Assigning annotations to assets using annotation rules on page 52
Creating a query on page 57

Modifying annotations

You might want to change the name, description, or values for an annotation, or delete an annotation that you no longer want to use.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage** and select **Annotations**.

   The Annotations page displays.

3. Position your cursor over the annotation you want to edit and click .

   The **Edit Annotation** dialog box displays.

4. You can make the following modifications to an annotation:

   a. Change the name, description, or both.

      However, note that you can enter a maximum of 255 characters for both the name and description, and you cannot change the type of any annotation. Additionally, for system-level annotations, you cannot change the name or description; however, you can add or remove values if the annotation is a list type.

Note: If a custom annotation is published to the Data Warehouse and you rename it, you will lose historical data.

b. To add another value to an annotation of list type, click

.

c. To remove a value from an annotation of list type, click .

You cannot delete an annotation value if that value is associated with an annotation contained in an annotation rule, query, or performance policy.

5. Click **Save** when you finish.

After you finish

If you are going to use annotations in the Data Warehouse, you need to force an update of annotations in the Data Warehouse. Refer to the *OnCommand Insight Data Warehouse Administration Guide*.

Deleting annotations

You might want to delete an annotation that you no longer want to use. You cannot delete a system-level annotation or an annotation that is used in an annotation rule, query, or performance policy.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage** and select **Annotations**.

   The Annotations page displays.
3. Position your cursor over the annotation you want to delete, and click .

   A confirmation dialog box displays.
4. Click **OK**.

Assigning annotations to assets using annotation rules

To automatically assign annotations to assets based on criteria that you define, you configure annotation rules. OnCommand Insight assigns the annotations to assets based on these rules. Insight also provides two default annotation rules, which you can modify to suit your needs or remove if you do not want to use them.

Default storage annotation rules

To expedite the assignment of storage annotations to your resources, OnCommand Insight includes 21 default annotation rules, which associate a tier level with a storage tier model. All of your storage resources are automatically associated with a tier upon acquisition of the assets in your environment.

The default annotation rules apply a tier annotations in the following way:

- Tier 1, storage quality tier
  The Tier 1 annotation is applied to the following vendors and their specified families: EMC (Symmetrix), HDS (HDS9500V, HDS9900, HDS9900V, R600, R700, USP r, USP V), IBM (DS8000), NetApp (FAS6000 or FAS6200), and Violin (Memory).
- Tier 2, storage quality tier
  The Tier 2 annotation is applied to the following vendors and their specified families: HP (3PAR StoreServ or EVA), EMC (CLARiiON), HDS (AMS or D800), IBM (XIV), and NetApp (FAS3000, FAS3100, and FAS3200).

You can edit the default settings of these rules to match your tier requirements, or you can remove them if you do not need them.

Creating annotation rules

As an alternative to manually applying annotations to individual assets, you can automatically apply annotations to multiple assets using annotation rules. Annotations set manually on an individual asset pages take precedence over rule-based annotations when Insight evaluates the annotation rules.

Before you begin

You must have created a query for the annotation rule.

About this task

Although you can edit the annotation types while you are creating the rules, you should have defined the types ahead of time.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage** and select **Annotation rules**.

   The Annotation Rules page displays the list of existing annotation rules.
3. Click

.

   The Add Rule dialog box displays.
4. Do the following:
   a. In the **Name** box, enter a unique name that describes the rule.

      This name will appear in the Annotation Rules page.
   b. Click **Query** and select the query that OnCommand Insight should use to apply the annotation to assets.
   c. Click **Annotation** and select the annotation you want to apply.
   d. Click **Value** and select a value for the annotation.

      For example, if you choose Birthday as the annotation, you specify a date for the value.
5. Click **Save**.
6. Click **Run all rules** if you want to run all the rules immediately; otherwise, the rules are run at a regularly scheduled interval.

Related tasks

Creating a query on page 57

Setting annotation rule precedence

By default, OnCommand Insight evaluates annotation rules sequentially; however, you can configure the order in which OnCommand Insight evaluates annotation rules if you want Insight to evaluate rules in a specific order.

Steps

1. Log in to the Insight web UI.

2. Click **Manage** and select **Annotation rules**.

   The Annotation Rules page displays the list of existing annotation rules.

3. Position your cursor over an annotation rule.

   The precedence arrows appear to the right of the rule.

4. To move a rule up in the list, click ; to move a rule down in the list, click .

   By default, new rules are added sequentially to the list of rules. Annotations set manually on an individual asset pages take precedence over rule-based annotations when Insight evaluates the annotation rules.

Modifying annotation rules

You can modify an annotation rule to change the rule's name, its annotation, the annotation's value, or the query associated with the rule.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage** and select **Annotation rules**.

   The Annotation Rules page displays the list of existing annotation rules.

3. Locate the rule that you want to modify:
   - On the Annotation Rules page, you can filter the annotation rules by entering a value in the filter box.
   - Click a page number to browse through the annotation rules by page if there are more rules than fit on a page.
4. Perform one of the following to display the **Edit Rule** dialog box:
   - If you are on the Annotation Rules page, position your cursor over the annotation rule and click .
   - If you are on an asset page, position your cursor over the annotation associated with the rule, position your cursor over the rule name when it displays, and then click the rule name.
5. Make the required changes and click **Save**.

Deleting annotation rules

You can delete an annotation rule when the rule is no longer required to monitor the objects in your network.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage**, and select **Annotation rules**.

   The Annotation Rules page displays the list of existing annotation rules.

3. Locate the rule that you want to delete:
   - On the Annotation Rules page, you can filter the annotation rules by entering a value in the filter box.

- Click a page number to browse through the annotation rules by page if there are more rules than fit on a single page.

4. Point the cursor over the rule that you want to delete, and then click .

   A confirmation message is displayed, prompting whether you want to delete the rule.

5. Click **OK**.

Importing annotation values

If you maintain annotations on SAN objects (such as storage, hosts, and virtual machines) in a CSV file, you can import that information into OnCommand Insight. You can import applications, business entities, or annotations such as tier and building.

About this task

The following rules apply:

- If an annotation value is empty, that annotation is removed from the object.
- When annotating volumes or internal volumes, the object name is a combination of storage name and volume name using the dash and arrow (->) separator:

  ```
  <storage_name>-><volume_name>
  ```

- When storage, switches, or ports are annotated, the Application column is ignored.
- The columns of Tenant, Line_of_Business, Business_Unit, and Project make up a business entity. Any of the values can be left empty. If an application is already related with a business entity different from the input values, the application is assigned to the new business entity.

The following object types are supported in the import utility:

Type	Name or Key
Host	*<Name>* or *<IP>*
VM	*<Name>*
Storage pool	*<Storage_name>-><Storage_Pool_name>*
Internal volume	*<Storage_name>-><Internal_volume_name>*
Volume	*<Storage_name>-><Volume_name>*
Storage	*<Name>* or *<IP>*
Switch	*<Name>* or *<IP>*
Port	*<WWN>*
Share	*<Storage Name>-><Internal Volume Name>-><Qtree Name>*
Qtree	*<Storage Name>-><Internal Volume Name>-><Share Name>-><Protocol>*[*-><Qtree* *<Qtree>* is optional in case there is a default Qtree.

The CSV file should use the following format:

```
, , <Annotation Type> [, <Annotation Type> ...]
[, Application] [, Tenant] [, Line_Of_Business] [,
Business_Unit] [, Project]

<Object Type Value 1>, <Object Name or Key 1>, <Annotation Value> [,
```

```
<Annotation Value> ...] [, <Application>] [, <Tenant>] [,
<Line_Of_Business>] [, <Business_Unit>] [, <Project>]

<Object Type Value 2>, <Object Name or Key 2>, <Annotation Value> [,
<Annotation Value> ...] [, <Application>] [, <Tenant>] [,
<Line_Of_Business>] [, <Business_Unit>] [, <Project>]

<Object Type Value 3>, <Object Name or Key 3>, <Annotation Value> [,
<Annotation Value> ...] [, <Application>] [, <Tenant>] [,
<Line_Of_Business>] [, <Business_Unit>] [, <Project>]

...

<Object Type Value N>, <Object Name or Key N>, <Annotation Value> [,
<Annotation Value> ...] [, <Application>] [, <Tenant>] [,
<Line_Of_Business>] [, <Business_Unit>] [, <Project>]
```

Steps

1. Log in to the Insight web UI.
2. Click **Admin** and select **Troubleshooting**.

   The Troubleshooting page displays.
3. In the **Other tasks section** of the page, click the **OnCommand Insight Portal** link.
4. Click **Insight Connect API**.
5. Log in to the portal.
6. Click **Annotation Import Utility**.
7. Save the `.zip` file, unzip it, and read the `readme.txt` file for additional information and samples.
8. Place the CSV file in same folder as the `.zip` file.
9. In the command line window, enter the following:

   ```
   java -jar import-utility.jar [-uusername] [-ppassword]
   [-aserver name or IP address] csv filename [-ltrue|false][-ctrue|
   false] [
   ```

   The -l option, which enables extra logging, and the -c option, which enables case sensitivity, are set to false by default; therefore, you need to specify them only when you want to use the features.

   Note: There are no spaces between the options and their values.

Querying assets

Queries enable you to monitor and troubleshoot your network by searching the assets in your environment at a granular level based on user-selected criteria (annotations and performance metrics). Additionally, annotation rules, which automatically assign annotations to assets, require a query.

Creating a query

You can create a query to enable you to search the assets in your environment at a granular level. Queries enable you to slice data by adding filters and then sorting the results to view inventory and performance data in one view.

About this task

For example, you can create a query for volumes, add a filter to find particular storages associated with the selected volume, add a filter to find a particular annotation, such as Tier 1, on the selected storages, and finally add another filter to find all storages with IOPS - Read (IO/s) greater than 25. When the results are displayed, you can then sort the columns of information associated with the query in ascending or descending order.

When a new data source is added which acquires assets or any annotation or application assignments are made, you can query for those assets, annotations, or applications after the queries are indexed, which occurs at a regularly scheduled interval.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Queries** and select **+ New Query**.
3. Click **Select Resource Type** and select a type of asset.

   When a resource is selected for a query, a number of default columns are automatically displayed; you can remove these columns or add new ones at any time.
4. In the **Name** text box, type the name of the asset or type a portion of text to filter through the asset names.

   You can use any of the following characters alone or combined to refine your search in any text box on the New Query page:

   - An asterisk enables you to search for everything. For example, “vol*rhel” displays all resources that start with “vol” and end with “rhel”.
   - The question mark enables you to search for a specific number of characters. For example, “BOS-PRD??-S12” displays BOS-PRD12-S12, BOS-PRD13-S12, and so on.
   - The OR operator enables you to specify multiple entities. For example, “FAS2240 OR CX600 OR FAS3270” finds multiple storage models.
5. Click to display the assets.
6. To add a criteria, click

, and do either of the following:

   - Type to search for a specific criteria and then select it.
   - Scroll down the list and select a criteria.
   - Enter a range of values if you choose a performance metric like IOPS - Read (IO/s).

Default annotations provided by Insight are indicated by ; it is possible to have annotations with duplicate names.

A column is added to the Query results list for the criteria and the results of the query in the list updates.

7. Optionally, you can click to remove an annotation or performance metric from the query results.

   For example, if your query shows maximum latency and maximum throughput for datastores and you want to show only maximum latency in the query results list, click this button, and clear the **Throughput - Max** check box. The Throughput - Max (MB/s) column is removed from the Query results list.

8. Click **Save**, enter a name for the query, and click **Save** again.

   If you have an account with an administrator role, you can create custom dashboards. A custom dashboard can comprise any of the widgets from Widget Library, several of which, let you represent query results in a custom dashboard. For more information about custom dashboards, see the *OnCommand Insight Getting Started Guide*.

Viewing queries

You can view your queries to monitor your assets and change how your queries display the data related to your assets.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Queries** and select **Show all queries**.
3. You can change how queries display by doing any of the following:
   - You can enter text in the **filter** box to search to display specific queries.
   - You can change the sort order of the columns in the table of queries to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header.
   - You can click a page number at the bottom of the table to browse through queries by page if there are more queries than fit on a single page.
   - You can change the time period (3h, 24h, 3d, 7d, all) for the data displayed, with Insight showing a maximum number of 1000 queries for the selected time period.

Exporting query results to a .CSV file

You might want to export the results of a query into a .CSV file to import the data into another application.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Queries** and select **Show all queries**.

   The Queries page is displayed.

3. Click a query.
4. Click to export query results to a `.CSV` file.
5. Do one of the following:

- Click **Open with** and then **OK** to open the file with Microsoft Excel and save the file to a specific location.
- Click **Save file** and then **OK** to save the file to your Downloads folder.

All of the objects attributes are exported to the file.

Modifying queries

You can change the criteria that are associated with a query when you want to change the search criteria for the assets that you are querying.

Steps

1. Log in to the Insight web UI.
2. Click **Queries** and select **Show all queries**.

   The Queries page is displayed.
3. Click the query name.
4. To remove a criterion from the query, click .
5. To add a criteria to the query, click , and select a criteria from the list.
6. Do one of the following:
   - Click **Save** to save the query with the name that was used initially.
   - Click **Save as** to save the query with another name.
   - Click **Rename** to change the query name that you had used initially.
   - Click **Revert** to change the query name back to the one that you had used initially.

Deleting queries

You can delete queries when they no longer gather useful information about your assets. You cannot delete a query if it is used in an annotation rule.

Steps

1. Log in to the Insight web UI.
2. Click **Queries** and select **Show all queries**.

   The Queries page displays.
3. Position your cursor over the query you want to delete and click .

   A confirmation message displays, asking if you want to delete the query.
4. Click **OK**.

Assigning multiple applications to or removing multiple applications from assets

You can assign multiple applications to or remove multiple application from assets by using a query instead of having to manually assign or remove them. You can perform this task on hosts, internal

volumes, and virtual machines by using the web UI; you must use the Java UI to perform this task on qtrees and shares.

Before you begin

You must have already created a query that finds all the assets that you to edit.

Steps

1. Click **Queries** and select **Show all queries**.

   The Queries page displays.

2. Click the name of the query that finds the assets.

   The list of assets associated with the query display.

3. Click , and select **All**, select **All on this page**, or individually select the assets in the list.

   The **Actions** button displays.

4. To add an application to the assets, click , and select **Edit Application**.

   a. Click **Application** and select one or more applications.

      You can select multiple applications for hosts, internal volumes, and virtual machines; however, you can select only one application for a volume.

   b. Click **Save**.

5. To remove an application assigned to the assets, click and select **Remove Application**.

   a. Select the application or applications you want to remove.

   b. Click **Delete**.

   Any new applications you assign override any applications on the asset that were derived from another asset. For example, volumes inherit applications from hosts, and when new applications are assigned to a volume, the new application takes precedence over the derived application.

Related tasks

Creating a query on page 57

Editing or removing multiple annotations from assets

You can edit multiple annotations for assets or remove multiple annotations from assets by using a query instead of having to manually edit or remove them. You can perform this task on hosts, internal volumes, and virtual machines by using the web UI; you must use the Java UI to perform this task on qtrees and shares.

Before you begin

You must have already created a query that finds all the assets that you want to edit.

Steps

1. Click **Queries** and select **Show all queries**.

   The Queries page displays.

2. Click the name of the query that find the assets.

   The list of assets associated with the query display.

3. Click , and select **All**, select **All on this page**, or individually select the assets in the list.

   The **Actions** button displays.

4. To add an annotation to the assets or edit the value of an annotation assigned to the assets, click , and select **Edit Annotation**.

   a. Click **Annotation** and select an annotation you want to change the value for or select a new annotation to assign it to all the assets.

   b. Click **Value** and select a value for the annotation.

   c. Click **Save**.

5. To remove an annotation assigned to the assets, click , and select **Remove Annotation**.

   a. Click **Annotation** and select the annotation you want to remove from the assets.

   b. Click **Delete**.

Related tasks

Creating a query on page 57

Managing performance policies

OnCommand Insight enables you to create performance policies to monitor your network for various thresholds and to raise alerts when those thresholds are crossed. Using performance policies, you can detect a violation of a threshold immediately, identify the implication, and analyze the impact and root cause of the problem in a manner that enables rapid and effective correction.

A performance policy enables you to set thresholds on any objects (datastore, disk, hypervisor, internal volume, port, storage, storage node, storage pool, VMDK, virtual machine, and volume) with reported performance counters (for example, total IOPS). When a violation of a threshold occurs, Insight detects and reports it in the associated asset page, by displaying a red solid circle; by email alert, if configured; and in the Violations Dashboard or any custom dashboard that reports violations.

Insight provides some default performance policies, which you can modify or delete if they are not applicable to your environment, for the following objects:

- Hypervisor
  There are ESX swapping and ESX utilization policies.
- Internal volume and volume
  There are two latency policies for each resource, one annotated for Tier 1 and the other annotated for Tier 2.
- Port
  There is a policy for BB credit zero.
- Storage node
  There is a policy for node utilization.
- Virtual machine
  There are VM swapping and ESX CPU and memory policies.

Creating performance policies

You create performance policies to set thresholds that trigger alerts to notify you about issues related to the resources in your network. For example, you can create a performance policy to alert you when total utilization for storage pools is greater than 60%.

Steps

1. Open OnCommandInsight in your browser.
2. From the **Manage** menu, select **Performance Policies**.

   The Performance Policies page is displayed:

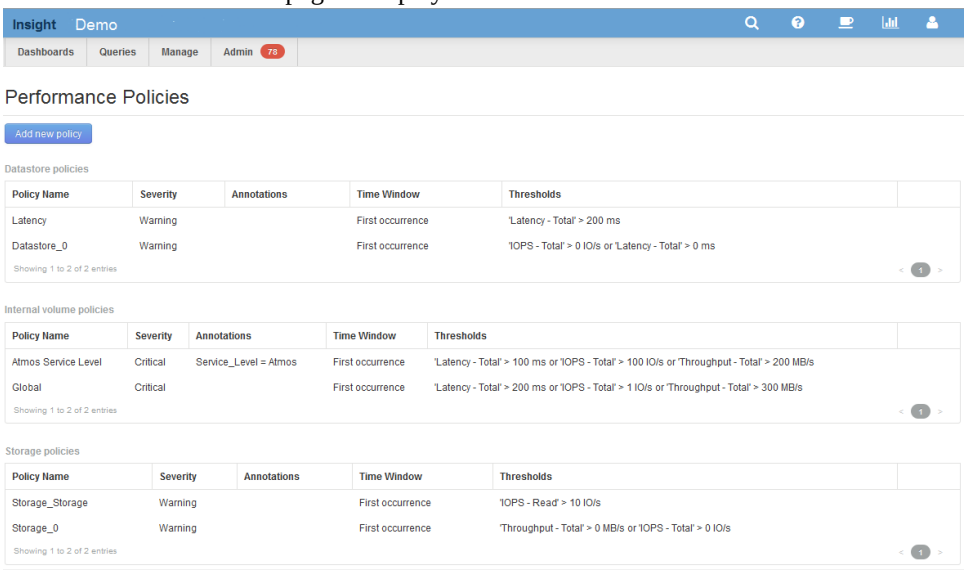

   Policies are organized by object and are evaluated in the order in which they appear in the list for that object.

3. Click **Add new policy**.

   The Add Policy dialog box displays:

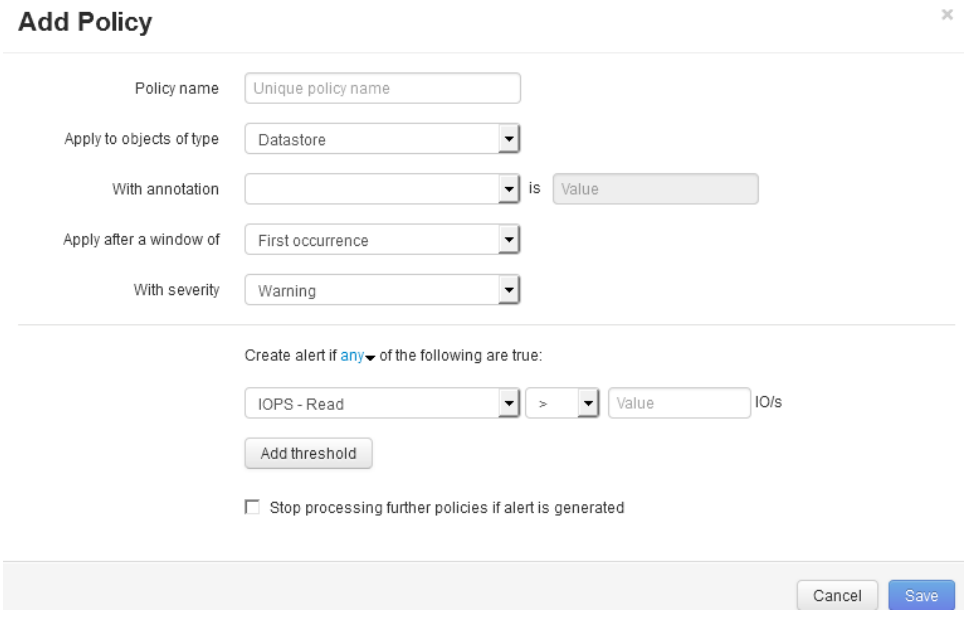

4. In the **Name** box, enter a name for the policy.

   You must use a name that is different from any of the other policy names for the object. For example, you cannot have two policies named “Latency” for an internal volume; however, you can have a “Latency” policy for an internal volume and another one for a different volume. The best practice is to always use a unique name for any policy, regardless of the object type.

5. In the **Apply to objects of type** list, select the type of object to which the policy applies.

6. In the **With annotation** list, select an annotation type, if applicable, and enter a value for the annotation in the **Value** box to have the policy apply only to objects that have this particular annotation set.

7. If you selected **Port** as the object type, select what the port is connected to in the **Connected to** list.

8. In the **Apply after a window of** list, select when an alert is raised to indicate a threshold violation.

   The **First occurrence** option triggers an alert when a violation of a threshold occurs on the first sample of data. All other options trigger an alert when if the threshold is crossed once and is continuously crossed for at least the specified amount of time.

9. In the **With severity** list, select the severity for the violation, depending on how important the violation is to the data center environment.

10. Click the **any** link in the **Create alert if any of the following are true** section to control how alerts are triggered for thresholds by selecting one of the following options:

    - **any**
      This is the default, which creates an alert when any of the thresholds related to a policy are crossed. For example, if a policy has four thresholds and two are crossed, Insight creates two alerts.

    - **all**
      This creates an alert when only all of the thresholds for a policy are crossed. When you select **all**, the first threshold you create for a performance policy is referred to as the primary rule. Ensure that the primary rule threshold is the violation that you are most concerned about for the performance policy.

11. In the **Create alert if** section, select a performance counter and an operator, and then enter a value to create a threshold.

12. Click **Add threshold** to add more thresholds.

13. To remove a threshold, click the trash can icon.

14. Check the **Stop processing further policies if alert is generated** box if you want the policy to stop processing when an alert occurs.

    #### Example

    For example, if you have four policies for datastores and the second policy is configured to stop processing when an alert occurs, the third and fourth policies are not processed while a violation of the second policy is active.

15. Click **Save**.

    The Performance Policies page displays, and the performance policy appears in the list of policies for the object type.

Performance policy evaluation precedence

The Performance Policies page groups policies by object type and Insight evaluates the policies in the order in which they appear in the object's performance policy list. You can change the order in which Insight evaluates policies in order to show the information that is most important to you in your network.

Insight evaluates all policies that are applicable to an object sequentially when performance data samples are taken into the system for that object; however, depending on annotations, not all policies apply to one group of objects. For example, suppose that internal volume has the following policies:

- Policy 1 (the Insight-supplied default policy)
- Policy 2 (with an annotation of “Service Level = Silver” with the **Stop processing further policies if alert is generated** option
- Policy 3 (with an annotation of“Service Level = Gold”)
- Policy 4

For an internal volume tier with a Gold annotation, Insight evaluates Policy 1, ignores Policy 2, and then evaluates Policy 3 and Policy 4. For an unannotated tier, Insight evaluates by the order of the policies; thus, Insight evaluates only Policy 1 and Policy 4. For an internal volume tier with a Silver annotation, Insight evaluates Policy 1 and Policy 2; however, if an alert is triggered when the policy's threshold is crossed once and is continuously crossed for the window of time specified in the policy, then Insight no longer evaluates the other policies in the list while it evaluates the current counters for the object. When Insight captures the next set of performance samples for the object, it again begins to evaluate the performance policies for the object by filter and then order.

Changing the precedence of a performance policy

By default, Insight evaluates an object's policies sequentially. You can configure the order in which Insight evaluates performance policies. For example, if you have a policy configured to stop processing when a violation occurs for Gold Tier storage, you can place that policy first in the list and avoid seeing more generic violations for the same storage asset.

Steps

1. Open Insight in your browser.
2. From the **Manage** menu, select **Performance Policies**.

   The Performance Policies page displays.
3. Hover your cursor over a policy name in an object type's performance policy list.

   The precedence arrows appear to the right of the policy.
4. To move a policy up in the list, click the up arrow; to move a policy down in the list, click the down arrow.

   By default, new policies are added sequentially to an object's list of policies.

Editing performance policies

You can edit existing and default performance policies to change how Insight monitors the conditions of interest to you in your network. For example, you might want to change a policy's threshold.

Steps

1. Open Insight in your browser.

2. From the **Manage** menu, select **Performance Policies**.

   The Performance Policies page displays.

3. Hover your cursor over a policy name in an object's performance policy list.

4. Click .

   The Edit Policy dialog box displays.

5. Make the required changes.

   If you change any option other than the policy name, Insight deletes all existing violations for that policy.

6. Click **Save.**

Deleting performance policies

You can delete a performance policy if you feel that it is no longer applicable to monitoring the objects in your network.

Steps

1. Open Insight in your browser.

2. From the **Manage** menu, select **Performance Policies**.

   The Performance Policies page displays.

3. Hover your cursor over the name of a policy in an object's performance policy list.

4. Click .

   A message appears, asking if you want to delete the policy.

5. Click **OK**.

Insight data source management

Data sources are the most critical component used to maintain an OnCommand Insight environment. Because they are the primary source of information for Insight, it is imperative that you maintain data sources in a running state.

You can monitor the data sources in your network by selecting a data source to check the events related to its status and noting any changes that might have caused problems.

In addition to examining an individual data source, you can perform these operations:

- Clone a data source to create many similar data sources in Insight
- Edit data source information
- Change credentials
- Control polling
- Delete the data source
- Install data source patches
- Install a new data source from a patch
- Prepare an error report for NetApp Customer Support

Related concepts

What data source patches are on page 132
Changes causing data collection problems on page 126

Setting up your data sources in Insight

Data sources are the most critical component when trying to maintain a Insight environment. Data sources discover network information that is used for analysis and validation. You need to configure your data sources within Insight so that they can be monitored within your network.

For each data source, the specific requirements to define that data source depend on the vendor and model of the corresponding devices. Before adding the data sources, you need network addresses, account information, and passwords for all devices and possibly these additional details:

- Switches
- Device management stations
- Storage systems that have IP connectivity
- Storage management stations
- Host servers running management software for storage devices that do not have IP connectivity

For more information about your data source definitions, see the "Vendor-specific data source reference" information in this section.

Related concepts

Insight data source management on page 66

Related tasks

Cloning a data source on page 69

Related references

Vendor-specific data source reference on page 71

Data source support information

As part of your configuration planning, you should ensure that the devices in your environment can be monitored by Insight. To do so, you can check the Data source support matrix for details about operating systems, specific devices, and protocols. Some data sources might not be available on all operating systems.

To access the data source support matrix from the Insight web UI, click the **Help** icon on the toolbar and then click the **Data source support matrix** option.

Adding data sources

You can add data sources quickly, using the Add data source dialog box.

Steps

1. Open OnCommand Insight in your browser.
2. If the Data sources table is not displayed, click the **Admin** icon on the toolbar to open it.
3. Click the **Actions** button and select the **Add data source** option.

   The Add data source dialog box is displayed.

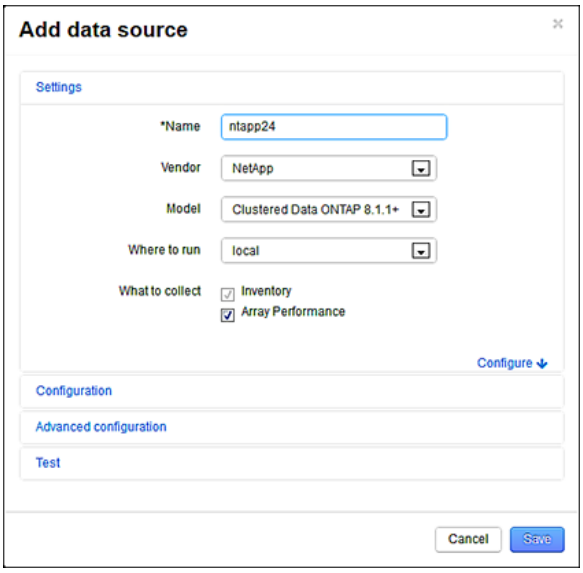

4. Enter your network **Name** for the new data source, as shown in this example.
5. In the **Add data source** dialog box, select the **Vendor** and **Model** of the data source from the drop-down lists.
6. Select any **What to collect** additional options.

   These selections change the available fields in the Configuration and Advanced configuration sections.

In this example, these options are Inventory and Array Performance. Inventory is always selected and cannot be changed. Other possible data collection options are Switch Performance, VM Performance, and Replication (for data synchronization).

7. Click the **Configuration** link and enter the basic setup information required for the data source with your selected data collection type.

   If this is all of the information needed for the data source, you might click the **Test** link at this point.

8. If this type of data source usually requires more detailed information to set it up in your network, click the **Advanced configuration** link to enter additional information.

   If you do not have the information required for the advanced configuration, see the "Vendor-specific data source reference" information for that data source.

9. Click the **Test** link to be certain that the data source is properly configured.

10. Click **Save**.

Related concepts

Data source support information on page 67

Related tasks

Cloning a data source on page 69

Related references

Vendor-specific data source reference on page 71
Testing the data source configuration on page 70

Importing data sources from a spreadsheet

You can import multiple data sources into OnCommand Insight from a spreadsheet. This might be helpful if you already maintain your discovery devices in a spreadsheet. This process adds new data sources, but cannot be used to update existing data sources.

About this task

OnCommand Insight includes a spreadsheet to help you create data sources. This spreadsheet has the following attributes:

- The spreadsheet can be used with Microsoft Excel 2003 or later.
- Each tab holds one data source type, for example, Brocade SSH/CLI.
- Each row represents an instance of a new data source to be created.

The spreadsheet includes a macro that creates a new data source in OnCommand Insight.

Steps

1. Locate the spreadsheet in the `<install_directory>/SANscreen/acq/bin/acqcli/SiteSurvey_DataSourceImporter_w_Macro.zip`.
2. In the spreadsheet, enter data source information in the cells with color.
3. Delete empty rows.
4. From the spreadsheet, run the `CreateDataSources` macro to create the data sources.

5. When prompted for credentials, enter the OnCommand Insight Server administration user name and password.

   The results are logged in the acquisition log.

6. A prompt asks if the machine currently running the macro has OnCommand Insight installed.

   Select one of the following:

   - No: Select "No" if a batch file will be created that must be run on the OnCommand Insight machine. Run this batch file from the install directory.
   - Yes: Select "Yes" if OnCommand Insight is already installed and no additional steps are required to generate the data source information.

7. To verify the addition of the data sources, open Insight in your browser.
8. Click the **Admin** icon on the toolbar.
9. Check the Data sources list for the data sources you imported.

Adding a new data source by patch

New data sources are released as patch files that can be loaded onto the system using the patch process. This process enables new data sources to be available between scheduled releases of OnCommand Insight.

Before you begin

You must have uploaded the patch file that you want to install.

Steps

1. On the Insight toolbar, click the **Admin** icon.
2. Select **Patches**.
3. Select **Actions > Install service pack or patches**.
4. In the **Install Service Pack or Patch** dialog box, click **Browse** to locate the uploaded patch file.
5. Click **Next** in the **Patch Summary** dialog box if you have the correct data source patch.
6. Review the **Read Me** information, and click **Next** to continue.
7. In the **Install** dialog box, click **Finish**.

Cloning a data source

Using the clone facility, you can quickly add a data source that has the same credentials and attributes as another data source. Cloning allows you to easily configure multiple instances of the same device type.

Steps

1. Click the **Admin** icon on the toolbar.

   The Data sources list opens.

2. Highlight the data source that has the setup information you want to use for your new data source.
3. To the right of the highlighted data source, click the **Clone** icon.

   The Clone this data source dialog box lists the information you must supply for the selected data source, as shown in this example for a NetApp data source:

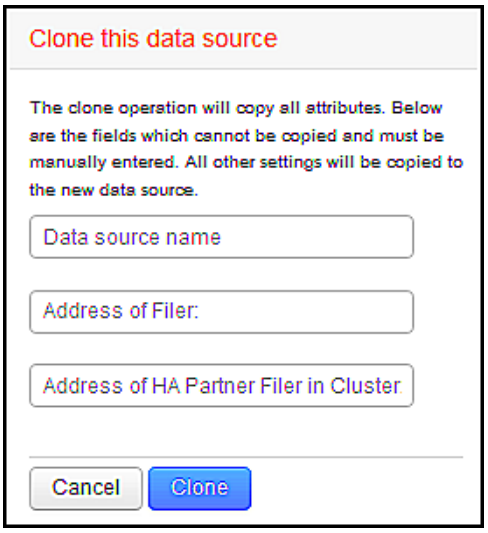

4. Enter the required information in the fields; those details cannot be copied from the existing data source.
5. Click **Clone**.

Result

The clone operation copies all other attributes and settings to create the new data source.

Related concepts

Data source support information on page 67

Related references

Vendor-specific data source reference on page 71

Testing the data source configuration

When you are adding a data source, you can verify the correctness of configuration to communicate with the device before saving or updating that data source.

When you click **Test** in the data source wizard, Insight tests the important inventory and performance configuration details for the specific data source. The test produces one of these results:

- PASSED: the data source is configured correctly for Insight.
- WARNING: the testing was incomplete, probably due to timing out during processing or acquisition not running.
- FAILED: the data source, as configured, cannot communicate with the specified device.

Related tasks

Adding data sources on page 67
Editing data source information on page 129

Vendor-specific data source reference

The configuration details vary depending on the vendor and model of the data source being added.

If a vendor's data source requires advanced Insight configuration instructions, such as special requirements and specific commands, that information is included in this section.

Related concepts

Data source support information on page 67

Related tasks

Adding data sources on page 67
Cloning a data source on page 69

3PAR InServ Firmware 2.2.2+ (SSH) data source

OnCommand Insight uses the 3PAR InServ Firmware 2.2.2+ (SSH) data source to discover inventory HP 3PAR StoreServ storage arrays.

Configuration

Field	Description
Cluster IP(s)	IP address of the storage.
User Name	Name used to log in to the 3PAR device.
Password	Password used to log in to the 3PAR device.
SMI-S Host IP	IP address of the SMI-S Provider Host.
SMI-S User Name Username	Name used to log in to the SMI-S Provider Host.
SMI-S Password	Password used to log in to the SMI-S Provider Host.

Requirements

- Read-only user name and password to the InServ Server
- Validated access by using `telnet` to port 22

Note: No third-party software is required.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device for inventory information
Exclude Devices	A comma-separated list of device IP addresses that should not be polled
SSH Process Wait Timeout (sec)	The time in seconds waiting for the command to finish. If not completed when timeout expires, the run fails.
Number of SSH Retries	Maximum number of retry attempts for a given command

Field	Description
SSH Banner Wait Timeout (sec)	Timeout for obtaining the first prompt after connection to the device
SMI-S Port	Port number of the SMI-S Provider Host
Protocol	The protocol used to connect to SMI-S provider
SMI-S Namespace	Interoperability namespace that the SMI-S provider is configured to use
Performance Poll Interval (sec)	Interval, in seconds, at which the data source should poll the device for performance information
Number of SMI-S Connection Retries	Maximum number of retry attempts for a given command

Commands

- showsys -csvtable
- showversion -a
- showlicense
- showport -csvtable
- showport -par -csvtable
- showld -csvtable
- showld -d -csvtable
- showpd -i -csvtable
- showpd -csvtable
- showvv -r -csvtable
- showld -d -csvtable
- showvv -csvtable
- showvlun -a -csvtable
- showhost -d -csvtable
- showcpg -r -csvtable
- showrcopy -csvtable

Brocade EFCM (9.5, 9.6, 9.7) data sources

For configuration, this data source requires the network address for the Enterprise Fabric Connectivity Manager (EFCM) server.

Configuration

Field	Description
EFC server	Network address for the Enterprise Fabric Connectivity (EFC) server, specified as either an IP (*nnn.nnn.nnn.nnn* format) or a hostname that can be resolved through DNS

Field	Description
User Name	Name used to log in to the switch
Password	Password used to log in to the McData device

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval (in minutes) at which the data source should poll the EFC server The default setting is 15 minutes.
Fabric Name	Fabric name to be reported by the EFCM data source (blank to report the fabric name as WWN)
Communication Port	The default setting is -1.
Enable Trapping	Check box that is used to trigger acquisition upon receiving an SNMP trap from the device If you select this check box, you must also activate SNMP notifications.
Minimum Time Between Traps (sec)	Minimum elapsed time (in seconds) between acquisition attempts triggered by traps The default setting is 15 seconds.
Inactive Zonesets	List of inactive zone sets on which to perform acquisition (in addition to performing acquisition on the active zone sets) You must separate multiple zone sets using a comma.
NIC to Use	Applicable when the data source reports to a Remote Acquisition Unit (RAU) that has two or more network cards: one that is used to acquire information about SAN devices, and the other that is used to connect to the Insight server The setting indicates which network interface should be used on the RAU when reporting on SAN devices.
Exclude Devices	Comma-separated list of devices that will not be discovered by the EFCM data source
Use the EFCM switch nickname as the Insight switch name	Check box that defaults to checked (true)
Performance Poll Interval (sec)	For Switch Performance, the recommended setting for the poll interval is 300 seconds, which is the default setting

Requirements

- IP address of the EFCM server
- Read-only username and password for the EFCM data source
- Access validation: Access to the switch by Telnet from the Insight server, using the read-only username and password over port 51512

Note: No third-party software is required.

Brocade FC Switch Firmware (SSH) data source

The Brocade FC Switch Firmware (SSH) data source supports Brocade or rebranded devices running Factored Operating System (FOS) firmware 4.2 and later. Devices in both FC switch and Access Gateway modes are supported.

Configuration

Field	Description
Switch IP	Network address for the device, specified as an IP (*nnn.nnn.nnn.nnn* format) or as a host name that can be resolved through DNS.
User Name	Name used to log in to the Brocade device. The default is admin.
Password	Password used to log in to the device.
SNMP Version	For Switch Performance, the default is V2.
SNMP Community:	For Switch Performance, SNMP *read-only* community string used to access the device.
SNMP User Name	For Switch Performance, the SNMP version protocol user name. This applies only to SNMP V3.
SNMP Password	For Switch Performance, the SNMP version protocol password. This applies only to SNMP V3.

Advanced configuration

Field	Description
Fabric Name	Fabric name to be reported by this data source (blank to report the fabric WWN).
Exclude Devices	Comma-separated list of Brocade devices, identified by IP address, that are not discovered by this data source.
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. The default is 15.
Timeout (sec)	The period of time that the data source pauses before it declares that the connection has failed (timed out). The default is 30.
Banner Wait Timeout (sec)	Timeout for obtaining the first prompt after connection to the device.
Admin Domains Active	Select to indicate that you implement Admin Domains.
Retrieve MPR Data	Select to acquire routing data for your multiprotocol router (MPR).
Enable Trapping	Check box used to trigger acquisition upon receiving an SNMP trap from the device. If you select this box, you must also activate SNMP notifications.
Minimum Time Between Traps (sec)	Minimum elapsed time between acquisition attempts triggered by traps, in seconds.

Field	Description
Discover all switches in the fabric	Check box used to trigger identification of all of the switches on this fabric.
Choose Favoring HBA vs. Zone Aliases	Determines which alias is chosen if multiple aliases exist for the same WWN.
Performance Poll Interval (sec)	For Switch Performance, the recommended setting and default for the poll interval is 300 seconds.
SNMP Auth Protocol	For Switch Performance, encryption type for SNMP V3 configurations. The default is MD5.
SNMP Privacy Protocol	For Switch Performance, selections include NONE, AES, DES, or TRIPLEDES for SNMP V3.
SNMP Privacy Password	For Switch Performance, provide a password for SNMP V3 encryption usage.
SNMP Retries	For Switch Performance, the maximum number of retry attempts for a given SNMP command.
SNMP Timeout (ms)	For Switch Performance, the timeout setting for connecting to the device using SNMP.

Requirements

- There must be IP connectivity to all switches in the fabric.
  If you select the **Discover all switches in the fabric** check box, OCI identifies all the switches in the fabric; however, it needs IP connectivity to these additional switches to discover them.
- The same account is needed globally across all switches in the fabric.
  You can use PuTTY (open source terminal emulator) to confirm access.
- If the Perform license is installed, ports 161 and 162 must be open to all switches in the fabric for SNMP performance polling.

Note: No third-party software is required.

Commands

- `ag - -modeshow`
- `agshow`
- `alishow`
- `fabricshow`
- `fosconfig - -show`
- `ipaddrshow`
- `lscfg - -show`
- `nbrstateshow`
- `nsshow -t`
- `portshow`
- `sfpshow`

- `snmpConfig - -show accessControl`
- `snmpConfig - -show mibCapability`
- `snmpConfig - -show snmpv3`
- `switchshow`
- `switchstatusshow`
- `trunkshow`
- `version`
- `wwn -sn`

Brocade Sphereon/Intrepid Switch (SNMP) data source

For configuration, this data source requires an IP address for every switch, and an SNMP *read-only* community string to access the device.

Configuration

Field	Description
Sphereon Switch	Network address for the switch, specified as an IP (*nnn.nnn.nnn.nnn* format) or as a host name that can be resolved through DNS.
SNMP Version	V2 is the default.
SNMP Community	Specifies the SNMP *read-only* community string used to access the device.
User Name	Specifies the SMIS user name used to log in to the switch. Used only for V3.
Password	Specifies the SMIS password used to log in to the switch. Used only for V3.

Requirements

- The IP address of every switch
- Read-only community string if using SNMP V1 or SNMP V2
- HTTP access to the switch to obtain zoning information
- Access validation
  Run the `snmpwalk` utility to the switch (see `<install_path>\bin\`).

Note: No third-party software is required.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Specifies the interval, in minutes, at which the data source should poll the device. The default is 15.
SNMP Auth Protocol	Specifies the encryption type for SNMP V3. The default is MD5.

Field	Description
SNMP Privacy Protocol	Selections include NONE, AES, DES, or TRIPLEDES for SNMP V3.
SNMP Privacy Password	Provide a password for SNMP V3 encryption usage.
SNMP Number of Retries	Specifies the maximum number of retry attempts for a given SNMP command.
SNMP Timeout (ms)	Specifies the timeout for connecting to the device using SNMP.
Fabric Name	Specifies the fabric name to be reported by this data source (blank to report the fabric WWN).
Enable Trapping	Use this check box to trigger acquisition upon receiving an SNMP trap from the device. If you select this box, you must also activate SNMP notifications.
Minimum Time Between Ttraps (seconds)	Specifies the minimum elapsed time between acquisition attempts triggered by traps, in seconds.
Performance Poll Interval (sec)	For Switch Performance, the recommended setting and default for the poll interval is 300 seconds.

SNMP MIBs

- EF-6000-MIB
- EOSN-FCPING-MIB
- EOSN-FCTRACEROUTE-MIB
- EOSN-MIB
- EOSN-PARTITIONING-MIB
- EOSN-REG
- EOSN-SECURITY-MIB
- EOSN-TC
- EOSN-TRAP-MIB
- FCFABRIC-ELEMENT-MIB
- FCMGMT-MIB
- FIBRE-CHANNEL-FE-MIB
- FIBRE-CHANNEL-MGMT-MIB
- IANAifType-MIB
- IF-MIB
- INET-ADDRESS-MIB
- RFC1155-SMI
- RFC1213-MIB
- RFC1215-MIB
- SNMP-COMMUNITY-MIB

- SNMP-FRAMEWORK-MIB
- SNMP-NOTIFICATION-MIB
- SNMP-TARGET-MIB
- SNMP-USER-BASED-SM-MIB
- SNMPv2-CONF
- SNMPv2-MIB
- SNMPv2-SMI
- SNMPv2-TC
- SNMP-VIEW-BASED-ACM-MIB
- TRANSPORT-ADDRESS-MIB

Cisco FC Switch Firmware (SNMP) data source

The Cisco FC Switch Firmware 2.0+ (SNMP) data source is applicable to all Cisco MDS Fibre Channel switches and also to a variety of Cisco Nexus FCoE switches on which the FC service is enabled. Additionally, you can discover many models of Cisco devices running in NPV mode with this data source.

Configuration

Field	Description
Cisco Switch IP	Network address for the switch, specified as an IP (*nnn.nnn.nnn.nnn* format) or as a host name that can be resolved through DNS.
SNMP Version	V2 is the default. V2 or later is required for performance; performance fails if it is configured with SNMP V1. V3 is supported since 6.0.
SNMP Community	SNMP *read-only* community string used to access the device. Community strings are not applicable for SNMP V3 configurations.
User Name	Name used to log in to the switch. Used for SNMP V3 configurations only.
Password	Password used to log in to the switch. Used for SNMP V3 configurations only.

Requirements

The following items are required:

- An IP address of one switch in the fabric or individual switches
- Chassis discovery, to enable fabric discovery
- If using SNMP V2, read-only community string
- Access validation using the `snmpwalk` utility to the switch (see `<install_path>\bin\`)

Note: No third-party software is required.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. The default is 40.
SNMP Auth Protocol	Encryption type for SNMP V3 configurations. The default is MD5.
SNMP Privacy Protocol	Selections include NONE, AES, DES, and TRIPLEDES for SNMP V3.
SNMP Privacy Password	Provide a password for SNMP V3 encryption usage.
SNMP Retries	Maximum number of retry attempts for a given SNMP command.
SNMP Timeout (ms)	Timeout for connecting to the device using SNMP.
Enable Trapping	Check box used to trigger acquisition upon receiving an SNMP trap from the device. If you mark this box, you must also activate SNMP notifications.
Minimum Time Between Traps (sec)	Minimum elapsed time between acquisition attempts triggered by traps, in seconds. 10 is the default.
Discover All Fabric Switches	If this check box is selected, the data source attempts to discover not only the switch configured in the data source, but also all other Cisco switches in the VSANs and fabrics of which this switch is a member.
Exclude Devices	This is a comma-separated list of devices, identified by switch IP address, that this data source will not discover. It supports wildcards and CIDR addresses (for example, 10.20.* and 192.168.1.0/24).
Include Devices	This is a comma-separated list of devices, identified by switch IP address, that this data source will discover. It supports wildcards and CIDR addresses (for example, 10.20.* and 192.168.1.0/24).
Check Device Type	If set, the data source accepts only those devices that explicitly advertise themselves as Cisco devices.

Field	Description
Primary Alias Type	Select the first preference for resolution of the alias from these options: **Device Alias** This is a user-friendly name for a port WWN (pWWN) that can be used in all configuration commands, as required. All switches in the Cisco MDS 9000 Family support Distributed Device Alias Services (device aliases). **None** Do not report any alias **Port Description** A description to help identify the port in a list of ports **Zone Alias (all)** A user-friendly name for a port that can be used only for zoning configuration **Zone Alias (only active)** A user-friendly name for a port that can be used only for the active configuration. This is the default.
Secondary Alias Type	Provide a second preference for resolution of the alias using the same options as the Primary Alias Type.
Tertiary Alias Type	Provide a third preference for resolution of the alias using the same options as the Primary Alias Type.
Enable SANTap Proxy Mode Support	Select this check box if your Cisco switch is using SANTap in proxy mode. If you are using EMC RecoverPoint, then you are probably using SANTap.
Performance Poll Interval (sec)	For Switch Performance, the recommended setting and default for the poll interval is 300 seconds.

SNMP MIBs

- CISCO-CDP-MIB
- CISCO-DM-MIB
- CISCO-DYNAMIC-PORT-VSAN-MIB
- CISCO-ENTITY-DISPLAY-MIB
- CISCO-FC-DEVICE-ALIAS-MIB

- CISCO-FC-FE-MIB
- CISCO-FCS-MIB
- CISCO-FDMI-MIB
- CISCO-IVR-MIB
- CISCO-LICENSE-MGR-MIB
- CISCO-NS-MIB
- CISCO-PSM-MIB
- CISCO-SMI
- CISCO-ST-TC
- CISCO-TC
- CISCO-VSAN-MIB
- CISCO-VTP-MIB
- CISCO-ZS-MIB
- ENTITY-MIB
- IANAifType-MIB
- IF-MIB
- INET-ADDRESS-MIB
- SNMP-FRAMEWORK-MIB
- SNMPv2-CONF
- SNMPv2-MIB
- SNMPv2-SMI
- SNMPv2-TC

Dell Compellent Storage Center data source

The Dell Compellent Storage Center data source collects network information for analysis and validation from Dell Compellent Storage Center storage arrays using the Dell Compellent Enterprise Manager. To add the data source, you must use specific configuration instructions and commands and understand data source requirements and usage recommendations.

Configuration

Field	Description
Address of Enterprise Manager	The IP address of the Dell Compellent Enterprise Manager server
Port	The port number of the Dell Compellent Enterprise Manager
User Name	The user name to log in to the Dell Compellent Enterprise Manager
Password	The password to log in to the Dell Compellent Enterprise Manager

General requirements

You need administrator credentials for the Dell Compellent Storage Center.

Advanced configuration

Field	Description
Poll Interval (min)	Poll interval in minutes
Performance Poll Interval (up to 3600 seconds)	Poll interval in seconds

Commands

This data source uses the Dell Java API.

EMC Atmos (HTTP) data source

This data source supports the EMC cloud offering, which includes a single pool of storage.

Implementation considerations

The EMC Atmos data source provides support to the EMC cloud offering, which includes a single pool of storage. This data source is implemented in the following ways:

- Atmos data appears only in the Reporting Storage and Storage Pool Capacity data model and the Chargeback data model.
- Atmos has a physical hierarchy and a logical hierarchy that impact what you see in reports.
- Atmos resource management group (RMG) capacity is tracked in Insight reporting as storage raw capacity.
- Atmos data is not grouped by and not reported on by tier.
- Atmos data can be reported on by an Atmos service level.

Atmos logical hierarchy and OnCommand Insight business entity divisions

Insight reports on Atmos logical hierarchy in its business entities as follows:

Business entity levels	Atmos logical hierarchy division
Tenant	Tenant
Line of business (LOB)	Subtenant
Business Unit	n/a
Project	n/a

Note: Business entities are created and maintained in the Java UI.

Configuration

Field	Description
Address of Atmos	Network address for the Atmos node that is currently running the EMC Capacity Reporting Tool. Specified as an IP (*nnn.nnn.nnn.nnn* format) or as a host name that can be resolved through DNS.
User Name	Name used for authentication into the EMC Capacity Reporting Tool API. The default is admin.
Password	Password used for authentication into the EMC Capacity Reporting Tool API.

Requirements

User name and password are not needed for EMC Capacity Reporting Tool authentication; however, they are needed for the reporting API.

- EMC Capacity Reporting Tool must be installed.
- An IP address of the Atmos node that is currently running the EMC Capacity Reporting Tool must be available.
- Access can be verified using `telnet` to the user-defined port.

Advanced configuration

Configuration

Field	Description
Port	Port used internally by the PC, where the Insight Data Warehouse (DWH) is running, to serve as an access point to the database.
HTTPS Enabled	Select to use the HTTPS protocol and connect through port 443.
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. The recommended setting and default is 20 minutes.

EMC Celerra data source

The Celerra data source collects inventory information from Celerra storage. For configuration, this data source requires the IP address of the storage processors and a *read-only* user name and password.

Configuration

Field	Description
Address of Celerra	IP address of storage processors
User Name	Name used to log in to the Celerra device
Password	Password used to log in to the Celerra device

Requirements

- The IP address of the storage processor
- Read-only user name and password
- SSH access via PuTTY

Advanced configuration

Field	Description
Inventory Poll Interval (minutes)	Interval, in minutes, at which the data source should poll the device The default is 20.
SSH Process Wait Timeout (sec)	The time, in seconds, to wait for the command to finish If not completed when the timeout expires, the SSH Client disconnects.
Number of Retries	Maximum number of retry attempts for a given Celerra command
SSH Banner Wait Timeout (sec)	Length of time to wait for a connection to a host before timing out

Commands

The following commands are invoked:

- `fs_dedupe -info -all`
- `nas_disk -info id=<disk id>"`
- `nas_disk -list`
- `nas_fs -info -size -all`
- `nas_pool -info -all`
- `nas_pool -size -all`
- `nas_quotas -list -tree -fs <file system name>"`
- `nas_quotas -report -tree -fs <file system name`
- `nas_storage -list`
- `nas_storage -info <storageName>`
- `nas_version -l`
- `nas_volume -info -size -all`
- `nas_xml -info:ALL -level:3`
- `/nas/sbin/t2tty -c <> "fcp topology"`

EMC CLARiiON (NaviCLI) data source

Before configuring this data source, make sure that the EMC Navisphere CLI is installed on the target device and on the Insight server. The Navisphere CLI version must match the firmware version on the controller. For performance data collection, statistics logging must be turned on.

NaviSphere Command Line Interface syntax

```
naviseccli.exe -h <IP address> -user <user> -password <password> -scope
<scope,use 0 for global scope> -port <use 443 by default> command
```

Configuration

Field	Description
CLARiiON storage	Network address for the storage device, specified as an IP (*nnn.nnn.nnn.nnn* format), or as a hostname that can be resolved through DNS.
User Name	Name used to log into the CLARiiON storage device.
Password	Password used to log into the CLARiiON storage device.
CLI Path to navicli.exe path or naviseccli.exe path	Full path to the executable.

Requirements

- An IP address of each CLARiiON storage processor
- Read-only Navisphere username and password to the CLARiiON arrays
- Access validation: Run NaviCLI from the Insight server to each array using the above username and password.
- NaviCLI version should correspond with the newest FLARE code on your array
- 3rd party software installed on Insight server/RAU: navicli

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. The default is 40.
Use Secure Client (naviseccli)	Provides a secure management facility for activating the required CLARiiON array replication support functionality.
Scope	The secure client scope. This is an enumeration with the following values: Global, Local, LDAP. The default is Global.
CLARiiON CLI Port	The port on the CLARiiON used for NaviCLI connections.
Inventory External Process Timeout (sec)	Maximum time allowed for inventory commands to complete.
Performance Poll Interval (sec)	For Array Performance, sets the Perform poll interval. The recommended setting is 300.

Field	Description
Retry Count to Retrieve Archive File	For Array Performance, the maximum number of retry attempts to retrieve an archive file.
Performance External process timeout (sec)	For Array Performance, the maximum time allowed for performance commands to complete.

Performance requirement

- Statistics logging must be turned on.

NaviSec CLI commands

The following commands are invoked:

- naviseccli.exe –help
- naviseccli.exe -h {ip} -User {user} -Password {password} -Scope 0 -port {port} getall
- naviseccli.exe -h {ip} -User {user} -Password {password} -Scope 0 -port {port} metalun -list
- naviseccli.exe -h {ip} -User {user} -Password {password} -Scope 0 -port {port} mirror sync -list
- naviseccli.exe -h {ip} -User {user} -Password {password} -Scope 0 -port {port} mirror async –list
- naviseccli.exe -h {ip} -User {user} -Password {password} -Scope 0 -port {port} storagepool -list
- naviseccli.exe -h {ip} -User {user} -Password {password} -Scope 0 -port {port} thinlun –list
- naviseccli.exe -h {ip} -User {user} -Password {password} -Scope 0 -port {port} storagegroup -list -gname {storage group name} (This can be run multiple times with different storage group names.)

NaviCLI (Classic) commands

- navicli.exe –help
- navicli.exe -h {ip} getall
- C:\Program Files\sanscreen\java\bin\java.exe –jar C:\Program Files\EMC\Navisphere CLI \navicli.jar –user {user} -password {password} -port {port} –address {ip} metalun -list
- navicli.exe -h {ip} mirror sync -list
- navicli.exe -h {ip} mirror async –list
- navicli.exe -h {ip} storagepool -list
- navicli.exe -h {ip} thinlun -list
- navicli.exe -h {ip} storagegroup -list -gname {storage group name} (This can be run multiple times with different storage group names.)

Performance commands

The following commands are run for acquiring performance data:

- naviseccli.exe -h {ip} -User {user} -Password {password} -Scope 0 -port {port} getall
- naviseccli.exe -h {ip} -User {user} -Password {password} -Scope 0 -port {port} getDisk
- naviseccli.exe -h {ip} -User {user} -Password {password} -Scope 0 -port {port} metalun -list

- naviseccli.exe -h {ip} -User {user} -Password {password} -Scope 0 -port {port} thinlun -list

EMC Data Domain data source

This data source collects storage and configuration information from EMC Data Domain deduplication storage systems. To add the data source, you must use specific configuration instructions and commands and be aware of data source requirements and usage recommendations.

Configuration

Field	Description
IP address	The IP address of the Data Domain storage array
User name	The user name you use to log in to the Data Domain storage array
Password	The password you use to log in to the Data Domain storage array

General requirements

- Read-only user name and password to the Data Domain storage
- Validated access by using `telnet` to port 22

Note: No non-NetApp software is required.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device for inventory information
SSH Process Wait Timeout (sec)	The time, in seconds, to wait for the command to finish. If not completed when the timeout expires, the SSH Client disconnects.
SSH Port	TCP port for SSH connections

Commands

- `cifs show clients`
- `disk show hardware`
- `disk show state`
- `enclosure show all`
- `filesys show compression summary`
- `filesys show space`
- `mtree list`
- `net show all`
- `nfs show clients`
- `quota show all`

- `storage show summary`
- `storage show tier active`
- `storage show tier archive`
- `system show all`
- `system show ports`
- `system show serialno`
- `system show version`

EMC ECC StorageScope data sources

This EMC device has three types of data sources: 5.x, 6.0, and 6.1.

Configuration

Field	Description
ECC server	Network address for the Enterprise Control Center (ECC) server, specified as an IP (*nnn.nnn.nnn.nnn* format), or as a hostname that can be resolved through DNS.
User Name	Name used to log into the ECC.
Password	Password used to log into the ECC.

Advanced configuration

Field	Description
ECC Port	Number of the ECC server port through which the data source acquires information. 1521=default. NOTE: If you are using version 6.1 of the ECC, the default port is now 1575.
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. The default is 30.
Protocol to Connect to Database	TCP is cleartext communication and TCPS is encrypted communication.
Query File System Information	Select to retrieve details for WWN Aliases and File Systems.

EMC XtremIO (HTTP) data source

To configure the EMC XtremIO data source, you must have the XMS Host address and an account with administrator privileges.

Configuration

Field	Description
XMS Host	IP address of FQDN of the XtremIO management server
User name	User name. The default is admin.
Password	Password used to log in to the machine.

Advanced configuration

Field	Description
TCP port	443
Inventory poll interval (min)	60
Connection timeout (sec)	60
Performance poll interval(sec)	300

EMC Isilon data source

The Isilon SSH data source collects inventory and performance from EMC Isilon scale-out NAS storage.

Configuration

Field	Description
IP address	IP address of the Isilon cluster
User name	User name you use to log in to the Isilon cluster
Password	Password you use to log in to the Isilon cluster

General requirements

- Root permissions to the Isilon storage
- Validated access by using `telnet` to port 22

Note: No non-NetApp software is required.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device for inventory information
Performance Poll Interval (sec)	For Array Performance, Interval, in seconds, at which the data source should poll the device for performance information
SSH Process Wait Timeout	Time, in seconds, to wait for the command to finish If not completed when the timeout expires, the SSH Client disconnects.
SSH Port	TCP port for SSH connections

Commands

- `isi devices -d {device id}`
- `isi license status`
- `isi nfs exports list -v`
- `isi quota quotas list --verbose --format=csv`
- `isi smb share list -v`
- `isi snapshot list`

- `isi snapshot settings view`
- `isi snapshot snapshots list --verbose --format=table`
- `isi status -n`
- `isi status -q`
- `isi version osrelease`
- `python -c "exec(\"import isi.app.lib.cluster as cluster;chassis = cluster.Cluster().device({node id}).chassis();print chassis[0][4]\")"`
- `python -c "exec(\"import isi.app.lib.cluster as cluster;device = cluster.Cluster().device({node id});print device.ipaddresses()\")"`
- `python -c "import isi.config.arrcfg as arrcfg; print arrcfg.config_load().guid;"`

EMC Solutions Enabler (CLI) with SMI-S Performance data source

OnCommand Insight discovers Symmetrix storage arrays by using Solutions Enabler `symcli` commands in conjunction with an existing Solutions Enabler server in your environment. The existing Solutions Enabler server has connectivity to the Symmetrix storage array through access to gatekeeper volumes. Administrator permissions are required to access this device.

Prerequisites

Before configuring this data source, you should ensure that the OnCommand Insight server has TCP connectivity to port 2707 on the existing Solutions Enabler server. OnCommand Insight discovers all the Symmetrix arrays that are "Local" to this server, as seen in "symcfg list" output from that server.

- On acquiring the OnCommand Insight host:
  - The EMC Solutions Enabler (CLI) with SMI-S provider application must be installed and the version must match or be earlier than the version running on the Solutions Enabler Server.
  - A properly configured `{installdir}\EMC\SYMAPI\config\netcnfg` file is required. This file defines service names for Solutions Enabler servers, as well as the access method (SECURE / NOSECURE /ANY).
- If you require read/write latency at the storage node level, the SMI-S Provider must communicate with a running instance of the UNISPHERE for VMAX application.

Note: If SMI-S user authentication is not enabled, the default values in the OnCommand Insight data source are ignored.

Having symauth enabled on Symmetrix arrays might inhibit the ability of OnCommand Insight to discover them. OnCommand Insight Acquisition runs as the SYSTEM user on the OnCommand Insight / Remote Acquisition Unit server that is communicating with the Solutions Enabler server. If hostname\SYSTEM does not have symauth privileges, OnCommand Insight fails to discover.

Note: See Knowledge Base article 2015297 on the NetApp Support Site at *mysupport.netapp.com*.

Configuration

The EMC Solutions Enabler Symmetrix CLI data source includes support for device configuration for thin provisioning and Symmetrix Remote Data Facility (SRDF).

Definitions are supplied for Fibre Channel and Switch Performance packages.

Field	Description
Service Name	Name of the Solution Enabler server as described in the local netcnfg file (Solutions Enabler install).
Full path to CLI	Full path to the SYMCLI commands. The default is `C:\Program Files\EMC\SYMCLI\bin`.

Requirements

- An IP address of the Solutions Enabler (SE) Server
- Administrator permissions
- Read-only user name and password to the SE software
- Solutions Enabler Server 6.5X requirements:
  - SMI-S provider 3.3.1 for SMIS-S V1.2 installed
  - After install, run `\Program Files\EMC\SYMCLI\bin>stordaemon start storsrvd`
- Access validation: Verify that the SMI-S provider is running: `telnet <se_server> 5988`
- The UNISPHERE for VMAX application is running and collecting statistic for the Symmetrix VMAX storage arrays that are managed by the SMI-S Provider installation

Note: No third-party software is required.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. The default is 40 minutes.
Inventory Exclude Devices	Comma-separated list of devices, identified by serial number, that are not discovered by this data source. For this data source, an example would be Sym-000187461207.

Field	Description
Connection Caching	Type of remote connection to the Solutions Enabler. Could be LOCAL, REMOTE_CACHED, or REMOTE: • LOCAL means that the OnCommand Insight Acquisition service is running on the Solutions Enabler server, which has Fibre Channel connectivity to the Symmetrix arrays you seek to discover, and has access to Gatekeeper volumes. This might be seen in some Remote Acquisition Unit (RAU) configurations. • REMOTE_CACHED is the default and should be used in most cases. This uses the NETCNFG file settings to connect using IP to the Solutions Enabler server, which must have Fibre Channel connectivity to the Symmetrix arrays you seek to discover, and has access to Gatekeeper volumes. • Sometimes REMOTE_CACHED options make CLI commands fail. In these cases, the REMOTE option should be used. Keep in mind that it will slow down the acquisition process (possibly to hours or even days in extreme cases). The NETCNFG file settings are still used for an IP connection to the Solutions Enabler server that has Fibre Channel connectivity to the Symmetrix arrays being discovered. **Note:** This setting does not change OnCommand Insight behavior with respect to the arrays listed as REMOTE by the "symcfg list" output. OnCommand Insight gathers data only on devices shown as LOCAL by this command.
CLI Timeout (sec)	The default is 7200.
SMI-S Host IP	For Array Performance, host name or IP address of the EMC SMI-S provider host.
SMI-S Port	For Array Performance, the SMI-S port. The default is 5988.
Protocol	For Array Performance, HTTP (if unsecured access) or HTTPS (if secured access).
SMI-S Namespace	For Array Performance, SMI-S namespace used to query the data.
SMI-S User Name	For Array Performance, name used for authentication by the SMI-S host for performance data.
SMI-S Password	For Array Performance, password used for authentication by the SMI-S host for performance data.
Performance Polling Interval (sec)	For Array Performance, the recommended setting and default for the poll interval is 1000 seconds.
Performance Filter Type	For Array Performance, set to exclude or include for acquisition by selecting the filter type as "exclude" or "include." The default is exclude.

Field	Description
Performance Filter Device List	For Array Performance, the comma-separated list of Solutions Enabler IP addresses to be included or excluded, depending upon the filter type selected. For this data source, an example would be `Sym-000187461207`.
RPO Polling Interval (sec)	For Replication, the interval, in seconds, at which the data source should poll the device. The default is 300.

EMC VNX (SSH) data source

For configuration, this data source requires the IP address of the Control Station and a *read-only* username and password.

Configuration

Field	Description
VNX IP	Celerra IP address
VNX User Name	Name used to log into the VNX device.
VNX Password	Password used to log into the VNX device.

Requirements

- An IP address of the Control Station
- Read-only username and password.
- Access validation: Verify SSH access via PuTTY.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. The default is 40.
VNX SSH Process Wait Timeout (sec)	Timeout for obtaining the first prompt once connected to the device. The default is 600.
Celerra Command Retry Attempts	Maximum number of retry attempts for a given Celerra command.
CLARiiON External Process Timeout for Inventory (sec)	Maximum time allowed for CLARiiON commands to complete.
Performance Poll Interval (sec)	For Array Performance, the interval, in seconds, at which the data source should poll the device.
CLARiiON External Process Timeout for Performance (sec)	For Array Performance, the maximum time allowed for CLARiiON commands to complete.

Commands

This data source runs the same commands as the EMC Celerra and EMC CLARiiON.

- /nas/opt/Navisphere/bin/navicli -h {ip} getall

- /nas/opt/Navisphere/bin/navicli -h {ip} metalun -list
- /nas/opt/Navisphere/bin/navicli -h {ip} mirror async –list
- /nas/opt/Navisphere/bin/navicli -h {ip} mirror sync -list
- /nas/opt/Navisphere/bin/navicli -h {ip} storagegroup -list -gname <storageGroupName>
- /nas/opt/Navisphere/bin/navicli -h {ip} storagepool -list
- /nas/opt/Navisphere/bin/navicli -h {ip} thinlun -list
- /nas/opt/Navisphere/bin/navicli -help
- /nas/sbin/t2tty -c <dataMoverId> "fcp topology"
- cat /etc/hosts
- fs_dedupe -info -all
- nas_disk -info id=<disk id>
- nas_disk -list
- nas_fs -info -size -all
- nas_pool -info -all
- nas_pool -size -all
- nas_quotas -list -tree -fs <file system name>
- nas_quotas -report -tree -fs <file system name>
- nas_storage -info <storageName>
- nas_storage -list
- nas_version -l
- nas_volume -info -size -all
- nas_xml -info:ALL -level:3

EMC VPLEX data source

For configuration, this data source requires an IP address of the VPLEX server and a *read-only* username and password to the VPLEX server.

Configuration

Field	Description
IP address of VPLEX Management Console	Host name or IP address of the VPLEX.
User Name	Name used for authentication.
Password	Password used for authentication.

Requirements

- An IP address of the VPLEX server
- Read-only username and password to the VPLEX server

- Validate access: Verify by using `telnet` to port 443. For a port other than the default port, with any browser use `HTTPS://<ip>:<port>`

Note: No third-party software is required.

Advanced configuration

Field	Description
Communication Port	The default is 443.
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. Recommended and default setting is 20 minutes.
Connection timeout (sec)	Time in seconds waiting for command to finish. If not completed when timeout expires, the run fails. The default is 60 seconds.
Number of Retries	Number of times a request is going to be sent after a connection timeout.

Commands

- ls -laf /clusters/*
- ls -laf /clusters/*
- ls -laf /devices/*
- ls -laf /engines/engine-*/directors/director-*
- ls -laf /distributed-storage/distributed-devices/*
- ls -laf /engines/engine-*
- ls -laf /clusters/*/storage-elements/extents/*
- ls -laf /clusters/*/exports/initiator-ports/*
- ls -laf /engines/engine-*/directors/director-*/hardware/ports/*
- ls -laf /clusters/*/storage-elements/storage-arrays/*
- ls -laf /clusters/*/exports/storage-views/*
- ls -laf /clusters/*/storage-elements/storage-volumes/*
- ls -laf /clusters/*/virtual-volumes
- version -a --verbose

Fujitsu Eternus data source

For configuration, this data source requires the IP address of the storage. It cannot be comma delimited.

Configuration

Field	Description
IP Address of Eternus Storage	IP address of the storage

Field	Description
User Name	The administration-level name used to log in to the Fujitsu device
Password	Password used to log in to the device

Requirements

- An IP address of the storage, which cannot be comma delimited
- Administration-level user name and password
- Access validation by using SSH to the configured IP and logging in
- Ensure that the page scroll is disabled.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the foundation package.
SSH Process Wait Timeout (sec)	Time, in seconds, to wait for the command to finish. If not completed when the timeout expires, the SSH Client disconnects.

Commands

- `show affinity-groups`
- `show affinity-groups -ag-number {affinity group id}`
- `show disks -disks all`
- `show enclosure-status`
- `show fru-ce -type {controller module name}`
- `show hardware-information`
- `show host-affinity -ag-number {affinity group id}`
- `show host-groups -all`
- `show mapping`
- `show network`
- `show port-groups -all`
- `show raid-groups`
- `show raid-groups -rg-number {raid group id}`
- `show snap-data-volume`
- `show thin-pro-pools`
- `show thin-pro-pools -pool-number {thin provisioning pool id}`
- `show volumes -mode expansion`

- `show volumes -type tpv`

HDS HiCommand Device Manager data source

The HDS data source supports the HiCommand Device Manager server. OnCommand Insight communicates with the HiCommand Device Manager server using the standard HiCommand API. This data source requires the IP address of the HiCommand Device Manager server and a *read-only* username and password for the HiCommand Device Manager software with peer privileges.

Configuration

Field	Description
HiCommand Server	Network address for the HiCommand Device Manager server, specified as an IP (*nnn.nnn.nnn.nnn* format), or as a host name that can be resolved through DNS.
User Name	Name used to log in to the HiCommand.
Password	Password used to log into the HiCommand.
Devices - USP, USP V, VSP/R600 Storages	Adds a Data Source Property Group information.
WMS/SMS/AMS Storages	Click **Add** to enter Data Source Property Group information.

Note: In HDS USP, USP V, VSP, any disk can belong to more than one array group.

Requirements

- IP address of the HiCommand Device Manager server
- Read-only user name and password for the HiCommand Device Manager software and peer privileges
- Validate access:
  - Log in to the HiCommand Device Manager software using peer user name and password.
  - Verify access to the HiCommand Device Manager API: `telnet <HiCommand Device_Manager_server_ip> 2001`

Note: No third-party software is required.

Advanced configuration

Field	Description
HiCommand Server Port	2001 Number of the HiCommand server port through which the data source acquires information.
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. The default is 40 minutes.
Choose 'Exclude' or 'Include' to specify a list	Comma-separated list of devices, identified by device ID, that are not discovered by this data source. For this data source, an example would be ARRAY.HDS9570V.65010524.

Field	Description
Exclude or Include Devices	Lists the devices to include or exclude depending on your previous selection.
Query Host Manager	Provides host information on aliases from HiCommand.
HTTP Timeout (sec)	The timeout values are used for the connect and read processes. A value of 0 indicates that there is no timeout. The default value is 0.
Performance Polling Interval (sec)	For Array Performance, the recommended setting and default for the poll interval is 300 seconds.
Report External Volumes	For Array Performance, this is checked to enable performance data to appear for backend LUNs on performance-related views in the client, if the Hitachi Universal Storage Platform Virtualization layer (USP V) is used.

Commands

- `<get target = "StorageArray">`
- `<get target = "StorageArray">`
  Components include logical unit, path, WWN, WWN group, host information, path, volume connection, port controller, array group, port, LUN group, and replication information.
- `<get target = "Host">`

Performance requirements

- HDS USP, USP V, and VSP performance
  - Performance Monitor must be licensed.
  - Monitoring switch must be enabled.
  - The Export Tool (`Export.exe`) must be copied to the OnCommand Insight Server.
  - The Export Tool version must match the microcode version of the target array.
- HDS AMS performance
  - Performance Monitor needs to be licensed.
  - The Storage Navigator Modular 2 (SNM2) CLI utility needs to be installed on the OnCommand Insight Server.
  - You must register all AMS, WMS, SMS storage arrays whose performance needs to be acquired by OnCommand Insight by using the following command:
    `auunitaddauto.exe -ip <IP address of Controller0> IP address of Controller1>`
  - You must ensure that all the arrays that you registered are listed in the output of this command: `auunitref.exe`.

HDS HiCommand Device Manager Lite data source

The HDS data source supports the HiCommand Device Manager server. OnCommand Insight communicates with the HiCommand Device Manager server using the standard HiCommand API.

Configuration

Field	Description
HiCommand Server	Network address for the HiCommand Device Manager server, specified as an IP (*nnn.nnn.nnn.nnn* format), or as a host name that can be resolved through DNS.
User Name	Name used to log in to the HiCommand API.
Password	Password used to log in to the HiCommand API.
Devices - USP, USP V, VSP/R600 Storages	Adds the Data Source Property Group information.
WMS/SMS/AMS Storages	Adds the Data Source Property Group information.

Advanced configuration

Field	Description
HiCommand Server Port	Number of the HiCommand server port through which the data source acquires information.
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device.
Choose 'Exclude' or 'Include' to specify a list	Comma-separated list of devices, identified by device ID, that are not discovered by this data source. For this data source, an example would be ARRAY.HDS9570V.65010524.
Exclude or Include Devices	Lists the devices to include or exclude, depending on your previous selection.
Query Host Manager	Provides host information about aliases from HiCommand.
HTTP Timeout (sec)	The default timeout is 0 seconds.
Performance Polling Interval (sec)	For Array Performance, the recommended setting and default for the poll interval is 300 seconds.

HP CommandView AE data source

This data source supports the HiCommand Device Manager server. OnCommand Insight communicates with the HiCommand Device Manager server using the standard HiCommand API. The HP CommandView AE is also referred to as *HDS HiCommand*.

Configuration

Field	Description
HiCommand Server	Network address for the HiCommand Device Manager server, specified as an IP (*nnn.nnn.nnn.nnn* format), or as a host name that can be resolved through DNS.

Field	Description
User Name	Name used to log in to the HiCommand Device Manager server.
Password	Password used to log in to the HiCommand Device Manager server.
Devices - USP, USP V, VSP/ R600 Storages	For Array Performance, click **Add** to enter Data Source Property Group information.
SNM2 Devices - WMS/SMS/AMS Storages	For Array Performance, click **Add** to enter Data Source Property Group information.

Requirements

- IP address of the Command View (AE) server
- Read-only user name and password to the CVAE software
- Access validation:
  - Log in to the CVAE software using a read-only user name and password.
  - Verify access to the CVAE API by entering `telnet <cvae_server_ip> 2001`

Note: No third-party software needs to be installed.

Advanced configuration

Field	Description
HiCommand Server Port	2001 Number of the HiCommand server port through which the data source acquires information.
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device.
Exclude Devices	Comma-separated list of devices, identified by device ID, that is not discovered by this data source. For this data source, an example would be ARRAY.HDS9570V.65010524.
Query Host Manager	Provides host information about aliases from HiCommand.
HTTP Timeout	The default value is 0, which indicates that there is no timeout.
Performance Polling Interval (sec)	For Array Performance, the recommended setting and default for the poll interval is 300 seconds.
Report External Volumes	For Array Performance, check this to enable performance data to appear for backend LUNs on performance-related views in the client, if Hitachi Universal Storage Platform Virtualization layer (USP V) is used.

HP CommandView AE Lite data source

The HP CommandView AE Lite data source supports the HP CommandView AE Lite edition, the CLI, and the SMI-S-only version of HP CommandView AE.

Configuration

Field	Description
HiCommand Server	Network address for the HiCommand Device Manager server, specified as either an IP (*nnn.nnn.nnn.nnn* format) or a host name that can be resolved through DNS
HiCommand Device Manager port	The HiCommand server port through which the CommandView AE Lite data source acquires information The default port is 2001.
User Name	Name used to log in to the HiCommand server
Password	Password used to log in to the HiCommand server
Devices - USP, USP V, VSP/ R600 Storages	For Array Performance, you should click **Add** to enter Data Source Property Group information
SNM2 Devices - WMS/SMS/AMS Storages	For Array Performance, you should click **Add** to enter Data Source Property Group information

Requirements

- IP address of the Command View (AE) server
- This version of the device manager has only the CLI licensed
- Read-only user name and password to the CVAE software
- Access validation: Log in to the CVAE software using the read-only user name and password

Note: You do not have to install any third-party software.

Advanced configuration

Query Host Manager	Description
HiCommand Server Port	The HiCommand server port through which the CommandView AE Lite data source acquires information The default port is 2001.
Inventory Poll Interval (min)	Interval (in minutes) at which the CommandView AE Lite data source should poll the device The default setting is 40 minutes.
Exclude Devices	Comma-separated list of devices—identified by device ID—that are not discovered by the CommandView AE Lite data source
Query Host Manager	Provides host information about aliases from the HiCommand server
HTTP Timeout (sec)	The default value is 0, which indicates that there is no timeout

Query Host Manager	Description
Performance Polling Interval (sec)	For Array Performance, the recommended setting for the poll interval is 300 seconds, which is the default setting

HP EVA Storage (SSSU) data source

For configuration, this data source requires the IP address of the Command View (CV) server and a *read-only* username and password to the CV software. The user must be defined in the CV software.

Configuration

Field	Description
CommandView Server	Name of the Command View (CV) manager.
User Name	Name used to log into the Command View manager. The name must be defined in Command View.
Password	Password used to log into the Command View.
Performance User Name	For Array Performance, the name used to log into the Command View manager. The name must be defined in Command View.
Performance Password	For Array Performance, the password used to log into the Command View.

Requirements

- IP address of the CV server
- Read-only username and password to the CV software. The user must be defined in the CV software.
- Third-party software installed on the OnCommand Insight Server/RAU: `sssu.exe`. The `sssu.exe` version should correspond to the CV version.
- Access validation: Run `sssu.exe` commands using username and password.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device.
CLI Home	Full pathname to the CLI home directory where sssu.exe is located. To validate access, run sssu.exe.
Inventory Exclude Devices	Comma-separated list of devices, identified by device ID, that will not be discovered by this data source for inventory.
Performance Poll Interval (sec)	For Array Performance, the interval, in minutes, at which the data source should poll the device for performance data.
Performance CLI Home	For Array Performance, full pathname to the CLI home directory where sssu.exe is located. To validate access, run `sssu.exe`.

Field	Description
Command Timeout (sec)	For Array Performance, the period of time that data source waits for the command to execute before it declares that command execution has failed (timed out). The command is a CLI source issues to the NSP administrating server (console). The command that the data default is 600 seconds.
Performance Exclude Devices	For Array Performance, comma-separated list of devices, identified by device ID, that will not be discovered by this data source.

Commands

- `LS CONTROLLER FULL XML`
- `LS DISK FULL XML`
- `LS DR_GROUP FULL XML`
- `LS HOST FULL XML`
- `LS SYSTEM FULL XML`
- `LS VDISK FULL XML`
- `SHOW VDISK FULL`

Performance requirements

The HP StorageWorks Command View EVA software suite must be installed on the OnCommand Insight Server, or install a Remote Acquisition Unit (RAU) on the EVA server.

1. Install HP StorageWorks Command View EVA Software Suite on the OnCommand Insight Server, or install a Remote Acquisition Unit on the Command View EVA server.
2. Locate the `evaperf.exe` command. For example, `c:\Program Files\Hewlett-Packard \EVA Performance Monitor\`
3. Using the IP of the Command View server, perform these steps:
   a. Run this command where 860 is the default port

      `Evaperf.exe server <Command View Server IP> 860 <username>`

   b. Enter the Command View server password at the password prompt.
      This should return a command line prompt and nothing else.
4. Verify the setup by running `evaperf.exe ls`.
   You should see a list of arrays or controllers managed by the Command View server. Each line shows a controller on an EVA array.

Performance commands

- **`Evaperf.exe all -csv -nfn -KB -cont <sampleDuration> -dur <sampleDuration>`**

IBM DS (CLI) data source

The IBM DS CLI data source supports DS6xxx and DS8xxx devices only. DS3xxx, DS4xxx, and DS5xxx devices are supported by the NetApp E-Series data source. You should refer to the Insight data source support matrix for supported models and firmware versions.

Configuration

Field	Description
DS storage	Network address for the DS device, specified as an IP (*nnn.nnn.nnn.nnn* format), or as a hostname that can be resolved through DNS.
User Name	Name used to log into the DS.
Password	Password used to log into the DS.
Executable dscli.exe Path	Full path to the dscli utility.

Requirements

- IP address of each DS array
- Storage Display Name is optional and cosmetic only
- Read-only username and password on each DS array
- Third-party software installed on the Insight server: IBM dscli
- Access validation: Run `dscli` commands using the username and password

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. The default is 40.
Storage Display Name	User-specified alias used as the storage name.
Inventory Exclude Devices	Comma-separated list of devices, identified by serial number, that will not be discovered by this data source.
Performance Poll Interval (sec)	For Array Performance, the recommended setting and default for the poll interval is 300 seconds.
Performance Filter Type	For Array Performance, can be either excluded or included for acquisition by selecting the filter type as "exclude" or "include."
Performance Filter Device List	For Array Performance, the comma-separated list of Solutions Enabler IP addresses to be included or excluded, depending upon the filter type selected.

Commands

- `dscli.exe -hmc <ip address> -user <user> -passwd <password> lsarray -v off -hdr off -fmt xml -dev IBM.XXX`
- `dscli.exe -hmc <ip address> -user <user> -passwd <password> lsddm -l -v off -hdr off -fmt xml IBM.XXX`

- `dscli.exe -hmc <ip address> -user <user> -passwd <password> lsfbvol -v off -hdr off -fmt xml -dev IBM.XXX`
- `dscli.exe -hmc <ip address> -user <user> -passwd <password> lshostconnect -v off -hdr off -fmt xml -dev IBM.XXX`
- `dscli.exe -hmc <ip address> -user <user> -passwd <password> lsioport -v off -hdr off -fmt xml -dev IBM.XXX`
- `dscli.exe -hmc <ip address> -user <user> -passwd <password> lsrank -v off -hdr off -fmt xml -dev IBM.XXX`
- `dscli.exe -hmc <ip address> -user <user> -passwd <password> lssi -v off -hdr off -fmt xml`
- `dscli.exe -hmc <ip address> -user <user> -passwd <password> lsvolgrp -v off -hdr off -fmt xml -dev IBM.XXX`
- `dscli.exe -hmc <ip address> -user <user> -passwd <password> -script`

Performance commands

- `dscli.exe -hmc <ip address> -user <user> -passwd <password> -script "diskdata"`
- `dscli.exe -hmc <ip address> -user <user> -passwd <password> -script "RankPerfScript.txt"`
- `dscli.exe -hmc <ip address> -user <user> -passwd <password> -script "VolPerfScript.txt"`
- `dscli.exe -hmc <ip address> -user <user> -passwd <password> -script "volumedata.txt"`

IBM PowerVM data source configuration and commands

The IBM PowerVM (SSH) data source collects information about virtual partitions running on IBM POWER hardware instances managed by a hardware management console (HMC). For configuration, this data source requires the user name, to log in to the HMC through SSH, and the view-level permission on HMC configurations.

Configuration

Field	Description
Hardware Management Console (HMC) Address	Network IP address for the management port of HMC, or a host name that can be resolved through DNS
HMC User	The user name to use to log in to the HMC through SSH The configured user must have view-level permission on all management system and logical partition security domains. The user must also have view permission on HMC configurations and the ability to collect VPD information for the HMC console security grouping. Finally, the user must have Virtual IO Server Command access under the Logical Partition security grouping.
Password	Password of the HMC user

Requirements

- User name and password that provide access to HMC through SSH
- User View permission on all management systems and logical partition security domains
  The user must also have View permission on HMC configurations and the ability to collect VPD information for the HMC console security grouping. The user must also be allowed Virtual IO Server Command access under the Logical Partition security grouping. It is a best practice to start from a role of an operator and then remove all roles. Read-only users on the HMC do not have privileges to run proxied commands on AIX hosts.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the HMC; defaults to 20 minutes
SSH Port	TCP port for SSH connections
SSH Process Wait Timeout (sec)	Timeout for any one SSH command, in seconds
Number of Retries	Number of times an SSH request is going to be sent after a timeout
Exclude Devices	List of managed POWER hardware systems to skip when retrieving data, specified as either the full system ID or as system display name strings Multiple values should be comma separated.

Commands

Some commands are run multiple times and with different values, depending on the values returned from previous commands.

- `lshmc -V -F --header`
- `lssyscfg -r sys -F --header`
- `lshwres -r mem -m {managed system id} --level sys -F --header`
- `lshwres -r proc -m {managed system id} --level sys -F --header`
- `lshwres -r io --rsubtype slot -m {managed system id} -F --header`
- `lspartition -dlpar | grep Partition`
- `lssyscfg -r lpar -m {managed system id} -F --header`
- `lshwres -r mem --level lpar -m {managed system id} -F --header`
- `lshwres -r proc --level lpar -m {managed system id} -F --header`
- `viosvrcmd -m {managed system id} --id {virtual IO server id} -c "ioslevel"`
- `viosvrcmd -m {managed system id} --id {virtual IO server id} -c "lsmap -all"`
- `viosvrcmd -m {managed system id} --id {virtual IO server id} -c "lsdev -dev fcs* -vpd"`
- `viosvrcmd -m {managed system id} --id {virtual IO server id} -c "lsdev"`

-
- `viosvrcmd -m {managed system id} --id {virtual IO server id} -c "lsdev -dev {physical volume name} -vpd"`
- `viosvrcmd -m {managed system id} --id {virtual IO server id} -c "lsvg"`
- `viosvrcmd -m {managed system id} --id {virtual IO server id} -c "lsvg {volume group name} -fmt ,"`
- `viosvrcmd -m {managed system id} --id {virtual IO server id} -c "lsvg -lv {volume group name} -fmt ,"`
- `viosvrcmd -m {managed system id} --id {virtual IO server id} -c "lsdev -dev {physical volume id} -attr"`
- `viosvrcmd -m {managed system id} --id {virtual IO server id} -c "lsdev -vpd"`
- `viosvrcmd -m {managed system id} --id {virtual IO server id} -c "lspv -size -fmt ,"`

IBM SVC (SSH) data source

The IBM SVC data source collects inventory and performance data using SSH. It supports a variety of devices that run the SVC operating system. The list of supported devices includes models such as the SVC, the v7000, the v5000, and the v3700. Refer to the Insight data source support matrix for supported models and firmware versions.

Configuration

Field	Description
Cluster/s IP	Network address for the cluster/s, specified as an IP (*nnn.nnn.nnn.nnn* format), or as a hostname that can be resolved through DNS
User Name	Name used to SSH into the device
Performance User Name	The name used to log in (for array performance)

Requirements

- IP address of each SVC cluster
- Public and private key pair that you either generate withInsight or reuse a keypair already in use on your SVC
  If you are reusing an existing keypair, you must convert them from Putty format to OpenSSH format.
- Public key installed on the SVC cluster
- Private key needs to be identified in the data source
- Access validation: Open `ssh` session to the SVC cluster using the private key

Note: No third-party software needs to be installed.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. The default is 40.
Exclude Devices	Comma-separated list of devices, identified by serial number, that will not be discovered by this data source.
Full Path to Private Key	The best practice is to store private key outside of the .. \sanscreen path to avoid losing the key during an upgrade.
SSH Process Wait Timeout (sec)	Time in seconds waiting for command to finish. If not completed when timeout expires, the run fails. The default is 10.
Performance Poll Interval (sec)	For Array Performance, the recommended setting and default for the poll interval is 300 seconds.
Performance Exclude Devices	For Array Performance, comma-separated list of devices, identified by device ID, that will not supply performance data.
Performance Full Path to Private Key	For Array Performance, the full path to the SSH key.
Performance SSH Process Wait Timeout (sec)	For Array Performance, the maximum time allowed for commands to complete.
To clean up dumped stats files	For Array Performance, remove dump files copied to the primary node during performance collection.

Commands

- `svcconfig backup -quiet`
- `svcinfo lsportfc -nohdr -delim :`
- Retrieves `/tmp/svc.config.backup.xml` via SCP

Performance requirements

- SVC Console, which is mandatory for any SVC cluster and required for the SVC discovery foundation package.
- Administrative access level required only for copying performance data files from cluster nodes to the config node

  Note: Because this access level is not required for the SVC foundation discovery package, the SVC foundation user might not work successfully.

- A private and public SSH key must be generated for this user, and the private key stored so that it is accessible from the Acquisition unit. If the SVC foundation user has the proper permissions, then the same user and key works. The same SSH key can be used for inventory and performance data.
- Enable data collection by connecting to the SVC cluster by SSH and running: `svctask startstats -interval 1`

  Note:

  Alternatively, enable data collection using the SVC management user interface.

IBM Tivoli Monitoring data source

This data source is used solely for File System Utilization. It communicates directly with the Tivoli Monitoring Database, also known as the Tivoli Monitoring Data Warehouse. Oracle and DB2 databases are supported.

Oracle error message

If the specified SID results in the error message containing "ORA-12154" on attempting to connect, double-check your Oracle DB network service configuration. If the access configuration specifies a fully qualified hostname (for example, "NAMES.DEFAULT_DOMAIN"), try inserting the fully qualified service name in the SID field. A simple example would be that the connection to SID `testdb` is failing and your Oracle configuration specifies a domain of `company.com`. The following string can be used instead of the base SID to try and connect: `testdb.company.com`.

Configuration

Field	Description
Tivoli Monitoring Database IP	IP address of server.
User Name	User name for login.
Password	Password for login.

Advanced configuration

Field	Description
Tivoli Monitoring Database Port	50001 should be configured for DB2.
DB2 Database Name	The default is `QTIVOLI`.
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. The default is 60.
Database Driver to Use	For DB2 environments, set it to DB2.
Protocol Used to Connect to the Database	TCP=default (TCPS is the secure option and depends on the database configuration.)
Database Schema	The name of the database schema to use in queries.

IBM TotalStorage DS4000 data source

This data source collects inventory and performance information. There are two possible configurations (firmware 6.x and 7.x+), and they both have the same values. The API collects the volume data statistics.

Configuration

Field	Description
Comma Separated List of Array SANtricity Controller IPs	Network address for the management controller of each array from which this data source will acquire information, specified as an IP (nnn.nnn.nnn.nnn format), or as a hostname that can be resolved through DNS. Separate multiple network addresses using a semicolon.

Requirements

- IP address of each DS5 or FAStT array
- Access validation: Ping the IP address of both controllers on each array.

Note: No third-party software needs to be installed.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device for performance data. The default is 30.
Performance Poll Interval (up to 3600 seconds)	For Array Performance, the recommended setting and default for the poll interval is 300 seconds.

IBM XIV (CLI) data source

IBM XIV data source inventory is performed by using the XIV CLI. XIV performance is accomplished by making SMI-S calls to the XIV array, which runs a SMI-S provider on port 5989.

Configuration

Field	Description
IP Address	IP address for the storage system
User Name	Name used to log in to the storage system
Password	Password used to log in to the storage system
Full path to XIV CLI directory	The default is `C:\Program Files\GUI10`.

Requirements

- Port requirement: TCP port 5989
- IP address of the XIV management interface
- Read-only user name and password
- Third-party software installed on the Insight server or RAU: XIV CLI
- Access validation: Log in to the XIV user interface from the Insight server using the user name and password.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. The default is 40 minutes.
CLI Process Wait Timeout (ms)	The default is 7200000 milliseconds.
SMI-S Host IP	For Array Performance, the host name or IP address of the SMI-S provider host.
SMI-S Port	For Array Performance, the SMI-S port. The default is 5989.

Field	Description
SMI-S Protocol	For Array Performance, HTTP (if unsecured access) or HTTPS (if secured access).
SMI-S Namespace	For Array Performance, the SMI-S namespace used to query the data.
Username	For Array Performance, the name used for authentication by the SMI-S host for performance data.
Password	For Array Performance, the password used for authentication by the SMI-S host for performance data.
Performance Poll Interval (sec)	For Array Performance, the recommended setting and default for the poll interval is 600 seconds.
Number of SMI-S Connection Retries	For Array Performance, the maximum number of connection retry attempts.

Commands

Insight invokes the following commands.

- `xcli -x -m <IP Address> -u <user> -p <password> config_get`
- `xcli -x -m <IP Address> -u <user> -p <password> disk_list`
- `xcli -x -m <IP Address> -u <user> -p <password> fc_port_list`
- `xcli -x -m <IP Address> -u <user> -p <password> fc_connectivity_list`
- `xcli -x -m <IP Address> -u <user> -p <password> host_list`
- `xcli -x -m <IP Address> -u <user> -p <password> ipinterface_list`
- `xcli -x -m <IP Address> -u <user> -p <password> mapping_list host = <hostname>`
- `xcli -x -m <IP Address> -u <user> -p <password> pool_list`
- `xcli -x -m <IP Address> -u <user> -p <password> system_capacity_list`
- `xcli -x -m <IP Address> -u <user> -p <password> version_get`
- `xcli -x -m <IP Address> -u <user> -p <password> vol_list`
- `xcli -x -m<IP Address> -u <user> -p <password> host_connectivity_list`
- `xcli -x -m <IP Address> -u <user> -p <password> pool_list`

Microsoft Hyper-V (WMI) data source

For configuration, the Microsoft Hyper-V (WMI) data source requires the IP address or the resolvable DNS name for the physical host (hypervisor).

Configuration

Field	Description
Physical Host IP Address	The IP address or the resolvable DNS name for the physical host (hypervisor)
User Name	The administrator name that is used to log in to the server

Field	Description
Password	The administrator password that is used to log in to the server
NT Domain	The NT domain for which the physical host (hypervisor) is configured

Requirements

- Administrative-level user account
- Windows Management Instrumentation (WMI) command, which is the default that is installed by Windows

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval (in minutes) at which the data source should poll the device The default setting is 20 minutes.
Connection Timeout (ms)	Time (in milliseconds) to wait for a WMI command to complete The default setting is 60000 milliseconds.

Commands

All commands are run with the following additional arguments: `/user:"user" /password:"password" /node:"IP"`

Cluster-level commands

- `wmic /namespace:\\root\mscluster /authlevel:pktprivacy path MSCluster_Cluster get Name,QuorumPath,SharedVolumesRoot /format:csv`
- `wmic /namespace:\\root\mscluster /authlevel:pktprivacy path MSCluster_ClusterToNode get Antecedent,Dependent`
- `wmic /namespace:\\root\mscluster /authlevel:pktprivacy path MSCluster_Node get Name,State`
- `wmic /namespace:\\root\mscluster /authlevel:pktprivacy path MSCluster_ClusterToClusterSharedVolume get GroupComponent,PartComponent`
- `wmic /namespace:\\root\mscluster /authlevel:pktprivacy path MSCluster_ClusterToQuorumResource get GroupComponent,PartComponent`
- `wmic /namespace:\\root\mscluster /authlevel:pktprivacy path MSCluster_ResourceToDisk get GroupComponent,PartComponent`
- `wmic /namespace:\\root\mscluster /authlevel:pktprivacy path MSCluster_DiskToDiskPartition get GroupComponent,PartComponent`
- `wmic /namespace:\\root\mscluster /authlevel:pktprivacy path MSCluster_ClusterSharedVolumeToPartition get GroupComponent,PartComponent`
- `wmic /namespace:\\root\mscluster /authlevel:pktprivacy path MSCluster_ClusterSharedVolume get FaultState,Name,VolumeName`

- `wmic /namespace:\\root\mscluster /authlevel:pktprivacy path MSCluster_Disk get ID,Name,ScsiBus,ScsiLUN,ScsiPort,ScsiTargetID,Signature`

Physical host-level commands

- `wmic /namespace:\\root\cimv2 path Win32_ComputerSystem get DNSHostName,Domain,Manufacturer,Model,NumberOfLogicalProcessors,SystemType,TotalPhysicalMemory`
- `wmic /namespace:\\root\cimv2 path Win32_OperatingSystem get BuildNumber,Caption,OSArchitecture,ServicePackMajorVersion,ServicePackMinorVersion,Version`
- `wmic /namespace:\\root\cimv2 path Win32_Processor get MaxClockSpeed`
- `wmic /namespace:\\root\cimv2 path Win32_SystemEnclosure get SerialNumber`
- `wmic /namespace:\\root\cimv2 path Win32_NetworkAdapterConfiguration get IPAddress,IPEnabled`
- `wmic /namespace:\\root\cimv2 path Win32_NetworkAdapter get NetConnectionStatus,PhysicalAdapter,Speed`
- `wmic /user:<user> /password:<password> /node:<IP>/namespace:\\root\wmi path MSFC_FCAdapterHBAAttributes get Active,DriverName,DriverVersion,FirmwareVersion,HardwareVersion,HBAStatus,InstanceName,Manufacturer,Model,ModelDescription,NodeWWN,NodeSymbolicName,NumberOfPorts,OptionROMVersion,SerialNumber,VendorSpecificID`
- `wmic /namespace:\\root\cimv2 path Win32_DiskDrive where InterfaceType='SCSI' get DeviceID,Manufacturer,Model,Name,Partitions,PNPDeviceID,SCSIBus,SCSILogicalUnit,SCSIPort,SCSITargetId,SerialNumber,Signature,Size,Status`
- `wmic /namespace:\\root\wmi path MSiSCSI_HBAInformation get InstanceName,UniqueAdapterId`
- `wmic /namespace:\\root\wmi path MSiSCSIInitiator_SessionClass get Devices,InitiatorName,SessionId,TargetName`
- `wmic /namespace:\\root\cimv2 path Win32_LogicalDiskToPartition get Antecedent,Dependent`
- `wmic /namespace:\\root\wmi path MSFC_FibrePortHBAAttributes get Active,Attributes,HBAStatus,InstanceName`
- `wmic /namespace:\\root\cimv2 path Win32_DiskDriveToDiskPartition get Antecedent,Dependent`
- `wmic /namespace:\\root\cimv2 path Win32_Volume get Capacity,DeviceID,DriveLetter,DriveType,FreeSpace,Label,Name`
- `wmic /namespace:\\root\cimv2 path StdRegProv call GetBinaryValue hDefKey="&H80000002" sSubKeyName="SYSTEM\MountedDevices" sValueName="\??\Volume{d71db670-8c30-11e2-b1c9-806e6f6e6963}"`

Virtual machine-level commands

- `wmic /namespace:\\root\virtualization path Msvm_VirtualSystemSettingDataComponent get GroupComponent,PartComponent`

- `wmic /namespace:\\root\virtualization path Msvm_ResourceAllocationSettingData where (ResourceSubType='Microsoft Virtual Hard Disk' and Parent != null or ResourceSubType='Microsoft Physical Disk Drive' and Parent != null) get Connection,HostResource,Parent,ResourceSubType`
- `wmic /namespace:\\root\virtualization path Msvm_ImageManagementService where name='vhdsvc' call GetVirtualHardDiskInfo 'C:\ClusterStorage\Volume1\ds2-client4\ds2-client1.vhd'`

NetApp Clustered Data ONTAP data source

This data source should be used for storage systems using Clustered Data ONTAP, and requires an administrator account used for read-only API calls.

Configuration

Field	Description
NetApp Management IP	IP address of a cluster management LIF.
User Name	Name used to log into the cluster.
Password	Password used to log into the NetApp cluster.

Requirements

- Administrator account used for read-only API calls
- Target IP is the cluster management LIF
- FCP license and mapped/masked volumes required for Insight discovery
- CIFS and/or NFS license and CIFS shares or NFS exports required for Insight NAS discovery
- Access validation: Telnet or SSH to the cluster management LIF from the Insight server using the username and password.

Note: Third-party software is not required.

Advanced configuration

Field	Description
Connection Type	The default is HTTP.
Communication Port	80 for HTTP or 443 for HTTPS. The default is 80.
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device.
Performance Poll Interval (sec)	For Array Performance, the recommended setting and default for the poll interval is 900 seconds.

NetApp Clustered Data ONTAP 8.1.x for Unified Manager data source

This data source collects clustered Data ONTAP 8.1.x data from the Unified Manager (UM) 6.0+ database. Using this data source, Insight discovers all clusters configured and populated in UM. For

efficiency, Insight does not call ZAPIs on the cluster itself. Performance is not supported in this data source.

Configuration

Field	Description
Unified Manager IP	IP address of the Unified Manager. The data source connects to the Unified Manager to collect Clustered Data ONTAP data.
User Name	Unified Manager's database user name.
Password	Unified Manager's database password.
Port	Unified Manager's database port. The default is 3306.

Advanced configuration

Field	Description
Inventory Poll Interval (min) Interval	Interval, in minutes, at which the data source should poll the device.
Exclude Clusters	To exclude the clusters from being collected in the poll. Provide IP addresses of the clusters being managed by the Unified Manager.

NetApp Data ONTAP operating in 7-Mode data source

For storage systems using Data ONTAP software operating in 7-Mode, you should use the ONTAPI data source, which uses the CLI to obtain capacity numbers.

Configuration

Field	Description
Address of Filer	Specifies the IP address.
User Name	Specifies the name used to log in to the device.
Password	Specifies the password used to log in to the NetApp device.
Address of HA Partner Filer in Cluster	If storage systems are configured as failover clusters, add the second storage system's IP address here.
User Name of HA Partner Filer in Cluster	If storage systems are configured as failover clusters, add the user name for the second storage system here.
Password of HA Partner Filer in Cluster	If storage systems are configured as failover clusters, add the second storage system password here.

Requirements

- IP address of the FAS storage controller and partner
- User name and password for the controller and the partner
- FCP license and mapped or masked volumes required for OnCommand Insight discovery
- CIFS or NFS licenses and CIFS shares or NFS exports required for OnCommand Insight NAS discovery

- Access validation
  Use telnet to access the FAS controller from the OnCommand Insight Server using the user name and password.

Note: Third-party software is not required.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Specifies the interval, in minutes, at which the data source should poll the device.
Connection Type	The default is HTTP.
Connection Port	The default is 80.
Performance Poll Interval (sec)	For Array Performance, the recommended setting and default for the poll interval is 300.

Storage systems connection requirements

As an alternative to using the default administrative user for this data source, you can configure a user with administrative rights directly on the NetApp storage systems so that this data source can acquire data from NetApp storage systems.

Connecting to NetApp storage systems requires that the user, who is specified when acquiring the main pfiler (on which the storage system exist), meet the following conditions:

- The user must be on vfiler0 (root filer/pfiler).
  Storage systems are acquired when acquiring the main pfiler.
- The following commands define the user role capabilities:
  - "api-*": Use this to allow OnCommand Insight to execute all NetApp storage API commands. This command is required to use the ZAPI.
  - "login-http-admin": Use this to allow OnCommand Insight to connect to the NetApp storage via HTTP. This command is required to use the ZAPI.
  - "security-api-vfiler": Use this to allow OnCommand Insight to execute NetApp storage API commands to retrieve vFiler unit information.
  - "cli-options": For "options" command and used for partner IP and enabled licenses.
  - " cli-lun": Access these command for managing LUNs. Displays the status (LUN path, size, online/offline state, and shared state) of the given LUN or class of LUNs.
  - "cli-df": For "df -s", "df -r", "df -A -r" commands and used to display free space.
  - "cli-ifconfig": For "ifconfig -a" command and used for getting filer IP address.
  - "cli-rdfile": For "rdfile /etc/netgroup" command and used for getting netgroups.
  - "cli-date": For "date" command and used to get full date for getting Snapshot copies.
  - "cli-snap": For "snap list" command and used for getting Snapshot copies.

If cli-date or cli-snap permissions are not provided, acquisition can finish, but Snapshot copies are not reported.

To acquire a 7-Mode data source successfully and generate no warnings on the storage system, you should use one of the following command strings to define your user roles. The second string listed here is a streamlined version of the first:

```
login-http-admin,api-*,security-api-vfile,cli-rdfile,cli-options,cli-
df,cli-lun,cli-ifconfig,cli-date,cli-snap,
or
login-http-admin,api-*,security-api-vfile,cli-*
```

OnCommand Insight executes the following calls made by the NetApp data source:

- aggr-list-info
- aggr-space-list-info - Get aggregate capacity details
- cifs-share-list-iter-end
- cifs-share-list-iter-next
- cifs-share-list-iter-start
- cifs-status
- cli-df
- cli-ifconfig - Get IP configuration (when net-ifconfig-get not supported)
- cli-lun
- cli-options
- disk-list-info
- fcp-adapter-list-info
- fcp-service-status
- iscsi-initiator-auth-list-info
- iscsi-initiator-get-default-auth
- iscsi-node-get-name
- iscsi-portal-list-info
- iscsi-service-status
- iscsi-session-list-info
- iscsi-tpgroup-list-info
- license-list-info
- lun-list-info
- lun-map-list-info
- net-ifconfig-get - Get IP configuration
- nfs-exportfs-list-rules
- nfs-exportfs-list-rules-2 - NFS exports for later version 7-mode
- nfs-status
- perf-object-counter-list-info

- perf-object-get-instances-iter-end
- perf-object-get-instances-iter-next
- perf-object-get-instances-iter-start
- perf-object-list-info
- portset-list-info
- rdfile
- qtree-list
- quota-report
- security-api-vfiler
- snapmirror-get-status - Get storage synchronization data
- snapmirror-list-connections
- snapmirror-list-destinations
- snapshot-reserve-list-info
- storage-adapter-get-adapter-list - Discover back end and initiator FC ports
- system-cli
- system-get-info
- system-get-ontapi-version
- system-get-vendor-info
- system-get-version
- volume-list-info

NetApp E-Series data source

This data source collects inventory and performance information. There are two possible configurations (firmware 6.x and firmware 7.x+), and they both have the same values.

Configuration

Field	Description
Comma-separated list of Array SANtricity Controller IPs	Network address for the management controller of each array from which this data source acquires information, specified as an IP (*nnn.nnn.nnn.nnn* format) or as a host name that can be resolved through DNS. Separate multiple network addresses using a comma.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	The default is 30.
Performance Poll Interval (up to 3600 seconds)	For Array Performance, the default is 300.

Requirements

- The IP address of each controller on the array
- Access validation by pinging the IP address of both controllers on the array

Note: No third-party software is required.

NetApp Host and VM File Systems data source

You can use the NetApp Host and VM File Systems data source to retrieve file system details and storage resource mappings for all Microsoft Windows host and VM (virtual machine) file systems and for all supported Linux VMs (those that are virtually mapped only) existing in the Insight server that are annotated with the configured Compute Resource Group (CRG).

Configuration

Field	Description
User Name	Operating system user with appropriate rights to retrieve file system data For Windows operating system users, this must include the domain prefix.
Password	Password for the operating system user
Compute Resource Group	Annotation value used to flag host and virtual machines for the data source discovers file systems A blank value indicates that the data source discovers file systems for all hosts and virtual machines not currently annotated with any Compute Resource Group.

General requirements

- This feature must be purchased separately.
  You can contact your Insight representative for assistance.
- You should check the Insight support matrix to verify that your host or virtual machine operating system is supported.
  To verify that links from file systems to storage resources are created, check that the relevant storage or virtualization vendor type and version report the volume or virtual disk identification data required.

Note: No non-NetApp software is required.

Microsoft Windows requirements

- This data source uses Window Management Instrumentation (WMI) data structures to retrieve data.
  This service must be operational and available remotely. In particular, port 135 must be accessible and must be opened if behind a firewall.
- Windows domain users must have the appropriate permissions to access WMI structures.
- Administrator permissions are required.

Linux requirements

- This data source uses a Secure Shell (SSH) connection to execute commands on Linux VMs.

The SSH service must be operational and available remotely. In particular, port 22 must be accessible and must be opened if behind a firewall.

- SSH users must have sudo permissions to execute read-only commands on Linux VMs.
  You must use the same password to log in to SSH and to answer any sudo password challenge.

Advanced configuration

Field	Description
Inventory poll interval (min)	Interval, in minutes, at which the data source should poll the device.

Microsoft Windows commands

- `root\cimv2 SELECT * FROM Win32_BIOS`
- `root\cimv2 SELECT * FROM Win32_ComputerSystem`
- `root\cimv2 SELECT * FROM Win32_OperatingSystem`
- `root\cimv2 SELECT * FROM Win32_NetworkAdapter WHERE AdapterTypeId=0`
- `root\cimv2 SELECT * FROM Win32_DiskDrive`
- `root\cimv2 SELECT * FROM Win32_DiskDriveToDiskPartition`
- `root\cimv2 SELECT * FROM Win32_LogicalDiskToPartition`
- `root\cimv2 SELECT * FROM Win32_LogicalDisk`
- `root\cimv2 SELECT * FROM Win32_Volume`
- `root\default StdRegProv:\HKEY_LOCAL_MACHINE\SYSTEM\MountedDevices`
- `root\default StdRegProv:\HKEY_LOCAL_MACHINE\HARDWARE\DEVICEMAP\Scsi\Scsi Port <port_number>\Scsi Bus <bus_number>\Target Id <target_id>\Logical Unit Id <logical_id>`

The last command runs for each discovered block device, as in the following example:

```
root\default StdRegProv:\HKEY_LOCAL_MACHINE\HARDWARE\DEVICEMAP\Scsi\Scsi
Port 2\Scsi Bus 0\Target Id 0\Logical Unit Id 0
```

```
root\default StdRegProv:\HKEY_LOCAL_MACHINE\HARDWARE\DEVICEMAP\Scsi\Scsi
Port 2\Scsi Bus 0\Target Id 0\Logical Unit Id 1
```

Linux commands

- `/bin/hostname`
- `/bin/domainname`
- `/bin/dnsdomainname`
- `/bin/uname -m`
- `/bin/cat /etc/os-release`
- `/bin/cat /sys/devices/virtual/dmi/id/product_name`
- `/bin/cat /sys/devices/virtual/dmi/id/chassis_vendor`

- `/bin/cat /sys/devices/virtual/dmi/id/sys_vendor`
- `/bin/cat /proc/cpuinfo`
- `/bin/cat /proc/meminfo`
- `/sbin/ifconfig -a`
- `/usr/bin/sudo /sbin/mii-tool <interface_name>`
  This command runs for each interface.
- `/bin/df -T -BK`
- `/bin/ls -l /sys/class/scsi_device`
- `/bin/cat /proc/swaps`
- `/bin/cat /proc/mdstat`
- `/bin/cat /proc/partitions`
- `/bin/ls -l /dev/mapper`
- `/bin/ls -l /dev/block`
- `/bin/ls -l /sys/block`
- `/bin/ls -l /dev/disk/by-id`
- `/bin/ls -l /dev/disk/by-uuid`
- `/bin/ls -l /dev/disk/by-path`
- `/bin/ls -l /dev/disk/by-label`
- `/usr/bin/sudo /sbin/pvdisplay -v`
- `/usr/bin/sudo /sbin/vgdisplay -v`
- `/usr/bin/sudo /sbin/lvdisplay -v`
- `/sbin/udevadm info --query=property --name=sda # for each SCSI device`
  This command is run for each SCSI device.
- `/bin/cat /sys/devices/pci0000:00/0000:00:07.1/host1/target1:0:0/1:0:0:0/vendor`
- `/bin/cat /sys/devices/pci0000:00/0000:00:07.1/host1/target1:0:0/1:0:0:0/model`
  This command is run for each SCSI device.

Usage recommendations

- You should annotate a group of hosts and virtual machines that have common operating system credentials using the same Compute Resource Group annotation.
  Each group has an instance of this data source discovering file system details from those hosts and virtual machines.
- If you have an instance of this data source for which the success rate is low (for example, OnCommand Insight is discovering file system details for only 50 of 1000 hosts and virtual machines in a group), you should move the hosts and virtual machines for which discovery is successful into a separate Compute Resource Group.

QLogic FC Switch (SNMP) data source

For configuration, this data source requires the network address for the FC Switch device, specified as an IP address, and an SNMP *read-only* community string used to access the device.

Configuration

Field	Description
SANSurfer Switch	Network address for the FC Switch (SANsurfer) device, specified as an IP (*nnn.nnn.nnn.nnn* format), or as a hostname that can be resolved through DNS.
SNMP version	The default is V2.
SNMP community	SNMP *read-only* community string used to access the device.
User Name	SNMP name used to log into the switch. Applies when using SNMP V3.
Password	SNMP password used to log into the switch. Applies when using SNMP V3.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device.
SNMP Auth Protocol	Encryption type. The default is MD5.
SNMP Retries	Maximum number of retry attempts for a given command.
SNMP Timeout (ms)	Timeout for connecting to the device using SNMP.
Enable Trapping	Check box used to trigger acquisition upon receiving an SNMP trap from the device. If you mark this box, you must also activate SNMP notifications.
Minimum Time Between Traps (sec)	Minimum elapsed time between acquisition attempts triggered by traps, in seconds.
Fabric Name	Fabric name to be reported by this data source (blank to report the fabric WWN).
Performance Poll Interval (sec)	For Switch Performance, the recommended setting and default for the poll interval is 300 seconds.

Violin Memory 6000 Series Flash Memory Array (HTTP) data source

This data source collects network information for analysis and validation from Violin 6000 series flash memory arrays. To add the data source, you must use specific configuration instructions and commands and understand data source requirements and usage recommendations.

Configuration

Field	Description
IP address or FQDN of Violin Memory Array Master Gateway	IP address or FQDN of the Violin Memory Array Master Gateway

Field	Description
User Name	User name you use to log in to the Violin Memory Array Master Gateway
Password	Password you use to log in to the Violin Memory Array Master Gateway

General requirements

- You need a read-only user name and password to the storage.
- Validate access with a web browser using the storage IP address.

Note: No non-NetApp software is required.

Advanced configuration

Field	Description
Communication Port	The communication port on the storage
HTTPS Enabled	Whether or not HTTPS is enabled on the storage
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device for inventory information
Connection timeout (sec)	Number of seconds Insight should wait before it stops trying to connect to the storage
Performance Poll Interval (sec)	Interval, in seconds, at which the data source should poll the device for performance information

Commands

- `POST http://{ip address}:{port}/admin/launch?script=rh&template=login&action=login HTTP/1.1`
- `GET http://{ip address}:{port}/xtree/system/smbios/system_serial_number/*** HTTP/1.1`
- `GET http://{ip address}:{port}/xtree/system/model/*** HTTP/1.1`
- `GET http://{ip address}:{port}/xtree/system/motd/*** HTTP/1.1`
- `GET http://{ip address}:{port}/xtree/system/num_cpus/*** HTTP/1.1`
- `GET http://{ip address}:{port}/xtree/system/uptime/*** HTTP/1.1`
- `GET http://{ip address}:{port}/xtree/system/version/composite/*** HTTP/1.1`
- `GET http://{ip address}:{port}/xtree/system/version/release/*** HTTP/1.1`
- `GET http://{ip address}:{port}/xtree/cluster/*** HTTP/1.1`
- `GET http://{ip address}:{port}/xtree/vshare/state/local/container/*** HTTP/1.1`
- `GET http://{ip address}:{port}/xtree/media/state/array/*** HTTP/1.1`
- `GET http://{ip address}:{port}/xtree/vshare/config/igroup/*** HTTP/1.1`

- `GET http://{ip address}:{port}/xtree/vshare/config/export/container/*** HTTP/1.1`
- `GET http://{ip address}:{port}/xtree/vshare/state/global/*** HTTP/1.1`

VMware vSphere (Web Services) data source

This data source collects ESX Hosts information and requires *read-only* privileges on all objects within the Virtual Center.

Configuration

Field	Description
Virtual Center Address	Network address for the Virtual Center or vSphere server, specified as an IP *(nnn.nnn.nnn.nnn* format) address or as a host name that can be resolved through DNS.
User Name	Name used to log in to the VMware server. You need an account with read-only privileges within Virtual Center or vSphere.
Password	Password used to log in to the VMware server.

Requirements

- Read-only privileges on all objects within Virtual Center.
- SDK access on the Virtual Center server
- Access validation: Log in to Virtual Center Client using your user name and password. Verify that the SDK is enabled by entering

  `telnet <vc_ip> 443`

  .

Note: Third-party software is not required.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval, in minutes, at which the data source should poll the device. Defaults to 20 minutes for VM Inventory.
Connection Timeout (ms)	The timeout default is 60,000 milliseconds.
Exclude VMs by	A comma-separated list of ESX host or data center IP addresses that should not be polled. This is optional.
Exclusion list (Comma Separated, or Semicolon Separated if Comma Is Used in the Value)	Comma-separated list of devices, identified by a generic device ID, that are not discovered by this data source.
Number of Retries for Requests to vCenter	Maximum number of retry attempts for a given command.
Communication Port	The possible communication ports are http-80 and https-443. The default is 443.

Field	Description
Performance Poll Interval (sec)	For VM Performance, the recommended setting and default for the poll interval is 300 seconds.

Commands

This data source uses the VMware Infrastructure (VI) SDK 2.5.0.

Changing data source credentials

If multiple data sources of the same type are sharing a username and password, you can change the password for all devices in the group at the same time.

Steps

1. On the Insight toolbar, click the **Admin** icon.

   The **Data sources** list opens.

2. Click the **Actions** button and select the **Change credentials** option.

3. In the Credentials Management dialog box, select one of the data source groups from the list.

   The Edit icon, a pen on a sheet of paper, becomes active to the right.

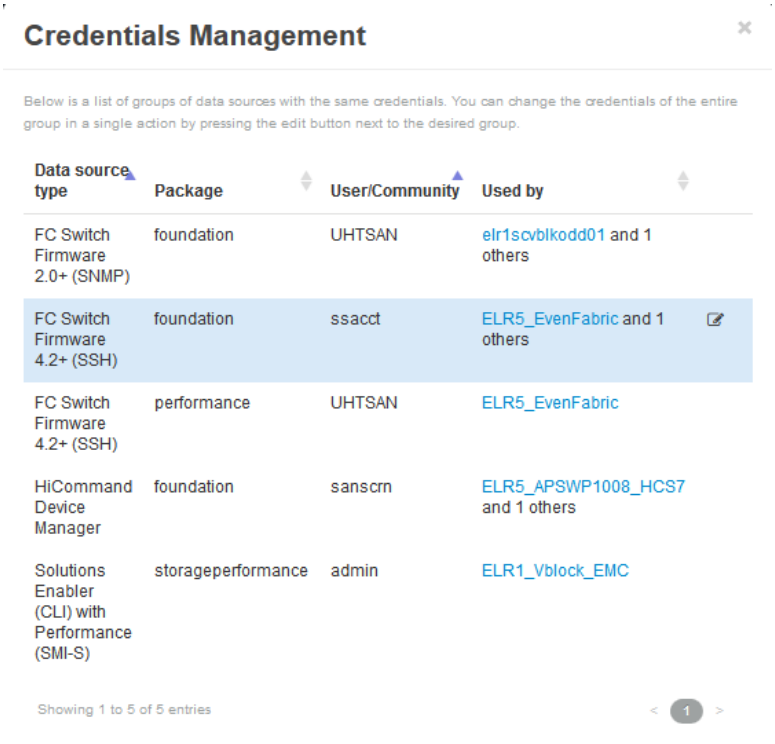

4. Click **Edit**.

5. Enter the new password and confirm it.

Changes causing data collection problems

If you are experiencing data collection problems in OnCommand Insight, changes in your environment are a likely cause. As a general maintenance rule, you should accommodate any changes in your environment in Insight as well.

You can use this checklist to identify changes to your network that might be causing problems:

- Have you changed any passwords? Were those passwords changed in Insight?
- Did you remove a device from your network? You must also remove the device from OnCommand Insight to prevent it from being rediscovered and reintroduced.
- Did you upgrade infrastructure software (such as HP CommandView EVA or EMC Solutions Enabler)?
  Ensure that the appropriate versions of the client tools are installed on the acquisition unit. If data source failures persist, you need to contact technical supportto request assistance and possibly a data source patch.
- Are all of your OnCommand Insight acquisition units using the same OnCommand Insight version? If the Remote Acquisition Units and local acquisition unit are running different OnCommand Insight versions, install the same version on all units to correct the data collection problem.
  If you need to install a new version of OnCommand Insight on all of the acquisition units, go to the support site and download the correct version.
- Have you changed any domain names or added a new domain? You must reconfigure the Auto Resolution (AR) methods in the Java UI.

Related concepts

What data source patches are on page 132

Related tasks

Deleting a data source on page 131

Examining one data source in detail

If you see that a data source has failed or slowed, you might want to examine a detailed summary of information for that data source to determine the cause of the problem. Data sources with conditions requiring your attention are marked with a solid red circle.

Steps

1. On the Insight toolbar, click the **Admin** icon.

   The **Data sources** list opens. Any listed data sources with potential problems are marked with a solid red circle. The most serious problems are at the top of the list.

2. Select the data source that is causing concern.
3. Click the data source name link.
4. On the data source summary page, check the information in any of these sections:
   - **Event timeline**

Lists events tied to the current status shown in the Data sources list. Events in this summary are displayed per device. Errors are shown in red. You can position your mouse pointer on timeline items to display additional information.

- **Devices reported by this data source**
  Lists the types of devices, their IP addresses, and links to more detailed information for each device.
- **Changes reported by this data source (last 3 weeks)**
  Lists any devices that were added or removed or had a change to the configuration.

5. After examining the data source information, you might want to perform one of these operations using the buttons at the top of the page:
   - **Edit** the description of the data source to correct the problem.
   - **Poll again** forces polling to reveal if the problem was persistent or intermittent.
   - **Postpone** data source polling for 1, 3, or 5 days to give you time to research the problem and stop the warning messages.
   - **Install a patch** on the data source to correct the problem.
   - Prepare an **Error report** for technical support.
   - **Delete** the data source from your Insight monitoring environment.

Researching a failed data source

If a data source has the "**Inventory failed !**" or "**Performance failed !**" message and a High or Medium Impact, you need to research this problem using the data source summary page with its linked information.

Steps

1. Click the linked **Name** of the data source to open the Summary page.

   This example shows a wide variety of information that might be provided for a failed data source.

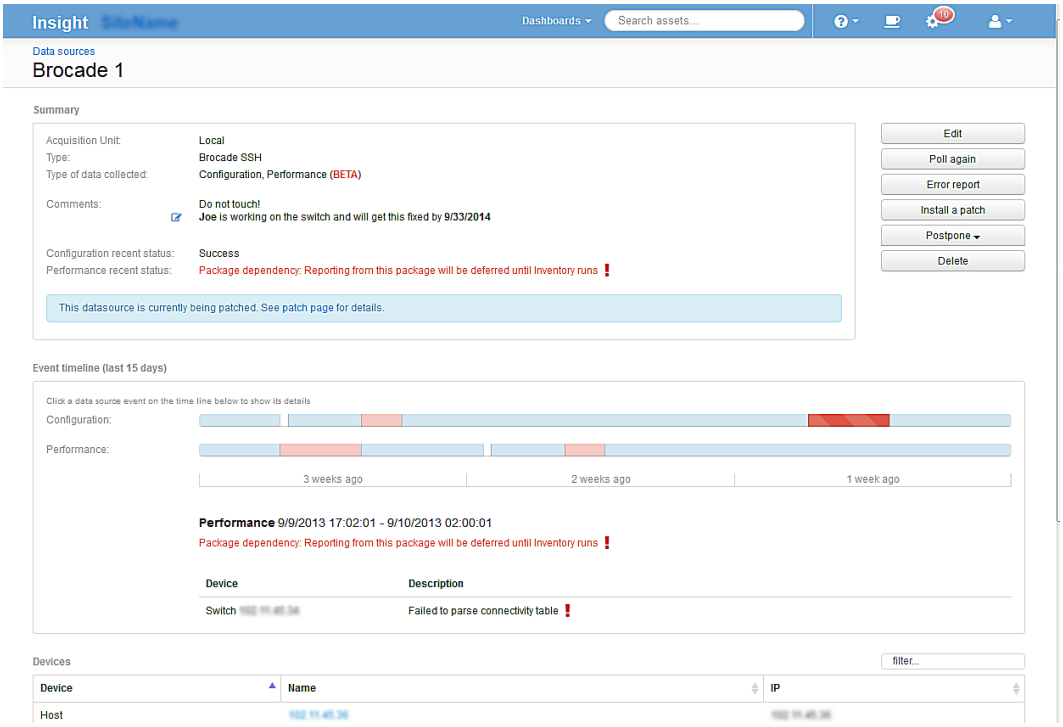

2. Check the **Comments** area to read any notes left by another engineer who might also be investigating this failure.
3. Note any performance messages.
4. If there is a patch being applied to this data source, click link to check the **patch page** to see if that has caused the problem.
5. Move your mouse pointer over the segments of the **Event timeline** graph to display additional information.
6. Select an error message for a Device and displayed below the Event timeline and click the **Error details** icon that displays to the right of the message.

   The Error details include the text of the error message, most likely causes, information in use, and suggestions of what can be tried to correct the problem.
7. In the Devices reported by this data source area, you might filter the list to display only devices of interest, and you can click the linked **Name** of a device to display the *asset page* for that device.
8. To return to previously displayed pages, use one of these techniques:
   - Click the browser back arrow.
   - Right-click the back arrow to display a list of the pages and select the page you want.
9. To display detailed information about other resources, click other linked names.
10. When you return to the data source summary page, check the **Changes** area at the bottom of the page to see if recent changes caused the problem.

Related concepts

Changes causing data collection problems on page 126
What data source patches are on page 132
Data source support information on page 67

Related tasks

Examining one data source in detail on page 126

Controlling data source polling

After making a change to a data source, you might want it to poll immediately to check your changes, or you might want to postpone the data collection on a data source for one, three, or five days while you work on a problem.

Steps

1. Click **Admin** on the toolbar.

   The **Data sources** list opens.

2. Select the data source on which you want to control the polling.
3. Click the data source name link.
4. On the data source summary page, check the information and click one of these two polling options:

   - **Poll again** to force the data source to collect data immediately.
   - **Postpone** and select the length of the polling delay from 1, 3, or 5 days.

After you finish

If you postponed the data collection on a data source and want to restart collection, click **Resume** on the summary page.

Editing data source information

You can quickly edit data source setup information.

Steps

1. Click **Admin** on the toolbar.

   The **Data sources** list opens.

2. Locate the data source that you want to edit.
3. Use one of these methods to begin the changes:

   - Click **Edit data source** to the right of the selected data source.
   - Click the linked name of the selected data source and click **Edit**.

   Either method opens the Edit data source dialog box.

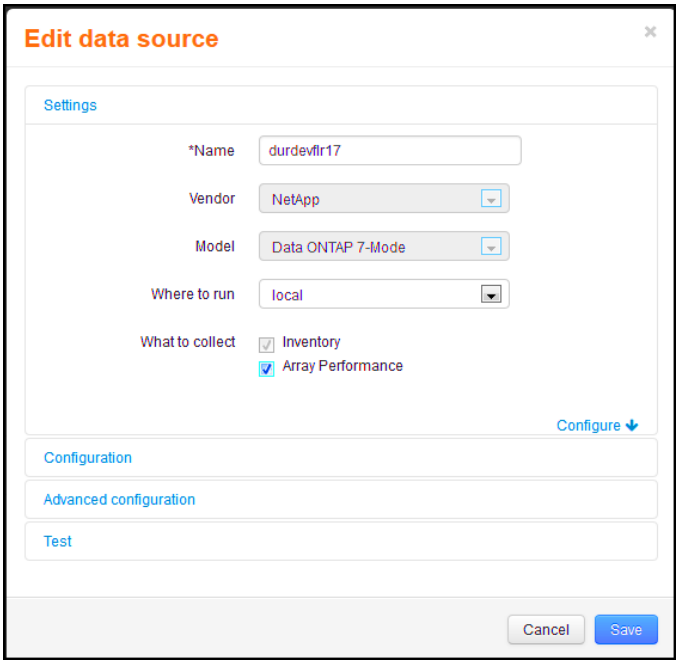

4. Select the required changes.
5. Click **Save**.

Editing information for multiple data sources

You can edit most of the information for multiple data sources of the same vendor and model at one time. For example, if these data sources share a user name and password, you can change the password in one place and thereby update the password for all the selected data sources.

About this task

Options that you cannot edit for the selected data sources appear dimmed or are not displayed in the Edit data source dialog box. Additionally, when an option displays a value of **Mixed**, it indicates that the value for the option varies between the selected data sources. For example, if the **Timeout (sec)** option for two selected data sources is **Mixed**, one data source could have a timeout value of 60 and the other could have a value of 90; therefore, if you change this value to 120 and save the changes to the data sources, the timeout setting for both data sources becomes 120.

Steps

1. Click **Admin** on the toolbar.

   The **Data Sources** list opens.
2. Click the **Actions** button and select the **Bulk edit** option.

   The Bulk Edit Multiple Data Sources dialog box displays:

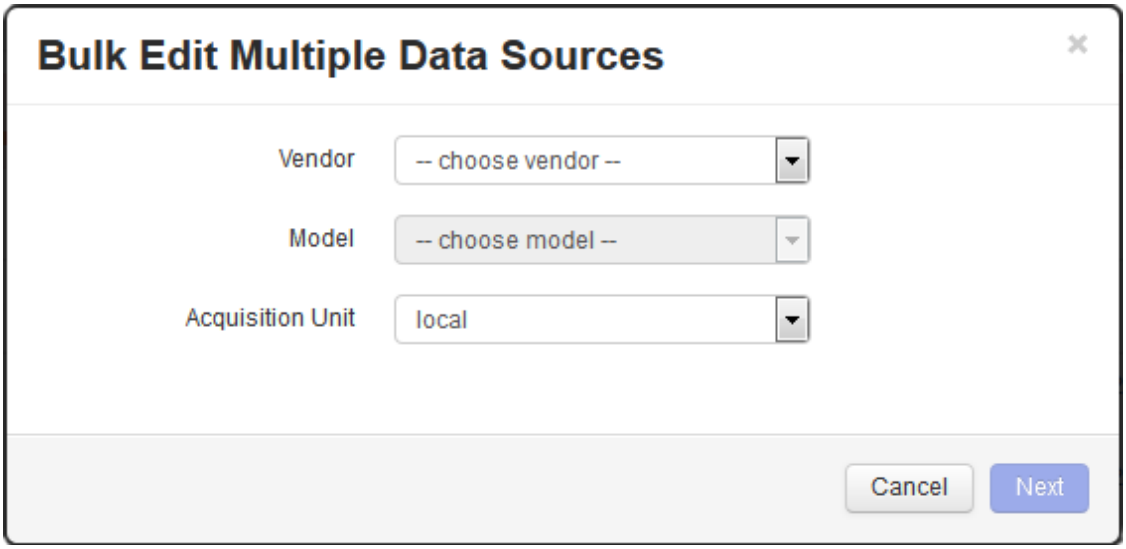

3. Select the vendor of the data sources from the Vendor list.

   If there is only one model for the vendor you selected, the dialog box lists all the data sources available.

4. Select the model of the data sources from the Model list.

   The dialog box lists all the data sources that match the vendor and model that you selected.

5. Select the location of the acquisition unit for the data sources from the Acquisition Unit list.

6. Click **Next**.

   The Edit data source dialog box displays.

7. Change any of the **Settings** options.

8. Click the **Configuration** link to change any of the basic options for the data sources.

9. Click the **Advanced Configuration** link to change any of the advanced options for the data sources.

10. Click **Save**.

Deleting a data source

If you have removed a data source from your environment, you must also delete it from the OnCommand Insight monitoring environment.

Steps

1. On the Insight toolbar, click the **Admin** icon.

   The Data sources list opens.

2. Select the data source that you want to delete.

3. Click the linked data source name.

4. Check the information for the selected data source on the summary page to be certain that it is the one you want to delete.

5. Click **Delete**.
6. Click **OK** to confirm the operation.

What data source patches are

Data source patches fix issues with existing patches and also enable you to easily add new data source types (vendors and models). For each data source type in your network, you can upload data source patches. You can also install, test, and manage the patching process. However, only one patch can be active for a data source type at a time.

For each patch, you can perform these tasks:

- Check the before and after comparison of each data source receiving the patch.
- Write comments to explain decisions or summarize research.
- Make changes to a data source that is not responding well to the patch.
- Approve the patch to be committed to your Insight server.
- Roll back a patch that is not operating as you intended.
- Replace a failing patch with a different one.

Applying a data source patch

Data source patches are periodically available and enable you to fix issues with an existing data source, add a data source for a new vendor, or add a new model for a vendor.

Before you begin

You must have obtained the `.zip` file that contains the latest data source `.patch` files from technical support.

Steps

1. Click the **Admin** icon on the toolbar.
2. Click **Patches**.
3. From the Actions button, select **Apply patch**.
4. In the **Apply data source patch** dialog box, click **Browse** to locate the `.patch` file.
5. Inspect the **Patch name**, **Description**, and **Impacted data source types**.
6. If the selected patch is correct, click **Apply Patch**.

   If you are applying a patch that fixes issues with a data source, all data sources of the same type are updated with the patch and you must approve the patch. Patches that do not affect any configured data sources are automatically approved.

After you finish

If you are applying a patch that adds a data source for a new vendor or a new model, you must add the data source after applying the patch.

Related tasks

Adding data sources on page 67

Installing a patch on one type of data source

After uploading a data source patch, you can install it on all of the data sources of the same type.

Before you begin

You must have uploaded a patch file that you want to install on one type of data source.

Steps

1. Click the **Admin** icon on the toolbar.
2. Click **Patches**.
3. From the Actions button, select **Apply patch**.
4. In the **Apply data source patch** dialog box, click **Browse** to locate the uploaded patch file.
5. Check the **Patch name**, **Description**, and **Impacted data source types**.
6. If the selected patch is correct, click **Apply Patch**.

   All data sources of the same type are updated with this patch.

Managing patches

You can review the current status of all of the data source patches being applied to your network. If you want to perform an action on a patch, you can click the linked name in the Patches currently under review table.

Before you begin

You must have already uploaded and be installing at least one patch.

Steps

1. On the Insight toolbar, click the **Admin** icon.
2. Click **Patches**.

   If no patches are being installed, the table of Patches currently under review is empty.
3. In **Patches currently under review**, check the status, as shown in this example, of the data source patches currently being applied.

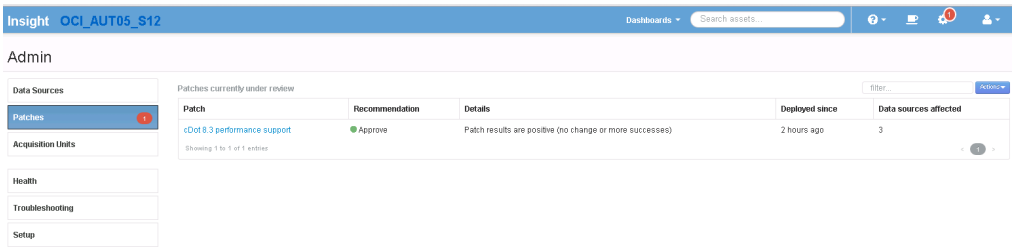

4. To examine the details associated with a specific patch, click the linked name of the patch.
5. For the selected patch, you might click any of these options to perform the next action on the patch:

   - **Approve patch** commits the patch to the data sources.

- **Rollback** removes the patch.
- **Replace patch** enables you to select a different patch for those data sources.

Committing a data source patch

You use the information in the Patches summary to decide if the patch is performing as expected and then commit the patch to your network.

Before you begin

You have installed a patch and need to decide if the patch is successful and should be approved.

Steps

1. On the Insight toolbar, click the **Admin** icon.
2. Click **Patches**.

   If no patches are being installed, the Patches currently under review is empty.
3. In **Patches currently under review**, check the status of the data source patches currently being applied.
4. To examine the details associated with a specific patch, click the linked name of the patch.
5. In the Patches summary information, shown in this example, check the **Recommendation** and **Comments** to assess the progress on the patch.

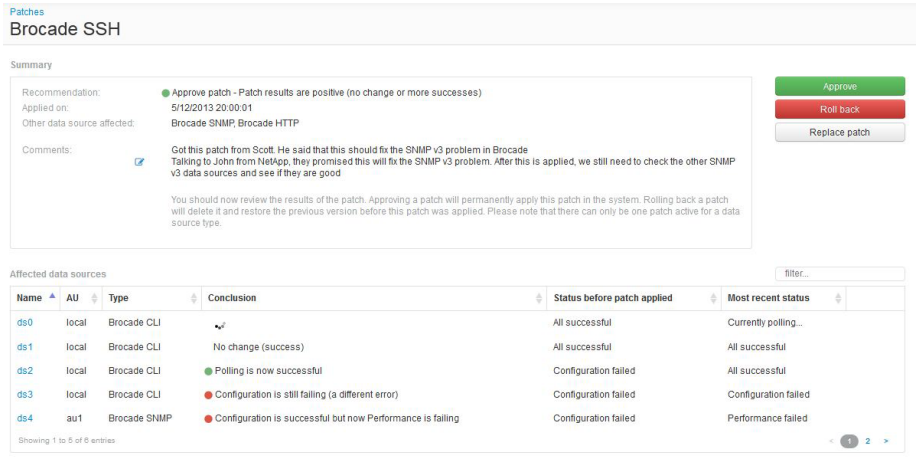

6. Check the **Data sources affected** table to see the status of each affected data source before and after the patch.

   If you are concerned that there is a problem with one of the data sources being patched, click the linked Name in the Data sources affected table.
7. If you conclude that the patch should be applied to that type of data source, click **Approve**.

   The data sources are changed and the patch is removed from Patches currently under review.

Rolling back a data source patch

If a data source patch is not working in the manner you expected, you can roll it back. Rolling back a patch deletes it, and restores the previous version as it was before this patch was applied.

Steps

1. On the Insight toolbar, click the **Admin** icon.
2. Click **Patches**.
3. In **Patches currently under review**, click the linked name of the patch that appears to be unsuccessful.
4. On the Patches page for the data source, examine this information:
   - **Summary** describes when the patch was applied, the affected data sources, and comments about the patch from you or other members of your team.
   - **Affected data sources** lists all of the data sources being patched and includes a comparison of the before and after patching status.
5. To display the details for a data source that is not successfully processing the patch, click the linked **Name**.
   a. Check the summary information.
   b. Check the **Event timeline** to see any configuration or performance data that might be affecting this data source.
6. If you conclude that the patch is not going to be successful, click the browser back arrow to return to the Patches summary page.
7. Click **Roll back** to remove that patch.

   If you know of a different patch that is more likely to be successful, click **Replace patch** and upload the new patch.

Maintaining Insight

Whether you are new to Insight and have a new system to set up, or your system has been operating for some time, you must take steps to maintain smooth operation of Insight and your network. The key maintenance concept is that changes in your network usually need to be accommodated in Insight.

These are the most common maintenance tasks:

- Maintaining Insight backups
- Updating expired Insight licenses
- Coordinating data source patches
- Updating the Insight version on all acquisition units
- Deleting removed data sources from Insight

Related concepts

What data source patches are on page 132

Managing Insight

OnCommand Insight monitors your environment, enabling you to research potential problems before a crisis is reported. The Assets Dashboard provides summary pie charts, heat maps for IOPs, and an interactive chart of the top 10 utilized storage pools.

Steps

1. Open the Insight **Assets Dashboard** and move your cursor over the pie charts to examine the asset distribution in these three charts:
   - Capacity by Vendor
   - Capacity by Tier
   - Switch Ports
2. View **Facts About Your Environment** to see information about your environment's used capacity, the capacity's efficiency, consumed FC resources, and virtual infrastructure statistics.
3. Position your cursor over a storage pool bar in the **Top 10 Utilized Pools** chart to view the used and unused capacity of the storage pool.
4. Click any asset name appearing in large text (which indicates that the asset has issues) in the **Storage IOP** heat map to display a page summarizing the current state of that asset.
5. In the lower right corner of the **Assets Dashboard**, click any asset name appearing in large text (which indicates the asset has issues) in the **Virtual Machine IOPS** heat map to display a page summarizing the current state of the asset.
6. On the toolbar, click **Admin** .
7. Find administration buttons with similar circled numbers and examine the details displayed in those areas.

If the Admin icon has a circled number, as shown in Step 6, the number indicates the number of possible problems detected. In the Admin area, potential problems are marked with a solid red circle in the web UI.

8. Click **Data Sources** to examine a list of all monitored data sources.

   Examine any data source with a **Status** column containing a message with a solid red circle and with an **Impact** listed as High or Medium. These are at the top of the table. The problems with those data sources affect a significant portion of your network, which you need to address.

9. Click **Acquisition Units** to note the status for each IP address running Insight and to restart an acquisition unit, if necessary

10. Click **Health** to see high-level instance monitoring of the Insight servers.

Monitoring OnCommand Insight system health

You can check the current status of your Insight system components by viewing the health page, which shows a component status and alerts you when there is an issue.

Steps

1. Log in to the Insight web UI.
2. Click **Admin** and select **Health**.

   The Health page displays.

3. View the summary of the current status of the components, paying particular attention to any attention status in the **Details** column that is preceded by a red circle, which indicates an issue that requires your immediate attention.

   The Health page displays information about the following Insight components:

   - Insight servers
   - Acquisition units
   - Data Warehouse servers (ETL)
   - Anomaly detection engine
     If the anomaly detection engine is installed and configured for monitored objects, see the `prelert.log` file in the following location: *`disk drive`*`:\`*`install directory`*`\SANscreen\jboss\server\onaro\log`.

Auditing system and user activities

If you want to locate unexpected changes, you can view an audit trail of the OnCommand Insight system and its user activities.

About this task

Insight generates audit entries for any user activities that affect the storage network or its management, including the following:

- Logging in
- Authorizing or unauthorizing a path
- Updating an authorized path

- Setting global policies or thresholds
- Adding or removing a data source
- Starting or stopping a data source
- Updating data source properties
- Adding, editing, or deleting a task
- Removing an application group
- Identifying or changing the identification for a device

Steps

1. Open Insight in your browser.
2. Click **Admin** and select **Audit**.

   The Audit page displays the audit entries in a table.
3. You can view the following details in the table:

   Time

   Date and time that the changes were made

   User

   Name of user associated with the audit entry

   Role

   User account's role, which is guest, user, or administrator

   IP

   IP address associated with the audit entry

   Action

   Type of activity in the audit entry

   Details

   Details of the audit entry

   If there is a user activity that affects a resource, such as a data source or an application, the details include a link to the resource's landing page.

   Note: When a data source is deleted, the user activity details related to the data source no longer contain a link to the data source's landing page.
4. You can display audit entries by choosing a particular time period (1 hour, 3 hours, 24 hours, 3 days, and 7 days), with Insight showing a maximum number of 1000 violations for the selected time period.

   You can click a page number below the table to browse through data by page if there is more data than fits on a single page.
5. You change the sort order of the columns in a table to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header; to return to the default sort order, click any other column header.

   By default, the table displays the entries in descending order.
6. You can use the **filter** box to show only the entries you want in the table.

Example

To see only the audit entries by the user `izzyk`, type `izzyk` in the **filter** box.

Monitoring the violations in your network

When Insight generates violations due to the thresholds set in performance policies, you can view them using the Violations Dashboard. The dashboard lists all the violations that occur in your network and enables you to locate and address issues.

Steps

1. Open OnCommand Insight in your browser.
2. On the Insight toolbar, click **Dashboards** and select **Violations Dashboard**.

   The Violations Dashboard displays.
3. You can use the **Violations By Policies** pie chart in the following ways:
   - You can position your cursor over any slice of a chart to display the percentage of the total violations that occurred for a particular policy or metric.
   - You can click a slice of a chart to “enlarge” it, which enables you to emphasize and study more carefully that slice by moving it away from the rest of the chart.
   - You can click the icon in the upper-right corner to display the pie chart in full screen mode, and click again to minimize the pie chart.

   A pie chart can contain a maximum of five slices; thus, if you have six policies that generate violations, Insight combines the fifth and sixth slices into an “Others” slice. Insight assigns the most violations to the first slice, the second most violations to the second slice, and so on.
4. You can use the **Violations History** chart in the following ways:
   - You can position your cursor over the chart to display the total number of violations that occurred at a particular time and the number that occurred out of the total for each specified metric.
   - You can click a legend label to remove the data associated with the legend from the chart. Click on the legend to display the data again.
   - You can click the icon in the upper-right corner to display the chart in full screen mode, and click again to minimize the pie chart.
5. You can use the**Violations Table** in the following ways:
   - You can click the icon in the upper-right corner to display the table in full screen mode, and click again to minimize the pie chart.
     If your window size is too small, then the Violations Table displays only three columns; however, when you click , additional columns (up to seven) display.
   - You can display violations for a particular time period (**1h**, **3h**, **24h**, **3d**, and **7d**), with Insight showing a maximum number of 1000 violations for the selected time period.
   - You can use the **filter** box to show only the violations you want.

- You can change the sort order of the columns in a table to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header; to return to the default sort order, click any other column header.
  By default, the table displays the violations in descending order.
- You can click a violation in the ID column to display the asset page for the duration of the violation.
- You can click the resource links (for example, storage pool and storage volume) in the Description column to display the asset pages associated with those resources.
- You can click the performance policy link in the Policy column to display the Edit Policy dialog box.
  You might want to adjust the thresholds for a policy if you feel it generates too few or too many violations.
- You can click a page number to browse through data by page if there is more data than fits on a single page.
- You can click to dismiss the violation.

Related tasks

Editing performance policies on page 64

Fibre Channel environment monitoring

Using OnCommand Insight's Fibre Channel asset pages, you can monitor the performance and inventory of the fabrics in your environment and be aware of any changes that might cause issues.

Fibre Channel asset pages

Insight's asset pages present summary information about the resource, its topology (the device and its connections), performance charts, and tables of associated resources. You can use the fabric, switch, and port asset pages to monitor your Fibre Channel environment. Particularly helpful when troubleshooting a Fibre Channel issue is the performance chart for each port asset, which shows the traffic for the selected top contributor port. Additionally, you can also show buffer-to-buffer credit metrics and port errors in this chart, with Insight displaying a separate performance chart for each metric.

Performance policies for port metrics

Insight enables you to create performance policies to monitor your network for various thresholds and to raise alerts when those thresholds are crossed. You can create performance polices for ports based on available port metrics. When a violation of a threshold occurs, Insight detects and reports it in the associated asset page by displaying a red solid circle; by email alert, if configured; and in the Violations Dashboard or any custom dashboard that reports violations.

Related tasks

Creating performance policies on page 62

Related references

Port metrics on page 141

Port metrics

Asset page performance charts display port metrics when the asset is associated with a port or when a port is added as a resource. Port metrics are useful when monitoring and troubleshooting your Fibre Channel environment, providing BB credit metrics and the type of port error (for example, link failures and Class 3 discards).

The following table lists the port metrics available in Insight and the Brocade and Cisco counters from which they are derived:

Insight Port Metrics	Brocade SNMP Counter	Cisco SNMP Counter
BB credit 0 -Tx Tx BB credit transitions to zero	**connUnitPortStatCountBBCreditZero** The count of transitions in to and out of BB credit zero state. The other side is not providing any credit.	**fcHCIfBBCreditTransistionFromZero** Increments if the transmit buffertobuffer credit is zero.
BB credit 0 - Rx Rx BB credit transitions to zero	**connUnitPortStatCountInputBuffersFull** The count of occurrences when all input buffers of a port were full and outbound BB credit transitioned to zero. There is no credit to provide to other side.	**fcHCIfBBCreditTransistionToZero** Increments if the receive buffertobuffer credit is zero.
BB credit zero duration - Tx Time in milliseconds during which Tx BB credit was zero during the sample interval	**connUnitPortStatCountBBCreditZero** An increment of this counter means that the frames could not be sent to the attached device for 2.5 µs because of zero Tx BB credit.	**fcIfTxWtAvgBBCreditTransitionToZero** Increments if the credit available is zero for 100 ms.

Insight Port Metrics	Brocade SNMP Counter	Cisco SNMP Counter
Timeout discard frames - Tx Discarded Tx frames caused by timeout	**swConnUnitC3Discard DueToTXTimeout** Number of Class 3 transmit frames discarded due to timeout	**fcIfTimeOutDiscards** Frames discarded because a timeout occurred within the switch
CRC errors CRC error count	**fcIfInvalidCrcs** The number of invalid CRCs detected by the FC-Port	**connUnitPortStatCount InvalidCRC** The number of frames received with invalid CRC
Link resets - Tx Tx Link Reset count	**connUnitPortStatCount TxLinkResets** Count of link resets. This is the number LRs transmitted.	**fcIfLinkResetOuts** The number of link reset protocol errors issued by the FCPort to the attached FCPort
Link resets - Rx Rx Link Reset count	**connUnitPortStatCount RxLinkResets** Count of link resets. This is the number of LRs received.	**fcIfLinkResetIns** The number of link reset protocol errors received from the attached FC-port
Link failures Link failure count	**connUnitPortStatCount LinkFailures** Count of link failures. This count is part of the LESB.	**fcIfLinkFailures** The number of link failures detected by the FCPort
Traffic - Tx Tx data rate in MB/s	**connUnitPortStatCount TxElements** The number of bytes that have been transmitted by this port	**fcIfC3OutOctets , fcIfCfOutOctets** The number of Class 3 frame octets, including the frame delimiters, transmitted by this port

Insight Port Metrics	Brocade SNMP Counter	Cisco SNMP Counter
Traffic - Rx Rx data rate in MB/s	**connUnitPortStatCount RxElements** The number of octets or bytes that have been received by this port	**fcIfC3InOctets , fcIfCfInOctets** The number of Class 3 frame octets, including the frame delimiters, received by this port
Frame rate - Tx Tx frame rate in fps	**connUnitPortStatCount TxObjects** The number of frames, packets, I/O instances, and so on transmitted by this port	**fcIfC3OutFrames, fcIfCfOutFrames** The number of frames, packets, I/O instances, and so on transmitted by this port
Frame rate - Rx Rx frame rate in fps	**connUnitPortStatCount RxObjects** The number of frames, packets, I/O instances, and so on received by this port	**fcIfC3InFrames, fcIfCfInFrames** The number of frames, packets, I/O instances, and so on received by this port
Frame size average - Tx Avg Tx frame size in B	**traffic.txRate * 2^{20} / traffic.txFrameRate** The ratio of transmitted traffic to transmitted frame rate	**traffic.txRate * 2^{20} / traffic.txFrameRate** The ratio of transmitted traffic to transmitted frame rate
Frame size average - Rx Avg Rx frame size in B	**traffic.rxRate * 2^{20} / traffic.rxFrameRate** The ratio of received traffic to received frame rate	**traffic.rxRate * 2^{20} / traffic.rxFrameRate** The ratio of received traffic to received frame rate

Related tasks

Creating performance policies on page 62

Acquisition unit status

The status of the Insight acquisition units connected to your server is displayed in the **Admin** > **Acquisition Units** table. If a unit in the list has a solid red circle in the description, that unit has a

serious problem that you must investigate. You can select an entry in this table and restart that acquisition unit.

The Acquisition Units table displays this information for each unit:

- **Name**
- **IP**
- **Status** is the operating status of the acquisition unit.
- **Last reported** displays last time a data source connected to the acquisition unit reported.

If you position your pointer over one of the units in the table, you can click the square icon to restart that unit.

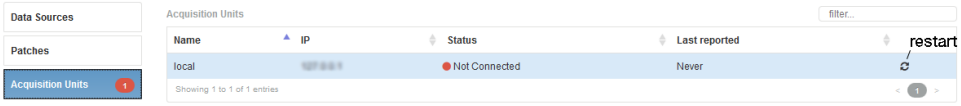

Restoring the Insight database

To restore your Insight database from a verified backup file, use the Troubleshooting options. This operation completely replaces your current OnCommand Insight data.

Before you begin

Best practice: Before restoring your OnCommand Insight database, use the manual backup process to create a copy of the current database. Check the backup file you plan to restore be certain that it was a successful backup containing the files you want to restore.

Steps

1. On the Insight toolbar, click the **Admin** icon.
2. Click **Troubleshooting**.

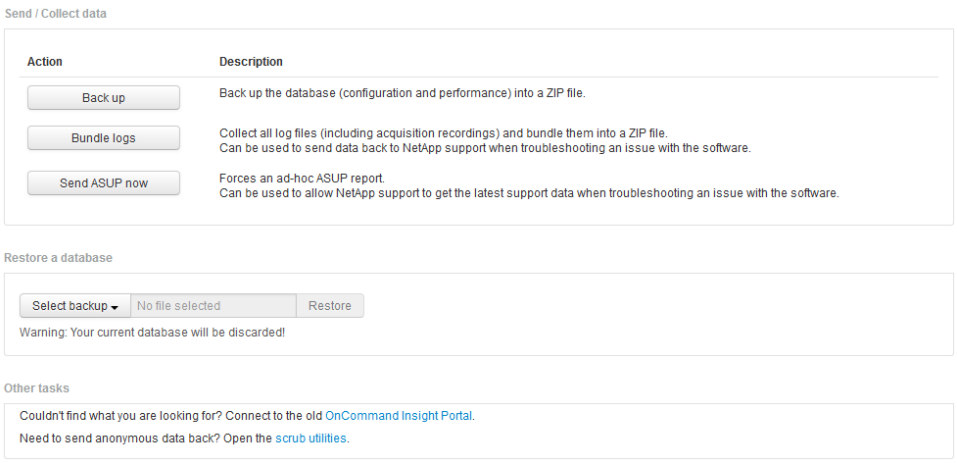

3. In the Restore a database section, select the backup file you want to restore from the **Select Backup** menu.

4. Click **Restore**.

Updating expired licenses

If one or more of your Insight licenses expired, you can update the licenses quickly using the same procedure as you did to install the licenses originally.

Steps

1. In a text editor, such as Notepad, open the new license file you received from NetApp Support and copy the license key text to your Windows Clipboard.
2. Open OnCommand Insight in your browser.
3. Click the **Admin** icon on the toolbar.
4. Click **Setup**.
5. Click the **Licenses** tab.
6. Click **Update License**.
7. Copy the license key text into the **License** text box.
8. Select the **Update (most common)** operation.

   This operation adds your new licenses to any currently active Insight licenses.
9. Click **Save**.

Licenses no longer compliant

If you notice the "Not Compliant" message on your Insight Licenses page, Insight is managing more terabytes than your company licensed.

The "Not Compliant" message means your company paid for fewer terabytes than Insight is currently managing. The difference between the managed terabytes and the licensed number of terabytes is shown beside the non-compliance message.

The operation of your Insight system is not affected, but you should contact your NetApp representative to increase your license coverage and update the appropriate license.

Replacing licenses for older Insight versions

If you have purchased a new Insight version that is not backward compatible with your older version of the product, you must replace the older licenses with the new licenses.

When you are installing the new licenses, you must select the **Replace** operation before you save the license key text.

Applying a service pack

Periodically, service packs are available, which you can apply to take advantage of fixes and enhancements to OnCommand Insight.

Before you begin

- You must have downloaded the service pack file (for example, `7.2service_pack_1.patch`) from the NOW site.

- You must have approved all patches.

Steps

1. Click the **Admin** icon on the toolbar.
2. Click **Patches**.
3. From the Actions button, select **Apply patch**.
4. In the **Apply data source patch** dialog box, click **Browse** to locate the service pack file.
5. Inspect the **Patch name**, **Description**, **Impacted data source types**, which shows if any data sources are affected, and **Details**, which describes the enhancements that the service pack contains.
6. If the selected service pack is correct, click **Apply Patch**.

   Service packs are approved automatically; no further action is required.

Preparing a special troubleshooting report

Insight sends information to NetApp Customer Support automatically through the ASUP system you set up after installing the software. However, you might want to create a troubleshooting report and open a case with the Support team for a specific problem.

You can use tools in Insight to perform a manual Insight backup, bundle the logs, and send that information to NetApp Customer Support.

Manually backing up the OnCommand Insight database

If you enabled weekly backups for the OnCommand Insight database, you are automatically generating copies that you can use to restore the database, if necessary. If you need to create a backup before a restore operation, or to send to NetApp technical support for assistance, you can create a backup `.zip` file manually.

Steps

1. On the Insight toolbar, click the **Admin** icon.
2. Click **Troubleshooting**.
3. In the Send/Collect data section, click **Backup**.
4. Click **Save File**.
5. Click **OK**.

Bundling logs for Support

When troubleshooting a problem with Insight software, you can quickly generate a zip file (using the "gz" format) of the logs and acquisition recordings to send to NetApp Customer Support.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Troubleshooting**.
3. In the Send / Collect data section, click **Bundle logs**.

4. Click **Save File**.
5. Click **OK**.

Sending information to NetApp Support

The NetApp automated support (ASUP) facility sends troubleshooting information directly to the NetApp Customer Support team. You can force a special report to be sent.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Setup**.
3. Click the **Backup/ASUP** tab.
4. In the Send/Collect data area, click **Send ASUP now** to submit your logs, recordings, and backup to NetApp Support.

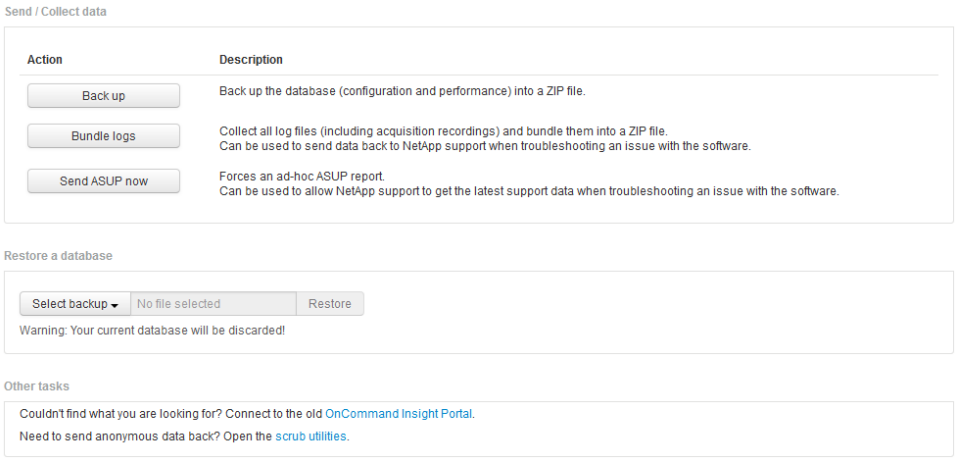

Advanced configuration tools

To complete your OnCommand Insight configuration, you must use the advanced configuration tools. These tools run in the browser and are opened from the **Admin > Troubleshooting** page.

To open the advanced configuration tools in the browser, click the **OnCommand Insight Portal** link at the bottom of this page:

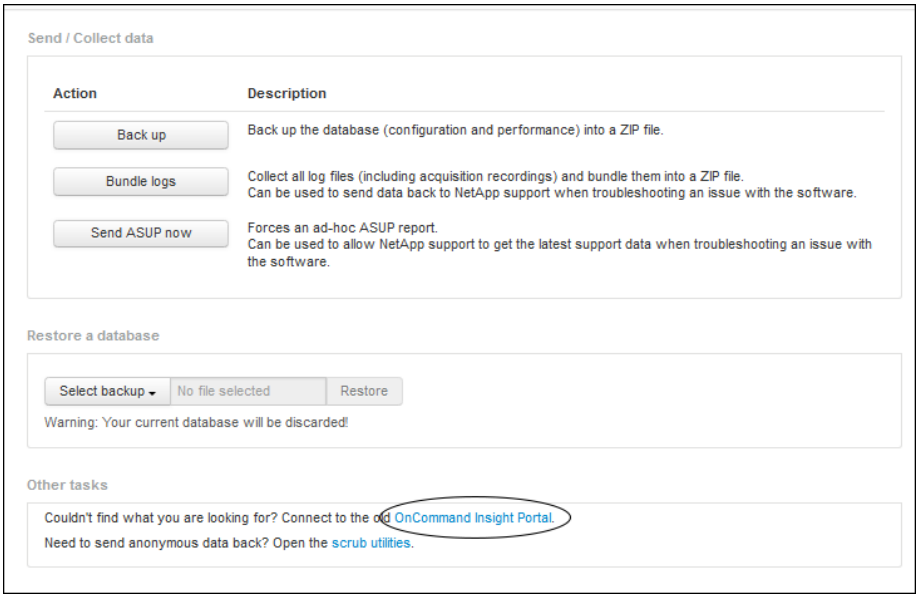

On the next page in the Advanced menu, click the **Troubleshooting** option to display these tools:

- Restart OnCommand Insight Server
- Acquisition
- Logs
- System Information
- Installed Software Packages
- Calculate Scale
- Migrate MySQL Data
- Force Update of Annotations for DWH
- Force Update of Annotations for DWH (include deleted objects)
- Configure Client Settings
- Check Tier Upgrade
- Server Resources Status
- Scrub Utilities
- Find Ghost Data Sources

In addition to these tools, which are available to all administrators, the chief administrator should enable the automatic disk model updates.

Restarting the OnCommand Insight Server

When you restart the OnCommand Insight Server, refresh the page and log into the OnCommand Insight Portal again.

About this task

Note: Both of these options should only be used upon request by NetApp Customer Support. There is no confirmation prior to restart.

Steps

1. In the web UI, click the **Admin** icon and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click the **Restart OnCommand Insight Server**.

Configuring the number of hours to ignore dynamic data

You can configure the number of hours during which OnCommand Insight ignores updating dynamic data, such as used capacity. If the default of six hours is used and no configuration changes occur, reports will not be updated with dynamic data until after the default number of hours. This option improves performance because this option defers updates when only the dynamic data changes.

About this task

If a value is set for this option, OnCommand Insight will update dynamic data based on the following rules:

- If no configuration changes occur but capacity data changes, data will not be updated.
- Dynamic data (other than configuration changes) will be updated only after the timeout specified in this option.
- If configuration changes occur, configuration and dynamic data is updated.

Dynamic data impacted by this option includes the following:

- Capacity violation data
- File Systems Allocated Capacity and Used Capacity
- Hypervisor
  - Virtual Disk Used Capacity
  - Virtual Machine Used Capacity
- Internal Volume
  - Data Allocated Capacity
  - Data Used Capacity
  - Deduplication Savings
  - Last Known Access Time
  - Last Snapshot Time
  - Other Used Capacity
  - Snapshot Count
  - Snapshot Used Capacity
  - Total Used Capacity
- iSCSI Session Initiator IPs, Target Session ID, and Initiator Session ID
- Qtree Quota Used Capacity
- Quota Used Files and Used Capacity

- Storage Efficiency Technology, Gain/Loss, and Potential Gain/Loss
- Storage Pool
  - Data Used Capacity
  - Deduplication Savings
  - Other Used Capacity
  - Snapshot Used Capacity
  - Total Used Capacity
- Volume
  - Deduplication Savings
  - Last Known Access Time
  - Used Capacity

Steps

1. In the web UI, click the **Admin** icon and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Select the **Ignore dynamic attribute changes for x hours** option and enter the number of hours that OnCommand Insight should ignore dynamic data.
5. Click **Save Settings**.
6. (Optional) To restart the acquisition unit, click the **Restart Acquisition Unit** link.

   Restating the local acquisition unit reloads all of the OnCommand Insight data source views. This change is applied during the next poll, so you do not have to restart the Acquisition Unit.

Generating logs for Customer Support

If requested by Customer Support, generate a server, acquisition, or remote log for troubleshooting purposes.

About this task

If NetApp Customer Support requests, use this option to generate the logs.

Steps

1. In the web UI, click the **Admin** icon and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click **OnCommand Insight Portal**.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click **Logs** and select:
   - **Server** - Server log files in {Home}\SANscreen\jboss\server\onaro\log
   - **Acq** - Acquisition log files in {Home} \SANscreen\acq\log
   - **Remote** - Remote logs from Remote Acquisition, OnCommand Insight Client, and Data Warehouse

Displaying system information

You can display the Microsoft Windows IP configuration information about the system on which OnCommand Insight server is deployed.

Steps

1. In the web UI, click the **Admin** icon and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Select **System Information**.

   The Windows IP configuration includes information such as the host name, DNS, IP address, subnet mask, OS information, memory, boot device, and connection name.

Listing installed OnCommand Insight components

You can display a list of the installed OnCommand Insight components including, among others, inventory, capacity, dimensions, and the Data Warehouse views. Customer Support might ask you for this information, or you might want to see what software versions were installed and when they were installed.

Steps

1. In the web UI, click the **Admin** icon and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click **Installed Software Packages**.

   OnCommand Insight lists the package names with the software version numbers and installation date and time of each.

Calculating the number of database objects

To determine the number of objects in the OnCommand Insight database, use the Calculate Scale feature.

Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click **Calculate Scale**.

   The system displays the number of discovered objects in the database.

Moving MySQL data using the migrate option

You can use migrate MySQL data directory to a different directory. You can retain the current data directory. You can use the migrate option on the Troubleshooting menu or you can use the command line. This procedure describes how to use the **Troubleshooting > Migrate MySQL data** option.

About this task

If you retain the current data directory, it will be kept as a backup and renamed.

Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Select **Migrate MySQL Data**.
5. Enter the path to which you want to migrate the data.
6. To retain the existing data directory, check **Keep existing data directory.**
7. Click **Migrate the data to selected directory**.

Moving MySQL data using the command line

You can use migrate MySQL data directory to a different directory. You can retain the current data directory. You can use the migrate option on the Troubleshooting menu or alternatively, you can use the command line. This procedure describes how to use the command line.

About this task

If you retain the current data directory, it will be kept as a backup and renamed.

You can use the Migrate MySQL Data utility or you can use a `java -jar mysqldatamigrator.jar` option in the OnCommand Insight path of `\bin\mysqldatamigrator` where the following parameters should be used:

- Mandatory parameters
  -path

  The new data path to which the data folder will be copied.

- Optional parameters
  -myCnf <my .cnf file>

  The path for the .cnf file. The default is `<install path>\mysql\my.cnf`. Use this flag only if a non-default MySQL is used.

  -doBackup

  If this flag is set, the current data folder will be renamed but not deleted.

Step

1. Access the command line tool here: <*`installation path`*> \bin\mysqldatamigrator \mysqldatamigrator.jar

Example usage

```
java -jar mysqldatamigrator.jar -path "C:\<new path>" -doBackup
```

Forcing annotation updates

If you have changed the annotations and want to use them in reports immediately, use one of the force annotation options.

Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Select one of these options:
   - **Force Update of Annotations for DWH** to force the update of annotations in data warehouse to be used for reports.
   - **Force Update of Annotations for DWH (including deleted objects)** to force an annotations update (including deleted objects) in the data warehouse to be used for reports.

Configuring Client settings

You can configure options for the OnCommand Insight Java UI Client on the server.

Steps

1. In the web UI, click the **Admin** icon and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. In the Advanced menu on the next page, click the **Troubleshooting** link.
4. Click **Configure Client Settings**.
5. Enter the following information:
   - **Initial Heap Size** is the size of the heap that is created when the VM (Virtual Machine) is started to run the client. This number is used to construct the `-Xms<size>` parameter to `java`. Decreasing this number means Java takes up less memory when it starts up. (However, it is then more likely that Java will need to reallocate the heap when it exceeds this number, which is an expensive operation, so you do not want to make this number too small.)
   - **Max Heap Size** is the largest that Java will allow the heap to grow in the VM that is running the client. This number is used to construct the `-Xmx<size>` parameter to `java`.
   - Log Level controls the amount of detail that is generated in the client logs. There are two possible values:
     - INFO - generates a "normal" amount of information.
     - DEBUG - generates more detailed information to the client log to help with debugging.
6. Click **Save Settings**.
7. Restart the client.

Note: If the client generates a "java.lang.OutOfMemoryError: Java heap space" error, you can increase the max heap size setting to see whether that fixes the problem.

Checking a tier upgrade

After upgrading from OnCommand Insight version 6.0 to a newer version, select the **Check Tier Upgrade** option to list any volume which presently had a different tier value than the tier value it has in the new version. The newer Insight versions improve the way tiers are calculated by using disk attributes instead of volume attributes. This change increases accuracy and reduces confusion when the volume is part of a storage pool with mixed disks.

About this task

Tier rule configuration no longer presents a disk attribute that is a combination of all the disk attributes of a storage pool, and tiers are calculated for a volume using the lowest performing disk attribute. This change might produce different values if your rules were setup to use the old volume data.

Steps

1. In the web UI, click the **Admin** icon and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click **Check Tier Upgrade**.
5. Locate the differences and decide if you need to adjust the tier rule.

Checking the status of server resources

This option displays the OnCommand Insight Server's information including server memory, disk space, OS, and CPU and OnCommand Insight database information including innoDB data size and the disk free space where the database resides.

Steps

1. In the web UI, click the **Admin** icon and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click **Server Resources Status**.

   For advanced OnCommand Insight users: The administrator can run some SQL tests to check the database and server's response time from the button at the end of the information summary. This option displays a warning if server resource is low.

Scrubbing data for transfer to support

Customers who have secure environments need to communicate with NetApp Customer Service to troubleshoot problems that arise without compromising their database information. The OnCommand Insight Scrub utilities allow you set up a comprehensive dictionary of keywords and patterns so that you can "cleanse" sensitive data and send scrubbed files to Customer Support.

Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.

2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click **Scrub Utilities**.

   There are three scrub utilities: Dictionary Admin, Scrub, and Lookup Dictionary.
5. Select the **Dictionary Admin** option.
6. Expand **Custom keywords**. Click **Add** to get the dialog box for entering new keywords (note that multiple entries must be line-separated
7. Enter a keyword and click **OK**. The keyword is added to the dictionary.
8. Expand **Patterns (regexp)**. Click **Add** to get the dialog box for entering a new pattern.
9. Enter a Name and a Pattern in the fields and click **OK**. The information has been added to the dictionary.

   Note: Patterns must be encompassed by round parentheses to identify a regular expression capturing group.
10. In the Database block, click **Build** to initiate compilation of the dictionary to the dictionary all words identified as sensitive from the OnCommand Insight database.

    On completion, you see a prompt informing you the revised dictionary is available. The Database description includes a line indicating how many keywords are in the dictionary. Check your keywords in the dictionary for accuracy. If you find problems and want to rebuild the dictionary, click **Reset** on the Database block to remove all keywords collected from the OnCommand Insight database from the dictionary. As the prompt advises, all other keywords will not be deleted. Return to the Scrub utilities and enter your Custom Keywords again.
11. Select **Scrub > Scrub log / XML files** again.
12. Browse to locate the log file and click **Scrub**.

    You can run and scrub signed MySQL queries.

    Note: The Scrub Utility will not accept an unsigned query. Queries can be signed only by NetApp personnel.

    While restrictions exist, the customer can scrub the query multiple times and send it in again to NetApp Customer Service if fine-tuning of the query is necessary.

Scrub dictionary administration

The scrub utility allows you to change and enhance the dictionary manually with words and patterns that were not picked up automatically by the system. The dictionary itself is never sent to NetApp and is entirely under your control. NetApp Customer Support can address the problems without needing to view sensitive data, so the keywords and patterns they receive do not interfere with their diagnostics.

To facilitate this interaction and to enable efficient troubleshooting:

- Every troubleshooting session must use the same dictionary. Failure to do so skews the data and prevents successful troubleshooting.
- A lookup dictionary displays the confirmed entries for keywords and patterns. The scrubbing process allows both you and the engineer to adjust and edit the scrubbed files until all possible patterns and keywords have been identified and entered into the dictionary. This ensures future exchanges will be expedited because the Lookup tables will be accurate, thorough and consistent.
- The dictionary remains valid and consistent unless and until you decide to clean (reset) the dictionary itself and build it again from scratch.

Lookup dictionary output

The Lookup dictionary is the output of the scrub utility that allows you to map back from the scrubbed fields to the real fields (for example, the name of a switch) in the original file.

The lookup utility displays the real values of what was scrubbed in the dictionary file.

Note: This is to be used only by the customer. This information is not available anywhere outside the customer's environment.

Information can be displayed by:

- Code assigned to a keyword.
- Keywords associated with a specific code

Therefore, you can perform a search using the {num} or by the entry that {num} replaced.

Troubleshooting with a scrubbed file

A customer in a secure environment has an unexpected slowdown occur in one of the OnCommand Insight applications. It is not a local problem. The customer contacts NetApp Customer Support. Both Customer Support and an engineer can use the Scrub Utilities to achieve a successful outcome. The process sequence for addressing the problem would resemble the steps shown below.

At the customer site

1. Customer Support requests a log file from the customer. (Types of data that may apply if you need to send NetApp include XML files, log files (for all components) and results of troubleshooting queries on the OnCommand Insight and the data warehouse database.
2. The customer has built a dictionary in the Scrub Utilities already (by clicking **Build** in the "Dictionary Admin" page). Typically the dictionary contains IP addresses, WWNs, VMs.)
3. The customer scrubs the file and checks its outcome. If there are keywords or patterns omitted, the customer edits the dictionary (using the "Dictionary Admin" and scrubs the log file again until it is correct. (These changes are made manually.)

At NetApp

1. NetApp Customer Support forwards the customer's scrubbed log file to an engineer.
2. The NetApp Customer Support representatives only see the results of the scrub. The engineer can identify and troubleshoot the problem without needing to know what the scrubbed material represents. Only the customer has access to the Lookup dictionary. Only the customer has the database!
3. If the engineer needs information about the scrubbed data, it can be discussed even by phone in a secure manner.

For example, if the engineer needs to know whether a device is a virtualizer, he can ask the customer this question and cite the code he sees in the log file for the address or name of the device. It could be "{123…}." When he gives the code to the customer, the customer can check using the "Lookup dictionary" by entering the code and receiving the real identity of the device. The customer is able to tell the support engineer that the device is or is not a virtualizer without ever needing to identify it to NetApp.

Searching in the lookup dictionary

If someone at the customer site wants to know the real identity of something references in the scrubbed file, the lookup dictionary contains this information.

Steps

1. Log in as an administrator.
2. In the Main Menu, click the **Troubleshooting** link.
3. Click **Scrub Utilities**.
4. Select the **Lookup dictionary** option.
5. Enter a code or keyword in the field.
6. Click **Lookup**.

   If you enter an incorrect format for the code, an error message displays. It could prompt you to check the server logs or correct the format of the scrubbed value.

Finding ghost data sources

If you have removed a device but the device data remains, you can locate any ghost data sources so that you can remove them.

Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click the **Find Ghost Datasources** link.

   OnCommand Insight produces a list of originators with their device information.

Adding a missing disk model

If acquisition fails due to a missing disk model, you can add the missing disk model to the `new_disk_models.txt` file and run acquisition again.

Before you begin

Understand the information needed to describe this data source because inaccurate information in the `new_disk_models.txt` file produces incorrect system data.

About this task

Because OnCommand Insight usually reads the internal disk model daily, the new data model information is updated automatically. Although you no longer need to restart the server for the changes to take effect, you can optionally restart the server to have the changes take effect before the next update.

Steps

1. Open Windows Notepad or another text editor to add the missing disk model information, .
2. Open the `new_disk_models.txt` file.

This file is located in your Insight installation at `C:\Program Files\SANscreen\jboss \server\onaro\deploy\datasources.war`.

3. Back up this file to prevent it from being overwritten during an upgrade.
4. Add the required information for the new data source.
5. Save the file in the `deploy\datasources.war` directory on your server.

Locating required information for new disk model

To locate the disk model information, identify the vendor and run an Internet search on the vendor and model number.

Steps

1. Vendor specifications do not always provide all of the necessary information based on the full model number. It is often useful to search for different parts of the model number string on the vendor's site to locate all of the information.
2. You need the disk vendor name, full model number, disk size and speed, and the interface type in order to define the new disk model in OnCommand Insight.

The following example shows the information used to define a Seagate disk:

```
ST3500320NS,Seagate,ST3500320NS,7200,500,SATA,,,1500,
http://h20000.www2.hp.com/bizsupport/
TechSupport/Document.jsp?
lang=en&cc=us&objectID=c00305257&prodTypeId=15351&prodSeriesId=397642
```

Content of the new_disk_models.txt file

The new_disk_models.txt file has required and optional fields. The fields are comma separated, so do not use commas *within* the fields.

Using a text editor, add the following information in this order:

1. **key**: use the model number (required)
2. **vendor**: name (required)
3. **model number**: full number (enter the same value as in "key")
4. **rpm of the disk**: for example 10000 or 15000 (required)
5. **size**: capacity in GB (required)
6. **interface type**: ATA, SATA, FC, SCSI, SAS, or FATA (required)
7. **seek time**: in ms (optional)
8. **potential transfer rate**: in MB/sec (optional)
9. **interface transfer rate**: the rate to and from the host in MB/sec (optional).
10. Additional useful information, such as a link to the vendor's specification page for the model (used only for Insight Perform) (optional).

Specialized administration toolkit

The Admin Tools menu provides specialized and infrequently-used options such as calculating RPO and customizing the Java UI Client Help menu.

The Admin Tools include these specialized options:

- Show OnCommand Insight logs
- Browse the SAN and list the devices in the SAN
- Calculate RPO
- Run the server
- Trigger a VM NAS job
- Customize the options in the Client Help menu

Opening the Admin Tools

To access the administrators' specialized maintenance toolkit, open the Admin Tools menu.

Steps

1. To use the Insight Admin Tools, open your browser and enter this URL customized for your environment:

   ```
   http://<server_host>:<server_port>/legacy/admin/all.do
   ```

2. Log in as an administrator. The Admin Tools menu opens.

Showing Insight logs

To research problems, you might analyze the contents of the OnCommand Insight logs. These logs contain errors and warnings and include information about the length of time that some actions took, which might also be helpful in troubleshooting.

Steps

1. To use the Insight Admin Tools, open your browser and enter this URL customized for your environment:

   ```
   http://<server_host>:<server_port>/legacy/admin/all.do
   ```

2. Select **Show Logs**.

   Note: This feature has the same function as the **Troubleshooting > Logs** option.

3. Select the server, acquisition, or remote log that you want to analyze.

   The last 100 lines from the selected log are displayed.

Browsing the SAN to locate devices

To locate a device or port in your environment based on its WWN, you can use the SANbrowse option.

Steps

1. To use the Insight Admin Tools, open your browser and enter this URL customized for your environment:

   ```
   http://<server_host>:<server_port>/legacy/admin/all.do
   ```

2. In the Admin Tools menu, select **SANbrowse**.

3. In the Search box, enter the WWN of the device or port you want to locate. Click **Search**.

   OnCommand Insight returns a report containing links to additional information, as shown in this example.

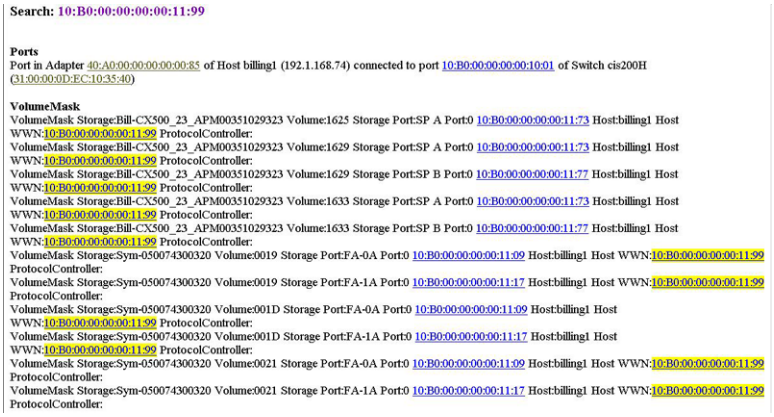

4. Click any of the links to see additional information related to a selected item. In this example, the first link in the search results was selected to produce this report with more links.

Calculating the Recovery Point Objective (RPO)

For users with the OnCommand Insight Replication license, you can calculate the Recovery Point Objective (RPO) time that would be used as the system recovery starting point. This is helpful if there is an outage and the system and data need to be restored.

Steps

1. To use the Insight Admin Tools, open your browser and enter this URL customized for your environment:

   ```
   http://<server_host>:<server_port>/legacy/admin/all.do
   ```

2. In the Admin Tools menu, select **Calculate RPO**.

   After the RPO is recalculated, a "successfully calculated" message displays above the Admin Tools menu options.

Running a paths and violations query

If you are concerned about a recent change in the environment, you can use the **Run server** option to report changes in the paths and any new violations.

Steps

1. To use the Insight Admin Tools, open your browser and enter this URL customized for your environment:

   ```
   http://<server_host>:<server_port>/legacy/admin/all.do
   ```

2. In the Admin Tools menu, select **Run server**.

   The results of this query list the changes found.

Triggering VM NAS jobs

To correlate the NFS share export IP address with the VM ESX IP address, you can use the Trigger VM NAS Job option in the Admin Tools. The NFS export IP relates to the NetApp data source, and the VM ESX IP relates to the Virtual Center data source.

Steps

1. To use the Insight Admin Tools, open your browser and enter this URL customized for your environment:

   ```
   http://<server_host>:<server_port>/legacy/admin/all.do
   ```

2. In the Admin Tools menu, select **Trigger VM NAS Job**.

   A report listing the results displays.

Copyright information

Copyright © 1994–2016 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

http://www.netapp.com/us/legal/netapptmlist.aspx

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

doccomments@netapp.com

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

Access Gateway
  Brocade data source support for *74*
accessing
  Violations Dashboard *139*
  web UI *12*
acquisition
  recordings for Support *146*
  restarting units *143*, *149*
acquisition failure
  adding missing disk to correct failure *157*
acquisition units
  monitoring health of *39*, *137*
adding missing
  disks *157*
administrators
  changing user passwords *20*
  experience level *9*
  permissions *18*
  specialized tools *159*
annotation rules
  deleting *54*
  modifying *54*
  predefined storage tiers *52*
  setting for automatic annotation assignment *52*
  setting precedence *53*
annotations
  assigning costs to *48*
  assignment of *48*
  configuring tier rules for automatic assignment *53*
  creating custom *49*
  default annotation rules *52*
  default types *46*
  defining *45*
  definition and purpose *45*
  deleting *52*
  importing *55*
  manually assigning to assets *50*
  modifying *51*
  predefined storage tiers *52*
  updating changes for reports *153*
Anomalies section
  on asset pages *36*
anomaly detection
  about *34*
  enabling *36*
anomaly detection engine
  monitoring health of *39*, *137*
anomaly detection software
  installing *35*
application
  enabling anomaly detection *36*
applications
  assigning to assets *43*
  defining *38*, *40*, *42*
  deleting *44*
  editing *44*
applying
  service packs *145*
architecture
  data warehouse *8*
  firewalls *8*
  OnCommand Insight *7*
architecture components, OnCommand Insight
  acquisition units *7*
  Data Warehouse *7*
  HTML5 web-based UI *7*
  Java UI *7*
  OnCommand Insight Server *7*
asset pages
  Anomalies section *36*
  opening from heat maps *136*
assets
  assigning applications to *43*
  manually assigning annotations to *50*
Assets Dashboard
  monitoring system *136*
assigning
  annotations *48*
  business entities to assets *44*
assigning multiple annotations
  using queries *60*
assurance violations
  configuring notifications *30*
ASUP
  configuring *30*
  process *147*
auditing
  user activities *137*
authentication
  LDAP *18*
  user *17*
automatic storage tiering
  configuring rules *53*
AutoSupport messages
  recommendation to use HTTPS *30*

B

backups
  OnCommand Insight *26*
  performing manually *146*
  restoring *144*
best practices
  backup directories *26*
  changing default passwords *12*
  defining annotation types before annotation rules *53*
  limited number of administrators *18*
  restoring OCI database *144*
Brocade
  EFCM data source *72*
  Enterprise Fabric Connectivity Manager data source *72*
  FC Switch Firmware (SSH) data source *74*
  Sphereon/Intrepid Switch (SNMP) *76*
Brocade EFCM data source
  requirements *72*

business entities
application data *38*, *40*, *42*
assigning to assets *44*
business unit *39*
creating *41*
hierarchy *40*
line of business *39*
project *39*
tenant *39*

C

CA
configuration for smart card *22*
CAC
browser requirements *23*
Cognos *25*
configuring for clients *23*
Data Warehouse configuration for *24*
certificate login
browser requirements *23*
client requirements *23*
configuration for Cognos *25*
configuration for Data Warehouse *24*
certificates
Cognos *25*
requirements for Smart Card and certificate login
Data Warehouse configuration for *24*
support *21*
Cisco
FC Switch Firmware data source *78*
CLARiiON
data source *85*
cloning
data sources *69*
cloud support
EMC data source *82*
clustered Data ONTAP
data source *114*
data source for Unified Manager *114*
Cognos
certificate configuration *25*
certificate login *25*
commands
3PAR *71*
7-Mode environments *115*
Brocade FC Switch Firmware *74*
Cisco *78*
EMC Celerra *83*
EMC Navisphere *85*
EMC VNX *93*
EMC VPLEX *94*
HDS HiCommand Device Manager *97*
HP EVA *102*
Hyper-V (WMI) *111*
IBM DS *104*
IBM SVC *107*
IBM XIV *110*
Solutions Enabler *90*
comments
how to send feedback about documentation *164*
committing
data source patches *134*
components, OnCommand Insight architecture
acquisition units *7*
Data Warehouse *7*
HTML5 web-based UI *7*
Java UI *7*
OnCommand Insight Server *7*
configuration
advanced *153*
advanced tools *147*
changes *126*
testing data sources *70*
configuration requirements
EMC XtremIO *88*
configuring
ASUP *30*
data sources *67*
email *28*
Insight *12*
Java UI Client *153*
performance policy evaluation precedence *64*
performance violation notifications *30*
SNMP notifications *28*
syslog *29*
users in LDAP *18*
cost
assigning to annotations *48*
creating
business entities *41*
custom annotations *49*
performance policies *62*
queries *57*
creating annotation rules
for automatic annotation assignment *52*
custom annotations
creating *49*
custom dashboards
deleting *34*
setting up *32*
Customer Support
configuring information sent to *30*
generate logs for *150*
product documentation *10*
scrubbed file example *156*
scrubbing report for *154*
sending data to *147*
customization
using annotations to monitor your environment *45*

D

Dashboard, Violations
accessing *139*
dashboards
setting up custom *32*
data collection
coordinating changes between your environment and Insight *126*
options *67*
data sources
adding *67*
adding new, by patch *69*
approving patches *134*
Brocade EFCM *72*

changing passwords for a group *125*
cloning to quickly add *69*
deleting *131*
editing multiple setup information *130*
editing setup information *129*
examining detailed summaries *126*
finding serious problems *136*
HP CommandView Lite *101*
importing *68*
installing a patch *132*, *133*
installing new *66*
locating ghost data *157*
managing *66*
managing patches *133*
NetApp *161*
patching
new data sources *69*
patching groups *132*
permissions to change *18*
polling *129*
researching failures *127*
rolling back patches *135*
setting up *66*
support matrix *67*
testing configurations *70*
vendor-specific configuration information *71*
data sources, supported
3PAR InServ Firmware (SSH) *71*
Brocade EFCM *72*
Brocade FC Switch Firmware (SSH) *74*
Brocade Sphereon/Intrepid Switch (SNMP) *76*
Cisco FC Switch Firmware (SNMP) *78*
Dell Compellent Storage Center *81*
EMC Atmos *82*
EMC Celerra SSH *83*
EMC CLARiiON (NaviCLI) *85*
EMC Data Domain *87*
EMC ECC StorageScope *88*
EMC Isilon *89*
EMC Solutions Enabler (CLI) with SMI-S Performance *90*
EMC VNX (SSH) *93*
EMC VPLEX *94*
EMC XtremIO *88*
Fujitsu Eternus *95*
HDS (HiCommand Device Manager Lite *99*
HDS (HiCommand Device Manager) *97*
HP (CommandView) *99*
HP (EVA) *102*
HP CommandView Lite *101*
IBM DS (CLI) *104*
IBM PowerVM *105*
IBM SVC (SSH) *107*
IBM Tivoli Monitoring *109*
IBM TotalStorage DS4000 *109*
IBM XIV (CLI) *110*
in 7-Mode environments *115*
Microsoft Hyper-V *111*
NetApp clustered Data ONTAP *114*
NetApp E-Series *118*
NetApp Host and VM File Systems *119*
QLogic FC Switch (SNMP) *122*
Violin Memory *122*
VMware vSphere *124*
Data Warehouse
architecture *8*
certificate configuration *24*
Data Warehouse server
monitoring health of *39*, *137*
databases
automatic weekly backups *26*
backing up *146*
OnCommand Insight *144*
SQL *152*
defining
applications *38*, *40*, *42*
deleting
annotation rules *54*
annotations *51*, *52*
applications *44*
custom dashboards *34*
data sources *131*
performance policies *65*
queries *59*
users *17*
Dell Compellent Storage Center
Insilon (SSH) data source *81*
devices
data source setup *66*
deleting inactive *126*
disk models
locating vendor information *158*
text file format *158*
documentation
how to receive automatic notification of changes to *164*
how to send feedback about *164*
locating on Support site *10*
online *10*
DWH
See Data Warehouse

E

editing
annotations *51*
applications *44*
data source setups *129*
multiple data source setups *130*
performance policies *64*
user accounts *21*
editing multiple applications
using queries *59*
email
changing user *21*
configuring for receiving Insight reports *28*
prerequisites *27*
to deliver reports *27*
EMC
Atmos data source *82*
Celerra data source *83*
CLARiiON data source *85*
ECC StorageScope data source *88*
Isilon SSH data source *89*
Solutions Enabler with SMI-S Performance *90*
VNX (SSH) data source *93*

VPLEX data source *94*
EMC Data Domain
data source *87*
EMC XtremeIO data source
configuration requirements *88*
error messages
data source *127*
non-compliant licenses *145*
evaluation precedence
performance policies *64*
events
data source *126*
exclamation mark
meaning of in user interface *126*
exporting to .CSV file
queries *58*

F

feedback
how to send comments about documentation *164*
Fibre Channel
monitoring *141*
port metrics *141*
Fibre Channel environment
monitoring *140*
files
backing up *146*
license *13*
restoring OCI database *144*
firewalls
architecture *8*
Fujitsu
Eternus data source *95*

H

HDS
data sources *97*, *99*
heat maps
showing high IOPs *136*
Help
access directly *10*
locating *10*
hosts
ESX *124*
HP
data sources *99*
EVA data source *102*
HP CommandView AE Lite data source
configuration *101*
requirements *101*
HTTP
protocol for ASUP *30*
hypervisor
data source *111*

I

IBM
data sources *104*, *105*, *107*, *109*, *110*
importing
data sources *68*
information
how to send feedback about improving
documentation *164*
Insight
backups *26*
component list *151*
data source patches *132*
data source setup requirements *66*
data source testing *70*
documentation *10*
license types *14*
licenses *14*, *145*
licenses, installing *13*
list of common maintenance tasks *136*
logging in *12*
maintenance *126*
number of database objects *151*
overview *7*
product information *10*
status *136*
Insight reports
configuring email for *28*
insight server
restart *147*
installation
of anomaly detection software *35*
installed software packages
displaying *147*
installing
licenses *13*
IP addresses
correlating export *161*
iSCSI
IBM data source support *110*

J

Java UI
access Help directly *10*
Help *10*
setting business entities *82*

L

LDAP
authentication *18*
rediscover user *21*
requirements for smart car and certificate login *22*
requirements for Smart Card and certificate login *25*
users *18*, *21*
licenses
for Insight *14*
installing *13*
not compliant *145*
replacing older versions *145*
types *14*
updating expired *14*, *145*
log files
accessing *147*
login
certificate *21*

Smart Card *21*
logs
audit *137*
creating a zip file for Support *146*
generate for Customer Support *150*
OnCommand Insight *159*
syslog *30*

M

machine-learning anomaly detection software
installing *35*
maintenance
coordinating changes between your environment and Insight *126*
specialized tools *159*
maintenance tasks
list of most common *136*
managing
performance policies *61*
Microsoft
data source *111*
missing disks
adding missing disk models *157*
modifying
annotation rules *54*
annotations *51*
data source setup *129*
multiple data source setups *130*
queries *59*
monitoring
anomalies *34*
applications *34*
data sources *66*
Fibre Channel *141*
Fibre Channel environment *140*
Insight component health *39*, *137*
violations *139*
monitoring your environment
using asset annotations *45*
moving
MySQL data *152*
MySQL data with command line *152*
MySQL
migrating data *152*

N

NetApp
data sources *114*, *115*, *119*
NetApp Customer Support
bundling logs for troubleshooting *146*
Insight documentation *10*
preparing troubleshooting report *146*
sending data to *147*
NetApp E-Series
data sources *118*
NetApp Host and VM File Systems
data source *119*
network
system information *151*
network maintenance
list of most common tasks *136*
notifications, SNMP
configuring *28*

O

OCI
See OnCommand Insight
OnCommand Insight
documentation *10*
Help *10*
managing data sources *66*
overview *7*
setup tasks *12*

P

passwords
changing for data source groups *125*
changing user *20*
coordinating changes *126*
encrypted for ASUP *30*
LDAP *18*
new users *17*
patches
approving *134*
coordinating changes between your environment and Insight *126*
data source, managing *133*
groups of data sources *132*
installing data source *132*, *133*
new data source, installing *66*
rolling back *135*
performance
alerts in syslog *29*
performance analysis
skills required *9*
performance policies
configuring evaluation precedence *64*
creating *62*
deleting *65*
editing *64*
evaluation precedence *64*
managing *61*
performance violations
configuring notifications *30*
planning
your data source support *67*
polling
data sources *129*
port metrics
Fibre Channel *141*
ports
locating *160*
precedence
setting for annotation rules *53*

Q

queries
assigning multiple annotations *60*
creating *57*

deleting *59*
editing multiple applications *59*
exporting to .CSV file *58*
modifying *59*
removing multiple annotations *60*
removing multiple applications *59*
viewing *58*

R

red circle
meaning of in the user interface *136*
removing
data sources *131*
user account *21*
removing multiple annotations
using queries *60*
removing multiple applications
using queries *59*
replication
calculate Recovery Point Objective *160*
reporting
Atmos data *82*
customized with annotations *46*
restart
insight server *147*
restarting
acquisition units *143*
Insight servers *148*
polling *129*
restoring
backups *144*
resume
polling *129*
roles
7-Mode environments *115*
users *17*, *18*

S

scaling
determine number of database objects *151*
security
firewall architecture *8*
user roles *18*
servers
backup *26*
EMC VPLEX *94*
LDAP *18*
monitoring health of *39*, *137*
proxy for ASUP *30*
restarting OnCommand Insight *148*
Service Level annotation
assigning cost to *48*
service packs
applying *145*
setting hours to ignore data updates
configuration *70*, *126*, *147*, *149*, *153*
setting up
custom dashboards *32*
Insight *12*
licenses *13*
Smart Card
support *21*
smart cards
host configuration for *22*
Smart Cards
Cognos, configuration for *25*
configuring for clients *23*
Data Warehouse configuration for *24*
SNMP
configuring notifications *28*
Qlogic data source *122*
software
anomaly detection *34*
software, anomaly detection
installing *35*
solid red circle
meaning *143*
storage
predefined tier annotations *52*
suggestions
how to send feedback about documentation *164*
Switch Level annotation
assigning cost to *48*
syslog
configuring *29*
contents *30*
system architecture
OnCommand Insight *7*

T

terabytes
licensed *145*
testing
added data sources *67*
data source configurations *70*
Tier annotation
assigning cost to *48*
tiers
checking changes after upgrading *154*
predefined annotations *52*
Tivoli
data source *109*
tracking data
using asset annotations *45*
troubleshooting
data collection problems *126*
database backup *146*
finding changes in paths *161*
locating a device or port *160*
locating ghost data *157*
logs *159*
preparing a special report *146*
restoring OnCommand Insight *144*
sending report to Support *147*
server resources status *154*
tools *147*
using audit log *137*
zipping logs for Support *146*
Twitter
how to receive automatic notification of documentation changes *164*

U

Unified Manager (UM)
    data source support *114*
upgrading
    check tiers after *154*
users
    auditing activities *137*
    authenticated from LDAP *18*
    changing passwords *20*
    defining *17*
    deleting *21*
    editing accounts *21*
    experience level *9*
    roles *18*

V

vendors
    finding disk model information *158*
    supported data sources *71*
vFiler unit
    connection requirements *115*
videos
    on Communities site *10*
viewing
    queries *58*
violations
    by policy *139*
    changes in syslog *29*
    configuring
        assurance violation notifications *30*
    history *139*
    in syslog *30*
    monitoring *139*
    SNMP notifications of *28*
Violations Dashboard
    accessing *139*
Violin Memory
    data source *122*
VM
    data sources *105*, *124*
VMI
    data sources *111*

W

web UI
    access data source support matrix *67*
    access Help directly *10*
    accessing *12*
    advanced configuration tools *147*
    documentation *10*
WWN
    locating a device or port using *160*