



Configuration and administration

OnCommand Insight

NetApp
June 10, 2024

This PDF was generated from <https://docs.netapp.com/us-en/oncommand-insight/config-admin/opening-insight.html> on June 10, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Configuration and administration 1
 - Setting up Insight 1
 - Insight Security 91
 - Smart Card and certificate login support 104
 - Configuring Data Warehouse for Smart Card and certificate login 116
 - Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.5 through 7.3.9) 117
 - Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.10 and later) 118
 - Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.5 to 7.3.9) 120
 - Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.10 and later) 122
 - Importing SSL certificates 124
 - Your business entities hierarchy 127
 - Defining annotations 130
 - Querying assets 144
 - Insight data source management 150
 - Device resolution 250
 - Maintaining Insight 269
 - Monitoring your environment 291

Configuration and administration

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.

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 web UI.

Steps

1. Do one of the following:

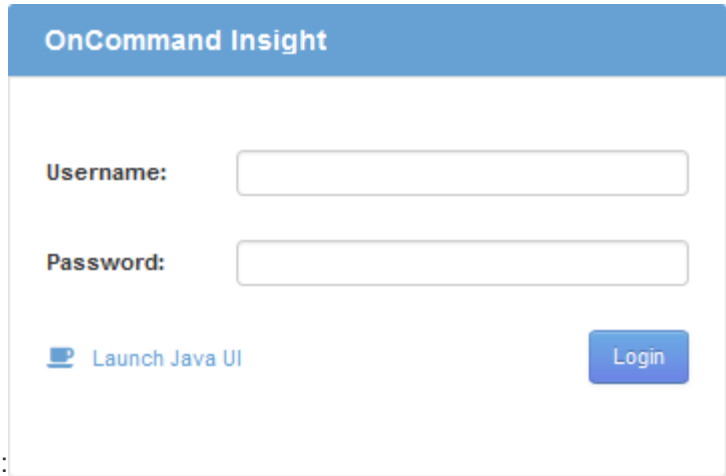
- Open Insight on the Insight server:

```
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 image shows a login dialog box for OnCommand Insight. It has a blue header bar with the text "OnCommand Insight". Below the header, there are two input fields: "Username:" and "Password:". To the left of the "Launch Java UI" link is a small blue icon of a laptop. The "Login" button is a blue rectangle with white text.

The OnCommand Insight dialog box displays:

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

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



An Insight browser session that is inactive for 30 minutes is timed out and you are automatically logged out of the system. For added security, it is recommended to close your browser after logging out of Insight.

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 .lic file.

Steps

1. Open the license file in a text editor and copy the text.
2. Open Insight in your browser.
3. On the Insight toolbar, click **Admin**.
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**.
10. If you are using the Insight consumption licensing model, you must check the box to **Enable sending usage information to NetApp** in the **Send usage information** section. Proxy must be properly configured and enabled for your environment.

After you finish

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

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

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
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
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
Qtrees	Yes	No performance data or charts
Quota	Yes	No performance data or charts
Path	Yes	No performance data or charts
Zone	Yes	No performance data or charts
Zone member	Yes	No performance data or charts
Generic device	Yes	No performance data or charts
Tape	Yes	No performance data or charts
Masking	Yes	No performance data or charts
ISCSI sessions	Yes	No performance data or charts
ICSI network portals	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)
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)

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.



Security best practices dictate that administrators configure the host operating system to prevent the interactive login of non-administrator/standard users.

Steps

1. Open Insight in your browser.
2. On the Insight toolbar, click **Admin**.
3. Click **Setup**.
4. Select the **Userstab**.
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.

Results

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 the 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.

- 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 permits you to 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 Insight for LDAP(s)

OnCommand Insight must be configured with Lightweight Directory Access Protocol (LDAP) settings as they are configured in your corporate LDAP domain.

Before configuring Insight for use with LDAP or secure LDAP (LDAPS), make note of the Active Directory configuration in your corporate environment. Insight settings must match those in your organization's LDAP domain configuration. Review the concepts below before configuring Insight for use with LDAP, and check with your LDAP domain administrator for the proper attributes to use in your environment.

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



OnCommand Insight supports LDAP and LDAPS via Microsoft Active Directory server or Azure AD. Additional LDAP implementations may work but have not been qualified with Insight. The procedures in these guides assume that you are using Microsoft Active Directory Version 2 or 3 LDAP (Lightweight Directory Access Protocol).

User Principal Name attribute:

The LDAP User Principal Name attribute (userPrincipalName) is what Insight uses as the username attribute. User Principal Name is guaranteed to be globally unique in an Active Directory (AD) forest, but in many large organizations, a user's principal name may not be immediately obvious or known to them. Your organization might use an alternative to the User Principal Name attribute for primary user name.

Following are some alternative values for the User Principal Name attribute field:

• sAMAccountName

This user attribute is the legacy pre-Windows 2000 NT username - this is what most users are accustomed to logging into their personal Windows machine. This is not guaranteed to be globally unique throughout an AD forest.



sAMAccountName is case-sensitive for the User Principal Name attribute.

• mail

In AD environments with MS Exchange, this attribute is the primary e-mail address for the end user. This should be globally unique throughout an AD forest, (and also familiar for end users), unlike their userPrincipalName attribute. The mail attribute will not exist in most non-MS Exchange environments.

• referral

An LDAP referral is a domain controller's way of indicating to a client application that it does not have a

copy of a requested object (or, more precisely, that it does not hold the section of the directory tree where that object would be, if in fact it exists) and giving the client a location that is more likely to hold the object. The client in turn uses the referral as the basis for a DNS search for a domain controller. Ideally, referrals always reference a domain controller that indeed holds the object. However, it is possible for the referred-to domain controller to generate yet another referral, although it usually does not take long to discover that the object does not exist and to inform the client.



sAMAccountName is generally preferred over User Principal Name. sAMAccountName is unique in the domain (though it may not be unique in the domain forest), but it is the string domain users typically use for login (For example, *netapp\username*). The Distinguished Name is the unique name in the forest, but is generally not known by the users.



On the Windows system part of the same domain, you can always open a command prompt and type SET to find the proper domain name (USERDOMAIN=). The OCI login name will then be USERDOMAIN\sAMAccountName.

For the domain name **mydomain.x.y.z.com**, use DC=x, DC=y, DC=z, DC=com in the Domain field in Insight.

Ports:

The default port for LDAP is 389, and the default port for LDAPs is 636

Typical URL for LDAPs: ldaps://<ldap_server_host_name>:636

Logs are at: \\<install_directory>\SANscreen\wildfly\standalone\log\ldap.log

By default, Insight expects the values noted in the following fields. If these change in your Active Directory environment, be sure to change them in the Insight LDAP configuration.

Role attribute
memberOf
Mail attribute
mail
Distinguished Name attribute
distinguishedName
Referral
follow

Groups:

To authenticate users with different access roles in the OnCommand Insight and DWH servers, you must create groups in Active Directory and enter those group names in OnCommand Insight and DWH servers. The

group names below are examples only; the names you configure for LDAP in Insight must match the ones set up for your Active Directory environment.

Insight Group	Example
Insight server administrator group	insight.server.admins
Insight administrators group	insight.admins
Insight users group	insight.users
Insight guests group	insight.guests
Reporting administrator group	insight.report.admins
Reporting pro authors group	insight.report.proauthors
Reporting authors group	insight.report.business.authors
Reporting consumers group	insight.report.business.consumers
Reporting recipients group	insight.report.recipients

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.

Before you begin

You must know the user and group attributes that have been configured for Insight in your LDAP domain.

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

About this task

OnCommand Insight supports LDAP and LDAPS via Microsoft Active Directory server. Additional LDAP implementations may work but have not been qualified with Insight. 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 **Admin > Setup > Users** list.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Setup**.
3. Click the **Users** tab.

4. Scroll to the LDAP section, as shown here.

LDAP

LDAP integration enables authentication of users via LDAP (or ActiveDirectory). This is done by assigning these users to LDAP groups. The groups are used to identify the user permissions.

☒ Enable LDAP

Please provide credentials for a user authorized for directory lookup queries.

LDAP servers:

User:

Password:

[Show more](#) ▼

5. Click **Enable LDAP** to allow the LDAP user authentication and authorization.

6. Fill in the fields:

- **LDAP servers:** Insight accepts a comma-separated list of LDAP URLs. Insight attempts to connect to the provided URLs without validating for LDAP protocol.



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

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

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

or, if using the default port:

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

When entering multiple LDAP servers in this field, ensure that the correct port number is used in each entry.

- **User name:** Enter the credentials for a user authorized for directory lookup queries on the LDAP servers.
- **Password:** Enter the password for the above user. To confirm this password on the LDAP server, click **Validate**.

7. If you want to define this LDAP user more precisely, click **Show more** and fill in the fields for the listed attributes.

These settings must match the attributes configured in your LDAP domain. Check with your Active Directory administrator if you are unsure of the values to enter for these fields.

- **Admins group**

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

- **Users group**

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

- **Guests group**

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

- **Server admins group**

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

- **Timeout**

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

- **Domain**

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

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

- **User principal name attribute**

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

- **Role attribute**

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

- **Mail attribute**

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



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

- **Distinguished name attribute**

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

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. On the Insight toolbar, click **Admin**.
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. On the Insight toolbar, click **Admin**.
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. On the Insight toolbar, click **Admin**.
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**.

Setting a Login Warning Message

OnCommand Insight allows administrators to set a custom text message that is displayed when users log in.

Steps

1. To set the message in the OnCommand Insight Server:
 - a. Navigate to **Admin > Troubleshooting > Advanced Troubleshooting > Advanced Settings**.
 - b. Enter your login message in the text area.
 - c. Click the **Client displays login warning message** checkbox.
 - d. Click **Save**.

The message will display upon login for all users.

2. To set the message in the Data Warehouse (DWH) and Reporting (Cognos):
 - a. Navigate to **System Information** and click the **Login Warning** tab.
 - b. Enter your login message in the text area.
 - c. Click **Save**.

The message will display upon DWH and Cognos Reporting login for all users.

Insight Security

The 7.3.1 release of OnCommand Insight introduced security features that allow Insight environments to operate with enhanced security. The features include improvements to

encryption, password hashing, and the ability to change internal user passwords and key pairs that encrypt and decrypt passwords. You can manage these features on all servers in the Insight environment.

The default installation of Insight includes a security configuration where all sites in your environment share the same keys and the same default passwords. To protect sensitive data, NetApp recommends you change the default keys and the Acquisition user password after an installation or upgrade.

Data source encrypted passwords are stored in the Insight Server database. The Server has a public key and encrypts passwords when a user enters them in a WebUI data source configuration page. The Server does not have the private keys required to decrypt the data source passwords stored in the Server database. Only Acquisition Units (LAU, RAU) have the data source private key required to decrypt data source passwords.

Rekeying servers

Using default keys introduces security vulnerability in your environment. By default, data source passwords are stored encrypted in the Insight database. They are encrypted using a key that is common to all Insight installations. In a default configuration, an Insight database sent to NetApp includes passwords that could theoretically be decrypted by NetApp.

Changing the Acquisition user password

Using the default 'Acquisition' user password introduces security vulnerability into your environment. All Acquisition Units use the "Acquisition" user to communicate with the Server. RAUs with default passwords can theoretically connect to any Insight server using default passwords.


Upgrade and installation considerations

When your Insight system contains non-default security configurations (you have rekeyed or changed passwords), you must back up your security configurations. Installing new software, or in some cases upgrading software, reverts your system to a default security configuration. When your system reverts to the default configuration, you must restore the non-default configuration in order for the system to operate correctly.

Managing keys in a complex service provider environment

A service provider can host multiple OnCommand Insight customers collecting data. The keys protect customer data from unauthorized access by multiple customers on the Insight server. Each customer's data is protected by their specific key pairs.

This implementation of Insight could be configured as shown in the following illustration.



You need to create individual keys for each customer in this configuration. Customer A requires identical keys for both RAUs. Customer B requires a single set of keys.

The steps you would take to change encryption keys for Customer A:

1. Perform a remote login to the server hosting RAU1.
2. Start the security admin tool.
3. Select Change Encryption Key to replace the default keys.
4. Select Backup to create a backup zip file of the security configuration.
5. Perform a remote login to the server hosting RAU2.
6. Copy the backup zip file of the security configuration to RAU2.
7. Start the security admin tool.
8. Restore the security backup from RAU1 to the current server.

The steps you would take to change encryption keys for Customer B:

1. Perform a remote login to the server hosting RAU3.
2. Start the security admin tool.
3. Select Change Encryption Key to replace the default keys.
4. Select Backup to create a backup zip file of the security configuration.

Managing security on the Insight server

The `securityadmin` tool allows you to manage security options on the Insight server. Security management includes changing passwords, generating new keys, saving and restoring security configurations you create, or restoring configurations to the default settings.

About this task

You use the `securityadmin` tool to manage security:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat`
- Linux - `/bin/oci-securityadmin.sh`

Steps

1. Perform a remote login to the Insight server.
2. Start the security admin tool in interactive mode:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat -i`
- Linux - `/bin/oci-securityadmin.sh -i`

The system requests login credentials.

3. Enter the user name and password for an account with “Admin” credentials.
4. Select **Server**.

The following server configuration options are available:

- **Backup**

Creates a backup zip file of the vault containing all passwords and keys and places the file in a location specified by the user, or in the following default locations:

- Windows - `C:\Program Files\SANscreen\backup\vault`
- Linux - `/var/log/netapp/oci/backup/vault`

- **Restore**

Restores the zip backup of the vault that was created. Once restored, all passwords and keys are reverted to values existing at the time of the backup creation.



Restore can be used to synchronize passwords and keys on multiple servers, for example: - Change the server encryption key on one server - Create a backup of the vault - Restore the vault backup to the second server

- **Change Encryption Key**

Change the server encryption key that is used to encrypt or decrypt proxy user passwords, SMTP user passwords, LDAP user passwords, and so on.



When you change encryption keys, you should backup your new security configuration so that you can restore it after an upgrade or installation.

- **Update Password**

Change password for the internal accounts that are used by Insight. The following options are displayed:

- `_internal`
- `acquisition`
- `cognos_admin`
- `dwh_internal`
- `hosts`
- `inventory`
- `root`



Some accounts need to be synchronized when passwords are changed. For example, if you change the password for the 'acquisition' user on the server, you need to change the password for the 'acquisition' user on the LAU, RAU, and DWH to match. Also, when you change passwords, you should backup your new security configuration so that you can restore it after an upgrade or installation.

• **Reset to Defaults**

Resets keys and passwords to default values. Default values are those provided during installation.

• **Exit**

Exit the `securityadmin` tool.

1. Chose the option you want to change and follow the prompts.

Managing security on the local acquisition unit

The `securityadmin` tool allows you to manage security options on the local acquisition user (LAU). Security management includes managing keys and passwords, saving and restoring security configurations you create or restoring configurations to the default settings.

Before you begin

You must have `admin` privileges to perform security configuration tasks.

About this task

You use the `securityadmin` tool to manage security:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat`
- Linux - `/bin/oci-securityadmin.sh`

Steps

1. Perform a remote login to the Insight server.
2. Start the security admin tool in interactive mode:
 - Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat -i`
 - Linux - `/bin/oci-securityadmin.sh -i`

The system requests login credentials.

3. Enter the user name and password for an account with “Admin” credentials.
4. Select **Local Acquisition Unit** to reconfigure the Local Acquisition Unit security configuration.

The following options are displayed:

- **Backup**

Creates a backup zip file of the vault containing all passwords and keys and places the file in a location specified by the user, or in the following default locations:

- Windows - C:\Program Files\SANscreen\backup\vault
- Linux - /var/log/netapp/oci/backup/vault

- **Restore**

Restores the zip backup of the vault that was created. Once restored, all passwords and keys are reverted to values existing at the time of the backup creation.



Restore can be used to synchronize passwords and keys on multiple servers, for example: - Change encryption keys on the LAU - Create a backup of the vault - Restore the vault backup to each of the RAUs

- **Change Encryption Keys**

Change the AU encryption keys used to encrypt or decrypt device passwords.



When you change encryption keys, you should backup your new security configuration so that you can restore it after an upgrade or installation.

- **Update Password**

Change password for 'acquisition' user account.



Some accounts need to be synchronized when passwords are changed. For example, if you change the password for the 'acquisition' user on the server, you need to change the password for the 'acquisition' user on the LAU, RAU, and DWH to match. Also, when you change passwords, you should backup your new security configuration so that you can restore it after an upgrade or installation.

- **Reset to Defaults**

Resets acquisition user password and acquisition user encryption keys to default values, Default values are those provided during installation.

- **Exit**

Exit the `securityadmin` tool.

5. Chose the option you want configure and follow the prompts.

Managing security on an RAU

The `securityadmin` tool allows you to manage security options on RAUs. You might need to backup or restore a vault configuration, change encryption keys, or update passwords for the acquisition units.

About this task

You use the `securityadmin` tool to manage security:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat`
- Linux - `/bin/oci-securityadmin.sh`

One scenario for updating the security configuration for the LAU, RAU is to update the 'acquisition' user password when the password for that user has been changed on the server. All of the RAUs, and the LAU use the same password as that of the server 'acquisition' user to communicate with the server.

The 'acquisition' user only exists on the Insight server. The RAU or LAU logs in as that user when they connect to the server.

Use the following steps to manage security options on an RAU:

Steps

1. Perform a remote login to the server running the RAU
2. Start the security admin tool in interactive mode:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat -i`
- Linux - `/bin/oci-securityadmin.sh -i`

The system requests login credentials.

3. Enter the user name and password for an account with “Admin” credentials.

The system displays the menu for the RAU.

- **Backup**

Creates a backup zip file of the vault containing all passwords and keys and places the file in a location specified by the user, or in the following default locations:

- Windows - `C:\Program Files\SANscreen\backup\vault`
- Linux - `/var/log/netapp/oci/backup/vault`

- **Restore**

Restores the zip backup of the vault that was created. Once restored, all passwords and keys are reverted to values existing at the time of the backup creation.



Restore can be used to synchronize passwords and keys on multiple servers, for example: - Change encryption keys on one server - Create a backup of the vault - Restore the vault backup to the second server

- **Change Encryption Keys**

Change the RAU encryption keys used to encrypt or decrypt device passwords.



When you change encryption keys, you should backup your new security configuration so that you can restore it after an upgrade or installation.

- **Update Password**

Change password for 'acquisition' user account.



Some accounts need to be synchronized when passwords are changed. For example, if you change the password for the 'acquisition' user on the server, you need to change the password for the 'acquisition' user on the LAU, RAU, and DWH to match. Also, when you change passwords, you should backup your new security configuration so that you can restore it after an upgrade or installation.

- **Reset to Defaults**

Resets encryption keys and passwords to default values. Default values are those provided during installation.

- **Exit**

Exit the `securityadmin` tool.

Managing security on the Data Warehouse

The `securityadmin` tool allows you to manage security options on the Data Warehouse server. Security management includes updating internal passwords for internal users on the DWH server, creating backups of the security configuration, or restoring configurations to the default settings.

About this task

You use the `securityadmin` tool to manage security:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat`
- Linux - `/bin/oci-securityadmin.sh`

Steps

1. Perform a remote login to the Data Warehouse server.
2. Start the security admin tool in interactive mode:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat -i`
- Linux - `/bin/oci-securityadmin.sh -i`

The system requests login credentials.

3. Enter the user name and password for an account with “Admin” credentials.

The system displays the security admin menu for the Data Warehouse:

- **Backup**

Creates a backup zip file of the vault containing all passwords and keys and places the file in a location specified by the user, or in the default location:

- Windows - C:\Program Files\SANscreen\backup\vault
- Linux - /var/log/netapp/oci/backup/vault

- **Restore**

Restores the zip backup of the vault that was created. Once restored, all passwords and keys are reverted to values existing at the time of the backup creation.



Restore can be used to synchronize passwords and keys on multiple servers, for example: - Change encryption keys on one server - Create a backup of the vault - Restore the vault backup to the second server

+

- **Change encryption keys**

Change the DWH encryption key used to encrypt or decrypt passwords such as connector passwords and SMPT passwords.

- **Update Password**

Change password for a specific user account.

- _internal
- acquisition
- cognos_admin
- dwh
- dwh_internal
- dwhuser
- hosts
- inventory
- root



When you change the dwhuser, hosts, inventory, or root passwords, you have the option to use SHA-256 password hashing. This options requires that all clients accessing the accounts use SSL connections.

- **Reset to Defaults**

Resets encryption keys and passwords to default values. Default values are those provided during installation.

- **Exit**

Exit the `securityadmin` tool.

Changing OnCommand Insight internal user passwords

Security policies might require you to change the passwords in your OnCommand Insight environment. Some of the passwords on one server exist on a different server in the environment, requiring that you change the password on both servers. For example, when you change the “inventory” user password on the Insight Server you must match the “inventory” user password on the Data Warehouse server Connector configured for that Insight Server.

Before you begin



You should understand the dependencies of the user accounts before you change passwords. Failing to update passwords on all required servers will result in communication failures between the Insight components.

About this task

The following table lists the internal user passwords for the Insight Server and lists the Insight components that have dependent passwords that need to match the new password.

Insight Server Passwords	Required changes
_internal	
acquisition	LAU, RAU
dwh_internal	Data Warehouse
hosts	
inventory	Data Warehouse
root	

The following table lists the internal user passwords for the Data Warehouse and lists the Insight components that have dependent passwords that need to match the new password.

Data Warehouse Passwords	Required changes
cognos_admin	
dwh	
dwh_internal (Changed using the Server Connector configuration UI)	Insight server

dwhuser	
hosts	
inventory (Changed using the Server Connector configuration UI)	Insight server
root	

Changing passwords in the DWH Server Connection Configuration UI

The following table lists the user password for the LAU and lists the Insight components that have dependent passwords that need to match the new password.

LAU Passwords	Required changes
acquisition	Insight Server, RAU

Changing the “inventory” and “dwh_internal” passwords using the Server Connection Configuration UI

If you need to change the “inventory” or “dwh_internal” passwords to match those on the Insight server you use the Data Warehouse UI.

Before you begin

You must be logged in as administrator to perform this task.

Steps

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

The **Edit Connector** screen is displayed.

Edit Connector

ID:	<input type="text" value="1"/>
Encryption:	<input type="text" value="Enabled"/>
Name:	<input type="text" value="Oci-stg06-s12r2.nane.netapp.com"/>
Host:	<input type="text" value="Oci-stg06-s12r2.nane.netapp.com"/>
Database user name:	<input type="text" value="inventory"/>
Database password:	<input type="password" value="••••••••"/>

[Advanced](#) ▼

3. Enter a new “inventory” password for the **Database password** field.
4. Click **Save**
5. To change the “dwh_internal” password, click **Advanced**.

The Edit Connector Advanced screen is displayed.

Edit Connector

ID:	<input type="text" value="1"/>
Encryption:	<input type="text" value="Enabled"/>
Name:	<input type="text" value="Oci-stg06-s12r2.nane.netapp.com"/>
Host:	<input type="text" value="Oci-stg06-s12r2.nane.netapp.com"/>
Database user name:	<input type="text" value="inventory"/>
Database password:	<input type="password" value="....."/>
Server user name:	<input type="text" value="dwh_internal"/>
Server password:	<input type="password" value="....."/>
HTTPS port:	<input type="text" value="443"/>
TCP port:	<input type="text" value="3306"/>

Basic ^

6. Enter the new password in the **Server password** field:

7. Click save.

Changing the dwh password using the ODBC Administration tool

When you change the password on for the dwh user on the Insight server, the password must also be changed on the Data Warehouse server. You use the ODBC Data Source Administrator tool to change the password on the Data Warehouse.


Before you begin

You must perform a remote login to the Data Warehouse server using an account with administrator privileges.

Steps


1. Perform a remote login to the server hosting that Data Warehouse.
2. Access the ODBC Administration tool at `C:\Windows\SysWOW64\odbcad32.exe`

The system displays the ODBC Data Source Administrator screen.



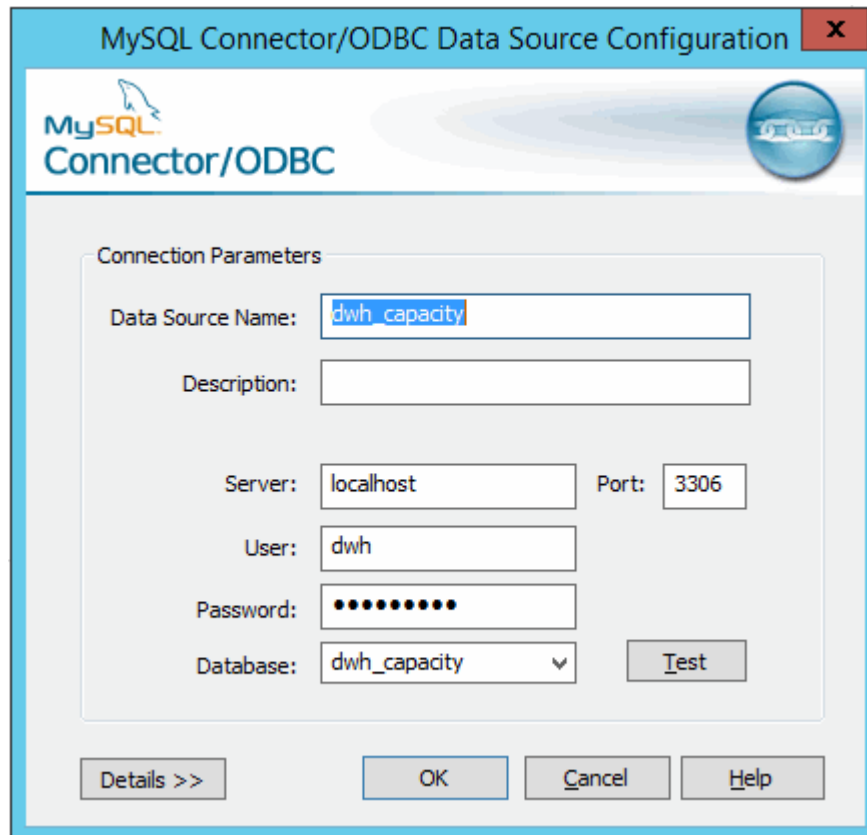
3. Click **System DSN**

The system data sources are displayed.



4. Select an OnCommand Insight Data Source from the list.
5. Click **Configure**

The Data Source Configuration screen is displayed.



The image shows a Windows-style dialog box titled "MySQL Connector/ODBC Data Source Configuration". The dialog has a blue header bar with the title and a close button (X). Below the header is a logo for "MySQL Connector/ODBC". The main area is titled "Connection Parameters" and contains several input fields: "Data Source Name" (text box with "dwh_capacity"), "Description" (empty text box), "Server" (text box with "localhost"), "Port" (text box with "3306"), "User" (text box with "dwh"), "Password" (password field with 10 dots), and "Database" (dropdown menu with "dwh_capacity" selected). There is a "Test" button next to the Database dropdown. At the bottom of the dialog are four buttons: "Details >>", "OK", "Cancel", and "Help".

6. Enter the new password in the **Password** field.

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 the system to enable these features.

After configuring the system 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 CAC users, smart cards contain multiple certificates, only one of which can match the trusted CA. The CAC certificate for identification should be used.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

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 ID.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Use the `regedit` utility to modify registry values in `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Apache Software Foundation\Procrun2.0\SANscreen Server\Parameters\Java:`

 - a. Change the `JVM_Option DclientAuth=false` to `DclientAuth=true`.

2. Back up the keystore file: `C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore`
3. Open a command prompt specifying `Run as administrator`

4. Delete the self-generated certificate: `C:\Program Files\SANscreen\java64\bin\keytool.exe -delete -alias "ssl certificate" -keystore C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore`
5. Generate a new certificate: `C:\Program Files\SANscreen\java64\bin\keytool.exe -genkey -alias "alias_name" -keyalg RSA -sigalg SHA1withRSA -keysize 2048 -validity 365 -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -dname "CN=commonName,OU=orgUnit,O=orgName,L=localityNameI,S=stateName,C=countryName"`
6. Generate a certificate signing request (CSR): `C:\Program Files\SANscreen\java64\bin\keytool.exe -certreq -sigalg SHA1withRSA -alias "alias_name" -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -file C:\temp\server.csr"`
7. After the CSR is returned in step 6, import the certificate, then export the certificate in Base-64 format and place it in "C:\temp" named servername.cer.
8. Extract the certificate from the keystore: `C:\Program Files\SANscreen\java64\bin\keytool.exe -v -importkeystore -srckeystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -srcalias "alias_name" -destkeystore "C:\temp\file.p12" -deststoretype PKCS12`
9. Extract a private key from the p12 file: `openssl pkcs12 -in "C:\temp\file.p12" -out "C:\temp\servername.private.pem"`
10. Merge the Base-64 certificate that you exported in step 7 with the private key: `openssl pkcs12 -export -in "<folder>\<certificate>.cer" -inkey "C:\temp\servername.private.pem" -out "C:\temp\servername.new.p12" -name "servername.abc.123.yyy.zzz"`
11. Import the merged certificate into the keystore: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -destkeystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -srckeystore "C:\temp\servername.new.p12" -srcstoretype PKCS12 -alias "alias_name"`
12. Import the root certificate: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -file "C:\<root_certificate>.cer" -trustcacerts -alias "alias_name"`
13. Import the root certificate into the server.trustore: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore" -file "C:\<email_certificate>.cer" -trustcacerts -alias "alias_name"`
14. Import the intermediate certificate: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore" -file "C:\<intermediate_certificate>.cer" -trustcacerts -alias "alias_name"`

Repeat this step for all intermediate certificates.

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

1. Restart the server.

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.

Before you begin



For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):

- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

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>)

Enabling CAC on a Linux server

Some modifications are required to enable CAC on a Linux OnCommand Insight server.

Steps

1. Navigate to `/opt/netapp/oci/conf/`
2. Edit `wildfly.properties` and change the value of `CLIENT_AUTH_ENABLED` to "True"
3. Import the "root certificate" that exists under `/opt/netapp/oci/wildfly/standalone/configuration/server.keystore`
4. Restart the server

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.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Use `regedit` to modify registry values in `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Apache Software Foundation\Procrun2.0\SANscreen Server\Parameters\Java`

- a. Change the `JVM_Option -DclientAuth=false` to `-DclientAuth=true`.

For Linux, modify the `clientAuth` parameter in `/opt/netapp/oci/scripts/wildfly.server`

2. Add certificate authorities (CAs) to the Data Warehouse trustore:

- a. In a command window, go to `..\SANscreen\wildfly\standalone\configuration`.

- b. Use the `keytool` utility to list the trusted CAs: `C:\Program Files\SANscreen\java64\bin\keytool.exe -list -keystore server.trustore -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\wildfly\standalone\configuration` and use the `keytool import` command: `C:\Program Files\SANscreen\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.

3. On the OnCommand Insight server, the `wildfly/standalone/configuration/standalone-full.xml` file needs to be modified by updating `verify-client` to "REQUESTED" in

/subsystem=undertow/server=default-server/https-listener=default-httpsto enable CAC. Log in to the Insight server and run the appropriate command:

OS	Script
Windows	<install dir>\SANscreen\wildfly\bin\enableCACforRemoteEJB.bat
Linux	/opt/netapp/oci/wildfly/bin/enableCACforRemoteEJB.sh

After executing the script, wait until the reload of the wildfly server is complete before proceeding to the next step.

- Restart the OnCommand Insight server.

Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.5 through 7.3.9)

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

Before you begin

This procedure is for systems running OnCommand Insight 7.3.5 through 7.3.9.



For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):

- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Add certificate authorities (CAs) to the Cognos truststore.
 - a. In a command window, go to `..\SANscreen\cognos\analytics\configuration\certs\`
 - b. Use the `keytool` utility to list the trusted CAs: `..\..\jre\bin\keytool.exe -list -keystore CAMKeystore.jks -storepass NoPassWordSet`

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.

- d. To include customer's CAs with OnCommand Insight trusted CAs, go to
`..\SANscreen\cognos\analytics\configuration\certs\.`
- e. Use the `keytool` utility to import the `.pem` file: `..\..\jre\bin\keytool.exe -importcert -keystore CAMKeystore.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 `NoPassWordSet`.
 - g. Answer `yes` when prompted to trust the certificate.
2. To enable CAC mode, execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`
 3. To disable CAC mode, execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`

Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.10 and later)

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

Before you begin

This procedure is for systems running OnCommand Insight 7.3.10 and later.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand Data Warehouse 7.3.3 and later](#)

Steps

1. Add certificate authorities (CAs) to the Cognos truststore.
 - a. In a command window, go to `..\SANscreen\cognos\analytics\configuration\certs\.`
 - b. Use the `keytool` utility to list the trusted CAs: `..\..\ibm-jre\jre\bin\keytool.exe -list -keystore CAMKeystore.jks -storepass NoPassWordSet`

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

e. Use the keytool utility to import the .pem file: `..\..\ibm-jre\jre\bin\keytool.exe -importcert -keystore CAMKeystore.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 `NoPassWordSet`.

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

2. To enable CAC mode, do the following:

a. Configure CAC logout page, using the following steps:

- Logon to Cognos portal (user must be part of System Administrators group i.e. `cognos_admin`)
- (Only for 7.3.10 and 7.3.11) Click Manage -> Configuration -> System -> Security
- (Only for 7.3.10 and 7.3.11) Enter `cacLogout.html` against Logout Redirect URL -> Apply
- Close browser.

b. Execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`

c. Start IBM Cognos service. Wait for Cognos service to start.

3. To disable CAC mode, do the following:

a. Execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`

b. Start IBM Cognos service. Wait for Cognos service to start.

c. (Only for 7.3.10 and 7.3.11) Unconfigure CAC logout page, using the following steps:

- Logon to Cognos portal (user must be part of System Administrators group i.e. `cognos_admin`)
- Click Manage -> Configuration -> System -> Security
- Enter `cacLogout.html` against Logout Redirect URL -> Apply
- Close browser.

Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.5 to 7.3.9)

You can add SSL certificates to enable enhanced authentication and encryption for your Data Warehouse and Cognos environment.

Before you begin

This procedure is for systems running OnCommmand Insight 7.3.5 through 7.3.9.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

About this task

You must have admin privileges to perform this procedure.

Steps

1. Create a backup of `..\SANSscreen\cognos\analytics\configuration\cogstartup.xml`.
2. Create a backup of the “certs” and “csk” folders under `..\SANSscreen\cognos\analytics\configuration`.
3. Generate a Certificate Encryption Request from Cognos. In an Admin CMD window, run:
 - a. `cd “\Program Files\sansscreen\cognos\analytics\bin”`
 - b. `ThirdPartyCertificateTool.bat -java:local -c -e -p NoPassWordSet -a RSA -d “CN=FQDN,O=orgname,C=US” -r c:\temp\encryptRequest.csr`
4. Open the `c:\temp\encryptRequest.csr` file and copy the generated content.
5. Send the `encryptRequest.csr` to the certificate authority (CA) to obtain an SSL certificate.

Make sure to add additional attributes such as “SAN:dns=FQDN (For example, hostname.netapp.com)” to add the SubjectAltName. Google Chrome version 58 and later complains if the SubjectAltName is missing from the certificate.

6. Download the chain certificates by including root certificate by using PKCS7 format

This will download `fqdn.p7b` file

7. Get a cert in `.p7b` format from your CA. Use a name that marks it as the certificate for the Cognos Webserver.
8. `ThirdPartyCertificateTool.bat` fails to import the entire chain, so multiple steps are required to export all certificates. Split the chain by exporting them individually as follows:
 - a. Open the `.p7b` certificate in “Crypto Shell Extensions”.
 - b. Browse in the left pane to “Certificates”.
 - c. Right-click on root CA > All Tasks > Export.
 - d. Select Base64 output.
 - e. Enter a file name identifying it as the root certificate.

- f. Repeat steps 8a through 8c to export all of the certificates separately into .cer files.
 - g. Name the files intermediateX.cer and cognos.cer.
9. Ignore this step if you have only one CA certificate, otherwise merge both root.cer and intermediateX.cer into one file.
 - a. Open intermediate.cer with NotePad and copy the content.
 - b. Open root.cer with NotePad and save the content from 9a.
 - c. Save the file as CA.cer.
10. Import the certificates into the Cognos keystore using the Admin CMD prompt:
 - a. cd "Program Files\sansscreen\cognos\analytics\bin"
 - b. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\CA.cer

This will set CA.cer as root Certificate Authority.

 - c. ThirdPartyCertificateTool.bat -java:local -i -e -r c:\temp\cognos.cer -t c:\temp\CA.cer

This will set Cognos.cer as encryption certificate which is signed by CA.cer.
11. Open the IBM Cognos Configuration.
 - a. Select Local Configuration → Security → Cryptography → Cognos
 - b. Change "Use third party CA?" to True.
 - c. Save the configuration.
 - d. Restart Cognos
12. Export the latest Cognos certificate into cognos.crt using the Admin CMD prompt:
 - a. "D:\Program Files\SANscreen\java\bin\keytool.exe" -exportcert -file "c:\temp\cognos.crt" -keystore "D:\Program Files\SANscreen\cognos\analytics\configuration\certs\CAMKeystore" -storetype PKCS12 -storepass NoPassWordSet -alias encryption
13. Import the "c:\temp\cognos.crt" into dwh trustore to establish SSL communication between Cognos and DWH, using the Admin CMD prompt window.
 - a. "D:\Program Files\SANscreen\java\bin\keytool.exe" -importcert -file "c:\temp\cognos.crt" -keystore "D:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore" -storepass changeit -alias cognoscert
14. Restart the SANscreen service.
15. Perform a backup of DWH to make sure DWH communicates with Cognos.

Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.10 and later)

You can add SSL certificates to enable enhanced authentication and encryption for your Data Warehouse and Cognos environment.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.10 and later.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

About this task

You must have admin privileges to perform this procedure.

Steps

1. Stop Cognos using the IBM Cognos Configuration tool. Close Cognos.
2. Create backups of the `..\SANSscreen\cognos\analytics\configuration` and `..\SANSscreen\cognos\analytics\temp\cam\freshness` folders.
3. Generate a Certificate Encryption Request from Cognos. In an Admin CMD window, run:
 - a. `cd "\Program Files\sansscreen\cognos\analytics\bin"`
 - b. `ThirdPartyCertificateTool.bat -java:local -c -e -p NoPassWordSet -a RSA -r c:\temp\encryptRequest.csr -d "CN=server.domain.com,O=NETAPP,C=US" -H "server.domain.com" -I "ipaddress"`. Note: here -H and -I are to add subjectAltNames like dns and ipaddress.
4. Open the `c:\temp\encryptRequest.csr` file and copy the generated content.
5. Input the `encryptRequest.csr` content and generate certificate using CA signing portal.
6. Download the chain certificates by including root certificate by using PKCS7 format

This will download `fqdn.p7b` file

7. Get a cert in `.p7b` format from your CA. Use a name that marks it as the certificate for the Cognos Webserver.
8. `ThirdPartyCertificateTool.bat` fails to import the entire chain, so multiple steps are required to export all certificates. Split the chain by exporting them individually as follows:
 - a. Open the `.p7b` certificate in "Crypto Shell Extensions".
 - b. Browse in the left pane to "Certificates".
 - c. Right-click on root CA > All Tasks > Export.
 - d. Select Base64 output.
 - e. Enter a file name identifying it as the root certificate.
 - f. Repeat steps 8a through 8e to export all of the certificates separately into `.cer` files.

- g. Name the files intermediateX.cer and cognos.cer.
9. Ignore this step if you have only one CA certificate, otherwise merge both root.cer and intermediateX.cer into one file.
 - a. Open root.cer with NotePad and copy the content.
 - b. Open intermediate.cer with NotePad and append the content from 9a (intermediate first and root next).
 - c. Save the file as chain.cer.
10. Import the certificates into the Cognos keystore using the Admin CMD prompt:
 - a. cd "Program Files\sansscreen\cognos\analytics\bin"
 - b. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\root.cer
 - c. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\intermediate.cer
 - d. ThirdPartyCertificateTool.bat -java:local -i -e -r c:\temp\cognos.cer -t c:\temp\chain.cer
11. Open the IBM Cognos Configuration.
 - a. Select Local Configuration → Security → Cryptography → Cognos
 - b. Change "Use third party CA?" to True.
 - c. Save the configuration.
 - d. Restart Cognos
12. Export the latest Cognos certificate into cognos.crt using the Admin CMD prompt:
 - a. cd "C:\Program Files\SANscreen"
 - b. java\bin\keytool.exe -exportcert -file c:\temp\cognos.crt -keystore cognos\analytics\configuration\certs\CAMKeystore -storetype PKCS12 -storepass NoPassWordSet -alias encryption
13. Back up the DWH server trustore


```
at.. \SANscreen\wildfly\standalone\configuration\server.trustore
```
14. Import the "c:\temp\cognos.crt" into DWH trustore to establish SSL communication between Cognos and DWH, using the Admin CMD prompt window.
 - a. cd "C:\Program Files\SANscreen"
 - b. java\bin\keytool.exe -importcert -file c:\temp\cognos.crt -keystore wildfly\standalone\configuration\server.trustore -storepass changeit -alias cognos3rdca
15. Restart the SANscreen service.
16. Perform a backup of DWH to make sure DWH communicates with Cognos.
17. The following steps should be performed even when only the "ssl certificate" is changed and the default Cognos certificates are left unchanged. Otherwise Cognos may complain about the new SANscreen certificate or be unable to create a DWH backup.
 - a. cd "%SANSSCREEN_HOME%cognos\analytics\bin\"
 - b. "%SANSSCREEN_HOME%java64\bin\keytool.exe" -exportcert -file "c:\temp\sansscreen.cer" -keystore "%SANSSCREEN_HOME%wildfly\standalone\configuration\server.keystore" -storepass changeit -alias "ssl certificate"
 - c. ThirdPartyCertificateTool.bat -java:local -i -T -r "c:\temp\sansscreen.cer"

Typically, these steps are performed as part of the Cognos certificate import process described in [How to](#)

import a Cognos Certificate Authority (CA) signed certificate into OnCommand DataWarehouse 7.3.3 and later

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.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):

- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)



Steps

1. Use `regedit` to modify registry values in `HKEY_LOCAL_MACHINE\SOFTWARE Wow6432Node\Apache Software Foundation\Procrun2.0\SANscreen Server\Parameters\Java`

- a. Change the `JVM_Option -DclientAuth=false` to `-DclientAuth=true`.

For Linux, modify the `clientAuth` parameter in `/opt/netapp/oci/scripts/wildfly.server`

2. Add certificate authorities (CAs) to the Data Warehouse trustore:

- a. In a command window, go to `..\SANscreen\wildfly\standalone\configuration`.
- b. Use the `keytool` utility to list the trusted CAs: `C:\Program Files\SANscreen\java64\bin\keytool.exe -list -keystore server.trustore -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\wildfly\standalone\configuration` and use the `keytool import` command: `C:\Program Files\SANscreen\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.

3. On the OnCommand Insight server, the `wildfly/standalone/configuration/standalone-full.xml` file needs to be modified by updating `verify-client` to "REQUESTED" in `/subsystem=undertow/server=default-server/https-listener=default-httpsto` enable CAC. Log in to the Insight server and run the appropriate command:

OS	Script
Windows	<install dir>\SANscreen\wildfly\bin\enableCACforRemoteEJB.bat
Linux	/opt/netapp/oci/wildfly/bin/enableCACforRemoteEJB.sh

After executing the script, wait until the reload of the wildfly server is complete before proceeding to the next step.

4. Restart the OnCommand Insight server.

Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.5 through 7.3.9)

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

Before you begin

This procedure is for systems running OnCommand Insight 7.3.5 through 7.3.9.



For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):

- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Add certificate authorities (CAs) to the Cognos truststore.
 - a. In a command window, go to `..\SANscreen\cognos\analytics\configuration\certs\`
 - b. Use the `keytool` utility to list the trusted CAs: `..\..\jre\bin\keytool.exe -list -keystore CAMKeystore.jks -storepass NoPassWordSet`

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.
 - d. To include customer's CAs with OnCommand Insight trusted CAs, go to `..\SANscreen\cognos\analytics\configuration\certs\`.
 - e. Use the `keytool` utility to import the `.pem` file: `..\..\jre\bin\keytool.exe -importcert -keystore CAMKeystore.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 `NoPassWordSet`.
 - g. Answer `yes` when prompted to trust the certificate.
2. To enable CAC mode, execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`
3. To disable CAC mode, execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`

Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.10 and later)

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

Before you begin

This procedure is for systems running OnCommand Insight 7.3.10 and later.



For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):

- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand Data Warehouse 7.3.3 and later](#)

Steps

1. Add certificate authorities (CAs) to the Cognos truststore.

- a. In a command window, go to `..\SANscreen\cognos\analytics\configuration\certs\`
- b. Use the `keytool` utility to list the trusted CAs: `..\..\ibm-jre\jre\bin\keytool.exe -list -keystore CAMKeystore.jks -storepass NoPassWordSet`

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.
- d. To include customer's CAs with OnCommand Insight trusted CAs, go to `..\SANscreen\cognos\analytics\configuration\certs\`.
- e. Use the `keytool` utility to import the `.pem` file: `..\..\ibm-jre\jre\bin\keytool.exe -importcert -keystore CAMKeystore.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 `NoPassWordSet`.
- g. Answer `yes` when prompted to trust the certificate.

2. To enable CAC mode, do the following:

- a. Configure CAC logout page, using the following steps:
 - Logon to Cognos portal (user must be part of System Administrators group i.e. `cognos_admin`)
 - (Only for 7.3.10 and 7.3.11) Click Manage -> Configuration -> System -> Security
 - (Only for 7.3.10 and 7.3.11) Enter `cacLogout.html` against Logout Redirect URL -> Apply
 - Close browser.
- b. Execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`
- c. Start IBM Cognos service. Wait for Cognos service to start.

3. To disable CAC mode, do the following:

- a. Execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`
- b. Start IBM Cognos service. Wait for Cognos service to start.
- c. (Only for 7.3.10 and 7.3.11) Unconfigure CAC logout page, using the following steps:
 - Logon to Cognos portal (user must be part of System Administrators group i.e. `cognos_admin`)
 - Click Manage -> Configuration -> System -> Security
 - Enter `cacLogout.html` against Logout Redirect URL -> Apply
 - Close browser.

Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.5 to 7.3.9)

You can add SSL certificates to enable enhanced authentication and encryption for your Data Warehouse and Cognos environment.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.5 through 7.3.9.



For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):

- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

About this task

You must have admin privileges to perform this procedure.

Steps

1. Create a backup of `..\SANSscreen\cognos\analytics\configuration\cogstartup.xml`.
2. Create a backup of the “certs” and “csk” folders under `..\SANSscreen\cognos\analytics\configuration`.
3. Generate a Certificate Encryption Request from Cognos. In an Admin CMD window, run:
 - a. `cd “\Program Files\sansscreen\cognos\analytics\bin”`
 - b. `ThirdPartyCertificateTool.bat -java:local -c -e -p NoPassWordSet -a RSA -d “CN=FQDN,O=orgname,C=US” -r c:\temp\encryptRequest.csr`
4. Open the `c:\temp\encryptRequest.csr` file and copy the generated content.
5. Send the `encryptRequest.csr` to the certificate authority (CA) to obtain an SSL certificate.

Make sure to add additional attributes such as “SAN:dns=FQDN (For example, hostname.netapp.com)” to add the SubjectAltName. Google Chrome version 58 and later complains if the SubjectAltName is missing from the certificate.

6. Download the chain certificates by including root certificate by using PKCS7 format

This will download `fqdn.p7b` file

7. Get a cert in `.p7b` format from your CA. Use a name that marks it as the certificate for the Cognos Webserver.
8. `ThirdPartyCertificateTool.bat` fails to import the entire chain, so multiple steps are required to export all certificates. Split the chain by exporting them individually as follows:
 - a. Open the `.p7b` certificate in “Crypto Shell Extensions”.
 - b. Browse in the left pane to “Certificates”.

- c. Right-click on root CA > All Tasks > Export.
 - d. Select Base64 output.
 - e. Enter a file name identifying it as the root certificate.
 - f. Repeat steps 8a through 8c to export all of the certificates separately into .cer files.
 - g. Name the files intermediateX.cer and cognos.cer.
9. Ignore this step if you have only one CA certificate, otherwise merge both root.cer and intermediateX.cer into one file.
- a. Open intermediate.cer with NotePad and copy the content.
 - b. Open root.cer with NotePad and save the content from 9a.
 - c. Save the file as CA.cer.
10. Import the certificates into the Cognos keystore using the Admin CMD prompt:
- a. cd "Program Files\sanscreen\cognos\analytics\bin"
 - b. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\CA.cer
- This will set CA.cer as root Certificate Authority.
- c. ThirdPartyCertificateTool.bat -java:local -i -e -r c:\temp\cognos.cer -t c:\temp\CA.cer
- This will set Cognos.cer as encryption certificate which is signed by CA.cer.
11. Open the IBM Cognos Configuration.
- a. Select Local Configuration → Security → Cryptography → Cognos
 - b. Change "Use third party CA?" to True.
 - c. Save the configuration.
 - d. Restart Cognos
12. Export the latest Cognos certificate into cognos.crt using the Admin CMD prompt:
- a. "D:\Program Files\SANscreen\java\bin\keytool.exe" -exportcert -file "c:\temp\cognos.crt" -keystore "D:\Program Files\SANscreen\cognos\analytics\configuration\certs\CAMKeystore" -storetype PKCS12 -storepass NoPassWordSet -alias encryption
13. Import the "c:\temp\cognos.crt" into dwh trustore to establish SSL communication between Cognos and DWH, using the Admin CMD prompt window.
- a. "D:\Program Files\SANscreen\java\bin\keytool.exe" -importcert -file "c:\temp\cognos.crt" -keystore "D:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore" -storepass changeit -alias cognoscert
14. Restart the SANscreen service.
15. Perform a backup of DWH to make sure DWH communicates with Cognos.

Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.10 and later)

You can add SSL certificates to enable enhanced authentication and encryption for your Data Warehouse and Cognos environment.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.10 and later.



For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):

- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

About this task

You must have admin privileges to perform this procedure.

Steps

1. Stop Cognos using the IBM Cognos Configuration tool. Close Cognos.
2. Create backups of the `..\SANSscreen\cognos\analytics\configuration` and `..\SANSscreen\cognos\analytics\temp\cam\freshness` folders.
3. Generate a Certificate Encryption Request from Cognos. In an Admin CMD window, run:
 - a. `cd "\\Program Files\sansscreen\cognos\analytics\bin"`
 - b. `ThirdPartyCertificateTool.bat -java:local -c -e -p NoPassWordSet -a RSA -r c:\temp\encryptRequest.csr -d "CN=server.domain.com,O=NETAPP,C=US" -H "server.domain.com" -I "ipaddress"`. Note: here -H and -I are to add subjectAltNames like dns and ipaddress.
4. Open the `c:\temp\encryptRequest.csr` file and copy the generated content.
5. Input the `encryptRequest.csr` content and generate certificate using CA signing portal.
6. Download the chain certificates by including root certificate by using PKCS7 format

This will download `fqdn.p7b` file

7. Get a cert in `.p7b` format from your CA. Use a name that marks it as the certificate for the Cognos Webserver.
8. `ThirdPartyCertificateTool.bat` fails to import the entire chain, so multiple steps are required to export all certificates. Split the chain by exporting them individually as follows:
 - a. Open the `.p7b` certificate in "Crypto Shell Extensions".
 - b. Browse in the left pane to "Certificates".
 - c. Right-click on root CA > All Tasks > Export.
 - d. Select Base64 output.

- e. Enter a file name identifying it as the root certificate.
- f. Repeat steps 8a through 8e to export all of the certificates separately into .cer files.
- g. Name the files intermediateX.cer and cognos.cer.
9. Ignore this step if you have only one CA certificate, otherwise merge both root.cer and intermediateX.cer into one file.
 - a. Open root.cer with NotePad and copy the content.
 - b. Open intermediate.cer with NotePad and append the content from 9a (intermediate first and root next).
 - c. Save the file as chain.cer.
10. Import the certificates into the Cognos keystore using the Admin CMD prompt:
 - a. cd "Program Files\sansscreen\cognos\analytics\bin"
 - b. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\root.cer
 - c. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\intermediate.cer
 - d. ThirdPartyCertificateTool.bat -java:local -i -e -r c:\temp\cognos.cer -t c:\temp\chain.cer
11. Open the IBM Cognos Configuration.
 - a. Select Local Configuration → Security → Cryptography → Cognos
 - b. Change "Use third party CA?" to True.
 - c. Save the configuration.
 - d. Restart Cognos
12. Export the latest Cognos certificate into cognos.crt using the Admin CMD prompt:
 - a. cd "C:\Program Files\SANscreen"
 - b. java\bin\keytool.exe -exportcert -file c:\temp\cognos.crt -keystore cognos\analytics\configuration\certs\CAMKeystore -storetype PKCS12 -storepass NoPassWordSet -alias encryption
13. Back up the DWH server trustore


```
at ..\SANscreen\wildfly\standalone\configuration\server.trustore
```
14. Import the "c:\temp\cognos.crt" into DWH trustore to establish SSL communication between Cognos and DWH, using the Admin CMD prompt window.
 - a. cd "C:\Program Files\SANscreen"
 - b. java\bin\keytool.exe -importcert -file c:\temp\cognos.crt -keystore wildfly\standalone\configuration\server.trustore -storepass changeit -alias cognos3rdca
15. Restart the SANscreen service.
16. Perform a backup of DWH to make sure DWH communicates with Cognos.
17. The following steps should be performed even when only the "ssl certificate" is changed and the default Cognos certificates are left unchanged. Otherwise Cognos may complain about the new SANscreen certificate or be unable to create a DWH backup.
 - a. cd "%SANSSCREEN_HOME%cognos\analytics\bin\"
 - b. "%SANSSCREEN_HOME%java64\bin\keytool.exe" -exportcert -file "c:\temp\sansscreen.cer" -keystore "%SANSSCREEN_HOME%wildfly\standalone\configuration\server.keystore" -storepass changeit -alias "ssl certificate"

```
c. ThirdPartyCertificateTool.bat -java:local -i -T -r "c:\temp\sansscreen.cer"
```

Typically, these steps are performed as part of the Cognos certificate import process described in [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Importing SSL certificates

You can add SSL certificates to enable enhanced authentication and encryption for enhancing the security of your OnCommand Insight environment.

Before you begin

You must ensure that your system meets the minimum required bit level (1024 bits).

About this task



Before you attempt to perform this procedure, you should back up the existing `server.keystore` file, and name the backup `server.keystore.old`. Corrupting or damaging the `server.keystore` file may result in an inoperable Insight server after the Insight server is restarted. If you create a backup, you can revert to the old file if problems occur.

Steps

1. Create a copy of the original keystore file:

```
cp c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore.old"
```
2. List the contents of the keystore:

```
C:\Program Files\SANscreen\java64\bin\keytool.exe -list -v -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore"
```

 - a. When prompted for a password, enter `changeit`.

The system displays the contents of the keystore. There should be at least one certificate in the keystore, "ssl certificate".
3. Delete the "ssl certificate":

```
keytool -delete -alias "ssl certificate" -keystore c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore
```
4. Generate a new key:

```
C:\Program Files\SANscreen\java64\bin\keytool.exe -genkey -alias "ssl certificate" -keyalg RSA -keysize 2048 -validity 365 -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore"
```

 - a. When prompted for first and last names, enter the fully qualified domain name (FQDN) that you intend to use.
 - b. Provide the following information about your organization and organizational structure:
 - Country: two-letter ISO abbreviation for your country (for example, US)
 - State or Province: name of the state or province where your organization's head office is located (for example, Massachusetts)
 - Locality: name of the city where your organization's head office is located (for example, Waltham)

- Organizational name: name of the organization that owns the domain name (for example, NetApp)
- Organizational unit name: name of the department or group that will use the certificate (for example, Support)
- Domain Name/ Common Name: the FQDN that is used for DNS lookups of your server (for example, www.example.com) The system responds with information similar to the following: Is CN=www.example.com, OU=support, O=NetApp, L=Waltham, ST=MA, C=US correct?

c. Enter Yes when the Common Name (CN) is equal to the FQDN.

d. When prompted for the key password, enter the password, or press the Enter key to use the existing keystore password.

5. Generate a certificate request file: `C:\Program Files\SANscreen\java64\bin\keytool.exe -certreq -alias "ssl certificate" -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -file c:\localhost.csr`

The `c:\localhost.csr` file is the certificate request file that is newly generated.

6. Submit the `c:\localhost.csr` file to your certificate authority (CA) for approval.

Once the certificate request file is approved, you want the certificate returned to you in `.der` format. The file might or might not be returned as a `.der` file. The default file format is `.cer` for Microsoft CA services.

Most organizations' CAs use a chain of trust model, including a root CA, which is often offline. It has signed the certificates for only a few child CAs, known as intermediate CAs.

You must obtain the public key (certificates) for the entire chain of trust—the certificate for the CA that signed the certificate for the OnCommand Insight server, and all the certificates between that signing CA up to and including the organizational root CA.

In some organizations, when you submit a signing request, you might receive one of the following:

- A PKCS12 file that contains your signed certificate and all the public certificates in the chain of trust
- A `.zip` file that contains individual files (including your signed certificate) and all the public certificates in the chain of trust
- Only your signed certificate

You must obtain the public certificates.

7. Import the approved certificate for `server.keystore`: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -alias OCI.hostname.com -file c:\localhost2.DER -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore"`

a. When prompted, enter the keystore password.

The following message is displayed: Certificate reply was installed in keystore

8. Import the approved certificate for `server.trustore`: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -alias OCI.hostname.com -file c:\localhost2.DER -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore"`

- a. When prompted, enter the trustore password.

The following message is displayed: Certificate reply was installed in trustore

9. Edit the SANscreen\wildfly\standalone\configuration\standalone-full.xml file:

Substitute the following alias string: alias="cbc-oci-02.muccbc.hq.netapp.com". For example:

```
<keystore path="server.keystore" relative-to="jboss.server.config.dir"
keystore-password="{VAULT::HttpsRealm::keystore_password:1}" alias="cbc-oci-
02.muccbc.hq.netapp.com" key-
password="{VAULT::HttpsRealm::key_password:1}"/>
```

10. Restart the SANscreen server service.

Once Insight is running, you can click the padlock icon to view the certificates that are installed on the system.

If you see a certificate containing "Issued To" information that matches "Issued By" information, you still have a self-signed certificate installed. The Insight installer-generated self-signed certificates have a 100-year expiration.

NetApp cannot guarantee that this procedure will remove digital certificate warnings. NetApp cannot control how your end user workstations are configured. Consider the following scenarios:

- Microsoft Internet Explorer and Google Chrome both utilize Microsoft's native certificate functionality on Windows.

This means that if your Active Directory administrators push your organization's CA certificates into the end user's certificate trustores, the users of these browsers will see certificate warnings disappear when the OnCommand Insight self-signed certificates have been replaced with the one signed by the internal CA infrastructure.

- Java and Mozilla Firefox have their own certificate stores.

If your system administrators do not automate ingesting the CA certificates into these applications' trusted certificates stores, using the Firefox browser might continue to generate certificate warnings because of an untrusted certificate, even when the self-signed certificate has been replaced. Getting your organization's certificate chain installed into the trustore is an additional requirement.

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.

The backup file will contain the following:

- Inventory data
- Up to 7 days of performance data

Steps

1. On the Insight toolbar, click **Admin > Setup**.
2. Click the **Backup & Archive** tab.
3. In the Weekly Backup section, select **Enable weekly backup**.
4. Enter the path to the **Backup location**. This can be on the on the local Insight server or on a remote server that is accessible from the Insight server.



The backup location setting is included in the backup itself, so if you restore the backup on another system, be aware that the backup folder location may be invalid on the new system. Double-check your backup location settings after restoring a backup.

5. Select the **Cleanup** option to keep either the last two or the last five backups.
6. Click **Save**.

Results

You can also go to **Admin > Troubleshooting** to create an on-demand backup.

What's included in the backup

Weekly and on-demand backups can be used for troubleshooting or migration.

The weekly or on-demand backup includes the following:

- Inventory data
- Performance data (if selected for inclusion in backup)
- Data sources and data source settings
- Integration packs
- Remote acquisition units
- ASUP/proxy settings
- Backup location settings
- Archive location settings
- Notification settings
- Users
- Performance policies
- Business entities and applications
- Device resolution rules and settings
- Dashboards and widgets
- Customized asset page dashboards and widgets
- Queries

- Annotations and annotation rules

The weekly backup does not include:

- Security tool settings / vault information (backed up via separate CLI process)
- Logs (can be saved to a .zip file on demand)
- Performance data (if not selected for inclusion in backup)
- Licenses



If you choose to include performance data in the backup, the most recent seven days of data is backed up. The remaining data will be in the archive if you have that feature enabled.

Performance data archiving

OnCommand Insight 7.3 introduces the ability to archive performance data on a daily basis. This supplements configuration and limited performance data backups.

OnCommand Insight retains up to 90 days of performance and violation data. However, when creating a backup of that data, only the most recent information is included in the backup. Archiving allows you to save the remainder of your performance data and load it as necessary.

Once the archive location is configured and archiving is activated, once a day Insight will archive the previous day's performance data for all objects into the archive location. Each day's archive is kept in the archive folder in a separate file. Archiving happens in the background and will continue as long as Insight is running.

The most recent 90 days of archives are retained; archive files older than 90 days are deleted as newer ones are created.

Enabling performance archive

To enable performance data archiving, follow these steps.

Steps

1. On the toolbar, click **Admin > Setup**.
2. Select the **Backup & Archive** tab.
3. In the Performance Archive section, ensure **Enable performance archive** is checked.
4. Specify a valid archive location.

You cannot specify a folder under the Insight installation folder.

Best Practice: Do not specify the same folder for archive as the Insight backup location.

5. Click **Save**.

The archive process is handled in the background and does not interfere with other Insight activities.

Loading performance archive

To load the performance data archive, follow these steps.

Before you begin

Before loading the performance data archive, you must restore a valid weekly or manual backup.

Steps

1. On the toolbar, click **Admin > Troubleshooting**.
2. In the Restore section, under **Load performance archive**, click **Load**.



Archive loading is handled in the background. Loading the full archive can take a long time as each day's archived performance data is populated into Insight. The status of the archive loading is displayed in the archive section of this page.

Configuring your email

You must configure OnCommand Insight to access your email system so that the OnCommand 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. On the Insight toolbar, 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 password-protected. 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 send a test email to specified recipients, click .
10. Click **Save**.

Configuring SNMP notifications

OnCommand Insight supports SNMP notifications for configuration and Global Path policy changes as well as 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 and can be found on the Insight server in `<install_dir>\SANscreen\MIBS\sanscreen.mib`.

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 as well as audit messages, and activate the logging process.

Before you begin

- You must have the IP address of the server on which 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 messages
- Performance alerts
- Optionally, Audit log messages

The following units are used in the syslog:

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

Steps

1. On the Insight toolbar, click **Admin** and select **Notifications**.
2. Scroll down to the **Syslog** section of the page.
3. Select the **Enable syslog** check box.
4. If desired, select the **Send audit** check box. New audit log messages will be sent to syslog in addition to being displayed on the Audit page. Note that already-existing audit log messages will not be sent to syslog; only newly-generated log messages will be sent.
5. In the **Server** field, enter the IP address of the log server.

You can specify a custom port by appending it following a colon at the end of the server IP (e.g. server:port). If port is not specified, the default syslog port of 514 is used.

6. In the **Facility** field, select the facility level that corresponds to the type of program that is logging the message.
7. 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 configured email, syslog, and SNMP sending methods for violations.

Steps

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

4. Click **Save**.

Configuring system-level event notifications

OnCommand Insight supports notifications for system-level events such as acquisition unit failures or data source errors. To receive notifications you must configure Insight to send email when one or more of these events occur.

Before you begin

You must have configured email recipients for receiving notifications in **Admin > Notifications > Sending Methods**.

Steps

1. Click **Admin > Notifications**.
2. Click **Events**.
3. In the **System Alert Events** Email section, select the severity level (**Warning and above** or **Critical**) for the notification, or choose **Do not send** if you do not wish to receive notifications of system-level events.
4. Click **Save**.
5. Click **Admin > System Alerts** to configure the alerts themselves.
6. To Add a new alert, click **+Add** and give the alert a unique **Name**. You can also click the right-side icon to **Edit** an existing alert.
7. Choose the **Event type** on which to alert, for example *Acquisition Unit Failure*.
8. Choose a **Snooze** interval to suppress notifications on duplicate events of the selected type for the selected time interval. If you select *Never*, you will receive repeat notifications once a minute until the event is no longer happening.
9. Choose a **Severity** (Warning or Critical) for the event notification.
10. Email notifications will be sent to the global email recipient list by default, or you can click the link provided to override the global list and send notifications to specific recipients.
11. Click **Save** to add the alert.

Configuring your ASUP processing

All NetApp products are equipped with automated capabilities to provide the best possible support for customers. The automated support (ASUP) periodically sends predefined and specific information to Customer Support. You can control the information to be 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

ASUP data is forwarded using the HTTPS protocol.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Setup**.
3. Click the **ASUP & Proxy** tab.
4. In the **ASUP** section, select **Enable ASUP** to activate the ASUP facility.
5. If you want to change your corporate information, update the following fields:
 - **Company name**
 - **Site name**
 - **What to send:** Logs, configuration data, performance data
6. Click **Test Connection** to ensure that the connection that you specified works.
7. Click **Save**.
8. In the **Proxy** section, choose whether to **Enable Proxy**, and specify your proxy **host**, **port**, and **user** information.
9. Click **Test Connection** to ensure that the proxy that you specified works.
10. Click **Save**.

What's included in the Autosupport (ASUP) package

The Autosupport package contains the database backup as well as extended information.

The Autosupport package includes the following:

- Inventory data
- Performance data (if selected for inclusion in ASUP)
- Data sources and data source settings
- Integration packs
- Remote acquisition units
- ASUP/proxy settings
- Backup location settings
- Archive location settings
- Notification settings
- Users
- Performance policies
- Business entities and applications
- Device resolution rules and settings
- Dashboards and widgets
- Customized asset page dashboards and widgets
- Queries
- Annotations and annotation rules
- Logs

- Licenses
- Acquisition / data source status
- MySQL status
- System information

The Autosupport package does not include:

- Security tool settings / vault information (backed up via separate CLI process)
- Performance data (if not selected for inclusion in ASUP)



If you choose to include performance data in the ASUP, the most recent seven days of data is included. The remaining data will be in the archive if you have that feature enabled. Archive data is not included in ASUP.

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: hosts, virtual machines, volumes, internal volumes, qtrees, shares, and hypervisors.

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.

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.

Assigning applications to assets

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

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.



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**.




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**.

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.

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  **Add** 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**.

Assigning business entities to assets

You can assign a business entity to an asset (host, port, storage, switch, virtual machine, qtree, share, 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.



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 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.

Assigning business entities to or removing business entities from multiple assets

You can assign business entities to or remove business entities from multiple assets by using a query instead of having to manually assign or remove them.

Before you begin

You must have already created the business entities you want to add to your desired assets.

Steps

1. Create a new query, or open an existing query.
2. If desired, filter for the assets to which you want to add business entities.
3. Select the desired assets in the list or click ☐ ▼ to select **All**.

The **Actions** button displays.

4. To add a business entity to the selected assets, click . If the selected asset type can have business entities assigned to it, you will see the menu choice to **Add Business Entity**. Select this.
5. Select the desired business entity from the list and click **Save**.

Any new business entity you assign overrides any business entities that were already assigned to the asset. Assigning applications to assets will also override the business entities assigned in the same way. Assigning business entities to an asset may also override any applications assigned to that asset.

6. To remove a business entity assigned to the assets, click and select **Remove Business Entity**.
7. Select the desired business entity from the list and click **Delete**.

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.
- 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
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
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



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 Insightweb 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 Object Storage 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.



Annotation names beginning or ending with a dot "." are not supported.

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.




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 **Listas** 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.

- c. Click  to add additional values.

- d. Click  to remove a value.

7. Click **Save**.

Your annotations appear in the list on the Annotations page.

Related information

[Importing and Exporting user data](#)


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**.
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.


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.



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  **Add**.

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.

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 Insightweb 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 or down in the list, click the up arrow or the down arrow.

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 Insightweb 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 Insightweb 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 and keys are supported in the import utility:

Type	Key
------	-----

Host	id-><id> or <Name> or <IP>
VM	id-><id> or <Name>
Storage pool	id-><id> or <Storage_name>-><Storage_Pool_name>
Internal volume	id-><id> or <Storage_name>-><Internal_volume_name>
Volume	id-><id> or <Storage_name>-><Volume_name>
Storage	id-><id> or <Name> or <IP>
Switch	id-><id> or <Name> or <IP>
Port	id-><id> or <WWN>
Share	id-><id> or <Storage Name>-><Internal Volume Name>-><Share Name>-><Protocol> <Qtree> is optional if there is a default qtree.
Qtree	id-><id> or <Storage Name>-><Internal Volume Name>-><Qtree Name>

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 Key 1>, <Annotation Value> [,
<Annotation Value> ...] [, <Application>] [, <Tenant>] [,
<Line_Of_Business>] [, <Business_Unit>] [, <Project>]

...

<Object Type Value N>, <Object 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 rest-import-utility.jar [-username] [-password]  
[-server name or IP address] [-batch size] [-ccase  
sensitive:true/false]  
[-lextra logging:true/false] csv filename
```

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



There are no spaces between the options and their values.



The following keywords are reserved and prevent users from specifying them as annotation names: - Application - Application_Priority - Tenant - Line_Of_Business - Business_Unit - Project Errors are generated if you attempt to import an annotation type using one of the reserved keywords. If you have created annotation names using these keywords, you must modify them so that the import utility tool can work correctly.



The Annotation Import utility requires Java 8 or Java 11. Ensure that one of those is installed prior to running the import utility. It is recommended to use the latest OpenJDK 11.

Assigning annotations to multiple assets using a query

Assigning an annotation to a group of assets helps you more easily identify or use those related assets in queries or dashboards.

Before you begin

Annotations that you wish to assign to assets must have previously been created.

About this task

You can simplify the task of assigning an annotation to multiple assets by using a query. For example, if you

want to assign a custom address annotation to all of your arrays at a specific data center location.

Steps

1. Create a new query to identify the assets on which you wish to assign an annotation. Click **Queries > +New Query**.
2. In the **Search for...** drop-down, choose **Storage**. You can set filters to further narrow down the list of storages displayed.
3. In the list of storages displayed, select one or more by clicking on the check box beside the storage name. You may also select all the displayed storages by clicking on the main check box at the top of the list.
4. When you have selected all of the desired storages, click **Actions > Edit Annotation**.

The system displays the Add Annotation dialog.

5. Select the **Annotation** and **Value** you want to assign to the storages and click **Save**.

If you are displaying the column for that annotation, it will appear on all the selected storages.

6. You can now use the annotation to filter for storages in a widget or query. In a widget, you can do the following:
 - a. Create a dashboard or open an existing one. Add a **Variable** and choose the annotation you set on the storages above. The variable is added to the dashboard.
 - b. In the variable field you just added, click on **Any** and enter the appropriate Value to filter on. Click on the check mark to save the variable value.
 - c. Add a widget. In the widget's Query, click on the **Filter by+** button and select the appropriate annotation from the list.
 - d. Click on **Any** and select the annotation variable you added above. Variables you have created start with "\$" and are displayed in the drop-down.
 - e. Set any other filters or fields you desire, then click **Save** when the widget is customized to your liking.

The widget on the dashboard displays the data for only the storages to which you assigned the annotation.

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.

Assets used in queries and dashboards

Insight queries and dashboard widgets can be used with a wide range of asset types

The following asset types can be used in queries, dashboard widgets, and custom asset pages. The fields and counters available for filters, expressions, and display will vary among asset types. Not all assets can be used in all widget types.

- Application
- Datastore

- Disk
- Fabric
- Generic Device
- Host
- Internal Volume
- iSCSI Session
- iSCSI Network Portal
- Path
- Port
- Qtree
- Quota
- Share
- Storage
- Storage Node
- Storage Pool
- Switch
- Tape
- VMDK
- Virtual Machine
- Volume
- Zone
- Zone Member

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 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.
- The NOT operator allows you to exclude text from the search results. For example, `NOT EMC*` finds everything that does not start with “EMC”. You can use `NOT *` to display fields that contain no value.

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.



Depending on the number of columns displayed in the query results table, you may not be able to view additional added columns. You can remove one or more columns until your desired columns become visible.

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*.

Related information

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.
 - To resize a column, hover the mouse over the column header until a blue bar appears. Place the mouse over the bar and drag it right or left.
 - To move a column, click on the column header and drag it right or left.
 - When scrolling through the query results, be aware that the results may change as Insight automatically polls your data sources. This may result in some items being missing, or some items appearing out of order depending on how they are sorted.

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. Only the attributes for the displayed columns will be exported. Some displayed columns, particularly those that are part of complex nested relationships, are not exported.



When a comma appears in an asset name, the export encloses the name in quotes, preserving the asset name and the proper .csv format.

+ When exporting query results, be aware that **all** rows in the results table will be exported, not just those selected or displayed on the screen, up to a maximum of 10,000 rows.

+

When opening an exported .CSV file with Excel, if you have an object name or other field that is in the format NN:NN (two digits followed by a colon followed by two more digits), Excel will sometimes interpret that name as a Time format, instead of Text format. This can result in Excel displaying incorrect values in those columns. For example, an object named "81:45" would show in Excel as "81:45:00". To work around this, import the .CSV into Excel using the following steps:

+



- Open a new sheet in Excel.
- On the "Data" tab, choose "From Text".
- Locate the desired .CSV file and click "Import".
- In the Import wizard, choose "Delimited" and click Next.
- Choose "Comma" for the delimiter and click Next.
- Select the desired columns and choose "Text" for the column data format.
- Click Finish.

Your objects should show in Excel in the proper format.

+


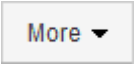
Modifying queries

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 Insightweb 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 Insightweb 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.

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 displays.

3. Select the desired assets in the list or click ☐ ▼ to select **All**.


The **Actions** button displays.

4. To add an application to the selected 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.

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.

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 displays.

3. Select the desired assets in the list or click ☐ ▼ to select **All**.

The **Actions** button displays.

4. To add an annotation to the assets or edit the value of an annotation assigned to the assets, click

Actions ▼,

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 **Actions** ▼,

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

Copying table values

You can copy values in tables for use in search boxes or other applications.

About this task

There are two methods you can use to copy values from tables or query results.

Steps

1. Method 1: Highlight the desired text with the mouse, copy it, and paste it into search fields or other applications.

2. Method 2: For single-value fields whose length exceeds the width of the table column, indicated by ellipses (...), hover over the field and click the clipboard icon. The value is copied to the clipboard for use in search fields or other applications.

Note that only values that are links to assets can be copied. Note also that only fields that include single values (i.e. non-lists) have the copy icon.

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.

- Volume

There are latency by tier and misaligned volume policies.

Creating performance policies

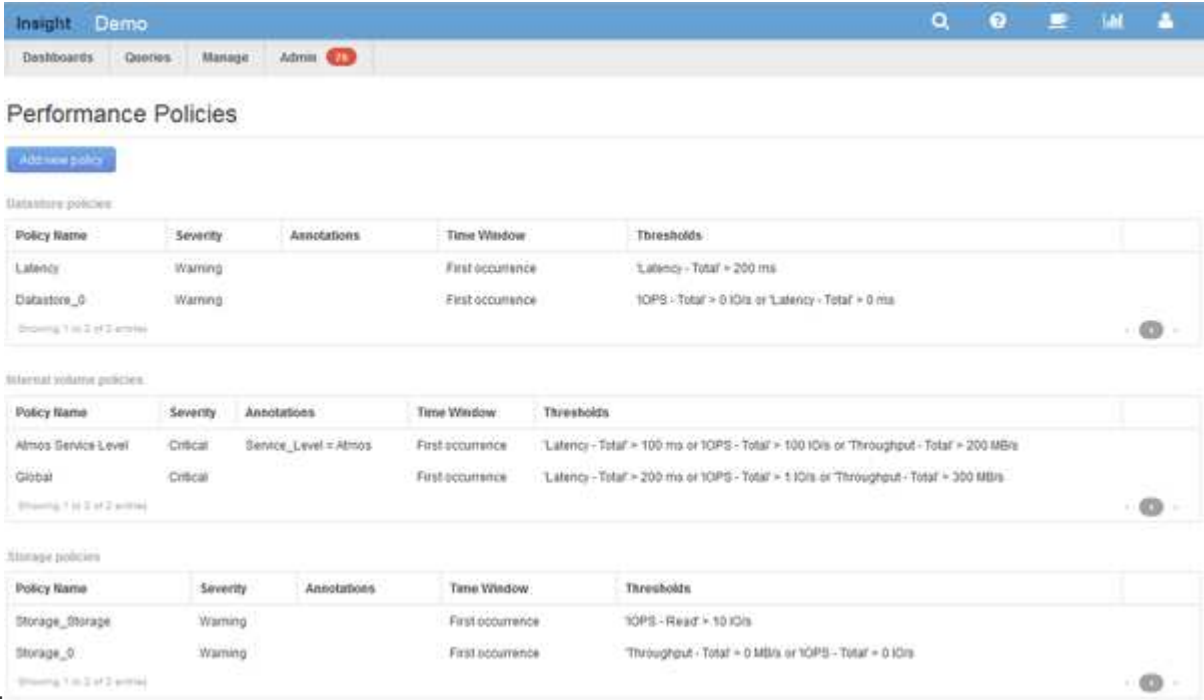
You create performance policies to set thresholds that trigger alerts to notify you about issues related to the resources in your network. For example, you can create a

performance policy to alert you when the total utilization for storage pools is greater than 60%.

Steps

- 1. Open OnCommand Insight in your browser.
- 2. Select **Manage > 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 is displayed.

- 4. In the **Policy name** field, enter a name for the policy.

You must use a name that is different from all 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 “Latency” policy for a different volume. The best practice is to always use a unique name for any policy, regardless of the object type.

- 5. From the **Apply to objects of type** list, select the type of object to which the policy applies.
- 6. From the **With annotation** list, select an annotation type, if applicable, and enter a value for the annotation in the **Value** box to apply the policy only to objects that have this particular annotation set.
- 7. If you selected **Port** as the object type, from the **Connected to** list, select what the port is connected to.
- 8. From 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 threshold is exceeded on the first sample of data. All other options trigger an alert when the threshold is crossed once and is continuously crossed for at least

the specified amount of time.

9. From the **With severity** list, select the severity for the violation.
10. By default, email alerts on policy violations will be sent to the recipients in the global email list. You can override these settings so that alerts for a particular policy are sent to specific recipients.
 - Click the link to open the recipients list, then click the **+** button to add recipients. Violation alerts for that policy will be sent to all recipients in the list.
11. Click the **any** link in the **Create alert if any of the following are true** section to control how alerts are triggered:
 - **any**

This is the default setting, which creates alerts when any of the thresholds related to a policy are crossed.
 - **all**

This setting creates an alert when all of the thresholds for a policy are crossed. When you select **all**, the first threshold that you create for a performance policy is referred to as the primary rule. You must ensure that the primary rule threshold is the violation that you are most concerned about for the performance policy.
12. In the **Create alert if** section, select a performance counter and an operator, and then enter a value to create a threshold.
13. Click **Add threshold** to add more thresholds.
14. To remove a threshold, click the trash can icon.
15. Select the **Stop processing further policies if alert is generated** check box if you want the policy to stop processing when an alert occurs.

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.
16. 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**.

Importing and Exporting user data

The import and export functions allow you to export annotations, annotation rules, queries, performance policies, and custom dashboards to one file. This file can then be imported into different OnCommand Insight servers.

The export and import functions are supported only between servers that are running the same version of OnCommand Insight.

To Export or Import user data, Click on **Admin** and select **Setup**, then choose the **Import/Export user data** tab.

During the import operation, data is added, merged, or replaced, depending on the objects and object types that are being imported.

- Annotation Types
 - Adds an annotation if no annotation with the same name exists in the target system.
 - Merges an annotation if the annotation type is a list, and an annotation with the same name exists in the target system.
 - Replaces an annotation if the annotation type is anything other than a list, and an annotation with the same name exists in the target system.



If an annotation with the same name but with a different type exists in the target system, the import fails. If objects depend on the failed annotation, those objects may show incorrect or unwanted information. You must check all annotation dependencies after the import operation is complete.

- Annotation Rules

- Adds an annotation rule if no annotation rule with the same name exists in the target system.
- Replaces an annotation rule if an annotation rule with the same name exists in the target system.



Annotation rules are dependent on both queries and annotations. You must check all the annotation rules for accuracy after the import operation is complete.

- Policies

- Adds a policy if no policy with the same name exists in the target system.
- Replaces a policy if a policy with the same name exists in the target system.



Policies may be out of order after the import operation is complete. You must check the policy order after the import. Policies that are dependent on annotations may fail if the annotations are incorrect. You must check all the annotation dependencies after the import.

+

- Queries

- Adds a query if no query with the same name exists in the target system.
- Replaces a query if a query with the same name exists in the target system, even if the resource type of the query is different.



If the resource type of a query is different, after the import, any dashboard widgets that use that query may display unwanted or incorrect results. You must check all query-based widgets for accuracy after the import. Queries that are dependent on annotations may fail if the annotations are incorrect. You must check all the annotation dependencies after the import.

+

- Dashboards

- Adds a dashboard if no dashboard with the same name exists in the target system.
- Replaces a dashboard if a dashboard with the same name exists in the target system, even if the resource type of the query is different.



You must check all query-based widgets in dashboards for accuracy after the import. If the source server has multiple dashboards with the same name, they are all exported. However, only the first one will be imported to the target server. To avoid errors during import, you should ensure that your dashboards have unique names before exporting them.

+

Insight Security

The 7.3.1 release of OnCommand Insight introduced security features that allow Insight environments to operate with enhanced security. The features include improvements to encryption, password hashing, and the ability to change internal user passwords and key pairs that encrypt and decrypt passwords. You can manage these features on all servers in the Insight environment.

The default installation of Insight includes a security configuration where all sites in your environment share the same keys and the same default passwords. To protect sensitive data, NetApp recommends you change the default keys and the Acquisition user password after an installation or upgrade.

Data source encrypted passwords are stored in the Insight Server database. The Server has a public key and encrypts passwords when a user enters them in a WebUI data source configuration page. The Server does not have the private keys required to decrypt the data source passwords stored in the Server database. Only Acquisition Units (LAU, RAU) have the data source private key required to decrypt data source passwords.

Rekeying servers

Using default keys introduces security vulnerability in your environment. By default, data source passwords are stored encrypted in the Insight database. They are encrypted using a key that is common to all Insight installations. In a default configuration, an Insight database sent to NetApp includes passwords that could theoretically be decrypted by NetApp.

Changing the Acquisition user password

Using the default 'Acquisition' user password introduces security vulnerability into your environment. All Acquisition Units use the "Acquisition" user to communicate with the Server. RAUs with default passwords can theoretically connect to any Insight server using default passwords.


Upgrade and installation considerations

When your Insight system contains non-default security configurations (you have rekeyed or changed passwords), you must back up your security configurations. Installing new software, or in some cases upgrading software, reverts your system to a default security configuration. When your system reverts to the default configuration, you must restore the non-default configuration in order for the system to operate correctly.

Managing keys in a complex service provider environment

A service provider can host multiple OnCommand Insight customers collecting data. The keys protect customer data from unauthorized access by multiple customers on the Insight server. Each customer's data is protected by their specific key pairs.

This implementation of Insight could be configured as shown in the following illustration.



You need to create individual keys for each customer in this configuration. Customer A requires identical keys for both RAUs. Customer B requires a single set of keys.

The steps you would take to change encryption keys for Customer A:

1. Perform a remote login to the server hosting RAU1.
2. Start the security admin tool.
3. Select Change Encryption Key to replace the default keys.
4. Select Backup to create a backup zip file of the security configuration.
5. Perform a remote login to the server hosting RAU2.
6. Copy the backup zip file of the security configuration to RAU2.
7. Start the security admin tool.
8. Restore the security backup from RAU1 to the current server.

The steps you would take to change encryption keys for Customer B:

1. Perform a remote login to the server hosting RAU3.
2. Start the security admin tool.
3. Select Change Encryption Key to replace the default keys.
4. Select Backup to create a backup zip file of the security configuration.

Managing security on the Insight server

The `securityadmin` tool allows you to manage security options on the Insight server. Security management includes changing passwords, generating new keys, saving and restoring security configurations you create, or restoring configurations to the default settings.

About this task

You use the `securityadmin` tool to manage security:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat`
- Linux - `/bin/oci-securityadmin.sh`

Steps

1. Perform a remote login to the Insight server.
2. Start the security admin tool in interactive mode:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat -i`
- Linux - `/bin/oci-securityadmin.sh -i`

The system requests login credentials.

3. Enter the user name and password for an account with “Admin” credentials.
4. Select **Server**.

The following server configuration options are available:

- **Backup**

Creates a backup zip file of the vault containing all passwords and keys and places the file in a location specified by the user, or in the following default locations:

- Windows - `C:\Program Files\SANscreen\backup\vault`
- Linux - `/var/log/netapp/oci/backup/vault`

- **Restore**

Restores the zip backup of the vault that was created. Once restored, all passwords and keys are reverted to values existing at the time of the backup creation.



Restore can be used to synchronize passwords and keys on multiple servers, for example: - Change the server encryption key on one server - Create a backup of the vault - Restore the vault backup to the second server

- **Change Encryption Key**

Change the server encryption key that is used to encrypt or decrypt proxy user passwords, SMTP user passwords, LDAP user passwords, and so on.



When you change encryption keys, you should backup your new security configuration so that you can restore it after an upgrade or installation.

- **Update Password**

Change password for the internal accounts that are used by Insight. The following options are displayed:

- `_internal`
- `acquisition`
- `cognos_admin`
- `dwh_internal`
- `hosts`
- `inventory`
- `root`



Some accounts need to be synchronized when passwords are changed. For example, if you change the password for the 'acquisition' user on the server, you need to change the password for the 'acquisition' user on the LAU, RAU, and DWH to match. Also, when you change passwords, you should backup your new security configuration so that you can restore it after an upgrade or installation.

- **Reset to Defaults**

Resets keys and passwords to default values. Default values are those provided during installation.

- **Exit**

Exit the `securityadmin` tool.

1. Chose the option you want to change and follow the prompts.

Managing security on the local acquisition unit

The `securityadmin` tool allows you to manage security options on the local acquisition user (LAU). Security management includes managing keys and passwords, saving and restoring security configurations you create or restoring configurations to the default settings.

Before you begin

You must have `admin` privileges to perform security configuration tasks.

About this task

You use the `securityadmin` tool to manage security:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat`
- Linux - `/bin/oci-securityadmin.sh`

Steps

1. Perform a remote login to the Insight server.
2. Start the security admin tool in interactive mode:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat -i`

- Linux - `/bin/oci-securityadmin.sh -i`

The system requests login credentials.

3. Enter the user name and password for an account with “Admin” credentials.
4. Select **Local Acquisition Unit** to reconfigure the Local Acquisition Unit security configuration.

The following options are displayed:

- **Backup**

Creates a backup zip file of the vault containing all passwords and keys and places the file in a location specified by the user, or in the following default locations:

- Windows - `C:\Program Files\SANscreen\backup\vault`
- Linux - `/var/log/netapp/oci/backup/vault`

- **Restore**

Restores the zip backup of the vault that was created. Once restored, all passwords and keys are reverted to values existing at the time of the backup creation.



Restore can be used to synchronize passwords and keys on multiple servers, for example: - Change encryption keys on the LAU - Create a backup of the vault - Restore the vault backup to each of the RAUs

- **Change Encryption Keys**

Change the AU encryption keys used to encrypt or decrypt device passwords.



When you change encryption keys, you should backup your new security configuration so that you can restore it after an upgrade or installation.

- **Update Password**

Change password for 'acquisition' user account.



Some accounts need to be synchronized when passwords are changed. For example, if you change the password for the 'acquisition' user on the server, you need to change the password for the 'acquisition' user on the LAU, RAU, and DWH to match. Also, when you change passwords, you should backup your new security configuration so that you can restore it after an upgrade or installation.

- **Reset to Defaults**

Resets acquisition user password and acquisition user encryption keys to default values, Default values are those provided during installation.

- **Exit**

Exit the `securityadmin` tool.

5. Chose the option you want configure and follow the prompts.

Managing security on an RAU

The `securityadmin` tool allows you to manage security options on RAUs. You might need to backup or restore a vault configuration, change encryption keys, or update passwords for the acquisition units.

About this task

You use the `securityadmin` tool to manage security:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat`
- Linux - `/bin/oci-securityadmin.sh`

One scenario for updating the security configuration for the LAU, RAU is to update the 'acquisition' user password when the password for that user has been changed on the server. All of the RAUs, and the LAU use the same password as that of the server 'acquisition' user to communicate with the server.

The 'acquisition' user only exists on the Insight server. The RAU or LAU logs in as that user when they connect to the server.

Use the following steps to manage security options on an RAU:

Steps

1. Perform a remote login to the server running the RAU

2. Start the security admin tool in interactive mode:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat -i`
- Linux - `/bin/oci-securityadmin.sh -i`

The system requests login credentials.

3. Enter the user name and password for an account with “Admin” credentials.

The system displays the menu for the RAU.

- **Backup**

Creates a backup zip file of the vault containing all passwords and keys and places the file in a location specified by the user, or in the following default locations:

- Windows - `C:\Program Files\SANscreen\backup\vault`
- Linux - `/var/log/netapp/oci/backup/vault`

- **Restore**

Restores the zip backup of the vault that was created. Once restored, all passwords and keys are reverted to values existing at the time of the backup creation.



Restore can be used to synchronize passwords and keys on multiple servers, for example: - Change encryption keys on one server - Create a backup of the vault - Restore the vault backup to the second server

- **Change Encryption Keys**

Change the RAU encryption keys used to encrypt or decrypt device passwords.



When you change encryption keys, you should backup your new security configuration so that you can restore it after an upgrade or installation.

- **Update Password**

Change password for 'acquisition' user account.



Some accounts need to be synchronized when passwords are changed. For example, if you change the password for the 'acquisition' user on the server, you need to change the password for the 'acquisition' user on the LAU, RAU, and DWH to match. Also, when you change passwords, you should backup your new security configuration so that you can restore it after an upgrade or installation.

- **Reset to Defaults**

Resets encryption keys and passwords to default values. Default values are those provided during installation.

- **Exit**

Exit the `securityadmin` tool.

Managing security on the Data Warehouse

The `securityadmin` tool allows you to manage security options on the Data Warehouse server. Security management includes updating internal passwords for internal users on the DWH server, creating backups of the security configuration, or restoring configurations to the default settings.

About this task

You use the `securityadmin` tool to manage security:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat`
- Linux - `/bin/oci-securityadmin.sh`

Steps

1. Perform a remote login to the Data Warehouse server.
2. Start the security admin tool in interactive mode:

- Windows - `C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat -i`
- Linux - `/bin/oci-securityadmin.sh -i`

The system requests login credentials.

3. Enter the user name and password for an account with “Admin” credentials.

The system displays the security admin menu for the Data Warehouse:

- **Backup**

Creates a backup zip file of the vault containing all passwords and keys and places the file in a location specified by the user, or in the default location:

- Windows - C:\Program Files\SANscreen\backup\vault
- Linux - /var/log/netapp/oci/backup/vault

- **Restore**

Restores the zip backup of the vault that was created. Once restored, all passwords and keys are reverted to values existing at the time of the backup creation.



Restore can be used to synchronize passwords and keys on multiple servers, for example: - Change encryption keys on one server - Create a backup of the vault - Restore the vault backup to the second server

+

- **Change encryption keys**

Change the DWH encryption key used to encrypt or decrypt passwords such as connector passwords and SMTP passwords.

- **Update Password**

Change password for a specific user account.

- _internal
- acquisition
- cognos_admin
- dwh
- dwh_internal
- dwhuser
- hosts
- inventory
- root



When you change the dwhuser, hosts, inventory, or root passwords, you have the option to use SHA-256 password hashing. This options requires that all clients accessing the accounts use SSL connections.

- **Reset to Defaults**

Resets encryption keys and passwords to default values. Default values are those provided during installation.

- **Exit**

Exit the `securityadmin` tool.

Changing OnCommand Insight internal user passwords

Security policies might require you to change the passwords in your OnCommand Insight environment. Some of the passwords on one server exist on a different server in the environment, requiring that you change the password on both servers. For example, when you change the “inventory” user password on the Insight Server you must match the “inventory” user password on the Data Warehouse server Connector configured for that Insight Server.

Before you begin



You should understand the dependencies of the user accounts before you change passwords. Failing to update passwords on all required servers will result in communication failures between the Insight components.

About this task

The following table lists the internal user passwords for the Insight Server and lists the Insight components that have dependent passwords that need to match the new password.

Insight Server Passwords	Required changes
_internal	
acquisition	LAU, RAU
dwh_internal	Data Warehouse
hosts	
inventory	Data Warehouse
root	

The following table lists the internal user passwords for the Data Warehouse and lists the Insight components that have dependent passwords that need to match the new password.

Data Warehouse Passwords	Required changes
cognos_admin	
dwh	

dwh_internal (Changed using the Server Connector configuration UI)	Insight server
dwhuser	
hosts	
inventory (Changed using the Server Connector configuration UI)	Insight server
root	

Changing passwords in the DWH Server Connection Configuration UI

The following table lists the user password for the LAU and lists the Insight components that have dependent passwords that need to match the new password.

LAU Passwords	Required changes
acquisition	Insight Server, RAU

Changing the “inventory” and “dwh_internal” passwords using the Server Connection Configuration UI

If you need to change the “inventory” or “dwh_internal” passwords to match those on the Insight server you use the Data Warehouse UI.

Before you begin

You must be logged in as administrator to perform this task.

Steps

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

The **Edit Connector** screen is displayed.

Edit Connector

ID:	<input type="text" value="1"/>
Encryption:	<input type="text" value="Enabled"/>
Name:	<input type="text" value="Oci-stg06-s12r2.nane.netapp.com"/>
Host:	<input type="text" value="Oci-stg06-s12r2.nane.netapp.com"/>
Database user name:	<input type="text" value="inventory"/>
Database password:	<input type="password" value="••••••••"/>

[Advanced](#) ▼

3. Enter a new “inventory” password for the **Database password** field.
4. Click **Save**
5. To change the “dwh_internal” password, click **Advanced**.

The Edit Connector Advanced screen is displayed.

Edit Connector

ID:	<input type="text" value="1"/>
Encryption:	<input type="text" value="Enabled"/>
Name:	<input type="text" value="Oci-stg06-s12r2.nane.netapp.com"/>
Host:	<input type="text" value="Oci-stg06-s12r2.nane.netapp.com"/>
Database user name:	<input type="text" value="inventory"/>
Database password:	<input type="password" value="....."/>
Server user name:	<input type="text" value="dwh_internal"/>
Server password:	<input type="password" value="....."/>
HTTPS port:	<input type="text" value="443"/>
TCP port:	<input type="text" value="3306"/>

Basic ^

6. Enter the new password in the **Server password** field:

7. Click save.

Changing the dwh password using the ODBC Administration tool

When you change the password on for the dwh user on the Insight server, the password must also be changed on the Data Warehouse server. You use the ODBC Data Source Administrator tool to change the password on the Data Warehouse.


Before you begin

You must perform a remote login to the Data Warehouse server using an account with administrator privileges.

Steps


1. Perform a remote login to the server hosting that Data Warehouse.
2. Access the ODBC Administration tool at C:\Windows\SysWOW64\odbcad32.exe

The system displays the ODBC Data Source Administrator screen.



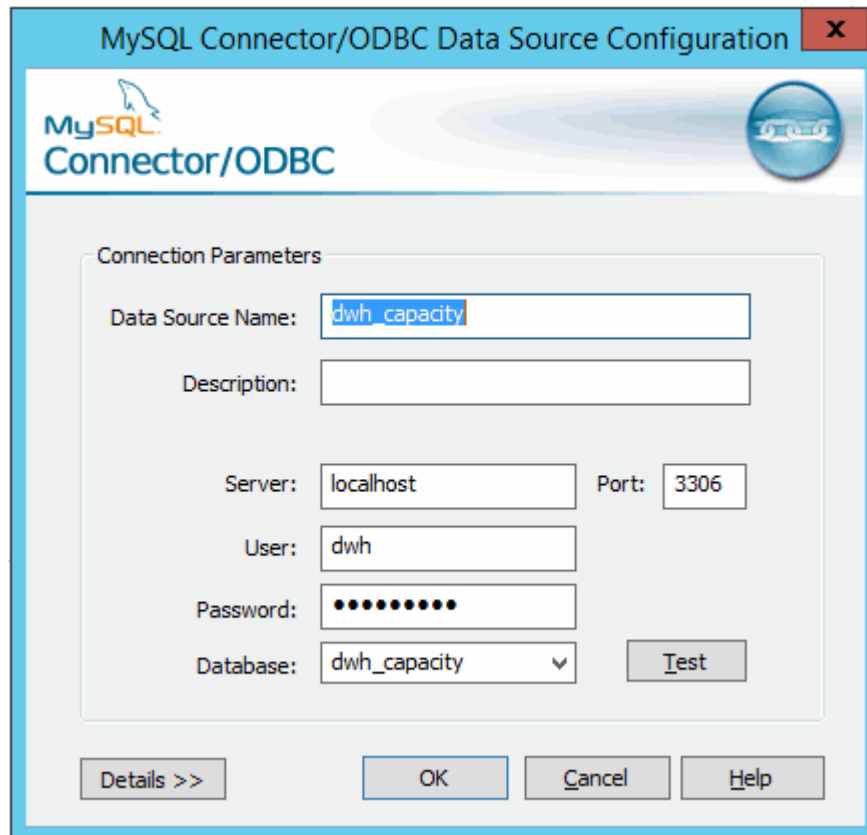
3. Click **System DSN**

The system data sources are displayed.



4. Select an OnCommand Insight Data Source from the list.
5. Click **Configure**

The Data Source Configuration screen is displayed.



The image shows a screenshot of the 'MySQL Connector/ODBC Data Source Configuration' dialog box. The title bar is blue with the text 'MySQL Connector/ODBC Data Source Configuration' and a red close button. The dialog has a light blue header with the MySQL logo and 'Connector/ODBC' text. The main area is titled 'Connection Parameters' and contains several input fields: 'Data Source Name' (containing 'dwh_capacity'), 'Description' (empty), 'Server' (containing 'localhost'), 'Port' (containing '3306'), 'User' (containing 'dwh'), 'Password' (containing masked characters), and 'Database' (a dropdown menu showing 'dwh_capacity'). There is a 'Test' button next to the Database field. At the bottom, there are four buttons: 'Details >>', 'OK', 'Cancel', and 'Help'.

6. Enter the new password in the **Password** field.

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 the system to enable these features.

After configuring the system 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 CAC users, smart cards contain multiple certificates, only one of which can match the trusted CA. The CAC certificate for identification should be used.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

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 ID.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Use the `regedit` utility to modify registry values in `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Apache Software Foundation\Procrun2.0\SANscreen Server\Parameters\Java:`

 - a. Change the `JVM_Option DclientAuth=false` to `DclientAuth=true`.

2. Back up the keystore file: `C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore`
3. Open a command prompt specifying `Run as administrator`

4. Delete the self-generated certificate: `C:\Program Files\SANscreen\java64\bin\keytool.exe -delete -alias "ssl certificate" -keystore C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore`
5. Generate a new certificate: `C:\Program Files\SANscreen\java64\bin\keytool.exe -genkey -alias "alias_name" -keyalg RSA -sigalg SHA1withRSA -keysize 2048 -validity 365 -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -dname "CN=commonName,OU=orgUnit,O=orgName,L=localityNameI,S=stateName,C=countryName"`
6. Generate a certificate signing request (CSR): `C:\Program Files\SANscreen\java64\bin\keytool.exe -certreq -sigalg SHA1withRSA -alias "alias_name" -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -file C:\temp\server.csr"`
7. After the CSR is returned in step 6, import the certificate, then export the certificate in Base-64 format and place it in "C:\temp" named servername.cer.
8. Extract the certificate from the keystore: `C:\Program Files\SANscreen\java64\bin\keytool.exe -v -importkeystore -srckeystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -srcalias "alias_name" -destkeystore "C:\temp\file.p12" -deststoretype PKCS12`
9. Extract a private key from the p12 file: `openssl pkcs12 -in "C:\temp\file.p12" -out "C:\temp\servername.private.pem"`
10. Merge the Base-64 certificate that you exported in step 7 with the private key: `openssl pkcs12 -export -in "<folder>\<certificate>.cer" -inkey "C:\temp\servername.private.pem" -out "C:\temp\servername.new.p12" -name "servername.abc.123.yyy.zzz"`
11. Import the merged certificate into the keystore: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -destkeystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -srckeystore "C:\temp\servername.new.p12" -srcstoretype PKCS12 -alias "alias_name"`
12. Import the root certificate: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -file "C:\<root_certificate>.cer" -trustcacerts -alias "alias_name"`
13. Import the root certificate into the server.trustore: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore" -file "C:\<email_certificate>.cer" -trustcacerts -alias "alias_name"`
14. Import the intermediate certificate: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore" -file "C:\<intermediate_certificate>.cer" -trustcacerts -alias "alias_name"`

Repeat this step for all intermediate certificates.

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

1. Restart the server.

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.

Before you begin

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

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>)

Enabling CAC on a Linux server

Some modifications are required to enable CAC on a Linux OnCommand Insight server.

Steps

1. Navigate to `/opt/netapp/oci/conf/`
2. Edit `wildfly.properties` and change the value of `CLIENT_AUTH_ENABLED` to "True"
3. Import the "root certificate" that exists under `/opt/netapp/oci/wildfly/standalone/configuration/server.keystore`
4. Restart the server

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.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Use `regedit` to modify registry values in `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Apache Software Foundation\Procrun2.0\SANscreen Server\Parameters\Java`

- a. Change the `JVM_Option -DclientAuth=false` to `-DclientAuth=true`.

For Linux, modify the `clientAuth` parameter in `/opt/netapp/oci/scripts/wildfly.server`

2. Add certificate authorities (CAs) to the Data Warehouse trustore:

- a. In a command window, go to `..\SANscreen\wildfly\standalone\configuration`.
- b. Use the `keytool` utility to list the trusted CAs: `C:\Program Files\SANscreen\java64\bin\keytool.exe -list -keystore server.trustore -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\wildfly\standalone\configuration` and use the `keytool import` command: `C:\Program Files\SANscreen\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.

3. On the OnCommand Insight server, the `wildfly/standalone/configuration/standalone-full.xml` file needs to be modified by updating `verify-client` to `"REQUESTED"` in

/subsystem=undertow/server=default-server/https-listener=default-httpsto enable CAC. Log in to the Insight server and run the appropriate command:

OS	Script
Windows	<install dir>\SANscreen\wildfly\bin\enableCACforRemoteEJB.bat
Linux	/opt/netapp/oci/wildfly/bin/enableCACforRemoteEJB.sh

After executing the script, wait until the reload of the wildfly server is complete before proceeding to the next step.

4. Restart the OnCommand Insight server.

Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.5 through 7.3.9)

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

Before you begin

This procedure is for systems running OnCommand Insight 7.3.5 through 7.3.9.



For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):

- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Add certificate authorities (CAs) to the Cognos truststore.
 - a. In a command window, go to `..\SANscreen\cognos\analytics\configuration\certs\`
 - b. Use the `keytool` utility to list the trusted CAs: `..\..\jre\bin\keytool.exe -list -keystore CAMKeystore.jks -storepass NoPassWordSet`

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.
- d. To include customer's CAs with OnCommand Insight trusted CAs, go to
`..\SANscreen\cognos\analytics\configuration\certs\`.
- e. Use the keytool utility to import the .pem file: `..\..\jre\bin\keytool.exe -importcert -keystore CAMKeystore.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 `NoPassWordSet`.
 - g. Answer `yes` when prompted to trust the certificate.
2. To enable CAC mode, execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`
 3. To disable CAC mode, execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`

Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.10 and later)

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

Before you begin

This procedure is for systems running OnCommand Insight 7.3.10 and later.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Add certificate authorities (CAs) to the Cognos truststore.
 - a. In a command window, go to `..\SANscreen\cognos\analytics\configuration\certs\`
 - b. Use the keytool utility to list the trusted CAs: `..\..\ibm-jre\jre\bin\keytool.exe -list -keystore CAMKeystore.jks -storepass NoPassWordSet`

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.
- d. To include customer's CAs with OnCommand Insight trusted CAs, go to
`..\SANscreen\cognos\analytics\configuration\certs\`.
- e. Use the keytool utility to import the .pem file: `..\..\ibm-jre\jre\bin\keytool.exe`
`-importcert -keystore CAMKeystore.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 `NoPassWordSet`.
 - g. Answer `yes` when prompted to trust the certificate.
2. To enable CAC mode, do the following:
 - a. Configure CAC logout page, using the following steps:
 - Logon to Cognos portal (user must be part of System Administrators group i.e. `cognos_admin`)
 - (Only for 7.3.10 and 7.3.11) Click Manage -> Configuration -> System -> Security
 - (Only for 7.3.10 and 7.3.11) Enter `cacLogout.html` against Logout Redirect URL -> Apply
 - Close browser.
 - b. Execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`
 - c. Start IBM Cognos service. Wait for Cognos service to start.
 3. To disable CAC mode, do the following:
 - a. Execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`
 - b. Start IBM Cognos service. Wait for Cognos service to start.
 - c. (Only for 7.3.10 and 7.3.11) Unconfigure CAC logout page, using the following steps:
 - Logon to Cognos portal (user must be part of System Administrators group i.e. `cognos_admin`)
 - Click Manage -> Configuration -> System -> Security
 - Enter `cacLogout.html` against Logout Redirect URL -> Apply
 - Close browser.

Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.5 to 7.3.9)

You can add SSL certificates to enable enhanced authentication and encryption for your Data Warehouse and Cognos environment.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.5 through 7.3.9.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

About this task

You must have admin privileges to perform this procedure.

Steps

1. Create a backup of `..\SANSscreen\cognos\analytics\configuration\cogstartup.xml`.
2. Create a backup of the “certs” and “csk” folders under `..\SANSscreen\cognos\analytics\configuration`.
3. Generate a Certificate Encryption Request from Cognos. In an Admin CMD window, run:
 - a. `cd “\Program Files\sansscreen\cognos\analytics\bin”`
 - b. `ThirdPartyCertificateTool.bat -java:local -c -e -p NoPassWordSet -a RSA -d “CN=FQDN,O=orgname,C=US” -r c:\temp\encryptRequest.csr`
4. Open the `c:\temp\encryptRequest.csr` file and copy the generated content.
5. Send the `encryptRequest.csr` to the certificate authority (CA) to obtain an SSL certificate.

Make sure to add additional attributes such as “SAN:dns=FQDN (For example, hostname.netapp.com)” to add the SubjectAltName. Google Chrome version 58 and later complains if the SubjectAltName is missing from the certificate.

6. Download the chain certificates by including root certificate by using PKCS7 format

This will download `fqdn.p7b` file

7. Get a cert in `.p7b` format from your CA. Use a name that marks it as the certificate for the Cognos Webserver.
8. `ThirdPartyCertificateTool.bat` fails to import the entire chain, so multiple steps are required to export all certificates. Split the chain by exporting them individually as follows:
 - a. Open the `.p7b` certificate in “Crypto Shell Extensions”.
 - b. Browse in the left pane to “Certificates”.
 - c. Right-click on root CA > All Tasks > Export.
 - d. Select Base64 output.

- e. Enter a file name identifying it as the root certificate.
- f. Repeat steps 8a through 8c to export all of the certificates separately into .cer files.
- g. Name the files intermediateX.cer and cognos.cer.
9. Ignore this step if you have only one CA certificate, otherwise merge both root.cer and intermediateX.cer into one file.
 - a. Open intermediate.cer with NotePad and copy the content.
 - b. Open root.cer with NotePad and save the content from 9a.
 - c. Save the file as CA.cer.
10. Import the certificates into the Cognos keystore using the Admin CMD prompt:
 - a. cd "Program Files\sansscreen\cognos\analytics\bin"
 - b. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\CA.cer

This will set CA.cer as root Certificate Authority.
 - c. ThirdPartyCertificateTool.bat -java:local -i -e -r c:\temp\cognos.cer -t c:\temp\CA.cer

This will set Cognos.cer as encryption certificate which is signed by CA.cer.
11. Open the IBM Cognos Configuration.
 - a. Select Local Configuration → Security → Cryptography → Cognos
 - b. Change "Use third party CA?" to True.
 - c. Save the configuration.
 - d. Restart Cognos
12. Export the latest Cognos certificate into cognos.crt using the Admin CMD prompt:
 - a. "D:\Program Files\SANscreen\java\bin\keytool.exe" -exportcert -file "c:\temp\cognos.crt" -keystore "D:\Program Files\SANscreen\cognos\analytics\configuration\certs\CAMKeystore" -storetype PKCS12 -storepass NoPassWordSet -alias encryption
13. Import the "c:\temp\cognos.crt" into dwh trustore to establish SSL communication between Cognos and DWH, using the Admin CMD prompt window.
 - a. "D:\Program Files\SANscreen\java\bin\keytool.exe" -importcert -file "c:\temp\cognos.crt" -keystore "D:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore" -storepass changeit -alias cognoscert
14. Restart the SANscreen service.
15. Perform a backup of DWH to make sure DWH communicates with Cognos.

Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.10 and later)

You can add SSL certificates to enable enhanced authentication and encryption for your Data Warehouse and Cognos environment.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.10 and later.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

About this task

You must have admin privileges to perform this procedure.

Steps

1. Stop Cognos using the IBM Cognos Configuration tool. Close Cognos.
2. Create backups of the `..\SANSscreen\cognos\analytics\configuration` and `..\SANSscreen\cognos\analytics\temp\cam\freshness` folders.
3. Generate a Certificate Encryption Request from Cognos. In an Admin CMD window, run:
 - a. `cd "\Program Files\sansscreen\cognos\analytics\bin"`
 - b. `ThirdPartyCertificateTool.bat -java:local -c -e -p NoPassWordSet -a RSA -r c:\temp\encryptRequest.csr -d "CN=server.domain.com,O=NETAPP,C=US" -H "server.domain.com" -I "ipaddress"`. Note: here -H and -I are to add subjectAltNames like dns and ipaddress.
4. Open the `c:\temp\encryptRequest.csr` file and copy the generated content.
5. Input the `encryptRequest.csr` content and generate certificate using CA signing portal.
6. Download the chain certificates by including root certificate by using PKCS7 format

This will download `fqdn.p7b` file

7. Get a cert in `.p7b` format from your CA. Use a name that marks it as the certificate for the Cognos Webserver.
8. `ThirdPartyCertificateTool.bat` fails to import the entire chain, so multiple steps are required to export all certificates. Split the chain by exporting them individually as follows:
 - a. Open the `.p7b` certificate in "Crypto Shell Extensions".
 - b. Browse in the left pane to "Certificates".
 - c. Right-click on root CA > All Tasks > Export.
 - d. Select Base64 output.
 - e. Enter a file name identifying it as the root certificate.
 - f. Repeat steps 8a through 8e to export all of the certificates separately into `.cer` files.

- g. Name the files intermediateX.cer and cognos.cer.
9. Ignore this step if you have only one CA certificate, otherwise merge both root.cer and intermediateX.cer into one file.
 - a. Open root.cer with NotePad and copy the content.
 - b. Open intermediate.cer with NotePad and append the content from 9a (intermediate first and root next).
 - c. Save the file as chain.cer.
10. Import the certificates into the Cognos keystore using the Admin CMD prompt:
 - a. cd "Program Files\sansscreen\cognos\analytics\bin"
 - b. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\root.cer
 - c. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\intermediate.cer
 - d. ThirdPartyCertificateTool.bat -java:local -i -e -r c:\temp\cognos.cer -t c:\temp\chain.cer
11. Open the IBM Cognos Configuration.
 - a. Select Local Configuration → Security → Cryptography → Cognos
 - b. Change "Use third party CA?" to True.
 - c. Save the configuration.
 - d. Restart Cognos
12. Export the latest Cognos certificate into cognos.crt using the Admin CMD prompt:
 - a. cd "C:\Program Files\SANscreen"
 - b. java\bin\keytool.exe -exportcert -file c:\temp\cognos.crt -keystore cognos\analytics\configuration\certs\CAMKeystore -storetype PKCS12 -storepass NoPassWordSet -alias encryption
13. Back up the DWH server trustore


```
at.. \SANscreen\wildfly\standalone\configuration\server.trustore
```
14. Import the "c:\temp\cognos.crt" into DWH trustore to establish SSL communication between Cognos and DWH, using the Admin CMD prompt window.
 - a. cd "C:\Program Files\SANscreen"
 - b. java\bin\keytool.exe -importcert -file c:\temp\cognos.crt -keystore wildfly\standalone\configuration\server.trustore -storepass changeit -alias cognos3rdca
15. Restart the SANscreen service.
16. Perform a backup of DWH to make sure DWH communicates with Cognos.
17. The following steps should be performed even when only the "ssl certificate" is changed and the default Cognos certificates are left unchanged. Otherwise Cognos may complain about the new SANscreen certificate or be unable to create a DWH backup.
 - a. cd "%SANSSCREEN_HOME%cognos\analytics\bin\"
 - b. "%SANSSCREEN_HOME%java64\bin\keytool.exe" -exportcert -file "c:\temp\sansscreen.cer" -keystore "%SANSSCREEN_HOME%wildfly\standalone\configuration\server.keystore" -storepass changeit -alias "ssl certificate"
 - c. ThirdPartyCertificateTool.bat -java:local -i -T -r "c:\temp\sansscreen.cer"

Typically, these steps are performed as part of the Cognos certificate import process described in [How to](#)

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.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Use regedit to modify registry values in `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Apache Software Foundation\Procrun2.0\SANscreen Server\Parameters\Java`

- a. Change the JVM_Option `-DclientAuth=false` to `-DclientAuth=true`.

For Linux, modify the `clientAuth` parameter in `/opt/netapp/oci/scripts/wildfly.server`

2. Add certificate authorities (CAs) to the Data Warehouse trustore:

- a. In a command window, go to `..\SANscreen\wildfly\standalone\configuration`.
- b. Use the `keytool` utility to list the trusted CAs: `C:\Program Files\SANscreen\java64\bin\keytool.exe -list -keystore server.trustore -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\wildfly\standalone\configuration and use the keytool import command: `C:\Program Files\SANscreen\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.

3. On the OnCommand Insight server, the `wildfly/standalone/configuration/standalone-full.xml` file needs to be modified by updating `verify-client` to "REQUESTED" in `/subsystem=undertow/server=default-server/https-listener=default-httpsto` to enable CAC. Log in to the Insight server and run the appropriate command:

OS	Script
Windows	<code><install dir>\SANscreen\wildfly\bin\enableCACforRemoteEJB.bat</code>
Linux	<code>/opt/netapp/oci/wildfly/bin/enableCACforRemoteEJB.sh</code>

After executing the script, wait until the reload of the wildfly server is complete before proceeding to the next step.

4. Restart the OnCommand Insight server.

Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.5 through 7.3.9)

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

Before you begin

This procedure is for systems running OnCommand Insight 7.3.5 through 7.3.9.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Add certificate authorities (CAs) to the Cognos truststore.

- In a command window, go to `..\SANscreen\cognos\analytics\configuration\certs\`
- Use the `keytool` utility to list the trusted CAs: `..\..\jre\bin\keytool.exe -list -keystore CAMKeystore.jks -storepass NoPassWordSet`

The first word in each line indicates the CA alias.

- 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\analytics\configuration\certs\`.
- Use the `keytool` utility to import the `.pem` file: `..\..\jre\bin\keytool.exe -importcert -keystore CAMKeystore.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.

- When prompted for a password, enter `NoPassWordSet`.
- Answer `yes` when prompted to trust the certificate.

2. To enable CAC mode, execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`

3. To disable CAC mode, execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`

Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.10 and later)

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

Before you begin

This procedure is for systems running OnCommand Insight 7.3.10 and later.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Add certificate authorities (CAs) to the Cognos truststore.

- In a command window, go to `..\SANscreen\cognos\analytics\configuration\certs\`
- Use the `keytool` utility to list the trusted CAs: `..\..\ibm-jre\jre\bin\keytool.exe -list -keystore CAMKeystore.jks -storepass NoPassWordSet`

The first word in each line indicates the CA alias.

- 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\analytics\configuration\certs\`.
- Use the `keytool` utility to import the `.pem` file: `..\..\ibm-jre\jre\bin\keytool.exe -importcert -keystore CAMKeystore.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.

- When prompted for a password, enter `NoPassWordSet`.
- Answer `yes` when prompted to trust the certificate.

2. To enable CAC mode, do the following:

- Configure CAC logout page, using the following steps:
 - Logon to Cognos portal (user must be part of System Administrators group i.e. `cognos_admin`)
 - (Only for 7.3.10 and 7.3.11) Click Manage -> Configuration -> System -> Security
 - (Only for 7.3.10 and 7.3.11) Enter `cacLogout.html` against Logout Redirect URL -> Apply
 - Close browser.
- Execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`
- Start IBM Cognos service. Wait for Cognos service to start.

3. To disable CAC mode, do the following:

- Execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`

- b. Start IBM Cognos service. Wait for Cognos service to start.
- c. (Only for 7.3.10 and 7.3.11) Unconfigure CAC logout page, using the following steps:
 - Logon to Cognos portal (user must be part of System Administrators group i.e. cognos_admin)
 - Click Manage -> Configuration -> System -> Security
 - Enter cacLogout.html against Logout Redirect URL -> Apply
 - Close browser.

Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.5 to 7.3.9)

You can add SSL certificates to enable enhanced authentication and encryption for your Data Warehouse and Cognos environment.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.5 through 7.3.9.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

About this task

You must have admin privileges to perform this procedure.

Steps

1. Create a backup of `..\SANSscreen\cognos\analytics\configuration\cogstartup.xml`.
2. Create a backup of the “certs” and “csk” folders under `..\SANSscreen\cognos\analytics\configuration`.
3. Generate a Certificate Encryption Request from Cognos. In an Admin CMD window, run:
 - a. `cd “\Program Files\sansscreen\cognos\analytics\bin”`
 - b. `ThirdPartyCertificateTool.bat -java:local -c -e -p NoPassWordSet -a RSA -d “CN=FQDN,O=orgname,C=US” -r c:\temp\encryptRequest.csr`
4. Open the `c:\temp\encryptRequest.csr` file and copy the generated content.

5. Send the encryptRequest.csr to the certificate authority (CA) to obtain an SSL certificate.

Make sure to add additional attributes such as "SAN:dns=FQDN (For example, hostname.netapp.com)" to add the SubjectAltName. Google Chrome version 58 and later complains if the SubjectAltName is missing from the certificate.

6. Download the chain certificates by including root certificate by using PKCS7 format

This will download fqdn.p7b file

7. Get a cert in .p7b format from your CA. Use a name that marks it as the certificate for the Cognos Webserver.
8. ThirdPartyCertificateTool.bat fails to import the entire chain, so multiple steps are required to export all certificates. Split the chain by exporting them individually as follows:
 - a. Open the .p7b certificate in "Crypto Shell Extensions".
 - b. Browse in the left pane to "Certificates".
 - c. Right-click on root CA > All Tasks > Export.
 - d. Select Base64 output.
 - e. Enter a file name identifying it as the root certificate.
 - f. Repeat steps 8a through 8c to export all of the certificates separately into .cer files.
 - g. Name the files intermediateX.cer and cognos.cer.
9. Ignore this step if you have only one CA certificate, otherwise merge both root.cer and intermediateX.cer into one file.
 - a. Open intermediate.cer with NotePad and copy the content.
 - b. Open root.cer with NotePad and save the content from 9a.
 - c. Save the file as CA.cer.
10. Import the certificates into the Cognos keystore using the Admin CMD prompt:
 - a. cd "Program Files\sanscreen\cognos\analytics\bin"
 - b. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\CA.cer

This will set CA.cer as root Certificate Authority.

- c. ThirdPartyCertificateTool.bat -java:local -i -e -r c:\temp\cognos.cer -t c:\temp\CA.cer

This will set Cognos.cer as encryption certificate which is signed by CA.cer.

11. Open the IBM Cognos Configuration.
 - a. Select Local Configuration → Security → Cryptography → Cognos
 - b. Change "Use third party CA?" to True.
 - c. Save the configuration.
 - d. Restart Cognos

12. Export the latest Cognos certificate into cognos.crt using the Admin CMD prompt:

- a. "D:\Program Files\SANscreen\java\bin\keytool.exe" -exportcert -file "c:\temp\cognos.crt" -keystore "D:\Program Files\SANscreen\cognos\analytics\configuration\certs\CAMKeystore" -storetype PKCS12 -storepass NoPassWordSet -alias encryption

13. Import the "c:\temp\cognos.crt" into dwh trustore to establish SSL communication between Cognos and DWH, using the Admin CMD prompt window.
 - a. "D:\Program Files\SANscreen\java\bin\keytool.exe" -importcert -file "c:\temp\cognos.crt" -keystore "D:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore" -storepass changeit -alias cognoscert
14. Restart the SANscreen service.
15. Perform a backup of DWH to make sure DWH communicates with Cognos.

Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.10 and later)

You can add SSL certificates to enable enhanced authentication and encryption for your Data Warehouse and Cognos environment.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.10 and later.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

About this task

You must have admin privileges to perform this procedure.

Steps

1. Stop Cognos using the IBM Cognos Configuration tool. Close Cognos.
2. Create backups of the ..\SANScreen\cognos\analytics\configuration and ..\SANScreen\cognos\analytics\temp\cam\freshness folders.
3. Generate a Certificate Encryption Request from Cognos. In an Admin CMD window, run:
 - a. cd "\\Program Files\sanscreen\cognos\analytics\bin"
 - b. ThirdPartyCertificateTool.bat -java:local -c -e -p NoPassWordSet -a RSA -r c:\temp\encryptRequest.csr -d "CN=server.domain.com,O=NETAPP,C=US" -H "server.domain.com" -I "ipaddress". Note: here -H and -I are to add subjectAltNames like dns and ipaddress.

4. Open the `c:\temp\encryptRequest.csr` file and copy the generated content.
5. Input the `encryptRequest.csr` content and generate certificate using CA signing portal.
6. Download the chain certificates by including root certificate by using PKCS7 format

This will download `fqdn.p7b` file

7. Get a cert in `.p7b` format from your CA. Use a name that marks it as the certificate for the Cognos Webserver.
8. `ThirdPartyCertificateTool.bat` fails to import the entire chain, so multiple steps are required to export all certificates. Split the chain by exporting them individually as follows:
 - a. Open the `.p7b` certificate in “Crypto Shell Extensions”.
 - b. Browse in the left pane to “Certificates”.
 - c. Right-click on root CA > All Tasks > Export.
 - d. Select Base64 output.
 - e. Enter a file name identifying it as the root certificate.
 - f. Repeat steps 8a through 8e to export all of the certificates separately into `.cer` files.
 - g. Name the files `intermediateX.cer` and `cognos.cer`.
9. Ignore this step if you have only one CA certificate, otherwise merge both `root.cer` and `intermediateX.cer` into one file.
 - a. Open `root.cer` with NotePad and copy the content.
 - b. Open `intermediate.cer` with NotePad and append the content from 9a (intermediate first and root next).
 - c. Save the file as `chain.cer`.
10. Import the certificates into the Cognos keystore using the Admin CMD prompt:
 - a. `cd "Program Files\sansscreen\cognos\analytics\bin"`
 - b. `ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\root.cer`
 - c. `ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\intermediate.cer`
 - d. `ThirdPartyCertificateTool.bat -java:local -i -e -r c:\temp\cognos.cer -t c:\temp\chain.cer`
11. Open the IBM Cognos Configuration.
 - a. Select Local Configuration → Security → Cryptography → Cognos
 - b. Change “Use third party CA?” to True.
 - c. Save the configuration.
 - d. Restart Cognos
12. Export the latest Cognos certificate into `cognos.crt` using the Admin CMD prompt:
 - a. `cd "C:\Program Files\SANscreen"`
 - b. `java\bin\keytool.exe -exportcert -file c:\temp\cognos.crt -keystore cognos\analytics\configuration\certs\CAMKeystore -storetype PKCS12 -storepass NoPassWordSet -alias encryption`
13. Back up the DWH server trustore


```
at.. \SANscreen\wildfly\standalone\configuration\server.trustore
```
14. Import the “`c:\temp\cognos.crt`” into DWH trustore to establish SSL communication between Cognos and DWH, using the Admin CMD prompt window.

- a. `cd "C:\Program Files\SANscreen"`
- b. `java\bin\keytool.exe -importcert -file c:\temp\cognos.crt -keystore wildfly\standalone\configuration\server.trustore -storepass changeit -alias cognos3rdca`

15. Restart the SANscreen service.

16. Perform a backup of DWH to make sure DWH communicates with Cognos.

17. The following steps should be performed even when only the “ssl certificate” is changed and the default Cognos certificates are left unchanged. Otherwise Cognos may complain about the new SANscreen certificate or be unable to create a DWH backup.

- a. `cd "%SANSSCREEN_HOME%cognos\analytics\bin\"`
- b. `"%SANSSCREEN_HOME%java64\bin\keytool.exe" -exportcert -file "c:\temp\sansscreen.cer" -keystore "%SANSSCREEN_HOME%wildfly\standalone\configuration\server.keystore" -storepass changeit -alias "ssl certificate"`
- c. `ThirdPartyCertificateTool.bat -java:local -i -T -r "c:\temp\sansscreen.cer"`

Typically, these steps are performed as part of the Cognos certificate import process described in [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Importing SSL certificates

You can add SSL certificates to enable enhanced authentication and encryption for enhancing the security of your OnCommand Insight environment.

Before you begin

You must ensure that your system meets the minimum required bit level (1024 bits).

About this task



Before you attempt to perform this procedure, you should back up the existing `server.keystore` file, and name the backup `server.keystore.old`. Corrupting or damaging the `server.keystore` file may result in an inoperable Insight server after the Insight server is restarted. If you create a backup, you can revert to the old file if problems occur.

Steps

1. Create a copy of the original keystore file: `cp c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore.old"`
2. List the contents of the keystore: `C:\Program Files\SANscreen\java64\bin\keytool.exe -list -v -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore"`

- a. When prompted for a password, enter `changeit`.

The system displays the contents of the keystore. There should be at least one certificate in the keystore,

"ssl certificate".

3. Delete the "ssl certificate": `keytool -delete -alias "ssl certificate" -keystore c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore`
4. Generate a new key: `C:\Program Files\SANscreen\java64\bin\keytool.exe -genkey -alias "ssl certificate" -keyalg RSA -keysize 2048 -validity 365 -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore"`
 - a. When prompted for first and last names, enter the fully qualified domain name (FQDN) that you intend to use.
 - b. Provide the following information about your organization and organizational structure:
 - Country: two-letter ISO abbreviation for your country (for example, US)
 - State or Province: name of the state or province where your organization's head office is located (for example, Massachusetts)
 - Locality: name of the city where your organization's head office is located (for example, Waltham)
 - Organizational name: name of the organization that owns the domain name (for example, NetApp)
 - Organizational unit name: name of the department or group that will use the certificate (for example, Support)
 - Domain Name/ Common Name: the FQDN that is used for DNS lookups of your server (for example, www.example.com) The system responds with information similar to the following: Is CN=www.example.com, OU=support, O=NetApp, L=Waltham, ST=MA, C=US correct?
 - c. Enter Yes when the Common Name (CN) is equal to the FQDN.
 - d. When prompted for the key password, enter the password, or press the Enter key to use the existing keystore password.
5. Generate a certificate request file: `C:\Program Files\SANscreen\java64\bin\keytool.exe -certreq -alias "ssl certificate" -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -file c:\localhost.csr`

The `c:\localhost.csr` file is the certificate request file that is newly generated.

6. Submit the `c:\localhost.csr` file to your certificate authority (CA) for approval.

Once the certificate request file is approved, you want the certificate returned to you in `.der` format. The file might or might not be returned as a `.der` file. The default file format is `.cer` for Microsoft CA services.

Most organizations' CAs use a chain of trust model, including a root CA, which is often offline. It has signed the certificates for only a few child CAs, known as intermediate CAs.

You must obtain the public key (certificates) for the entire chain of trust—the certificate for the CA that signed the certificate for the OnCommand Insight server, and all the certificates between that signing CA up to and including the organizational root CA.

In some organizations, when you submit a signing request, you might receive one of the following:

- A PKCS12 file that contains your signed certificate and all the public certificates in the chain of trust
- A `.zip` file that contains individual files (including your signed certificate) and all the public certificates in the chain of trust

- Only your signed certificate

You must obtain the public certificates.

7. Import the approved certificate for `server.keystore`: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -alias OCI.hostname.com -file c:\localhost2.DER -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore"`

- a. When prompted, enter the keystore password.

The following message is displayed: Certificate reply was installed in keystore

8. Import the approved certificate for `server.trustore`: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -alias OCI.hostname.com -file c:\localhost2.DER -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore"`

- a. When prompted, enter the trustore password.

The following message is displayed: Certificate reply was installed in trustore

9. Edit the `SANscreen\wildfly\standalone\configuration\standalone-full.xml` file:

Substitute the following alias string: `alias="cbc-oci-02.muccbc.hq.netapp.com"`. For example:

```
<keystore path="server.keystore" relative-to="jboss.server.config.dir"
keystore-password="{VAULT::HttpsRealm::keystore_password:1}" alias="cbc-oci-
02.muccbc.hq.netapp.com" key-
password="{VAULT::HttpsRealm::key_password:1}"/>
```

10. Restart the SANscreen server service.

Once Insight is running, you can click the padlock icon to view the certificates that are installed on the system.

If you see a certificate containing "Issued To" information that matches "Issued By" information, you still have a self-signed certificate installed. The Insight installer-generated self-signed certificates have a 100-year expiration.

NetApp cannot guarantee that this procedure will remove digital certificate warnings. NetApp cannot control how your end user workstations are configured. Consider the following scenarios:

- Microsoft Internet Explorer and Google Chrome both utilize Microsoft's native certificate functionality on Windows.

This means that if your Active Directory administrators push your organization's CA certificates into the end user's certificate trustores, the users of these browsers will see certificate warnings disappear when the OnCommand Insight self-signed certificates have been replaced with the one signed by the internal CA infrastructure.

- Java and Mozilla Firefox have their own certificate stores.

If your system administrators do not automate ingesting the CA certificates into these applications' trusted certificates stores, using the Firefox browser might continue to generate certificate warnings

because of an untrusted certificate, even when the self-signed certificate has been replaced. Getting your organization's certificate chain installed into the trustore is an additional requirement.

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.

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  **Add** 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**.

Assigning business entities to assets

You can assign a business entity to an asset (host, port, storage, switch, virtual machine, qtree, share, 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.

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 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.


Assigning business entities to or removing business entities from multiple assets

You can assign business entities to or remove business entities from multiple assets by using a query instead of having to manually assign or remove them.


Before you begin

You must have already created the business entities you want to add to your desired assets.

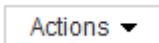
Steps

1. Create a new query, or open an existing query.
2. If desired, filter for the assets to which you want to add business entities.
3. Select the desired assets in the list or click  ▼ to select **All**.

The **Actions** button displays.

4. To add a business entity to the selected assets, click . If the selected asset type can have business entities assigned to it, you will see the menu choice to **Add Business Entity**. Select this.
5. Select the desired business entity from the list and click **Save**.

Any new business entity you assign overrides any business entities that were already assigned to the asset. Assigning applications to assets will also override the business entities assigned in the same way. Assigning business entities to as asset may also override any applications assigned to that asset.

6. To remove a business entity assigned to the assets, click  and select **Remove Business Entity**.
7. Select the desired business entity from the list and click **Delete**.

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.
- 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
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

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



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 Insightweb 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 Object Storage 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.



Annotation names beginning or ending with a dot "." are not supported.

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.




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.

c. Click  to add additional values.

d. Click  to remove a value.

7. Click **Save**.

Your annotations appear in the list on the Annotations page.

Related information

[Importing and Exporting user data](#)


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**.
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.

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 Insightweb 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.



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  **Add**.

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.

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 Insightweb 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 or down in the list, click the up arrow or the down arrow.

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 Insightweb 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 Insightweb 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 and keys are supported in the import utility:

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

Qtree	id-><id> or <Storage Name>-><Internal Volume Name>-><Qtree Name>
-------	--

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 Key 1>, <Annotation Value> [,
<Annotation Value> ...] [, <Application>] [, <Tenant>] [,
<Line_Of_Business>] [, <Business_Unit>] [, <Project>]

...

<Object Type Value N>, <Object 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 rest-import-utility.jar [-username] [-ppassword]
[-aserver name or IP address] [-bbatch size] [-ccase
sensitive:true/false]
[-lextra logging:true/false] csv filename
```

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



There are no spaces between the options and their values.



The following keywords are reserved and prevent users from specifying them as annotation names: - Application - Application_Priority - Tenant - Line_Of_Business - Business_Unit - Project Errors are generated if you attempt to import an annotation type using one of the reserved keywords. If you have created annotation names using these keywords, you must modify them so that the import utility tool can work correctly.



The Annotation Import utility requires Java 8 or Java 11. Ensure that one of those is installed prior to running the import utility. It is recommended to use the latest OpenJDK 11.

Assigning annotations to multiple assets using a query

Assigning an annotation to a group of assets helps you more easily identify or use those related assets in queries or dashboards.

Before you begin

Annotations that you wish to assign to assets must have previously been created.

About this task

You can simplify the task of assigning an annotation to multiple assets by using a query. For example, if you want to assign a custom address annotation to all of your arrays at a specific data center location.

Steps

1. Create a new query to identify the assets on which you wish to assign an annotation. Click **Queries > +New Query**.
2. In the **Search for...** drop-down, choose **Storage**. You can set filters to further narrow down the list of storages displayed.
3. In the list of storages displayed, select one or more by clicking on the check box beside the storage name. You may also select all the displayed storages by clicking on the main check box at the top of the list.
4. When you have selected all of the desired storages, click **Actions > Edit Annotation**.

The system displays the Add Annotation dialog.

5. Select the **Annotation** and **Value** you want to assign to the storages and click **Save**.

If you are displaying the column for that annotation, it will appear on all the selected storages.

6. You can now use the annotation to filter for storages in a widget or query. In a widget, you can do the following:
 - a. Create a dashboard or open an existing one. Add a **Variable** and choose the annotation you set on the storages above. The variable is added to the dashboard.
 - b. In the variable field you just added, click on **Any** and enter the appropriate Value to filter on. Click on the check mark to save the variable value.
 - c. Add a widget. In the widget's Query, click on the **Filter by+** button and select the appropriate annotation from the list.
 - d. Click on **Any** and select the annotation variable you added above. Variables you have created start with "\$" and are displayed in the drop-down.

- e. Set any other filters or fields you desire, then click **Save** when the widget is customized to your liking.
- The widget on the dashboard displays the data for only the storages to which you assigned the annotation.

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.

Assets used in queries and dashboards

Insight queries and dashboard widgets can be used with a wide range of asset types

The following asset types can be used in queries, dashboard widgets, and custom asset pages. The fields and counters available for filters, expressions, and display will vary among asset types. Not all assets can be used in all widget types.

- Application
- Datastore
- Disk
- Fabric
- Generic Device
- Host
- Internal Volume
- iSCSI Session
- iSCSI Network Portal
- Path
- Port
- Qtree
- Quota
- Share
- Storage
- Storage Node
- Storage Pool
- Switch
- Tape
- VMDK
- Virtual Machine
- Volume
- Zone

- Zone Member

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 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.
- The NOT operator allows you to exclude text from the search results. For example, `NOT EMC*` finds everything that does not start with “EMC”. You can use `NOT *` to display fields that contain no value.


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.



Depending on the number of columns displayed in the query results table, you may not be able to view additional added columns. You can remove one or more columns until your desired columns become visible.

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*.

Related information

[Importing and Exporting user data](#)

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.
 - To resize a column, hover the mouse over the column header until a blue bar appears. Place the mouse over the bar and drag it right or left.
 - To move a column, click on the column header and drag it right or left.
 - When scrolling through the query results, be aware that the results may change as Insight automatically polls your data sources. This may result in some items being missing, or some items appearing out of order depending on how they are sorted.


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. Only the attributes for the displayed columns will be exported. Some displayed columns, particularly those that are part of complex nested relationships, are not exported.



When a comma appears in an asset name, the export encloses the name in quotes, preserving the asset name and the proper .csv format.

+ When exporting query results, be aware that **all** rows in the results table will be exported, not just those selected or displayed on the screen, up to a maximum of 10,000 rows.

+

When opening an exported .CSV file with Excel, if you have an object name or other field that is in the format NN:NN (two digits followed by a colon followed by two more digits), Excel will sometimes interpret that name as a Time format, instead of Text format. This can result in Excel displaying incorrect values in those columns. For example, an object named "81:45" would show in Excel as "81:45:00". To work around this, import the .CSV into Excel using the following steps:

+



- Open a new sheet in Excel.
 - On the "Data" tab, choose "From Text".
 - Locate the desired .CSV file and click "Import".
 - In the Import wizard, choose "Delimited" and click Next.
 - Choose "Comma" for the delimiter and click Next.
 - Select the desired columns and choose "Text" for the column data format.
 - Click Finish.
- Your objects should show in Excel in the proper format.

+



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 Insightweb 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 Insightweb 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.

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 displays.

3. Select the desired assets in the list or click ☐ ▼ to select **All**.

The **Actions** button displays.

4. To add an application to the selected assets, click **Actions** ▼, 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 **Actions** ▼ 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.

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.

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 displays.

3. Select the desired assets in the list or click ☐ ▼ to select **All**.

The **Actions** button displays.

4. To add an annotation to the assets or edit the value of an annotation assigned to the assets, click

Actions ▼

, 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 **Actions** ▼, and select **Remove Annotation**.

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

Copying table values

You can copy values in tables for use in search boxes or other applications.

About this task

There are two methods you can use to copy values from tables or query results.

Steps

1. Method 1: Highlight the desired text with the mouse, copy it, and paste it into search fields or other applications.
2. Method 2: For single-value fields whose length exceeds the width of the table column, indicated by ellipses (...), hover over the field and click the clipboard icon. The value is copied to the clipboard for use in search fields or other applications.

Note that only values that are links to assets can be copied. Note also that only fields that include single values (i.e. non-lists) have the copy icon.

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

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.

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.

Location of the most up-to-date version of the Data Source Support Matrix

The OnCommand Insight Data Source Support Matrix is updated with each service pack release. The most current version of the document can be found at the [NetApp Support Site](#).

Adding data sources

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

Steps

1. Open OnCommand Insight in your browser and log in as a user with administrative permissions.
2. Select **Admin** and choose **Data sources**.
3. Click the **+Add** button.

The Add data source wizard opens.

4. In the **Settings** section, enter the following information:

Field	Description
Name	Enter a unique network name for this data source. NOTE: only letters, numbers and the underscore (_) character are allowed in the data source name.
Vendor	Choose the vendor of the data source from the drop-down.
Model	Choose the model of the data source from the drop-down.
Where to run	Choose Local, or you may choose a remote acquisition unit if RAU's are configured in your environment.
What to collect	For most data sources, these options will be Inventory and Performance. Inventory is always selected by default and cannot be un-selected. Note that some data sources may have different options. The collection options you select change the available fields in the Configuration and Advanced configuration sections.

- Click the **Configuration** link and enter the basic setup information required for the data source with your selected data collection type.
- 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.
- For details about configuration or advanced configuration information required or available for your specific data source, see the [Vendor-specific data source reference](#).
- Click the **Test** link to be certain that the data source is properly configured.
- Click **Save**.

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. On the Insight toolbar, click **Admin**.
 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 **Admin**.
2. Select **Patches**.
3. Select **Actions > Install service pack or patch**.
4. In the **Install Service Pack or Patch** dialog box, click **Browse** to locate and select the patch file that you uploaded.
5. Click **Next** in the **Patch Summary** dialog box.
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.

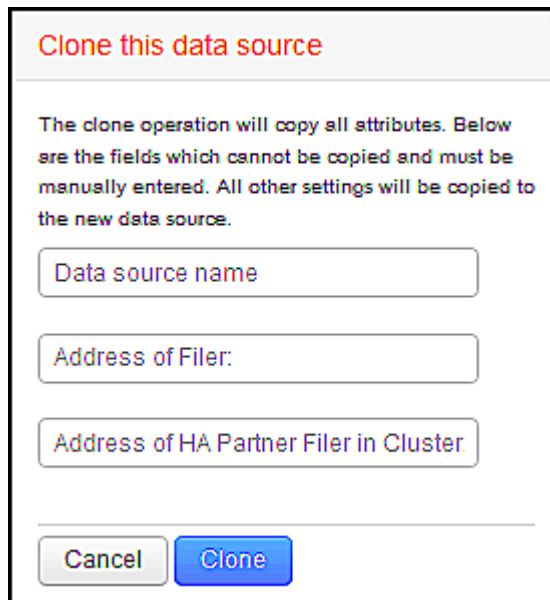
Steps

1. On the Insight toolbar, click **Admin**.

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:



Clone this data source

The clone operation will copy all attributes. Below are the fields which cannot be copied and must be manually entered. All other settings will be copied to the new data source.

Data source name

Address of Filer:

Address of HA Partner Filer in Cluster

Cancel Clone

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

Results

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

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 the **Test** button in the data source wizard, communication with the specified device is checked. The test produces one of these results:

- **PASSED:** the data source is configured correctly.

- **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. Check your configuration settings and re-test.

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.

3PAR InServ data source

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

Terminology

OnCommand Insight acquires the following inventory information from the 3PAR InServ data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Physical Disk	Disk
Storage System	Storage
Controller Node	Storage Node
Common Provisioning Group	Storage Pool
Virtual Volume	Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- IP address or FQDN of the InServ cluster
- For inventory, read-only user name and password to the InServ Server.
- For performance, read-write user name and password to the InServ Server.
- Port requirements: 22 (inventory collection), 5988 or 5989 (performance collection) [Note: 3PAR Performance is supported for InServ OS 3.x+]
- For performance collection confirm that SMI-S is enabled by logging into the 3PAR array via SSH.

Configuration

Field	Description
Cluster IP	IP address or fully-qualified domain name of the InServ cluster
User Name	User name for the InServ Server
Password	Password used for the InServ Server
SMI-S Host IP	IP address of the SMI-S Provider Host
SMI-S User Name	User name for the SMI-S Provider Host
SMI-S Password	Password used for the SMI-S Provider Host

Advanced Configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Exclude Devices	Comma-separated list of device IPs to exclude
SSH Process Wait Timeout (sec)	SSH process timeout (default 60 seconds)
Number of SSH Retries	Number of SSH retry attempts
SSH Banner Wait Timeout (sec)	SSH banner wait timeout (default 20 seconds)
SMI-S Port	Port used by SMI-S Provider Host
Protocol	Protocol used to connect to the SMI-S provider
SMI-S namespace	SMI-S namespace
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Number of SMI-S Connection Retries	Number of SMI-S connection retry attempts

Amazon AWS EC2 data source

OnCommand Insight uses this data source to discover inventory and performance for Amazon AWS EC2.

Pre-requisites:

In order to collect data from Amazon EC2 devices, you must have the following information:

- You must have the IAM Access Key ID
- You must have the Secret Access Key for your Amazon EC2 cloud account
- You must have the "list organization" privilege
- Port 433 HTTPS
- EC2 Instances can be reported as a Virtual Machine, or (less naturally) a Host. EBS Volumes can be reported as both a VirtualDisk used by the VM, as well as a DataStore providing the Capacity for the VirtualDisk.

Access keys consist of an access key ID (for example, AKIAIOSFODNN7EXAMPLE) and a secret access key (for example, wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY). You use access keys to sign programmatic requests that you make to EC@ if you use the Amazon EC2 SDKs, REST, or Query API operations. These keys are provided with your contract from Amazon.

How to configure this data source

To configure the Amazon AWS EC2 data source, you will need the AWS IAM Access Key ID and Secret Access Key for your AWS account.

Fill in the data source fields according to the tables below:

Configuration:

Field	Description
AWS Region	Choose AWS region
IAM Role	For use only when acquired on an AU in AWS. See below for more information on IAM Roles.
AWS IAM Access Key ID	Enter AWS IAM Access Key ID. Required if you do not use IAM Role.
AWS IAM Secret Access Key	Enter AWS IAM Secret Access Key. Required if you do not use IAM Role.
I understand AWS will bill me for API requests	Check this to verify your understanding that AWS bills you for API requests made by Insight polling

Advanced Configuration:

Field	Description
Include Extra Regions	Specify additional regions to include in polling.

Cross Account Role	Role for accessing resources in different AWS accounts.
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
HTTP connection and socket timeout (sec)	HTTP connection timeout (default 300 seconds)
Include AWS tags	Check this to enable support for AWS tags in Insight annotations
Performance Poll Interval (sec)	Interval between performance polls (default 1800 seconds)

Mapping AWS tags to Insight annotations

The AWS EC2 data source includes an option that allows you to populate Insight annotations with tags configured on AWS. The annotations must be named exactly as the AWS tags. Insight will always populate same-named text-type annotations, and will make a "best attempt" to populate annotations of other types (number, boolean, etc). If your annotation is of a different type and the data source fails to populate it, it may be necessary to remove the annotation and re-create it as a text type.

Note that AWS is case-sensitive, while Insight is case-insensitive. So if you create an annotation named "OWNER" in Insight, and tags named "OWNER", "Owner", and "owner" in AWS, all of the AWS variations of "owner" will map to Insight's "OWNER" annotation.

Related Information:

[Managing Access Keys for IAM Users](#)

Include Extra Regions

In the AWS Data Collector **Advanced Configuration** section, you can set the **Include extra regions** field to include additional regions, separated by comma or semi-colon. By default, this field is set to **us-.***, which collects on all US AWS regions. To collect on *all* regions, set this field to **.***.

If the **Include extra regions** field is empty, the data collector will collect on assets specified in the **AWS Region** field as specified in the **Configuration** section.

Collecting from AWS Child Accounts

Insight supports collection of child accounts for AWS within a single AWS data collector. Configuration for this collection is performed in the AWS environment:

- You must configure each child account to have an AWS Role that allows the primary account ID to access EC2 details from the children account.
- Each child account must have the role name configured as the same string
- Enter this role name string into the Insight AWS Data Collector **Advanced Configuration** section, in the **Cross Account Role** field.

Best Practice: It is highly recommended to assign the AWS predefined AmazonEC2ReadOnlyAccess policy to the ECS primary account. Also, the user configured in the data source should have at least the predefined

AWSOrganizationsReadOnlyAccess policy assigned, in order to query AWS.

Please see the following for information on configuring your environment to allow Insight to collect from AWS child accounts:

[Tutorial: Delegate Access Across AWS Accounts Using IAM Roles](#)

[AWS Setup: Providing Access to an IAM User in Another AWS Account That You Own](#)

[Creating a Role to Delegate Permissions to an IAM User](#)

IAM Roles

When using *IAM Role* security, you must ensure that the role you create or specify has the appropriate permissions needed to access your resources.

For example, if you create an IAM role named *InstanceEc2ReadOnly*, you must set up the policy to grant EC2 read-only list access permission to all EC2 resources for this IAM role. Additionally, you must grant STS (Security Token Service) access so that this role is allowed to assume roles cross accounts.

After you create an IAM role, you can attach it when you create a new EC2 instance or any existing EC2 instance.

After you attach the IAM role *InstanceEc2ReadOnly* to an EC2 instance, you will be able to retrieve the temporary credential through instance metadata by IAM role name and use it to access AWS resources by any application running on this EC2 instance.



IAM role can be used only when the Acquisition Unit is running in an AWS instance.

Brocade Enterprise Fabric Connectivity Manager data source

OnCommand Insight uses the Brocade Enterprise Fabric Connectivity Manager (EFCM) data source to discover inventory for Brocade EFCM switches. Insight supports EFCM versions 9.5, 9.6, and 9.7.

Requirements



This data collector is not available starting with OnCommand Insight 7.3.11.

- Network address or fully-qualified domain name for the EFCM server
- EFCM version must be 9.5, 9.6, or 9.7
- IP address of the EFCM server
- Read-only username and password for the EFCM server
- Validated access to the Connectrix switch by Telnet from the Insight server, using the read-only username and password over port 51512

Configuration

Field	Description
-------	-------------

EFC server	IP address or fully-qualified domain name of the EFC Server
User Name	User name for the switch
Password	Password used for the switch

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 15 minutes)
Fabric Name	Fabric name to be reported by the EFCM data source. Leave blank to report the fabric name as WWN.
Communication Port	Port used for communication with the switch
Enable Trapping	Select to enable acquisition upon receiving an SNMP trap from the device. If you select enable trapping, you must also activate SNMP.
Minimum Time Between Traps (sec)	Minimum time between acquisition attempts triggered by traps (default 15 seconds)
Inactive Zonesets	Comma-separated list of inactive Zonesets on which to perform acquisition, in addition to performing acquisition on the active zone sets
NIC to Use	Specify which network interface should be used on the RAU when reporting on SAN devices
Exclude Devices	Comma-separated list of unit names to include or exclude from polling
Use the EFCM switch nickname as the Insight switch name	Select to use the EFCM switch nickname as the Insight switch name
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Brocade FC Switch data source

OnCommand Insight uses the Brocade FC Switch (SSH) data source to discover inventory for Brocade or rebranded switch devices running Factored Operating System (FOS) firmware 4.2 and later. Devices in both FC switch and Access Gateway modes are supported.

Terminology

OnCommand Insight acquires the following inventory information from the Brocade FC Switch data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Switch	Switch
Port	Port
Virtual Fabric, Physical Fabric	Fabric
Zone	Zone
Logical Switch	Logical Switch
LSAN Zone	IVR Zone



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- The Acquisition Unit (local or remote) will initiate connections to TCP Port 22 on Brocade switches to collect inventory data. The AU will also initiate connections to UDP port 161 for collection of performance data.
- 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.
- SNMP read-only Community String

Configuration

Field	Description
Switch IP	IP address or fully-qualified domain name of the switch
User Name	User name for the switch
Password	Password used for the switch

SNMP Version	SNMP version
SNMP Community String	SNMP read-only community string used to access the switch
SNMP User Name	SNMP version protocol user name (applies only to SNMP v3)
SNMP Password	SNMP version protocol password (applies only to SNMP v3)

Advanced configuration

Field	Description
Fabric Name	Fabric name to be reported by the data source. Leave blank to report the fabric name as WWN.
Exclude Devices	Comma-separated list of device IDs to exclude from polling
Inventory Poll Interval (min)	Interval between inventory polls (default 15 minutes)
Timeout (sec)	Connection timeout (default 30 seconds)
Banner Wait Timeout (sec)	SSH banner wait timeout (default 5 seconds)
Admin Domains Active	Select if using Admin Domains
Retrieve MPR Data	Select to acquire routing data from your multiprotocol router (MPR)
Enable Trapping	Select to enable acquisition upon receiving an SNMP trap from the device. If you select enable trapping, you must also activate SNMP.
Minimum Time Between Traps (sec)	Minimum time between acquisition attempts triggered by traps (default 10 seconds)
Discover all switches in the fabric	Select to discover all switches in the fabric
Choose Favoring HBA vs. Zone Aliases	Choose whether to favor HBA or zone aliases
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

SNMP Auth Protocol	SNMP authentication protocol (SNMP v3 only)
SNMP Privacy Protocol	SNMP privacy protocol (SNMP v3 only)
SNMP Privacy Password	SNMP privacy password (SNMP v3 only)
SNMP Retries	Number of SNMP retry attempts
SNMP Timeout (ms)	SNMP timeout (default 5000 ms)

Brocade Sphereon/Intrepid Switch data source

OnCommand Insight uses the Brocade Sphereon/Intrepid Switch (SNMP) data source to discover inventory for Brocade Sphereon or Intrepid switches.

Requirements



This data collector not available starting with OnCommand Insight 7.3.11.

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

Configuration

Field	Description
Sphereon Switch	IP address or fully-qualified domain name of the switch
SNMP Version	SNMP version
SNMP Community	SNMP read-only community string used to access the switch
User Name	SMI-S user name for the switch (SNMP v3 only)
Password	SMI-S password for the switch (SNMP v3 only)

Advanced configuration

Field	Description
-------	-------------

Inventory Poll Interval (min)	Interval between inventory polls (default 15 minutes)
SNMP Auth Protocol	SNMP authentication protocol (SNMPv3 only)
SNMP Privacy Protocol	SNMP privacy protocol (SNMPv3 only)
SNMP Privacy Password	SNMP privacy password
SNMP Number of Retries	Number of SNMP retry attempts
SNMP Timeout (ms)	SNMP timeout (default 5000 ms)
Fabric Name	Fabric name to be reported by the data source. Leave blank to report the fabric name as WWN.
Enable Trapping	Select to enable acquisition upon receiving an SNMP trap from the device. If you select enable trapping, you must also activate SNMP.
Minimum Time Between Ttraps (seconds)	Minimum time between acquisition attempts triggered by traps (default 10 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Cisco FC Switch Firmware (SNMP) data source

OnCommand Insight uses the Cisco FC Switch Firmware 2.0+ (SNMP) data source to discover inventory for Cisco MDS Fibre Channel switches as well as 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.

Terminology

OnCommand Insight acquires the following inventory information from the Cisco FC Switch data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Switch	Switch
Port	Port
VSAN	Fabric

Zone	Zone
Logical Switch	Logical Switch
Name Server Entry	Name Server Entry
Inter-VSAN Routing (IVR) Zone	IVR Zone



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- 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
- Port 161 is used to access the device
- Access validation using the `snmpwalk` utility to the switch (see `<install_path>\bin\`)

Configuration

Field	Description
Cisco Switch IP	IP address or fully-qualified domain name of the switch
SNMP Version	SNMP version v2 or later is required for performance acquisition
SNMP Community String	SNMP read-only community string used to access the switch (not applicable for SNMP v3)
User Name	User name for the switch (SNMP v3 only)
Password	Password used for the switch (SNMPv3 only)

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
SNMP Auth Protocol	SNMP authentication protocol (SNMPv3 only)
SNMP Privacy Protocol	SNMP privacy protocol (SNMPv3 only)

SNMP Privacy Password	SNMP privacy password
SNMP Retries	Number of SNMP retry attempts
SNMP Timeout (ms)	SNMP timeout (default 5000 ms)
Enable Trapping	Select to enable trapping. If you enable trapping, you must also activate SNMP notifications.
Minimum Time Between Traps (sec)	Minimum time between acquisition attempts triggered by traps (default 10 seconds)
Discover All Fabric Switches	Select to discover all switches in the fabric
Exclude Devices	Comma-separated list of device IPs to exclude from polling
Include Devices	Comma-separated list of device IPs to include in polling
Check Device Type	Select to accept only those devices that explicitly advertise themselves as Cisco devices

Primary Alias Type	<p>Provide a first preference for resolution of the alias. Choose from the following:</p> <ul style="list-style-type: none"> • Device Alias <p>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).</p> <ul style="list-style-type: none"> • None <p>Do not report any alias</p> <ul style="list-style-type: none"> • Port Description <p>A description to help identify the port in a list of ports</p> <ul style="list-style-type: none"> • Zone Alias (all) <p>A user-friendly name for a port that can be used only for zoning configuration</p> <ul style="list-style-type: none"> • Zone Alias (only active) <p>A user-friendly name for a port that can be used only for the active configuration. This is the default.</p>
Secondary Alias Type	Provide a second preference for resolution of the alias
Tertiary Alias Type	Provide a third preference for resolution of the alias
Enable SANTap Proxy Mode Support	Select 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)	Interval between performance polls (default 300 seconds)

EMC Celerra data source

The Celerra (SSH) 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.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Celerra data source. For each

asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Celerra Network Server	Storage
Celerra Meta Volume / Celerra Storage Pool	Storage Pool
File System	Internal Volume
Data Mover	Controller
File System mounted on a data Mover	File Share
CIFS and NFS Exports	Share
Disk Volume	Backend LUN



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- The IP address of the storage processor
- Read-only user name and password
- SSH port 22

Configuration

Field	Description
Address of Celerra	IP address or fully-qualified domain name of the Celerra device
User Name	Name used to log in to the Celerra device
Password	Password used to log in to the Celerra device

Advanced configuration

Field	Description
Inventory Poll Interval (minutes)	Interval between inventory polls (default 20 minutes)
SSH Process Wait Timeout (sec)	SSH process timeout (default 600 seconds)

Number of Retries	Number of inventory retry attempts
SSH Banner Wait Timeout (sec)	SSH banner wait timeout (default 20 seconds)

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
```

Terminology

OnCommand Insight acquires the following inventory information from the EMC CLARiiON data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Storage	Storage
Storage Processor	Storage Node
Thin Pool, RAID Group	Storage Pool
LUN	Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- An IP address of each CLARiiON storage processor
- Read-only Navisphere username and password to the CLARiiON arrays
- NaviCLI must be installed on the Insight server/RAU
- 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

- For performance, statistics logging must be turned on.
- Port requirements: 80, 443

Configuration

Field	Description
CLARiiON storage	IP address or fully-qualified domain name of the CLARiiON Storage
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 <code>navicli.exe</code> OR <code>naviseccli.exe</code> executable

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Use Secure Client (naviseccli)	Select to use secure client (naviseccli)
Scope	The secure client scope. The default is Global.
CLARiiON CLI Port	Port used for CLARiiON CLI
Inventory External Process Timeout (sec)	External process timeout (default 1800 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Performance External process timeout (sec)	External process timeout (default 1800 seconds)

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.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Data Domain data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing

or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Array	Storage
Port	Port
Filesys	Internal Volume
Mtree	QTree
Quota	Quota
NFS and CIFS share	FileShare



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- IP address of the Data Domain device
- Read-only user name and password to the Data Domain storage
- SSH port 22

Configuration

Field	Description
IP address	The IP address or fully-qualified domain name of the Data Domain storage array
User name	The user name for the Data Domain storage array
Password	The password for the Data Domain storage array

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
SSH Process Wait Timeout (sec)	SSH process timeout (default 180 seconds)
SSH Port	SSH service port

EMC ECC StorageScope data source

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

Configuration



This data collector is no longer available starting with OnCommand Insight 7.3.11.

Field	Description
ECC server	IP address or fully-qualified domain name of the ECC Server
User Name	User name for the ECC server
Password	Password for the ECC server

Advanced configuration

Field	Description
ECC Port	Port used for the ECC server
Inventory Poll Interval (min)	Interval between inventory polls (default 30 minutes)
Protocol to Connect to Database	Protocol Used to Connect to the Database
Query File System Information	Select to retrieve details for WWN Aliases and File Systems.

Dell EMC ECS data source

This data collector acquires inventory and performance data from EMC ECS storage systems. For configuration, the data collector requires an IP address of the ECS server and an administrative level domain account..

Terminology

OnCommand Insight acquires the following inventory information from the EMC ECS data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Cluser	Storage
Tenant	Storage Pool

Bucket	Internal Volume
Disk	Disk



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- An IP address of the ECS Management Console
- Administrative level domain account for the ECS system
- Port 443 (HTTPS). Requires outbound connectivity to TCP port 443 on the ECS system.
- For performance, read-only username and password for ssh/scp access.
- For performance, port 22 is required.

Configuration

Field	Description
ECS Host	IP addresses or fully-qualified domain names of the ECS system
ECS Host Port	Port used for communication with ECS Host
ECS Vendor ID	Vendor ID for ECS
Password	Password used for ECS

Advanced configuration

Field	Description
Inventory Poll Interval (minutes)	Interval between inventory polls. The default is 360 minutes.

EMC Isilon data source

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

Terminology

OnCommand Insight acquires the following inventory information from the EMC Isilon data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive	Disk
Cluster	Storage
Node	Storage Node
File System	Internal Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

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

Configuration

Field	Description
IP address	The IP address or fully-qualified domain name of the Isilon cluster
User name	The user name for the Isilon cluster
Password	The password for the Isilon cluster

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
SSH Process Wait Timeout	SSH process timeout (default 60 seconds)
SSH Port	SSH service port

Running CLI Commands

Starting with OnCommand Insight version 7.3.11 and Service Pack 9, the EMC Isilon data source contains an enhancement that will result in Insight running more CLI commands. If you are using a non-root user within your data source, you will likely have configured a "sudoers" file to grant that user account the ability to run specific CLI commands via SSH.

In order for Insight to understand EMC's "Access Zones" feature, Insight will now additionally run the following new CLI commands:

- `sudo isi zone zones list --format json -verbose`
- `sudo isi zone zones list`

Insight parses the output of these commands, and runs more instances of existing commands, to obtain the logical configuration of objects like qtrees, quotas and NAS shares/exports that reside in non-default Access Zones. Insight now reports those objects for non-default Access Zones as the result of this enhancement. As Insight obtains that data by running existing commands (with different options) no sudoers file change is required in order for those to work; it is only with the introduction of the new commands above that the change is required.

Please update your sudoers file to allow your Insight service account to run those commands before upgrading to this Insight release. Failure to do so will result in your Isilon data sources failing.

"File System" statistics

Beginning with OnCommand Insight 7.3.12, the EMC Isilon data collector introduces "File System" statistics on the node object for EMC Isilon. The existing node statistics reported by OnCommand Insight are "disk" based - i.e, for IOPs and throughput of a storage node, what are the disks in this node doing in aggregate? But for workloads where reads are cached in memory and/or compression is in use, the file system workload may be substantially higher than what actually hits the disks - a data set that compresses 5:1 could therefore have a "File System Read Throughput" value 5x the storage node Read Throughput, as the latter measures the reads off of disk, which expand 5x when the node uncompresses the data to service the client's read request.

Dell EMC PowerStore data source

The Dell EMC PowerStore data collector gathers inventory information from Dell EMC PowerStore storage. For configuration, the data collector requires the IP address of the storage processors and a read-only user name and password.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Data Domain data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
host	host
host_volume_mapping	host_volume_mapping
hardware (it has Drives under "extra_details" object): Drives	Disk
Appliance	StoragePool
Cluster	Storage Array

Node	StorageNode
fc_port	Port
volume	Volume
InternalVolume	file_system
Filesys	Internal Volume
Mtree	QTree
Quota	Quota
NFS and CIFS share	FileShare



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- IP address or fully-qualified domain name of storage processor
- Read-only user name and password

Parent Serial Number explained

Traditionally Insight is capable of reporting the storage array serial number, or the individual storage node serial numbers. However, some storage array architectures do not cleanly align to this. A PowerStore cluster can be comprised of 1-4 appliances, and each appliance has 2 nodes. If the appliance itself has a serial number, that serial number is neither the serial number for the cluster nor the nodes.

The attribute "Parent Serial Number" on the storage node object is populated appropriately for Dell/EMC PowerStore arrays when the individual nodes sit inside an intermediate appliance/enclosure that is just part of a larger cluster.

Configuration

Field	Description
PowerStore gateway(s)	IP addresses or fully-qualified domain names of PowerStore storage
User Name	User name for PowerStore
Password	Password used for PowerStore

Advanced configuration

Field	Description
HTTPS Port	Default is 443
Inventory Poll Interval (minutes)	Interval between inventory polls. The default is 60 minutes.

OnCommand Insight's PowerStore performance collection makes use of PowerStore's 5-minute granularity source data. As such, Insight polls for that data every five minutes, and this is not configurable.

EMC RecoverPoint data source

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

The EMC RecoverPoint data source gathers the volume-to-volume replication relationships that RecoverPoint coordinates across other storage arrays. OnCommand Insight shows a storage array for each RecoverPoint cluster, and collects inventory data for nodes and storage ports on that cluster. No storage pool or volume data is collected.

Requirements

- IP address or fully-qualified domain name of storage processor
- Read-only user name and password
- REST API access via port 443
- SSH access via PuTTY

Configuration

Field	Description
Address of RecoverPoint	IP address or fully-qualified domain name of RecoverPoint cluster
User Name	User name for the RecoverPoint cluster
Password	Password for the RecoverPoint cluster

Advanced configuration

Field	Description
TCP Port	TCP Port used to connect to Recoverpoint cluster
Inventory Poll Interval (minutes)	Interval between inventory polls (default 20 minutes)

Excluded Clusters	Comma-separated list of cluster IDs or names to exclude when polling
-------------------	--

EMC Solutions Enabler 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.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Solutions Enabler data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Disk Group	Disk Group
Storage Array	Storage
Director	Storage Node
Device Pool, Storage Resource Pool (SRP)	Storage Pool
Device, TDev	Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

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.

- 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.

- Administrator permissions on the Solutions Enabler (SE) Server
- 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>stordaeomon start storsrvd`
- The UNISPHERE for VMAX application must be running and collecting statistics for the Symmetrix VMAX storage arrays that are managed by the SMI-S Provider installation
- Access validation: Verify that the SMI-S provider is running: `telnet <se_server> 5988`

Configuration



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 the array.


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	Service name as specified in netcnfg file
Full path to CLI	Full path to the Symmetrix CLI

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data
Inventory Exclude Devices	Comma-separated list of device IDs to include or exclude

Connection Caching	<p>Choose connection caching method:</p> <ul style="list-style-type: none"> • 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. • In the event that REMOTE_CACHED options make CLI commands fail, use the REMOTE option. 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. <div>  <p>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.</p> </div>
CLI Timeout (sec)	CLI process timeout (default 7200 seconds)
SMI-S Host IP	IP address of the SMI-S Provider Host
SMI-S Port	Port used by SMI-S Provider Host
Protocol	Protocol used to connect to the SMI-S provider
SMI-S Namespace	Interoperability namespace that the SMI-S provider is configured to use
SMI-S User Name	User name for the SMI-S Provider Host
SMI-S Password	User name for the SMI-S Provider Host

Performance Polling Interval (sec)	Interval between performance polls (default 1000 seconds)
Performance Filter Type	Specify whether to include or exclude the array list below when collecting performance data
Performance Filter Device List	Comma-separated list of device IDs to include or exclude
RPO Polling Interval (sec)	Interval between RPO polls (default 300 seconds)

EMC VNX data source

For configuration, the EMC VNX (SSH) data source requires the IP address of the Control Station and a *read-only* username and password.

Configuration

Field	Description
VNX IP	IP address or fully-qualified domain name of the VNX Control Station
VNX User Name	User name for the VNX Control Station
VNX Password	Password for the VNX Control Station

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 between inventory polls (default 40 minutes)
VNX SSH Process Wait Timeout (sec)	VNX SSH process timeout (default 600 seconds)
Celerra Command Retry Attempts	Number of Celerra command retry attempts
CLARiiON External Process Timeout for Inventory (sec)	CLARiiON external process timeout for inventory (default 1800 seconds)

Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
CLARiiON External Process Timeout for Performance (sec)	CLARiiON external process timeout for performance (default 1800 seconds)

EMC VNXe data source

The EMC VNXe data source provides inventory support for EMC VNXe and Unity unified storage arrays.

This data source is CLI-based and requires that you install the Unisphere for VNXe CLI (uemcli.exe) on the acquisition unit that the VNXe data source resides on. uemcli.exe uses HTTPS as the transport protocol, so the acquisition unit must be able to initiate HTTPS connections to the VNXe/Unity arrays. You must have at least a read-only user for use by the data source.

Terminology

OnCommand Insight acquires the following inventory information from the EMC VNXe data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Storage Array	Storage
Processor	Storage Node
Storage Pool	Storage Pool
General iSCSI Block info, VMWare VMFS	Volume
Shared Folder	Internal Volume
CIFS Share, NFS Share, Share from VMWare NFS datastore	Share
Replication Remote System	Synchronization
iSCSI Node	iSCSI Target Node
iSCSI Initiator	iSCSI Target Initiator



These are common terminology mappings only and might not represent every case for this data source.

Requirements

The following are requirements to configure and use this data source:

- The VNXe data collector is CLI based; you must install the Unisphere for VNXe CLI, (uemcli.exe) onto the acquisition unit where your VNXe data collector resides.
- uemcli.exe uses HTTPS as the transport protocol, so the acquisition unit will need to be able to initiate HTTPS connections to the VNXe.
- You must have at least a read-only user for use by the data source.
- IP address of the managing Solutions enabler server.
- HTTPS on Port 443 is required
- The EMC VNXe data collector provides NAS and iSCSI support for inventory; fibre channel volumes will be discovered, but Insight does not report on FC mapping, masking, or storage ports.

Configuration

Field	Description
VNXe Storage	IP address or fully-qualified domain name of the VNXe device
User Name	User name for the VNXe device
Password	Password for the VNXe device
Full path to the uemcli executable	Full path to the uemcli.exe executable

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
VNXe CLI Port	Port used for the VNXe CLI
Inventory External Process Timeout (sec)	External process timeout (default 1800 seconds)

EMC VPLEX data source

For configuration, this data source requires an IP address of the VPLEX server and an administrative level domain account.

Terminology

OnCommand Insight acquires the following inventory information from the EMC VPLEX data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Cluster	Storage
Engine	Storage Node
Device, System Extend	Backend Storage Pool
Virtual Volume	Volume
Front-End Port, Back-End Port	Port
Distributed Device	Storage Synchronization
Storage View	Volume Map, Volume Mask
Storage Volume	Backend LUN
ITLs	Backend Path



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- An IP address of the VPLEX server
- Administrative level domain account for the VPLEX server
- Port 443 (HTTPS). Requires outbound connectivity to TCP port 443 on the VPLEX management station.
- For performance, read-only username and password for ssh/scp access.
- For performance, port 22 is required.
- Validate access: Verify by using `telnet` to port 443. For a port other than the default port, with any browser use `HTTPS://<ip>:<port>`

Configuration

Field	Description
IP address of VPLEX Management Console	IP address or fully-qualified domain name of the VPLEX Management Console
User Name	User name for VPLEX CLI
Password	Password used for VPLEX CLI

Performance Remote IP Address of VPLEX Management Console	Performance Remote IP address of the VPLEX Management Console
Performance Remote User Name	Performance Remote user name of VPLEX Management Console
Performance Remote Password	Performance Remote Password of VPLEX Management Console

Advanced configuration

Field	Description
Communication Port	Port used for VPLEX CLI
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection timeout (sec)	Connection timeout (default 60 seconds)
Number of Retries	Number of inventory retry attempts
Performance Poll Interval (sec)	Interval between performance polls (default 600 seconds)
Performance SSH Process Wait Timeout (sec)	SSH process timeout (default 600 seconds)
SSH Banner Wait Timeout (sec)	SSH banner wait timeout (default 20 seconds)
Number of Retries	Number of performance retry attempts

EMC XtremIO data source

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

Terminology

OnCommand Insight acquires the following inventory information from the EMC XtremIO data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk (SSD)	Disk
Cluster	Storage

Controller	Storage Node
Volume	Volume
LUN Map	Volume Map
Initiator, Target	Volume Mask



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- An IP address of each XtremIO Management Server
- An account with Administrator privileges
- Access to port 443 (HTTPS)

Configuration

Field	Description
XMS Host	IP address or fully-qualified domain name of the XtremIO Management Server
User name	User name for the XtremIO Management Server
Password	Password for the XtremIO Management Server

Advanced configuration

Field	Description
TCP port	TCP Port used to connect to XTremIO Management Server (default 443)
Inventory poll interval (min)	Interval between inventory polls (default 60 minutes)
Connection timeout (sec)	Connection timeout (default 60 seconds)
Performance poll interval(sec)	Interval between performance polls (default 300 seconds)

Fujitsu Eternus data source

The Fujitsu Eternus data source requires the IP address of the storage. It cannot be comma delimited.

Terminology

OnCommand Insight acquires the following inventory information from the Fujitsu Eternus data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Storage	Storage
Thin Pool, Flexible Tier Pool, Raid Group	Storage Pool
Standard Volume, Snap Data Volume (SDV), Snap Data Pool Volume (SDPV), Thin Provisioning Volume (TPV)	Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- An IP address of the Eternus storage, which cannot be comma delimited
- SSH Administration-level user name and password
- Port 22
- Ensure that the page scroll is disabled. (clienv-show-more-scroll disable)

Configuration

Field	Description
IP Address of Eternus Storage	IP address of the Eternus storage
User Name	User name for Eternus storage
Password	Password used for the sternus

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)

SSH Process Wait Timeout (sec)	SSH process timeout (default 600 seconds)
--------------------------------	---

Hitachi Content Platform (HCP) data source

This data collector supports the Hitachi Content Platform (HCP) using the HCP Management API.

Terminology

OnCommand Insight acquires the following inventory information from the HCP data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
HCP Cluster	Storage
Tenant	Storage Pool
Namespace	Internal Volume
Node	Node



These are common terminology mappings only and might not represent every case for this data source.

Inventory Requirements

- IP address of the HCP server
- Read-only user name and password for the HCP software and peer privileges

Configuration

Field	Description
HCP Host	IP address or fully-qualified domain name of the HCP host
HCP Port	Default is 9090
HCP user ID	User name for the HCP host
HCP Password	Password used for the HCP host
HCP Authentication Type	Choose HCP_LOCAL or ACTIVE_DIRECTORY

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Performance Polling Interval (sec)	Interval between performance polls (default 900 seconds)

HDS HiCommand Device Manager data source

The HDS HiCommand and HiCommand Lite data sources support the HiCommand Device Manager server. OnCommand Insight communicates with the HiCommand Device Manager server using the standard HiCommand API.

Terminology

OnCommand Insight acquires the following inventory information from the HDS HiCommand and HiCommand Lite data sources. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
PDEV	Disk
Journal Pool	Disk Group
Storage Array	Storage
Port Controller	Storage Node
Array Group, DP Pool	Storage Pool
Logical Unit, LDEV	Volume



These are common terminology mappings only and might not represent every case for this data source.

Inventory Requirements

- IP address of the HiCommand Device Manager server
- Read-only user name and password for the HiCommand Device Manager software and peer privileges
- Port requirements: 2001 (http) or 2443 (https)
- 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`

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`.

Configuration

Field	Description
HiCommand Server	IP address or fully-qualified domain name of the HiCommand Device Manager server
User Name	User name for the HiCommand Device Manager server.
Password	Password used for the HiCommand Device Manager server.
Devices - VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages	<p>Device list for VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages. Each storage requires:</p> <ul style="list-style-type: none">• Array's IP: IP address of the storage• User Name: User name for the storage• Password: Password for the storage• Folder Containing Export Utility JAR Files: The folder containing the Export utility <code>.jar</code> files

SNM2Devices - WMS/SMS/AMS Storages	Device list for WMS/SMS/AMS storages. Each storage requires: <ul style="list-style-type: none"> • Array's IP: IP address of the storage • Storage Navigator CLI Path: SNM2 CLI path • Account Authentication Valid: Select to choose valid account authentication • User Name: User name for the storage • Password: Password for the storage
Choose Tuning Manager for Performance	Choose Tuning Manager for performance and override other performance options
Tuning Manager Host	IP address or fully-qualified domain name of tuning manager
Tuning Manager Port	Port used for Tuning Manager
Tuning Manager Username	User name for Tuning Manager
Tuning Manager Password	password for Tuning Manager



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

Advanced configuration

Field	Description
HiCommand Server Port	Port used for the HiCommand Device Manager
HTTPs Enabled	Select to enable HTTPs
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data
Exclude or Include Devices	Comma-separated list of device ID's or array names to include or exclude
Query Host Manager	Select to query host manager
HTTP Timeout (sec)	HTTP connection timeout (default 60 seconds)

Performance Polling Interval (sec)	Interval between performance polls (default 300 seconds)
Export timeout in seconds	Export utility timeout (default 300 seconds)

Hitachi Ops Center data collector

This data collector uses Hitachi Ops Center's integrated suite of applications to access inventory and performance data of multiple storage devices. For inventory and capacity discovery, your Ops Center installation must include both the "Common Services" and "Administrator" components. For performance collection, you must additionally have "Analyzer" deployed.

Terminology

OnCommand Insight acquires the following inventory information from this data collector. For each asset type acquired, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

Vendor/Model Term	OnCommand Insight Term
Storage Systems	Storage
Volume	Volume
Parity Groups	Storage Pool(RAID), Disk Groups
Disk	Disk
Storage Pool	Storage Pool(Thin, SNAP)
External Parity Groups	Storage Pool(Backend), Disk Groups
Port	Storage Node → Controller Node → Port
Host Groups	Volume Mapping and Masking
Volume Pairs	Storage Synchronization

Note: These are common terminology mappings only and might not represent every case for this data collector.

Inventory Requirements

You must have the following in order to collect inventory data:

- IP address or hostname of the Ops Center server hosting the "Common Services" component
- Root/sysadmin user account and password that exist on all servers hosting Ops Center components. HDS did not implement REST API support for usage by LDAP/SSO users until Ops Center 10.8+

Performance requirements

The following requirements must be met in order to collect performance data:

- The HDS Ops Center "Analyzer" module must be installed
- Storage arrays must be feeding the Ops Center "Analyzer" module

Configuration

Field	Description
Hitachi Ops Center IP Address	IP address or fully-qualified domain name of the Ops Center server hosting the "Common Services" component
User Name	User name for the Ops Center server.
Password	Password used for the Ops Center server.

Advanced configuration

Field	Description
Connection Type	HTTPS (port 443) is the default
Override TCP Port	Specify the port to use if not the default
Inventory Poll Interval (min)	Interval between inventory polls. The default is 40.
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data.
Filter device List	Comma-separated list of device serial numbers to include or exclude
Performance Poll Interval (sec)	Interval between performance polls. The default is 300.

HDS Storage

Terms applying to objects or references that you might find on HDS storage asset landing pages.

HDS Storage Terminology

The following terms apply to objects or references that you might find on HDS storage asset landing pages. Many of these terms apply to other data collectors as well.

- Name — comes directly from HDS HiCommand Device Manager's "name" attribute via the GetStorageArray XML API call
- Model - comes directly from HDS HiCommand Device Manager's "arrayType" attribute via the GetStorageArray XML API call
- Vendor — HDS
- Family - comes directly from HDS HiCommand Device Manager's "arrayFamily" attribute via the GetStorageArray XML API call
- IP — this is the management IP address of the array, not an exhaustive list of all IP addresses on the array
- Raw Capacity — a base2 value representing the sum of the total capacity of all disks in this system, regardless of disk role.

HDS Storage Pool

Terms applying to objects or references that you might find on HDS storage pool asset landing pages.

HDS Storage Pool Terminology

The following terms apply to objects or references that you might find on HDS storage pool asset landing pages. Many of these terms apply to other data collectors as well.

- **Type:** The value here will be one of:
 - **RESERVED** — if this pool is dedicated for purposes other than data volumes, i.e, journaling, snapshots
 - **Thin Provisioning** — if this is a HDP pool
 - **Raid Group** — you will not likely see these for a few reasons:

OCI takes a strong stance to avoid double counting capacity at all costs. On HDS, one typically needs to build Raid Groups from disks, create pool volumes on those Raid Groups, and construct pools (often HDP, but could be special purpose) from those pool volumes. If OCI reported both the underlying Raid Groups as is, as well as the Pools, the sum of their raw capacity would vastly exceed the sum of the disks.

Instead, OCI's HDS HiCommand data collector arbitrarily shrinks the size of Raid Groups by the capacity of pool volumes. This may result in OCI not reporting the Raid Group at all. Additionally, any resulting Raid Groups are flagged in a way such that they are not visible in the OCI WebUI, but they do flow into the OCI Data Warehouse (DWH). The purpose of these decisions is to avoid UI clutter for things that most users do not care about — if your HDS array has Raid Groups with 50MB free, you probably cannot use that free space for any meaningful outcome.

- **Node** - N/A, as HDS pools are not tied to any one specific node
- **Redundancy** - the RAID level of the pool. Possibly multiple values for a HDP pool comprised of multiple RAID types
- **Capacity %** - the percent used of the pool for data usage, with the used GB and total logical GB size of the pool
- **Over-committed Capacity** - a derived value, stating “the logical capacity of this pool is oversubscribed by this percentage by virtue of the sum of the logical volumes exceeding the logical capacity of the pool by this percentage”
- **Snapshot** - shows the capacity reserved for snapshot usage on this pool

HDS Storage Node

Terms applying to objects or references that you might find on HDS storage node asset landing pages.

HDS Storage Node Terminology

The following terms apply to objects or references that you might find on HDS storage node asset landing pages. Many of these terms apply to other data collectors as well.

- **Name** — The name of the Front-end director (FED) or Channel Adapter on monolithic arrays, or the name of the controller on a modular array. A given HDS array will have 2 or more Storage Nodes

- Volumes — The Volume table will show any volume mapped to any port owned by this storage node

Hitachi Ops Center data collector

This data collector uses Hitachi Ops Center's integrated suite of applications to access inventory and performance data of multiple storage devices. For inventory and capacity discovery, your Ops Center installation must include both the "Common Services" and "Administrator" components. For performance collection, you must additionally have "Analyzer" deployed.

Terminology

OnCommand Insight acquires the following inventory information from this data collector. For each asset type acquired, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

Vendor/Model Term	OnCommand Insight Term
Storage Systems	Storage
Volume	Volume
Parity Groups	Storage Pool(RAID), Disk Groups
Disk	Disk
Storage Pool	Storage Pool(Thin, SNAP)
External Parity Groups	Storage Pool(Backend), Disk Groups
Port	Storage Node → Controller Node → Port
Host Groups	Volume Mapping and Masking
Volume Pairs	Storage Synchronization

Note: These are common terminology mappings only and might not represent every case for this data collector.

Inventory Requirements

You must have the following in order to collect inventory data:

- IP address or hostname of the Ops Center server hosting the "Common Services" component
- Root/sysadmin user account and password that exist on all servers hosting Ops Center components. HDS did not implement REST API support for usage by LDAP/SSO users until Ops Center 10.8+

Performance requirements

The following requirements must be met in order to collect performance data:

- The HDS Ops Center "Analyzer" module must be installed
- Storage arrays must be feeding the Ops Center "Analyzer" module

Configuration

Field	Description
Hitachi Ops Center IP Address	IP address or fully-qualified domain name of the Ops Center server hosting the "Common Services" component
User Name	User name for the Ops Center server.
Password	Password used for the Ops Center server.

Advanced configuration

Field	Description
Connection Type	HTTPS (port 443) is the default
Override TCP Port	Specify the port to use if not the default
Inventory Poll Interval (min)	Interval between inventory polls. The default is 40.
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data.
Filter device List	Comma-separated list of device serial numbers to include or exclude
Performance Poll Interval (sec)	Interval between performance polls. The default is 300.

HDS NAS (HNAS) data source

The HDS NAS (HNAS) data source is an inventory and configuration data source to support discovery of HDS NAS clusters. Insight supports discovering NFS and CIFS shares, file systems (Insight Internal Volumes), and spans (Insight Storage Pools).

This data source is SSH based, so the acquisition unit that will host it needs to be able to initiate SSH sessions to TCP 22 on the HNAS itself, or the Systems Management Unit (SMU) that the cluster is connected to.

Terminology

OnCommand Insight acquires the following inventory information from the HNAS data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Tier	Disk Group
Cluster	Storage
Node	Storage Node
Span	Storage Pool

File System	Internal Volume
-------------	-----------------



These are common terminology mappings only and might not represent every case for this data source.

Requirements

The following are requirements to configure and use this data source:

- Device IP address
- Port 22, SSH protocol
- Username and password - privilege level: Supervisor
- NOTE: This data collector is SSH based, so the AU that hosts it must be able to initiate SSH sessions to TCP 22 on the HNAS itself, or the Systems Management Unit (SMU) that the cluster is connected to.



This data collector is SSH based, so the AU that hosts it must be able to initiate SSH sessions to TCP 22 on the HNAS itself, or the Systems Management Unit (SMU) that the cluster is connected to.

Configuration

Field	Description
HNAS Host	IP address or fully-qualified domain name of HNAS Management Host
User Name	User name for HNAS CLI
Password	Password used for HNAS CLI

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 30 minutes)
SSH Banner Wait Timeout (sec)	SSH banner wait timeout (default 15 seconds)
SSH Command Timeout (sec)	SSH command timeout (default 30 seconds)

HP CommandView AE data source

The HP CommandView Advanced Edition (AE) and CommandView AE CLI/SMI (AE Lite) data sources support inventory and performance from a CommandView (also referred to as HiCommand) Device Manager server.

Terminology

OnCommand Insight acquires the following inventory information from the HP CommandView AE and AE Lite data sources. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
PDEV	Disk
Journal Pool	Disk Group
Storage Array	Storage
Port Controller	Storage Node
Array Group, DP Pool	Storage Pool
Logical Unit, LDEV	Volume



These are common terminology mappings only and might not represent every case for this data source.

Inventory Requirements

- IP address of the HiCommand Device Manager server
- Read-only user name and password for the CommandView AE software and peer privileges
- The CommandView AE Lite version of the device manager has only the CLI licensed
- Port requirement: 2001

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`.

Configuration

Field	Description
HiCommand Server	IP address or fully-qualified domain name of the HiCommand Device Manager server
User Name	User name for the HiCommand Device Manager server.
Password	Password used for the HiCommand Device Manager server.
Devices - USP, USP V, VSP/R600 Storages	Device list for VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages. Each storage requires: <ul style="list-style-type: none"> • Array's IP: IP address of the storage • User Name: User name for the storage • Password: Password for the storage • Folder Containing Export Utility JAR Files: The folder containing the Export utility <code>.jar</code> files
SNM2Devices - WMS/SMS/AMS Storages	Device list for WMS/SMS/AMS storages. Each storage requires: <ul style="list-style-type: none"> • Array's IP: IP address of the storage • Storage Navigator CLI Path: SNM2 CLI path • Account Authentication Valid: Select to choose valid account authentication • User Name: User name for the storage • Password: Password for the storage
Choose Tuning Manager for Performance	Choose Tuning Manager for performance and override other performance options
Tuning Manager Host	IP address or fully-qualified domain name of tuning manager
Tuning Manager Port	Port used for Tuning Manager
Tuning Manager Username	User name for Tuning Manager
Tuning Manager Password	password for Tuning Manager



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

Advanced configuration

Field	Description
HiCommand Server Port	Port used for the HiCommand Device Manager
HTTPs Enabled	Select to enable HTTPs
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data
Exclude or Include Devices	Comma-separated list of device ID's or array names to include or exclude
Query Host Manager	Select to query host manager
HTTP Timeout (sec)	HTTP connection timeout (default 60 seconds)
Performance Polling Interval (sec)	Interval between performance polls (default 300 seconds)
Export timeout in seconds	Export utility timeout (default 300 seconds)

HP EVA Storage data source

For configuration, The EVA Storage (SSSU) 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.

Terminology

OnCommand Insight acquires the following inventory information from the HP EVA data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Disk Group	Disk Group (not modeled)
Storage Cell	Storage

Virtual Disk	Storage Pool
Virtual Disk	Volume



These are common terminology mappings only and might not represent every case for this data source.

Inventory 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.

Performance Requirements

The HP StorageWorks Command View EVA software suite must be installed on the OnCommand Insight Server. Alternatively, you can 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.

Configuration

Field	Description
CommandView Server	IP address or fully-qualified domain name of the EVA Storage Manager
User Name	User name for the Command View manager. The name must be defined in Command View.

Password	Password used for the Command View manager.
Performance User Name	For performance, the user name for the Command View manager. The name must be defined in Command View.
Performance Password	For performance, the password used for the Command View manager.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
CLI Home	Full pathname to the CLI home directory where <code>sssu.exe</code> is located
Inventory Exclude Devices	Comma-separated list of device names to include
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Performance CLI Home	For Array Performance, full pathname to the CLI home directory where <code>sssu.exe</code> is located. To validate access, run <code>sssu.exe</code>
Command Timeout (sec)	<code>evaperf</code> command wait timeout (default 600 seconds)
Performance Exclude Devices	Comma-separated list of device names to exclude from collecting performance data

HPE Nimble data source

The HPE Nimble data collector supports inventory and performance data for HPE Nimble storage arrays.

Terminology

OnCommand Insight acquires the following inventory information from the HPE Nimble data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Array	Storage

Disk	Disk
Pool	Storage Pool
Volume	Volume
Initiator	Storage Host Alias
Controller	Storage Node
Fibre Channel Interface	Controller



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- The array must be installed and configured, and reachable from the client through its fully qualified domain name (FQDN) or array management IP address.
- The array must be running NimbleOS 2.3.x or later.
- You must have a valid user name and password to the array.
- Port 5392 must be open on the array.

Configuration

Field	Description
Array Management IP Address	Fully qualified domain name (FQDN) or array management IP address.
User Name	User name for the Nimble array
Password	Password for the Nimble array

Advanced configuration

Field	Description
Port	Port used by Nimble REST API. The default is 5392.
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)

Note: The default performance poll interval is 300 seconds and can not be changed. This is the only interval supported by Nimble.

Huawei OceanStor data source

OnCommand Insight uses the Huawei OceanStor (REST/HTTPS) data source to discover inventory for Huawei OceanStor storage.

Terminology

OnCommand Insight acquires the following inventory and performance information from the Huawei OceanStor. For each asset type acquired by OnCommand Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

Vendor/Model Term	OnCommand Insight Term
Storage Pool	Storage Pool
File System	Internal Volume
Controller	Storage Node
FC Port (Mapped)	Volume Map
Host FC Initiator (Mapped)	Volume Mask
NFS/CIFS Share	Share
Share	iSCSI Target Node
iSCSI Link Initiator	iSCSI Initiator Node
Disk	Disk
LUN	Volume

Requirements

The following are requirements to configure and use this data collector:

- Device IP
- Credentials to access OceanStor device manager
- Port 8088 must be available

Configuration

Field	Description
OceanStor Host IP Address	IP address or fully-qualified domain name of the OceanStor Device Manager

User Name	Name used to log into the OceanStor Device Manager
Password	Password used to log into the OceanStor Device Manager

Advanced configuration

Field	Description
TCP Port	TCP Port used to connect to OceanStor Device Manager (default 8088)
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Connection Timeout (sec)	Connection timeout (default 60 seconds)

IBM Cleversafe data source

This data source collects inventory and performance data for IBM Cleversafe.

Requirements

The following are requirements for configuring this data source:

- Manager IP Address or Host Name
- A username and password for same
- Port 9440

Configuration

Field	Description
Cleversafe manager Host Name or IP Address	Host IP address of the CleverSafe device
User Name	Name used to log into the Cleversafe
Password	Password used to log into the Cleversafe

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Default is 60 minutes
HTTP Connection Timeout)	Default is 60 seconds

IBM DS 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.

Terminology

OnCommand Insight acquires the following inventory information from the IBM DS data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk Drive Module	Disk
Storage Image	Storage
Extent Pool	Storage Pool
Fixed Block Volume	Volume



These are common terminology mappings only and might not represent every case for this data source.

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
- Port requirements: 80, 443, & 1750

Configuration

Field	Description
DS storage	IP address or fully-qualified domain name of the DS Storage Host
User Name	Name used for the DS CLI
Password	Password used for the DS CLI
Executable dscli.exe Path	Full path to the <code>dscli.exe</code> utility.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Storage Display Name	Name of the IBM DS storage array
Inventory Exclude Devices	Comma-separated list of device serial numbers to exclude from inventory collection
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Performance Filter Type	Include: Data collected only from devices on list. Exclude: No data from these devices is collected
Performance Filter Device List	Comma-separated list of device IDs to include or exclude from performance collection

IBM PowerVM data source

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.

Terminology

OnCommand Insight acquires the following inventory information from the IBM PowerVM data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
hdisk	Virtual Disk
Managed System	Host
LPAR, VIO Server	Virtual Machine
Volume Group	Data Store
Physical Volume	LUN



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- IP address of the Hardware Management Console (HMC)
- User name and password that provide access to HMC through SSH
- Port requirement SSH-22
- 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.

- IBM best practice is to have the devices monitored by two or more HMCs. Be aware that this may cause OnCommand Insight to report duplicated devices, therefore it is highly recommended to add redundant devices to the "Exclude Devices" list in the Advanced Configuration for this data collector.

Configuration

Field	Description
Hardware Management Console (HMC) Address	IP address or fully-qualified domain name of the PowerVM Hardware Management Console
HMC User	User name for the Hardware Management Console
Password	Password used for the Hardware Management Console

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
SSH Port	Port used for SSH to the PowerVM
SSH Process Wait Timeout (sec)	SSH process timeout (default 600 seconds)
Number of Retries	Number of inventory retry attempts
Exclude Devices	Comma-separated list of device IDs or display names to exclude

IBM SVC data source

The IBM SVC data source collects inventory and performance data using SSH, supporting 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.

Terminology

OnCommand Insight acquires the following inventory information from the IBM SVC data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive	Disk
Cluster	Storage
Node	Storage Node
Mdisk Group	Storage Pool
Vdisk	Volume
Mdisk	Backend LUN



These are common terminology mappings only and might not represent every case for this data source.

Inventory Requirements

- IP address of each SVC cluster
- Port 22 available
- 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



No third-party software needs to be installed.

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.



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

- Port 22 required
- 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`



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

Parent Serial Number explained

Traditionally Insight is capable of reporting the storage array serial number, or the individual storage node serial numbers. However, some storage array architectures do not cleanly align to this. An SVC cluster can be comprised of 1-4 appliances, and each appliance has 2 nodes. If the appliance itself has a serial number, that serial number is neither the serial number for the cluster nor the nodes.

The attribute "Parent Serial Number" on the storage node object is populated appropriately for IBM SVC arrays when the individual nodes sit inside an intermediate appliance/enclosure that is just part of a larger cluster.

Configuration

Field	Description
Cluster/s IP	IP address of fully-qualified domain name for the SVC storage
Choose 'Password' or 'OpenSSH Key File' to specify credential type	The credential type used to connect to the device via SSH
Inventory User Name	User name for the SVC CLI
Inventory Password	Password for the SVC CLI
Full Path to Inventory Private Key	Full path to the Inventory private key file
Performance User Name	User name for the SVC CLI for performance collection
Performance Password	Password for the SVC CLI for performance collection
Full Path to Performance Private Key	Full path to the Performance private key file

Advanced configuration

Field	Description
-------	-------------

Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Exclude Devices	Comma-separated list of device IDs to exclude from inventory collection
SSH Process Wait Timeout (sec)	SSH process timeout (default 200 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Performance Exclude Devices	Comma-separated list of device IDs to exclude from performance collection
Performance SSH Process Wait Timeout (sec)	SSH process timeout (default 200 seconds)
To clean up dumped stats files	Select to clean up dumped stats files

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



This data collector is no longer available starting with OnCommand Insight 7.3.11.

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 to connect: `testdb.company.com`.

Configuration

Field	Description
Tivoli Monitoring Database IP	IP address or fully-qualified domain name of the Tivoli Monitoring server
User Name	User name for the Tivoli Monitoring server
Password	Password for the Tivoli Monitoring server

Advanced configuration

Field	Description
Tivoli Monitoring Database Port	Port used for Tivoli monitoring database
Oracle SID or DB2 Database Name	Oracle listener service ID or DB2 database name
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Database Driver to Use	Choose Database Driver to use
Protocol Used to Connect to the Database	Protocol Used to Connect to the Database
Database Schema	Enter Database Schema

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	IP addresses or fully-qualified domain names of controllers, separated by commas

Requirements

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

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 30 minutes)
Performance Poll Interval (up to 3600 seconds)	Interval between performance polls (default 300 seconds)

IBM XIV data source

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

Terminology

OnCommand Insight acquires the following inventory information from the IBM XIV data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Storage System	Storage
Storage Pool	Storage Pool
Volume	Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- Port requirement: TCP port 7778
- IP address of the XIV management interface
- Read-only user name and password
- The XIV CLI must be installed on the Insight server or RAU
- Access validation: Log in to the XIV user interface from the Insight server using the user name and password.

Configuration

Field	Description
IP Address	IP address or fully-qualified domain name for the XIV storage
User Name	User name for the XIV storage
Password	Password for the XIV storage
Full path to XIV CLI directory	Full path to the XIV CLI directory

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)

CLI Process Wait Timeout (ms)	CLI process timeout (default 7200000 ms)
SMI-S Host IP	IP address of the SMI-S Provider Host
SMI-S Port	Port used by SMI-S Provider Host
SMI-S Protocol	Protocol used to connect to the SMI-S provider
SMI-S Namespace	SMI-S namespace
Username	User name for the SMI-S Provider Host
Password	Password for the SMI-S Provider Host
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Number of SMI-S Connection Retries	Number of SMI-S connection retry attempts

Infinidat InfiniBox data source

The Infinidat InfiniBox (HTTP) data source is used to collect information from the Infinidat InfiniBox storage. You must have access to the InfiniBox Management Node.

Terminology

OnCommand Insight acquires the following inventory information from the InfiniBox data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive	Disk
InfiniBox	Storage
Node	Storage Node
Pool	Storage Pool
Volume	Volume
FC Port	Port
Filesystem	Internal Volume

Filesystem	FileShare
Filesystem Exports	Share



These are common terminology mappings only and might not represent every case for this data source.

Configuration

Field	Description
InfiniBox Host	IP address or fully-qualified domain name of the InfiniBox Management Node
User Name	User name for InfiniBox Management Node
Password	Password for the InfiniBox Management Node

Advanced configuration

Field	Description
TCP Port	TCP Port used to connect to InfiniBox Server (default 443)
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Connection Timeout	Connection timeout (default 60 seconds)

Microsoft Azure compute data source

OnCommand Insights uses the Azure compute data collector to acquire inventory and performance data from Azure compute instances.

Requirements

You need the following information to configure this data collector:

- Port requirement: 443 HTTPS
- Azure Management Rest IP (management.azure.com)
- Azure Service Principal Application (Client) ID (user account)
- Azure Service Principal Authentication key (user password)

You need to set up an Azure account for Insight discovery. Once the account is properly configured and you register the application in Azure, you will have the credentials required to discover the Azure instance with Insight. The following link describes how to set up the account for discovery:<https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-create-service-principal-portal>

Configuration

Enter data into the data source fields according to the table below:

Field	Description
Azure Service Principal Application (Client) ID (Reader role required)	Sign-in ID to Azure. Requires Reader Role access.
Azure tenant ID	Microsoft tenant ID
Azure Service Principal Authentication Key	Login authentication key
I understand Microsoft bills me for API requests	Check this to verify your understanding that Microsoft bills you for API requests made by Insight polling.

Advanced Configuration

Enter data into the data source fields according to the table below:

Field	Description
Inventory Poll Interval (min)	The default is 60
Choose 'Exclude' or 'Include' to Apply to Filter VMs by Tags	Specify whether to include or exclude VM's by Tags when collecting data. If 'Include' is selected, the Tag Key field can not be empty.
Tag Keys and Values on which to Filter VMs	Click + Filter Tag to choose which VMs (and associated disks) to include/exclude by filtering for keys and values that match keys and values of tags on the VM. Tag Key is required, Tag Value is optional. When Tag Value is empty, the VM is filtered as long as it matches the Tag Key.
Performance Poll Interval (sec)	The default is 300

Azure NetApp Files data source

This data source acquires inventory and performance data for Azure NetApp Files (ANF).

Requirements

The following are requirements for configuring this data source:

- Port requirement: 443 HTTPS
- Azure Management Rest IP (management.azure.com)
- Azure Service Principal Application (Client) ID (user account)
- Azure Service Principal authentication key (user password)

- You need to set up an Azure account for Cloud Insights discovery.

Once the account is properly configured and you register the application in Azure, you will have the credentials required to discover the Azure instance with Cloud Insights. The following link describes how to set up the account for discovery:

<https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-create-service-principal-portal>

Configuration

Field	Description
Azure Service Principal Application (Client) ID	Sign-in ID to Azure
Azure Tenant ID	Azure Tenant ID
Azure Service Principal Authentication Key	Login authentication key
I understand Microsoft bills me for API requests	Check this to verify your understanding that Microsoft bills you for API requests made by Insight polling.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Default is 60 minutes

Microsoft Hyper-V data source

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

Terminology

OnCommand Insight acquires the following inventory information from the Hyper-V data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Virtual hard Disk	Virtual Disk
Host	Host
Virtual Machine	Virtual Machine
Cluster Shared Volumes (CSV), Partition Volume	Data Store

Internet SCSI Device, Multi Path SCSI LUN	LUN
Fiber Channel Port	Port



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- The Hyper-V requires port 5985 opened for data collection and remote access/management.
- IP address of Clustering group node
- Local Administrator user & password on the hypervisor
- Administrative-level user account
- Port requirements: Port 135 and Dynamic TCP ports assigned 1024-65535 for Windows 2003 and older and 49152-65535 for Windows 2008.
- DNS resolution must succeed, even if the data collector is pointed at only an IP address.
- Each Hyper-V hypervisor must have “Resource Metering” turned on for every VM, on every host. This allows each hypervisor to have more data available for Cloud Insights on each guest. If this is not set, fewer performance metrics are acquired for each guest. More information on Resource metering can be found in the microsoft documentation:

[Hyper-V Resource Metering Overview](#)

[Enable-VMResourceMetering](#)

Configuration

Field	Description
Physical Host IP Address	The IP address or fully-qualified domain name for the physical host (hypervisor)
User Name	Administrator user name doe the hypervisor
Password	Password for the hypervisor
NT Domain	The DNS name used by the nodes in the cluster

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection Timeout (ms)	Connection timeout (default 60000 ms)

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.

Terminology

OnCommand Insight acquires the following inventory information from the Clustered Data ONTAP data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Raid Group	Disk Group
Cluster	Storage
Node	Storage Node
Aggregate	Storage Pool
LUN	Volume
Volume	Internal Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- Administrator account used for read-only API calls
- Target IP is the cluster management LIF
- Username (with read-only role name to ontapi application to the default Vserver) and password to log into NetApp cluster
- Port requirements: 80 or 443
- License requirements: FCP license and mapped/masked volumes required for discovery

Configuration

Field	Description
NetApp Management IP	IP address or fully-qualified domain name of the NetApp cluster
User Name	User name for the NetApp cluster

Password	Password for the NetApp cluster
----------	---------------------------------

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Clustered Data ONTAP Storage

Terms applying to objects or references that you might find on NetApp Clustered Data ONTAP storage asset landing pages.

Clustered Data ONTAP Storage Terminology

The following terms apply to objects or references that you might find on NetApp Clustered Data ONTAP storage asset landing pages. Many of these terms apply to other data collectors as well.

- **Model** — A comma delimited list of the unique, discrete node model names within this cluster. If all the nodes in the clusters are the same model type, just one model name will appear.
- **Vendor** — same Vendor name you would see if you were configuring a new data source.
- **Serial number** — The array serial number. On cluster architecture storage systems like NetApp Clustered Data Ontap, this serial number may be less useful than the individual “Storage Nodes” serial numbers.
- **IP** — generally will be the IP(s) or hostname(s) as configured in the data source.
- **Microcode version** — firmware.
- **Raw Capacity** — base 2 summation of all the physical disks in the system, regardless of their role.
- **Latency** — a representation of what the host facing workloads are experiencing, across both reads and writes. Ideally, OCI is sourcing this value directly, but this is often not the case. In lieu of the array offering this up, OCI is generally performing an IOPs-weighted calculation derived from the individual internal volumes’ statistics.
- **Throughput** — aggregated from internal volumes.
- **Management** — this may contain a hyperlink for the management interface of the device. Created programmatically by the Insight data source as part of inventory reporting.

Clustered Data ONTAP Storage Pool

Terms applying to objects or references that you might find on NetApp Clustered Data ONTAP storage pool asset landing pages.

Clustered Data ONTAP Storage Pool Terminology

The following terms apply to objects or references that you might find on NetApp Clustered Data ONTAP storage pool asset landing pages. Many of these terms apply to other data collectors as well.

- **Storage** — what storage array this pool lives on. Mandatory.
- **Type** — a descriptive value from a list of an enumerated list of possibilities. Most commonly will be “Aggregate” or “RAID Group”.
- **Node** — if this storage array’s architecture is such that pools belong to a specific storage node, its name will be seen here as a hyperlink to its own landing page.
- **Uses Flash Pool** — Yes/No value — does this SATA/SAS based pool have SSDs used for caching acceleration?
- **Redundancy** — RAID level or protection scheme. RAID_DP is dual parity, RAID_TP is triple parity.
- **Capacity** — the values here are the logical used, usable capacity and the logical total capacity, and the percentage used across these.
- **Over-committed capacity** — If by using efficiency technologies you have allocated a sum total of volume or internal volume capacities larger than the logical capacity of the storage pool, the percentage value here will be greater than 0%.
- **Snapshot** — snapshot capacities used and total, if your storage pool architecture dedicates part of its capacity to segments areas exclusively for snapshots. Ontap in MetroCluster configurations are likely to exhibit this, while other Ontap configurations are less so.
- **Utilization** — a percentage value showing the highest disk busy percentage of any disk contributing capacity to this storage pool. Disk utilization does not necessarily have a strong correlation with array performance — utilization may be high due to disk rebuilds, deduplication activities, etc in the absence of host driven workloads. Also, many arrays’ replication implementations may drive disk utilization while not showing as internal volume or volume workload.
- **IOPS** — the sum IOPs of all the disks contributing capacity to this storage pool.
- **Throughput** — the sum throughput of all the disks contributing capacity to this storage pool.

Clustered Data ONTAP Storage Node

Terms applying to objects or references that you might find on NetApp Clustered Data ONTAPs storage node asset landing pages.

Clustered Data ONTAP Storage Node Terminology

The following terms apply to objects or references that you might find on NetApp Clustered Data ONTAP storage pool asset landing pages. Many of these terms apply to other data collectors as well.

- **Storage** — what storage array this node is part of. Mandatory.
- **HA Partner** — on platforms where a node will fail over to one and only one other node, it will generally be seen here.
- **State** — health of the node. Only available when the array is healthy enough to be inventoried by a data source.
- **Model** — model name of the node.
- **Version** — version name of the device.
- **Serial number** — The node serial number.
- **Memory** — base 2 memory if available.
- **Utilization** — On Ontap, this is a controller stress index from a proprietary algorithm. With every performance poll, a number between 0 and 100% will be reported that is the higher of either WAFL disk contention, or average CPU utilization. If you observe sustained values > 50%, that is indicative of

undersizing — potentially a controller/node not large enough or not enough spinning disks to absorb the write workload.

- IOPS — Derived directly from Ontap ZAPI calls on the node object.
- Latency — Derived directly from Ontap ZAPI calls on the node object.
- Throughput — Derived directly from Ontap ZAPI calls on the node object.
- Processors — CPU count.

NetApp Clustered Data ONTAP for Unified Manager data source

This data source collects 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



This data collector is no longer available starting with OnCommand Insight 7.3.11.

Field	Description
Unified Manager IP	IP address or fully-qualified domain name of the Unified Manager
User Name	User name for the Unified Manager
Password	Password for the Unified Manager
Port	Port used for communication with the Unified Manager (default 3306)

Advanced configuration

Field	Description
Inventory Poll Interval (min) Interval	Interval between inventory polls (default 15 minutes)
Exclude Clusters	Comma-separated list of cluster IPs to exclude

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.

Terminology

OnCommand Insight acquires the following inventory information from the NetApp Data ONTAP 7-Mode data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown.

When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Raid Group	Disk Group
Filer	Storage
Filer	Storage Node
Aggregate	Storage Pool
LUN	Volume
Volume	Internal Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- IP address of the FAS storage controller and partner
- Port 443
- User name and password for the controller and the partner
- A custom admin level username and password for controller and partner controller with the following role capabilities for 7-Mode:
 - "api-*": Use this to allow OnCommand Insight to execute all NetApp storage API commands.
 - "login-http-admin": Use this to allow OnCommand Insight to connect to the NetApp storage via HTTP.
 - "security-api-vfiler": Use this to allow OnCommand Insight to execute NetApp storage API commands to retrieve vFiler unit information.
 - "cli-options": Use this to read storage system options.
 - "cli-lun": Access these commands 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": Use this to display free disk space.
 - "cli-ifconfig": Use this to display interfaces and IP addresses.

Configuration

Field	Description
Address of Filer	IP address or fully-qualified domain name for the NetApp Filer

User Name	User name for the NetApp Filer
Password	password for the NetApp Filer
Address of HA Partner Filer in Cluster	IP address or fully-qualified domain name for the HA Partner Filer
User Name of HA Partner Filer in Cluster	User name for the NetApp HA Partner Filer
Password of HA Partner Filer in Cluster	password for the NetApp HA Partner Filer

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection Type	Choose connection type
Connection Port	Port used for NetApp API
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Storage systems connection

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-*
```

NetApp E-Series data source

The NetApp E-Series 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.

Terminology

OnCommand Insight acquires the following inventory information from the NetApp E-Series data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive	Disk
Volume Group	Disk Group
Storage Array	Storage
Controller	Storage Node
Volume Group	Storage Pool
Volume	Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- The IP address of each controller on the array
- Port requirement 2463

Configuration

Field	Description
Comma-separated list of Array SANtricity Controller IPs	IP addresses and/or fully-qualified domain names for the array controllers

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 30 minutes)
Performance Poll Interval (up to 3600 seconds)	Interval between performance polls (default 300 seconds)

E-Series Storage

Terms applying to objects or references that you might find on NetApp E-Series storage asset landing pages.

E-Series Storage Terminology

The following terms apply to objects or references that you might find on NetApp E-Series storage asset landing pages. Many of these terms apply to other data collectors as well.

- Model — model name of the device.
- Vendor — same Vendor name you would see if you were configuring a new data source.
- Serial number — The array serial number. On cluster architecture storage systems like NetApp Clustered Data Ontap, this serial number may be less useful than the individual “Storage Nodes” serial numbers.
- IP — generally will be the IP(s) or hostname(s) as configured in the data source.
- Microcode version — firmware.
- Raw Capacity — base 2 summation of all the physical disks in the system, regardless of their role.
- Latency — a representation of what the host facing workloads are experiencing, across both reads and writes. Insight calculates an IOPs-weighted average derived from the volumes in the storage.
- Throughput — the array’s total host facing throughput. Insight sums the volumes’ throughput to derive this value.
- Management — this may contain a hyperlink for the management interface of the device. Created programmatically by the Insight data source as part of inventory reporting.

E-Series Storage Pool

Terms applying to objects or references that you might find on NetApp E-Series storage

pool asset landing pages.

E-Series Storage Pool Terminology

The following terms apply to objects or references that you might find on NetApp E-Series storage pool asset landing pages. Many of these terms apply to other data collectors as well.

- **Storage** — what storage array this pool lives on. Mandatory.
- **Type** — a descriptive value from a list of an enumerated list of possibilities. Most commonly will be “Thin Provisioning” or “RAID Group”.
- **Node** — if this storage array’s architecture is such that pools belong to a specific storage node, its name will be seen here as a hyperlink to its own landing page.
- **Uses Flash Pool** — Yes/No value.
- **Redundancy** — RAID level or protection scheme. E-Series reports “RAID 7” for DDP pools.
- **Capacity** — the values here are the logical used, usable capacity and the logical total capacity, and the percentage used across these. These value both include E-Series “preservation” capacity, resulting both in numbers and the percentage being higher than what the E-Series own user interface may show.
- **Over-committed capacity** — If by using efficiency technologies you have allocated a sum total of volume capacities larger than the logical capacity of the storage pool, the percentage value here will be greater than 0%.
- **Snapshot** — snapshot capacities used and total, if your storage pool architecture dedicates part of its capacity to segments areas exclusively for snapshots.
- **Utilization** — a percentage value showing the highest disk-busy percentage of any disk contributing capacity to this storage pool. Disk utilization does not necessarily have a strong correlation with array performance — utilization may be high due to disk rebuilds, deduplication activities, etc in the absence of host-driven workloads. Also, many arrays’ replication implementations may drive disk utilization while not showing as volume workload.
- **IOPS** — the sum IOPs of all the disks contributing capacity to this storage pool.
- **Throughput** — the sum throughput of all the disks contributing capacity to this storage pool.

E-Series Storage Node

Terms applying to objects or references that you might find on NetApp E-Series storage node asset landing pages.

E-Series Storage Node Terminology

The following terms apply to objects or references that you might find on NetApp E-Series storage pool asset landing pages. Many of these terms apply to other data collectors as well.

- **Storage** — what storage array this node is part of. Mandatory.
- **HA Partner** — on platforms where a node will fail over to one and only one other node, it will generally be seen here.
- **State** — health of the node. Only available when the array is healthy enough to be inventoried by a data source.
- **Model** — model name of the node.
- **Version** — version name of the device.

- Serial number — The node serial number.
- Memory — base 2 memory if available.
- Utilization — Utilization is not currently available for NetApp E-Series.
- IOPS — Calculated by summing all the IOPs for volumes that belong exclusively to this node.
- Latency — a number representing the typical host latency or response time on this controller. Insights calculates an IOPs weighted average from volumes that belong exclusively to this node.
- Throughput — a number representing the host driven throughput on this controller. Calculated by summing all the throughput for volumes that belong exclusively to this node.
- Processors — CPU count.

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).

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.

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.
- Dynamic TCP ports assigned 1024-65535 for Windows 2003 and older
- Ports 49152—65535 for Windows 2008



As a general rule, when trying to use a firewall between Insight, an AU, and this data source, you should consult with your Microsoft team to identify the ports they believe will be 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.

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.

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.

Advanced configuration

Field	Description
Inventory poll interval (min)	Interval between inventory polls (default 360 minutes)

NetApp SolidFire data source

The NetApp SolidFire data source supports both iSCSI and Fibre Channel SolidFire configurations, for both inventory and performance collection.

The SolidFire data source utilizes the SolidFire REST API. The acquisition unit where the data source resides needs to be able to initiate HTTPS connections to TCP port 443 on the SolidFire cluster management IP address. The data source needs credentials capable of making REST API queries on the SolidFire cluster.

Terminology

OnCommand Insight acquires the following inventory information from the NetApp SolidFire data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive	Disk
Cluster	Storage
Node	Storage Node
Volume	Volume
Fibre Channel Port	Port
Volume Access Group, LUN Assignment	Volume Map
iSCSI Session	Volume Mask



These are common terminology mappings only and might not represent every case for this data source.

Requirements

The following are requirements for configuring this data source:

- Management Virtual IP Address
- Port 443

Configuration

Field	Description
Management Virtual IP Address (MVIP)	Management Virtual IP address of the SolidFire Cluster
User Name	Name used to log into the SolidFire cluster
Password	Password used to log into the SolidFire cluster

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)

TCP Port	TCP Port used to connect to SolidFire Server (default 443)
Connection Timeout (sec)	Connection timeout (default 60 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Troubleshooting

When SolidFire reports an error it is displayed in OnCommand Insight as follows:

An error message was received from a SolidFire device while trying to retrieve data. The call was <method> (<parameterString>). The error message from the device was (check the device manual): <message>

Where:

- The <method> is an HTTP method, such as GET or PUT.
- The <parameterString> is a comma separated list of parameters that were included in the REST call.
- The <message> is whatever the device returned as the error message.

NetApp StorageGRID data source

This data source collects inventory and performance data for StorageGRID.

Requirements

The following are requirements for configuring this data source:

- StorageGRID Host IP Address
- A username and password for a user that has had the Metric Query and Tenant Access roles assigned
- Port 443

Configuration

Field	Description
StorageGRID Host IP Address (MVIP)	Host IP address of the StorageGRID
User Name	Name used to log into the StorageGRID
Password	Password used to log into the StorageGRID

Advanced configuration

Field	Description
-------	-------------

Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Performance Poll Interval (sec)	Interval between performance polls (default 900 seconds)

OpenStack data source

The OpenStack (REST API / KVM) data source collects information about OpenStack hardware instances. This data source collects inventory data for all OpenStack instances, and optionally, VM performance data.

Requirements

The following are requirements for configuring the OpenStack data source.

- IP address of the OpenStack controller
- OpenStack admin role credentials and sudo access to the Linux KVM hypervisor are recommended.



If you are not using an admin account or admin equivalent privileges, you can still acquire data from the data source. You will need to modify the policy configuration file (i.e. `etc/nova/policy.json`) to allow users with non-admin role to call the API:

- `"os_compute_api:os-availability-zone:detail": ""`
- `"os_compute_api:os-hypervisors": ""`
- `os_compute_api:servers:detail:get_all_tenants": ""`
- For performance collection the OpenStack Ceilometer module must be installed and configured. Configuring the Ceilometer is done by editing the `nova.conf` file for each hypervisor and then restart the Nova Compute service on each hypervisor. The option name changes for different releases of OpenStack:
 - Icehouse
 - Juno
 - Kilo
 - Liberty
 - Mitaka
 - Newton
 - Ocata
- For CPU stats, `"compute_monitors=ComputeDriverCPUMonitor"` needs to be turned on in `/etc/nova/nova.conf` on compute nodes.
- Port requirements:
 - 5000 for http and 13000 for https, for the Keystone service
 - 22 for KVM SSH
 - 8774 for Nova Compute Service
 - 8776 for Cinder Block Service
 - 8777 for Ceilometer Performance Service

- 9292 for Glance Image Service



The port binds to the specific service, and the service may run on the controller or another host in larger environments.

Configuration

Field	Description
OpenStack Controller IP Address	IP address or fully-qualified domain name of the OpenStack Controller
OpenStack Administrator	User name for an OpenStack Admin
OpenStack Password	Password used for the OpenStack Admin
OpenStack Administrator Tenant	OpenStack Administrator Tenant
KVM Sudo User	KVM Sudo User name
Choose 'Password' or 'OpenSSH Key File' to specify credential type	The credential type used to connect to the device via SSH
Full Path to Inventory Private Key	Full Path to Inventory Private Key
KVM Sudo Password	KVM Sudo Password

Advanced configuration

Field	Description
Enable hypervisor inventory discovery through SSH	Check this to enable hypervisor inventory discovery through SSH
OpenStack Admin URL port	OpenStack Admin URL port
Use HTTPS	Check to use secure HTTP
HTTP Connection Timeout (sec)	Timeout for HTTP connection (default 300 seconds)
SSH Port	Port used for SSH
SSH Process Wait Timeout (sec)	SSH process timeout (default 30 seconds)
SSH Process Retries	Number of inventory retry attempts

Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
-------------------------------	---

Oracle ZFS data source

The Oracle ZFS data source supports inventory and performance collection.

Terminology

OnCommand Insight acquires the following inventory information from this data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk (SDD)	Disk
Cluster	Storage
Controller	Storage Node
LUN	Volume
LUN Map	Volume Map
Initiator, Target	Volume mask
Share	Internal Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

The following are requirements for configuring this data source:

- Host names for the ZFS Controller-1 and the ZFS Controller-2
- Administrator user name and credentials
- Port requirement: 215 HTTP/HTTPS

Configuration

ZFS Controller-1 Hostname	Host name for storage controller 1
ZFS Controller-2 Hostname	Host name for storage controller 2

User name	User name for the storage system administrator user account
Password	Password for the administrator user account

Advanced configuration

Field	Description
TCP port	TCP Port used to connect to ZFS (default 215)
Connection Type	HTTP or HTTPS
Inventory poll interval	Inventory poll interval (default 60 minutes)
Connection Timeout	Default is 60 seconds
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Troubleshooting

Some things to try if you encounter problems with this data collector:

Problem:	Try This:
"Invalid login credentials"	validate Zfs user account and password
"Configuration error" with error message "REST Service is disabled"	Verify REST service is enabled on this device.
"Configuration error " with error message "User unauthorized for command"	<p>Likely due to certain roles (for example, 'advanced_analytics') are not included for the configured user <userName>.Possible Solution:</p> <ul style="list-style-type: none"> • Correct the Analytics (statistic) scope for the user \${user} with the read only role:- From the Configuration → Users screen, put your mouse over the role and double click to allow editing • Select "Analytics" from the Scope drop down menu. A list of the possible properties appears. • Click the top most check box and it will select all three properties.- Click the Add button on the right side. • Click the Apply button at the top right of the pop-up window. The pop-up window will close.

Pure Storage FlashArray data source

The Pure Storage FlashArray (HTTP) data source is used to collect information from the Pure Storage Flash Array. Insight supports both inventory and performance collection.

Terminology

OnCommand Insight acquires the following inventory information from the Pure Storage FlashArray data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive (SSD)	Disk
Array	Storage
Controller	Storage Node
Volume	Volume
Port	Port
LUN Map (Host, Host Group, Target Port)	Volume Map, Volume Mask



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- Storage system IP address
- User name and password for the Administrator account of the Pure storage system.
- Port requirement: HTTP/HTTPS 80/443

Configuration

Field	Description
FlashArray Host	IIP address or fully-qualified domain name of FlashArray Management Server
User Name	User name for the FlashArray Management Server
Password	Password for the FlashArray Management Server

Advanced configuration

Field	Description
-------	-------------

Connection Type	Management Server
TCP Port	TCP Port used to connect to FlashArray Server (default 443)
Connection Timeout (sec)	Connection timeout (default 60 seconds)
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

QLogic FC Switch data source

For configuration, the QLogic FC Switch (SNMP) 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	IP address or fully-qualified domain name for the SANSurfer switch
SNMP version	SNMP version
SNMP community	SNMP Community String
User Name	User name for the SANSurfer switch
Password	Password for the SANSurfer switch

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 15 minutes)
SNMP Auth Protocol	SNMP authentication protocol (SNMPv3 only)
SNMP Retries	Number of SNMP retry attempts
SNMP Timeout (ms)	SNMP timeout (default 5000 ms)
Enable Trapping	Select to enable trapping

Minimum Time Between Traps (sec)	Minimum time between acquisition attempts triggered by traps (default 10 seconds)
Fabric Name	Fabric name to be reported by the data source. Leave blank to report the fabric name as WWN.
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Red Hat (RHEV) data source

The Red Hat Enterprise Virtualization (REST) data source collects information about RHEV instances via HTTPS.

Requirements

- IP address of the RHEV server over port 443 via REST API
- Read-only username and password
- RHEV Version 3.0+

Configuration

Field	Description
RHEV Server IP Address	IP address or fully-qualified domain name of the RHEV server
User Name	User name for the RHEV server
Password	Password used for the RHEV server

Advanced configuration

Field	Description
HTTPS Communication Port	Port used for HTTPS communication to RHEV
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection timeout (sec)	Connection timeout (default 60 seconds)

Violin Flash Memory Array data source

The Violin 6000-Series Flash Memory Array (HTTP) data source collects network information for analysis and validation from Violin 6000-series flash memory arrays.

Terminology



This data collector is no longer available starting with OnCommand Insight 7.3.11.

OnCommand Insight acquires the following inventory information from the Violin 6000-Series Flash Memory Array data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Violin Intelligent Memory Module (VIMM)	Disk
Container	Storage
Memory Gateway	Storage Node
LUN	Volume
Initiator, Initiator Group, Target	Volume Map, Volume Mask



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- You need a read-only user name and password to the storage.
- Validate access with a web browser using the storage IP address.

Configuration

Field	Description
IP address or FQDN of Violin Memory Array Main Gateway	IP address or fully-qualified domain name of the Violin Memory Array Main Gateway
User Name	User name for the Violin Memory Array Main Gateway
Password	Password for the Violin Memory Array Main Gateway

Advanced configuration

Field	Description
Communication Port	Port used for communication with Violin array
HTTPS Enabled	Select to use HTTPS

Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection timeout (sec)	Connection timeout (default 60 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

VMware vSphere data source

The VMware vSphere (Web Services) data source collects ESX Host information and requires *read-only* privileges on all objects within the Virtual Center.

Terminology

OnCommand Insight acquires the following inventory information from the VMware vSphere data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Virtual Disk	Disk
Host	Host
Virtual Machine	Virtual Machine
Data Store	Data Store
LUN	LUN
Fiber Channel Port	Port



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- IP address of the Virtual Center server
- Read-only username and password in Virtual Center
- Read-only privileges on all objects within the Virtual Center.
- SDK access on the Virtual Center server
- Port requirements: http-80 https-443
- Validate access by logging in to Virtual Center Client using your user name and password and verifying that the SDK is enabled by entering `telnet <vc_ip> 443`.

Configuration

Field
Description
Virtual Center Address
Network address for the Virtual Center or vSphere server, specified as an IP (<i>nnn.nnn.nnn.nnn</i> format) address or as a host name that can be resolved through DNS.
User Name
User name for the VMware server.
Password
Password for the VMware server.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection Timeout (ms)	Connection timeout (default 60000 ms)
Filter VMs by	Choose how to filter VMs
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the VM list below when collecting data
List of VMs to filter (Comma Separated, or Semicolon Separated If Comma Is Used in the Value)	Comma-separated or semicolon-separated list of VMs to include or exclude from polling
Number of Retries for Requests to vCenter	Number of vCenter Request retry attempts
Communication Port	Port used for VMware server
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

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 **Admin**.

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 support to 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 update your Device Resolution (formerly Auto Resolution) methods.

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 **Admin**.

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 3, 7, or 30 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.
2. On the Summary page, 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.

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** and navigate to the data source list view
2. Select the data source for 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 3, 7, or 30 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** and navigate to the data source list view
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. Make the desired changes and 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** and navigate to the data source list view
2. Select the data sources you want to modify. Selected data sources must belong to same vendor, model and acquisition unit.
3. Click the **Actions** button and select the **Edit** option.
4. In the edit dialog, change any of the **Settings** as needed.
5. Click the **Configuration** link to change any of the basic options for the data sources.
6. Click the **Advanced Configuration** link to change any of the advanced options for the data sources.

7. Click **Save**.

Mapping data source tags to annotations

When a data source is configured to poll tag data, Insight automatically sets annotation values for an existing Insight annotation with the same name as a tag.

When the Insight annotation exists before the tags are enabled in the data source, the data source tag data is automatically added to the Insight annotation.

When you create an annotation after the tag is enabled, initial polling of the data source does not automatically update the annotation. There is a delay in the time it takes to replace or populate the Insight annotation. To avoid the delay, you can force the tag to annotation update by postponing and then resuming the data source.

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 **Admin**.

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. On the Insight toolbar, click **Admin**.
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.

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. On the Insight toolbar, click **Admin**.
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 **Admin**.
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 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 **Admin**.
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.

Patches
Brocade SSH

Summary

Recommendation: ● Approve patch - Patch results are positive (no change or more successes)

Applied on: 5/12/2013 20:00:01

Other data source affected: Brocade SHMP; Brocade HTTP

Comments: ● Got this patch from Scott. He said that this should fix the SHMP v3 problem in Brocade. Talking to John from NetApp, they promised this will fix the SHMP v3 problem. After this is applied, we still need to check the other SHMP v3 data sources and see if they are good.

You should now review the results of the patch. Approving a patch will permanently apply this patch to the system. Rolling back a patch will delete it and restore the previous version before this patch was applied. Please note that there can only be one patch active for a data source type.

[Approve](#)
[Roll back](#)
[Replace patch](#)

Affected data sources

Name	Alt	Type	Conclusion	Status before patch applied	Most recent status
ds0	local	Brocade CLI	● No change (success)	All successful	Currently polling...
ds1	local	Brocade CLI	● No change (success)	All successful	All successful
ds2	local	Brocade CLI	● Polling is now successful	Configuration failed	All successful
ds3	local	Brocade CLI	● Configuration is still failing (a different error)	Configuration failed	Configuration failed
ds4	asr1	Brocade SHMP	● Configuration is successful but now Performance is failing	Configuration failed	Performance failed

Showing 1 to 5 of 5 entries

- 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.

- 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

- On the Insight toolbar, click **Admin**.
- Click **Patches**.
- In **Patches currently under review**, click the linked name of the patch that appears to be unsuccessful.
- 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.
- To display the details for a data source that is not successfully processing the patch, click the linked **Name**.
 - Check the summary information.
 - Check the **Event timeline** to see any configuration or performance data that might be affecting this data source.
- 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.

Device resolution

You need to discover all of the devices you want to monitor with OnCommand Insight. Discovery is necessary in order to accurately track performance and inventory in your environment. Typically the majority of devices in your environment are discovered through automatic device resolution.



If you are performing an upgrade and have inactive Auto Resolution rules in the system you are upgrading from, these rules will be deleted during the upgrade. To preserve inactive Auto Resolution rules, activate the rules (check the box) before the upgrade is performed.

After you install and configure data sources, devices in your environment, including switches, storage arrays and your virtual infrastructure of hypervisors and VMs are identified. However, this does not normally identify 100% of the devices in your environment.

After data source type devices have been configured, best practice is to leverage device resolution rules to help identify the remaining unknown devices in your environment. Device resolution can help you resolve unknown devices as the following device types:

- physical hosts
- storage arrays
- tapes
- switches

Devices remaining as “unknown” after device resolution are considered generic devices, which you can also show in queries and on dashboards.

The rules created in turn will automatically identify new devices with similar attributes as they are added to your environment. In some cases, Device resolution also allows for manual identification bypassing the device resolution rules for undiscovered devices within Insight.

Incomplete identification of devices can result in issues including:

- Incomplete paths
- Unidentified multipath connections
- The inability to group applications
- Inaccurate topology views
- Inaccurate data in the Data warehouse and reporting

The Device resolution feature (**Manage > Device resolution**) includes the following tabs, each of which plays a role in device resolution planning and viewing results:

- “FC identify” contains a list WWNs and port information of Fibre Channel devices that were not resolved

through automatic device resolution. The tab also identifies the percentage of devices that have been identified.

- “IP identify” contains a list of devices accessing CIFS shares and NFS shares that were not identified through automatic device resolution. The tab also identifies the percentage of devices that have been identified.
- “Auto resolution rules” contains the list of rules that are run when performing Fibre channel device resolution. These are rules you create to resolve unidentified Fibre channel devices.
- “Preferences” provides configuration options that you use to customize device resolution for your environment.


Before you begin

You need to know how your environment is configured before you define the rules for identifying devices. The more you know about your environment the easier it will be to identify devices.

You need to answer questions similar to the following to help you create accurate rules:

- Does your environment have naming standards for zones or hosts and what percentage of these are accurate?
- Does your environment use a switch alias or storage alias and do they match the host name?
- Does your environment use an SRM tool and can you use it to identify host names? What coverage does the SRM provide?
- How often do naming schemes change in your environment?
- Have there been any acquisitions or mergers that introduced different naming schemes?


After analyzing your environment, you should be able to identify what naming standards exist that you can expect to reliably encounter. The information you gathered might be represented graphically in a figure similar to the following:



In this example the largest number of devices are reliably represented by storage aliases. Rules that identify hosts using storage aliases should be written first, rules using switch aliases should be written next , and the last rules created should use zone aliases. Due to the overlap of the use of zone aliases and switch aliases, some storage alias rules might identify additional devices, leaving less rules required for zone aliases and switch aliases.

Steps to defining devices in your environment

Typically, you would use workflow similar to the following to identify devices in your environment. Identification is an iterative process and might require multiple steps of planning and refining rules.



If you have unidentified devices (otherwise known as “unknown” or generic devices) in your environment and you subsequently configure a data source that identifies those devices upon polling, they will no longer be displayed or counted as generic devices.

Planning device resolution rules for your environment

Using rules to identify devices in your environment is typically an iterative process that requires a thorough analysis of your environment and the creation of multiple rules to identify as many devices as possible. The best case scenario is to set a goal to identify 100% of the devices in your environment.

The most efficient order for rules is to place the most restrictive rules first, resulting in most entries not pattern matching, with the process proceeding to less restrictive rules. This allows Insight to apply more patterns to each entry increasing the possibility of patterns matching and of positive host identification.

When you create rules, your objective should be to create rules that address the largest number of unidentified devices possible. For example, creating rules that follow a pattern of coverage similar to the following is far more efficient than creating 30 rules with lower percentages of coverage:

Rule	Percentage of coverage
Rule 1	60%
Rule 2	25%
Rule 3	8%
Rule 4	4%
Rule 5	1%

Creating device resolution rules

You create device resolution rules to identify hosts, storage, and tapes that are not automatically identified currently by OnCommand Insight. The rules that you create identify devices currently in your environment and also identify similar devices as they are added to your environment.

About this task

When you create rules you start by identifying the source of information that the rule runs against, the method used to extract information, and whether DNS lookup is applied to the results of the rule.

Source that is used to identify the device
<ul style="list-style-type: none"> • SRM aliases for hosts • Storage alias containing an embedded host or tape name • Switch alias containing an embedded host or tape name • Zone names containing an embedded host name
Method that is used to extract the device name from the source
<ul style="list-style-type: none"> • As is (extract a name from an SRM) • Delimiters • Regular expressions
DNS lookup
Specifies if you use DNS to verify the host name.

You create rules in the Auto Resolution Rules tab. The following steps describe the rule creation process.

Steps

1. Click **Manage > Device resolution**
2. In the **Auto resolution rules** tab, click **+Add**

The New Rule screen is displayed.



The New Rule screen includes a ? icon, that provides help and examples for creating regular expressions.

3. In the **Type** list select the device you want to identify.

You can select Host or Tape.

4. In the **Source** list, select the source you want to use to identify the host.

Depending on the source you chose, Insight displays the following response:

- Zones lists the zones and WWN that need to be identified by Insight.
- SRM lists the unidentified aliases that need to be identified by Insight
- Storage alias lists storage aliases and WWN that need to be identified by Insight
- Switch alias lists the switch aliases that need to be identified by Insight

5. In the **Method** list select the method you want to employ to identify the host.

Source	Method
SRM	"As is", "Delimiters", "Regular expressions"
Storage alias	"Delimiters", or "Regular expressions"
Switch alias	"Delimiters", or "Regular expressions"
Zones	"Delimiters", or "Regular expressions"

- Rules using "Delimiters" require the delimiters and the minimum length of the host name.


The minimum length of the host name is number of characters that Insight should use to identify a host. Insight performs DNS lookups only for host names that are this long or longer.

For rules using Delimiters, the input string is tokenized by the delimiter and a list of host name candidates is created by making several combinations of the adjacent token. The list is then sorted, largest to smallest. For example, for vipsnq03_hba3_emc3_12ep0 the list would result in the following:

- vipsnq03_hba3_emc3_12ep0
- vipsnq03_hba3_emc3
- hba3 emc3_12ep0
- vipsnq03_hba3
- emc3_12ep0

- hba3_emc3
 - vipsnq03
 - 12ep0
 - emc3
 - hba3
- Rules using “Regular expression” require a regular expression, the format, and cases sensitivity selection.

6.

Click  to run all rules, or click the down-arrow in the button to run the rule you created (and any other rules that have been created since the last full run of AR.)

Results

The results of the rule run are displayed in the FC identify tab.

Starting an automatic device resolution update

A device resolution update commits manual changes that have been added since the last full automatic device resolution run. Running an update can be used to commit and run only the new manual entries made to the device resolution configuration. No full device resolution run is performed.

Steps

1. Log into the Insight web UI.
2. Click **Manage > Device Resolution**
3. In the **Device resolution** screen, click the down-arrow in the **Run AR** button.
4. Click **Update** to start the update.

Rule assisted manual identification

This feature is used for special cases where you want to run a specific rule or a list of rules (with or without a one-time reordering) to resolve unknown hosts, storage, and tape devices or group of them.

Before you begin

You have a number of devices that have not been identified and you also have multiple rules that successfully identified other devices.

About this task



If your source only contains part of a host or device name, use a regular expression rule and format it to add the missing text.

Steps

1. Log into the OnCommand Insight web UI.

2. Click **Manage > Device resolution**

3. Click the **FC Identify** tab.

The system displays the identified and unidentified devices.

4. Select multiple unidentified devices.

5. Click **Identify > Set host resolution** or **> Set tape resolution**

The system displays the Identify screen which contains a list of all of the rules that successfully identified devices.

6. Change the order of the rules to an order that meets your needs.

The order of the rules are changed in the Identify screen, but are not changed globally.

7. Select the method that that meets your needs.

OnCommand Insight executes the host resolution process in the order in which the methods appear, beginning with those at the top.

When rules that apply are encountered, rule names are shown in the rules column and identified as manual.

Fibre Channel device resolution

The FC Identify screen displays the WWN and WWPN of Fibre Channel devices whose hosts have not been identified by automatic device resolution. The screen also displays any devices that have been resolved by manual device resolution.

Devices that have been resolved by manual resolution contain a status of “OK” and identify the rule used to identify the device. Missing devices have a status of “Unidentified”. The total coverage for identification of devices is listed on this page.

+ Add

Total coverage
30% (3/10)

FC Identify (10)

Identify

Unidentify

filter...

↑

▾

<div><div>☐</div></div>	WWN	Port WWN	IP	Name	Type	Status	Rule	
<div><div>☐</div></div>	30:E0:00:00:00:00:00	10:B0:00:00:00:00:28:20	1.1.1.1	ResolvedHost1	Host	OK	Hosts by zone	
<div><div>☐</div></div>	30:E0:00:00:00:00:00:02	10:B0:00:00:00:00:28:22	2.2.2.2	ResolvedHost2	Host	OK	Rule deleted	
<div><div>☐</div></div>	30:E0:00:00:00:00:00:03	10:B0:00:00:00:00:28:23			Unknown	Unidentified		
<div><div>☐</div></div>	30:E0:00:00:00:00:00:04	10:B0:00:00:00:00:28:24			Unknown	Unidentified		
<div><div>☐</div></div>	30:E0:00:00:00:00:00:05	10:B0:00:00:00:00:28:25			Unknown	Unidentified		

Showing 1 to 5 of 10 entries

<

1

2

>

You perform bulk actions by selecting multiple devices on the left-hand side of the FC identify screen. Actions can be performed on a single device by hovering over a device and selecting the identify or unidentify buttons on the far right of the list.

The Total coverage link displays a list of the "number of devices identified/number of devices available" for your configuration:

- SRM alias

- Storage alias
- Switch alias
- Zones
- User defined

Adding a Fibre Channel device manually

You can manually add a Fibre Channel device to OnCommand Insight using the manual add feature available in the Device resolution FC Identify tab. This process might be used for pre-identification of a device that is expected to be discovered in the future.

Before you begin

To successfully add a device identification to the system you need to know the WWN or IP address and the device name.

About this task

You can add a Host, Storage, Tape or Unknown Fibre Channel device manually.

Steps

1. Log in to the Insight web UI
2. Click **Manage > Device resolution**
3. Click the **FC Identify** tab.
4. Click the add button.

The Add Device dialog is displayed

5. Enter the WWN or IP address, the device name, and select the device type.

Results

The device you enter is added to the list of devices in the FC Identify tab. The “Rule” is identified as Manual.

Importing Fibre Channel device identification from a CSV file

You can manually import Fibre Channel device identification into OnCommand Insight Device Resolution feature using a list of devices in a CSV file.

Before you begin

You must have a correctly formatted CSV file in order to import device identifications directly into the Device Resolution feature. The CSV file for Fibre Channel devices requires the following information:

WWN
IP

Name
Type



As a best practice, it is recommended to first export the FC Identify information to a CSV file, make your desired changes in that file, and then import the file back into FC Identify. This ensures that the expected columns are present and in the proper order.

To import FC Identify information:

Steps

1. Log into the Insight web UI.
2. Click **Manage > Device Resolution**
3. Select the **FC identify** tab.
4. Click **Identify > Identify from file**
5. a. Navigate to the folder containing your CSV files for import and select the desired file.

The devices you enter are added to the list of devices in the FC Identify tab. The “Rule” is identified as “Manual”.

Exporting Fibre Channel device identifications to a CSV file

You can export existing Fibre Channel device identifications to a CSV file from the OnCommand Insight device resolution feature. You might want to export a device identification so that you can modify it and then import it back into Insight where it is then used to identify devices that are similar to those originally matching the exported identification.


About this task

This scenario might be used when devices have similar attributes that can be easily edited in the CSV file and then imported back into the system.

When you export a Fibre Channel device identification to a CSV file, the file contains the following information in the order shown:

WWN
IP
Name
Type

Steps

1. Log into the Insight web UI.
2. Click **Manage > Device Resolution**
3. Select the **FC identify** tab.
4. Select the Fibre Channel device or devices whose identification you want to export.
5. Click the export  icon.
6. Chose if you want to open the CSV file or save the file.

IP device resolution

The IP Identify screen displays any iSCSI and CIFS or NFS shares that have been identified by automatic device resolution or by manual device resolution. Unidentified devices are also shown. The screen includes the IP address, Name, Status, iSCSI node, and share name for devices. The percentage of devices that have been successfully identified is also displayed.

+ Add

Total coverage
20% (2/10)

IP identify (10)

Identify

Unidentify

filter...

↑

⌵

<input type="checkbox"/>	Address	IP	Name	Status	iSCSI node	Share name
<input type="checkbox"/>	1.1.1.1	1.1.1.1	LA3-CNS-SQL-06A	OK		/vol/ServerLogs_STG/
<input type="checkbox"/>	0.0.0.0/0					/vol/ServerLogs_STG/
<input type="checkbox"/>	10.56.100.18				iqn.1991-05.com.microsoft:la3-cns-sql-06b.cns.comcastnets.com	
<input type="checkbox"/>	10.56.100.19				iqn.1991-05.com.microsoft:jec20643597717.tfyd.com	/vol/wc_sc_libraries_prod/libraries_qtree/
<input type="checkbox"/>	100.54.18.100	100.54.18.100	ushapip000961b	OK		

Showing 1 to 5 of 10 entries

< 1 2 >

Adding IP devices manually

You can manually add an IP device to OnCommand Insight using the manual add feature available in the IP Identify screen.

Steps

1. Log in to the Insight web UI.
2. Click **Manage > Device resolution**
3. Click the **IP Identify** tab.
4. Click the add button.

The Add Device dialog is displayed

5. Enter the address, IP address, and a unique device name.

Results

The device you enter is added to the list of devices in the IP Identify tab.

Importing IP device identification from a CSV file

You can manually import IP device identifications into the Device Resolution feature using a list of device identifications in a CSV file.

Before you begin

You must have a correctly formatted CSV file in order to import device identifications. The CSV file for IP devices requires the following information:

Address
IP
Name



As a best practice, it is recommended to first export the IP Identify information to a CSV file, make your desired changes in that file, and then import the file back into IP Identify. This ensures that the expected columns are present and in the proper order.

To import IP Identify information:

Steps

1. Log into the Insight web UI.
2. Click **Manage > Device Resolution**
3. Select the **IP identify** tab.
4. Click **Identify > Identify from file**
5. a. Navigate to the folder containing your CSV files for import and select the desired file.

The devices you enter are added to the list of devices in the IP Identify tab.

Exporting IP device identification to a CSV file

You can export existing IP device identifications from Insight using the Device Resolution feature. You might want to export a device identification so that you can modify it and then import it back into Insight so that it can be used to identify devices that are similar to those in the exported identification.


About this task

When you export an IP device identification to a CSV file, the file contains the following information in the order shown:

Address
IP

Name

Steps

1. Log into the Insight web UI.
2. Click **Manage > Device Resolution**
3. Select the **IP Identify** tab.
4. Select the IP device or devices whose identification you want to export.
5. Click the export  icon.
6. Chose if you want to open the CSV file or save the file.

Setting options in the Preferences tab

The device resolution preferences tab lets you create an auto resolution schedule, specify storage and tape venders to include or exclude from identification, and set DNS lookup options.

Auto resolution schedule

An auto resolution schedule can specify when automatic device resolution is run:

Option	Description
Every	Use this option to run automatic device resolution on intervals of days, hours, or minutes.
Every day	Use this option to run automatic device resolution daily at a specific time.
Manually	Use this option to only run automatic device resolution manually.
On every environment change	Use this option to run automatic device resolution whenever there is a change in the environment.

If you specify manually, nightly automatic device resolution is disabled.

DNS processing options

DNS processing options allow you to select the following features:

- When DNS lookup result processing is enabled, you can add a list of DNS names to append to resolved devices.
- You can select "Auto resolution of IPs:" to enables automatic host resolution for iSCSI initiators and hosts accessing NFS shares by using DNS lookup. If this is not specified, only FC-based resolution is performed.
- You can choose to allow underscores in host names and to use a "connected to" alias instead of the standard port alias in results.

Including or excluding specific storage and tape vendors

You can include or exclude specific storage and tape vendors for automatic resolution. You might want to exclude specific vendors if you know, for example, that a specific host will become a legacy host and should be excluded from your new environment. You can also re-add vendors that you earlier excluded but no longer want excluded.



Device resolution rules for tape only work for WWNs where the Vendor for that WWN is set to **Included as Tape only** in the Vendors preferences.

Regular expression examples

If you have selected the regular expression approach as your source naming strategy, you can use the regular expression examples as guides for your own expressions used in the OnCommand Insight automatic resolution methods.

Formatting regular expressions

When creating regular expressions for OnCommand Insight automatic resolution, you can configure output format by entering values in a field named `FORMAT`.

The default setting is `\1`, which means that a zone name that matches the regular expression is replaced by the contents of the first variable created by the regular expression. In a regular expression, variable values are created by parenthetical statements. If multiple parenthetical statements occur, the variables are referenced numerically, from left to right. The variables can be used in the output format in any order. Constant text can also be inserted in the output, by adding it to the `FORMAT` field.

For example, you might have the following zone names for this zone naming convention:

```
[Zone number]_[data center]_[hostname]_[device type]_[interface number]
```

- S123_Miami_hostname1_filer_FC1
- S14_Tampa_hostname2_switch_FC4
- S3991_Boston_hostname3_windows2K_FC0
- S44_Raleigh_hostname4_solaris_FC1

And you might want the output to be in the following format:

```
[hostname]-[data center]-[device type]
```

To do this, you need to capture the host name, data center, and device type fields in variables, and use them in the output. The following regular expression would do this:

```
. * ? _ ( [a-zA-Z0-9] + ) _ ( [a-zA-Z0-9] + ) _ ( [a-zA-Z0-9] + ) _ . *
```

Because there are three sets of parentheses, the variables `\1`, `\2` and `\3` would be populated.

You could then use the following format to receive output in your preferred format:

```
\2-\1-\3
```

Your output would be as follows:

```
hostname1-Miami-filer
hostname2-Tampa-switch
hostname3-Boston-windows2K
hostname4-Raleigh-solaris
```

The hyphens between the variables provide an example of constant text that is inserted in the formatted output.

Example 1 showing zone names

In this example, you use the regular expression to extract a host name from the zone name. You could create a regular expression if you have something similar to the following zone names:

- S0032_myComputer1Name-HBA0
- S0434_myComputer1Name-HBA1
- S0432_myComputer1Name-HBA3

The regular expression that you could use to capture the host name would be:

```
S[0-9]+_([a-zA-Z0-9]*)[_-]HBA[0-9]
```

The outcome is a match of all zones beginning with S that are followed by any combination of digits, followed by an underscore, the alphanumeric hostname (myComputer1Name), an underscore or hyphen, the capital letters HBA, and a single digit (0-9). The hostname alone is stored in the `\1` variable.

The regular expression can be broken into its components:

- "S" represents the zone name and begins the expression. This matches only an "S" at the beginning of the zone name.
- The characters [0-9] in brackets indicate that what follows "S" must be a digit between 0 and 9, inclusive.
- The + sign indicates that the occurrence of the information in the preceding brackets has to exist 1 or more times.
- The _ (underscore) means that the digits after S must be followed immediately by only an underscore character in the zone name. In this example, the zone naming convention uses the underscore to separate the zone name from the host name.
- After the required underscore, the parentheses indicate that the pattern contained within will be stored in the `\1` variable.
- The bracketed characters [a-zA-Z0-9] indicate that the characters being matched are all letters (regardless of case) and numbers.

- The * (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.
- The bracketed characters [_-] (underscore and dash) indicate that the alphanumeric pattern must be followed by an underscore or a dash.
- The letters HBA in the regular expression indicate that this exact sequence of characters must occur in the zone name.
- The final set of bracketed characters [0-9] match a single digit from 0 through 9, inclusive.

Example 2

In this example, skip up to the first underscore "", *then match E and everything after that up to the second ""*, and then skip everything after that.

Zone: Z_E2FHDBS01_E1NETAPP

Hostname: E2FHDBS01

RegExp: . ? (E. ?) . * ?

Example 3

The parentheses "(")" around the last section in the Regular Expression (below) identifies which part is the hostname. If you wanted VSAN3 to be the host name, it would be: _([a-zA-Z0-9]).*

Zone: A_VSAN3_SR48KENT_A_CX2578_SPA0

Hostname: SR48KENT

RegExp: _[a-zA-Z0-9]+_([a-zA-Z0-9]).*

Example 4 showing a more complicated naming pattern

You could create a regular expression if you have something similar to the following zone names:

- myComputerName123-HBA1_Symm1_FA3
- myComputerName123-HBA2_Symm1_FA5
- myComputerName123-HBA3_Symm1_FA7

The regular expression that you could use to capture these would be:

```
([a-zA-Z0-9]*)_.*
```

The \1 variable would contain only myComputerName123 after being evaluated by this expression.

The regular expression can be broken into its components:

- The parentheses indicate that the pattern contained within will be stored in the \1 variable.
- The bracketed characters [a-zA-Z0-9] mean that any letter (regardless of case) or digit will match.
- The * (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.

- The `_` (underscore) character in the regular expression means that the zone name must have an underscore immediately following the alphanumeric string matched by the preceding brackets.
- The `.` (period) matches any character (a wildcard).
- The `*` (asterisk) indicates that the preceding period wildcard may occur 0 or more times.

In other words, the combination `.*` indicates any character, any number of times.

Example 5 showing zone names without a pattern

You could create a regular expression if you have something similar to the following zone names:

- `myComputerName_HBA1_Symm1_FA1`
- `myComputerName123_HBA1_Symm1_FA1`

The regular expression that you could use to capture these would be:

```
(.*?)_.*
```

The `\1` variable would contain *myComputerName* (in the first zone name example) or *myComputerName123* (in the second zone name example). This regular expression would thus match everything prior to the first underscore.

The regular expression can be broken into its components:

- The parentheses indicate that the pattern contained within will be stored in the `\1` variable.
- The `.*` (period asterisk) match any character, any number of times.
- The `*` (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.
- The `?` character makes the match non-greedy. This forces it to stop matching at the first underscore, rather than the last.
- The characters `_.*` match the first underscore found and all characters that follow it.

Example 6 showing computer names with a pattern

You could create a regular expression if you have something similar to the following zone names:

- `Storage1_Switch1_myComputerName123A_A1_FC1`
- `Storage2_Switch2_myComputerName123B_A2_FC2`
- `Storage3_Switch3_myComputerName123T_A3_FC3`

The regular expression that you could use to capture these would be:

```
.*?_.*?_([a-zA-Z0-9]*[ABT])_.*
```

Because the zone naming convention has more of a pattern, we could use the above expression, which will match all instances of a hostname (*myComputerName* in the example) that ends with either an A, a B, or a T, placing that hostname in the `\1` variable.

The regular expression can be broken into its components:

- The .* (period asterisk) match any character, any number of times.
- The ? character makes the match non-greedy. This forces it to stop matching at the first underscore, rather than the last.
- The underscore character matches the first underscore in the zone name.
- Thus, the first .*?_ combination matches the characters *Storage1_* in the first zone name example.
- The second .*?_ combination behaves like the first, but matches *Switch1_* in the first zone name example.
- The parentheses indicate that the pattern contained within will be stored in the \1 variable.
- The bracketed characters [a-zA-Z0-9] mean that any letter (regardless of case) or digit will match.
- The * (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.
- The bracketed characters in the regular expression [ABT] match a single character in the zone name which must be A, B, or T.
- The _ (underscore) following the parentheses indicates that the [ABT] character match must be followed up an underscore.
- The .* (period asterisk) match any character, any number of times.

The result of this would therefore cause the \1 variable to contain any alphanumeric string which:

- was preceded by some number of alphanumeric characters and two underscores
- was followed by an underscore (and then any number of alphanumeric characters)
- had a final character of A, B or T, prior to the third underscore.

Example 7

Zone: myComputerName123_HBA1_Symm1_FA1

Hostname: myComputerName123

RegExp: ([a-zA-Z0-9]+)_.*

Example 8

This example finds everything before the first _.

Zone: MyComputerName_HBA1_Symm1_FA1

MyComputerName123_HBA1_Symm1_FA1

Hostname: MyComputerName

RegExp: (.*?)_.

Example 9

This example finds everything after the 1st _ and up to the second _.

Zone: Z_MyComputerName_StorageName

Hostname: MyComputerName

RegExp: .? (. ?) . * ?

Example 10

This example extracts "MyComputerName123" from the zone examples.

Zone: Storage1_Switch1_MyComputerName123A_A1_FC1

Storage2_Switch2_MyComputerName123B_A2_FC2

Storage3_Switch3_MyComputerName123T_A3_FC3

Hostname: MyComputerName123

RegExp: .? . ? ([a-zA-Z0-9]+) [ABT] _ .

Example 11

Zone: Storage1_Switch1_MyComputerName123A_A1_FC1

Hostname: MyComputerName123A

RegExp: .? . ? ([a-zA-Z0-9]+) . * ?

Example 12

The ^ (circumflex or caret) **inside square brackets** negates the expression, for example, [^Ff] means anything except uppercase or lowercase F, and [^a-z] means everything except lowercase a to z, and in the case above, anything except the _ . The format statement adds in the "-" to the output host name.

Zone: mhs_apps44_d_A_10a0_0429

Hostname: mhs-apps44-d

RegExp: ` ([^_]) _ ([AB]) . * ` `Format in OnCommand Insight: ` \1 - \2 () _

([^_]) _ () . * `Format in OnCommand Insight: \1 - \2 - \3

Example 13

In this example, the storage alias is delimited by "\" and the expression needs to use "\\" to define that there are actually "\" being used in the string, and that those are not part of the expression itself.

Storage Alias: \Hosts\E2DOC01C1\E2DOC01N1

Hostname: E2DOC01N1

RegExp: \\ . ? \\ . ? \\ (. * ?)

Example 14

This example extracts "PD-RV-W-AD-2" from the zone examples.

Zone: PD_D-PD-RV-W-AD-2_01

Hostname: PD-RV-W-AD-2

RegExp: [^_]- (.-\d+) .+

Example 15

The format setting in this case adds the "US-BV-" to the hostname.

Zone: SRV_USBVM11_F1

Hostname: US-BV-M11

RegExp: SRV_USBV([A-Za-z0-9]+)_F[12]

Format: US-BV-\1

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

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 shows the total raw capacity for storage by each vendor.
 - Capacity by Tier shows the total useable capacity for each storage tier.

- Switch Ports pie chart shows the manufacturers of ports and shows the percentage of ports used.
- 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 Insight toolbar, click **Admin**.
- 7. Note any areas showing solid red circles.

In the OnCommand Insightweb UI, potential problems are marked with a solid red circle.

- 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 should periodically check the current status of your Insight system components by viewing the health page, which shows the status of each component and alerts you when there is an issue.

Steps

- 1. Log in to the Insightweb UI.
- 2. Click **Admin** and select **Health**.

The Health page is displayed.

- 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 any or all of the following Insight components based on your system configuration:

Component	Test	Details	Displays
-----------	------	---------	----------

Acquisition	Inventory data processing	Status of local acquisition unit	<p>“OK” if number of concurrently-polling data sources is less than 75% of execution pool maximum (default maximum is 30).</p> <p>“Acquisition is busy” if usage is greater than 75%, and recommends increasing polling interval or adding more remote acquisition units.</p>
DWH	Backup	Status of Data Warehouse scheduled backup	<p>“OK” and the last successful DWH backup time if DWH scheduled backup is enabled.</p> <p>Otherwise, displays information about any error found.</p>
DWH	ETL	Status of Data Warehouse ETL	<p>“OK” and the last successful DWH build time if no errors.</p> <p>Otherwise, displays information about any error found.</p>
Server	ASUP	Status of ASUP	<p>“ASUP Enabled” and the last successful phonehome time if available. “ASUP Failed” if phonehome is enabled but encountered a problem.</p> <p>+ “Invalid backup location” if backup directory is not valid.</p> <p>+ Displays the last successful phonehome time as well as time of the last failed attempt if available.</p> <p>+ “ASUP Disabled” if phonehome is disabled.</p>

Server	Auto resolution	Status of automatic device resolution	<p>“OK” if no errors. “Auto resolution is blocked” if identification errors prevent resolution progress.</p> <p>+ “Low success rate” if less than 75% of generic devices could be identified.</p>
Server	Elasticsearch	Status of elastic search data store	<p>“OK” if no errors. “Service unavailable” if unable to connect to elastic search service.</p> <p>+ "Cluster mode detected" if more than one node is detected.</p> <p>+ "High memory utilization" if heap space used is more than 85%.</p> <p>+ "Status: RED" indicates an error reported by elastic search. Displays information about the error and recommends contacting customer support.</p>
Server	CPU	Insight CPU usage	<p>“OK” if CPU load is less than 65%. “System CPU load is high. Reduce your CPU load.” if CPU load is greater than 65%.</p>
Server	Disk space	Status of disk space	<p>Free disk space, disk space in use by Insight, and recommended disk space reserved for Insight. “Low Disk Space” if disk utilization is more than 80%.</p>

Server	EventBus	Status of EventBus	“EventBus is empty” if EventBus queue is empty, otherwise displays status of EventBus queue.
Server	Inventory data processing	Status of inventory data processing capability of Insight server	“OK” if Insight server is not busy. “Server is busy” if the server is busy at least 75% of the time for the last hour. Recommends not adding more data sources and recommends splitting the environment to several servers.
Server	MySQL	Status of MySQL database	“OK” if no problems are detected. “The database is having performance issues. Some queries are taking too long to run” if the number of slow queries is more than 5%. + “The database log file grew more than <size> in the past hour. Check MySQL log file” if the error log grows to more than 20 KB.
Server	Performance archive	Status of performance archive	“Performance archive is enabled” or “Performance archive is not enabled”.
Server	Physical memory	Status of physical memory	“OK” if memory usage is less than 85%. “Memory usage is high. Reduce your overall memory footprint for system stability” if memory usage is greater than 85%.

Server	Service pack	Service pack availability	Displays whether a service pack is available for Insight. If a service pack is available, displays instructions.
Server	Usage information	Status of sending of usage information	<p>Displays whether sending of usage information to NetApp is enabled or disabled. Recommends enabling if disabled. Displays last attempted or last successful send time.</p> <p>+ Displays information on any problems encountered.</p>
Server	Violation	Status of open violations	<p>“OK” if the number of open violations is less than 75% of the violations limit. “Maximum number of open violations allowed is <number>” if the number of open violations is greater than 75% of the violations limit. Recommends reviewing performance policy configuration.</p> <p>+ “Violation manager is blocked” if the number of open violations is at the violations limit.</p> <p>+ Note that the violation manager cannot create new violations and recommends reviewing performance policy configuration.</p>
Server	Weekly backup	Status of weekly backup	“OK” if weekly backup is enabled, otherwise displays “Weekly backup is not enabled”.

Deleting inactive devices

Deleting devices that are inactive helps keep your data cleaner and easier to navigate.

About this task

To delete inactive devices from Insight, do the following:

Steps

1. Create a new query or open an existing query.
2. Choose either the *generic device*, *host*, *storage*, *switch*, or *tape* asset type.
3. Add a filter for **Is active**, and set the filter to **No**.

The results table displays only assets that are not active.

4. Select the devices that you want to delete.
5. Click the **Actions** button and select **Delete Inactive Devices**.

Your inactive devices are deleted and will no longer be displayed in Insight.

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. Audit log messages can optionally be sent to syslog in addition to being displayed on the Audit page.

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
- Create a user
- Delete a user
- User role change

- Modify a user (Guest à Admin)
- Logout of a user (either forced logout or manual logout)
- Deleting an acquisition unit
- Update License
- Enabling backup
- Disabling Backup
- Enabling ASUP (Enabling Proxy on same page is reported in audit log)
- Disabling ASUP (Disabling Proxy on same page is reported in audit log)
- Security - re-key, change system passwords.
- Removing/adding annotations on assets
- CAC user logon / logoff
- CAC user session timeout

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.



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.

To see only the audit entries by the user `izzzyk`, type `izzzyk` 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, 7d, and 30d**), 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.

Acquisition unit status

The Acquisition Unit screen provides a view of all your acquisition units, including status and any errors present.

The status of the Insight acquisition units connected to your server is displayed in the **Admin > Acquisition Units** table. This table displays the following information for each acquisition unit:

- **Name**
- **IP**
- **Status** is the operating status of the acquisition unit.
- **Last reported** displays the last time a data source connected to the acquisition unit reported.
- **Note** displays a user-entered note related to the AU.

If an acquisition unit in the list has a problem, the Status field will show a red circle with brief information about the problem. You should investigate any acquisition unit problems, as they likely affect data collection.

To restart an acquisition unit, hover over the unit and click on the *Restart Acquisition Unit* button that appears..

To add a text note, hover over an acquisition unit and click the *Add Note* button that appears. Only the most recently entered note is displayed.

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 **Admin**.
2. Click **Troubleshooting**.

The screenshot shows the 'Send / Collect data' section with three buttons: 'Back up', 'Bundle logs', and 'Send ASUP now'. Below this is the 'Restore a database' section with a 'Select backup' dropdown menu, a 'No file selected' status, and a 'Restore' button. A warning message states: 'Warning: Your current database will be discarded!'. At the bottom is the 'Other tasks' section with two links: 'Connect to the old OnCommand Insight Portal' and 'Open the scrub utilities'.

Action	Description
Back up	Back up the database (configuration and performance) into a ZIP file.
Bundle logs	Collect all log files (including acquisition recordings) and bundle them into a ZIP file. Can be used to send data back to NetApp support when troubleshooting an issue with the software.
Send ASUP now	Forces an ad-hoc ASUP report. Can be used to allow NetApp support to get the latest support data when troubleshooting an issue with the software.

Restore a database

Select backup ▼ No file selected Restore

Warning: Your current database will be discarded!

Other tasks

Couldn't find what you are looking for? Connect to the old OnCommand Insight Portal

Need to send anonymous data back? Open the scrub utilities.

3. In the Restore a database section, select the backup file you want to restore from the **Select Backup** menu.
4. Click **Restore**.
5. On the warning that all data will be replaced, click **OK**

The status of the restore activity is displayed on the restore page.

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 on **Admin** 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**.
10. If you are using the Insight consumption licensing model, you must check the box to **Enable sending usage information to NetApp** in the usage section. Proxy must be properly configured and enabled for your environment.

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. On the Insight toolbar, click **Admin**.
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 **Admin**.
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.

Send / Collect data

Action	Description
<input type="button" value="Back up"/>	Back up the database (configuration and performance) into a ZIP file.
<input type="button" value="Bundle logs"/>	Collect all log files (including acquisition recordings) and bundle them into a ZIP file. Can be used to send data back to NetApp support when troubleshooting an issue with the software.
<input type="button" value="Send ASUP now"/>	Forces an ad-hoc ASUP report. Can be used to allow NetApp support to get the latest support data when troubleshooting an issue with the software.

Restore a database

Warning: Your current database will be discarded!

Other tasks

Couldn't find what you are looking for? Connect to the old [OnCommand Insight Portal](#).

Need to send anonymous data back? Open the [scrub utilities](#).

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 **Scrub utilities** link.

There are several scrub sections: Lookup in Dictionary, Scrub data, and Build dictionary, Custom keywords, and Regular expressions.

3. a. In the **Lookup in dictionary** section, Enter a code to display the value it replaces, or enter a value to see the code that replaces it. Note: before you can do a lookup, you must **Build** the dictionary to identify values to scrub from the support data.
4. To add your own keywords to scrub from the support data, in the **Custom keywords** section, click **Actions** › **Add custom keyword**. Enter a keyword and click **Save**. The keyword is added to the dictionary.
5. Expand **Patterns (regex)**. Click **Add** to get the dialog box for entering a new pattern.
6. To use a regular expression to identify words or phrases to scrub, enter a pattern or patterns in the **Regular expressions** section. Click **Actions** › **Add regular expression**, enter a Name for the pattern and the Regular expression in the fields and click **Save**. The information has been added to the dictionary.



Patterns must be encompassed by round parentheses to identify a regular expression capturing group.

7. In the **Build dictionary** section, click **Build** to initiate compilation of the dictionary of 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, no other keywords will be deleted. Return to the Scrub utilities and enter your Custom Keywords again.

8. After you create a Scrub dictionary, you can use it to scrub a log, XML, or other text file to make the data anonymous.
9. To scrub a log, XML, or other text file, in the **Scrub data** section, Browse to locate the file and click **Scrub file**.

Advanced troubleshooting

To complete your OnCommand Insight configuration, you must use the advanced troubleshooting tools. These tools run in the browser and are opened from the **Admin > Troubleshooting** page.

To open the advanced troubleshooting tools in the browser, click the **Advanced Troubleshooting** link at the bottom of the page.

The advanced troubleshooting tools allow you to view various reports, system information, installed packages, and logs, as well as perform numerous actions such as restarting the server or acquisition units, update DWH annotations, and import annotations.

See the Advanced Troubleshooting page for all available options.

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. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **Advanced Troubleshooting** link.
3. Click the **Advanced settings** tab, in the Acquisition Dynamic Attributes section enter the number of hours that OnCommand Insight should ignore dynamic data for Acquisition Dynamic Attributes.
4. Click **Save**.
5. (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. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click **Advanced Troubleshooting**.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click the **Logs** tab and select the log file to download.

A dialog box opens allowing you to open the log or save the log locally.

Displaying system information

You can display the Microsoft Windows IP configuration information about the system on which OnCommand Insight server is deployed.

Steps

1. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **Advanced Troubleshooting** link.
3. On the Advanced Troubleshooting page, click the **Reports** tab.
4. Click **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. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **Advanced Troubleshooting** link.
3. On the Advanced Troubleshooting page, click the **Reports** tab.
4. Click **Installed Software Packages**.

Calculating the number of database objects

To determine the number of objects in the OnCommand Insight database, use the Calculate Scale feature.

Steps

1. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **Advanced Troubleshooting** link.
3. On the Advanced Troubleshooting page, click the **Reports** tab.
4. Click **Calculated Scale**.

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



Both of these options should only be used upon request by NetApp Customer Support. There is no confirmation prior to restart.

Steps

1. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **Advanced Troubleshooting** link.
3. On the next page in the Advanced menu, click the **Actions** tab.
4. Click **Restart Server**.

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. Click **Advanced Troubleshooting**.
3. Select the **Actions** tab
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**.

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.

Steps

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. On the bottom of the page, click the **Advanced Troubleshooting** link.
3. Click the **Actions** tab.
4. Select one of these options:
 - **Update DWH Annotations** to force the update of annotations in data warehouse to be used for reports.
 - **Update DWH Annotations (incl. deleted)** to force an annotations update (including deleted objects) in the data warehouse to be used for reports.

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. On the Insight toolbar, 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 **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.

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 **Advanced Troubleshooting** link.

3. On the **Reports** tab, click the **Ghost Data Sources** link.

OnCommand Insight produces a list of originators with their device information.

Adding a missing disk model

If acquisition fails due to an unknown disk model, you can add the missing disk model to the `new_disk_models.txt` file and run acquisition again.

About this task

As part of a poll of a storage device by OnCommand Insight acquisition, the disk models on the storage device are read. If a vendor has added new disk models to their array that Insight doesn't know about, or if there is a mismatch between the model number Insight looks for and the one returned by the storage device, acquisition of that data source will fail with an error. In order to prevent these errors, it is necessary to update the disk model information known to Insight. New disk models are added to Insight with updates, patches and maintenance releases. However, you may decide to update this information manually instead of waiting for a patch or update.

Because OnCommand Insight reads the disk model file every five minutes, any new data model information you enter is updated automatically. You do not need to restart the server for the changes to take effect, but you can opt to restart the server and any remote acquisition units (RAUs) to have the changes take effect before the next update.

Disk model updates are added to the `new_disk_models.txt` file located in the `<SANScreenInstallDir>\wildfly\standalone\deployments\datasources.war` directory. Understand the information needed to describe your new disk model before updating the `new_disk_models.txt` file. Inaccurate information in the file produces incorrect system data and could result in failed acquisition.

Follow these instructions to manually update Insight disk models:

Steps

1. Locate the proper information for your disk model.
2. Using a text editor, open the `new_disk_models.txt` file.
3. Add the required information for the new data source.
4. Save the file in the `<SANScreenInstallDir>\wildfly\standalone\deployments\datasources.war` directory on your server.
5. Back up the `new_disk_models.txt` file to a safe location. During any subsequent OnCommand Insight upgrade, this file will be overwritten. If your disk model information is not present in the upgraded file, you will need to re-enter it.

Locating required information for new disk model

To locate the disk model information, identify the vendor and model number and run an Internet search.

About this task

Locating disk model information is as simple as running an internet search. Be sure to note the vendor name and disk model number before searching.

Steps

1. It is recommended to use an advanced internet search for the vendor, model, and document type “PDF” to find the vendor’s data sheet and/or installation guide for the drive. These data sheets are usually the best source for vendor disk information.
2. 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.
3. Locate 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 You can use the following table as a guide to help note this information as you find it:

For this field:	Which is:	Enter this:
Model number (aka Key)	Required	
Vendor	Required	
Disk speed (RPM)	Required	
Size (in GB)	Required	
Interface Type (select one)	Required	ATA, SATA, SATA2, SATA3, FC, SAS, FATA, SSD, OTHER
Seek time in ms	Optional	
Maximum transfer rate in MB/sec	Optional	
Interface transfer rate in MB/sec	Optional	
Link to vendor/model information	Optional but recommended	

4. Enter that information into the `new_disk_models.txt` file. See [Content of the new_disk_models.txt file](#) for format, order, and examples.

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.

All fields are required except for seek time, transfer rates and `additional_info`. If available, include the vendor/model website link in the `additional_info` field.

Using a text editor, enter the following information in this order, separated by commas, for each new disk model you wish to add:

1. **key**: use the model number (required)
2. **vendor**: name (required)
3. **model number**: full number (usually the same value as in "key") (required)
4. **rpm of the disk**: for example 10000 or 15000 (required)
5. **size**: capacity in GB (required)
6. **interface type**: ATA, SATA, FC, SAS, FATA, SSD, OTHER (required)
7. **seek time**: in ms (optional)
8. **potential transfer rate**: the potential transfer rate in MB/sec. Maximum transfer rate of the disk itself. (optional)
9. **interface transfer rate**: the rate to and from the host in MB/sec (optional).
10. **Additional Info**: Any additional information you want to capture. Best practice is to enter the link to the vendor page where the specs are found, for reference (optional)

For any optional fields left blank, be sure to include the comma.

Examples (each on one line with no spaces):

```
ST373405,Seagate,ST373405,10000,73,FC,5.3,64,160,http://www.seagate.com/staticfiles/support/disc/manuals/enterprise/cheetah/73(LP)/100109943e.pdf
```

```
SLR5B-M400SS,HITACHI,SLR5B-M400SS,1000000,400,SSD,,,,,
```

```
X477_THARX04TA07,TOSHIBA,X477_THARX04TA07,7200,4000,SATA,9.5,,,https://storage.toshiba.eu/export/sites/toshiba-sdd/media/products/datasheets/MG03ACAxxy.pdf
```

Monitoring your environment

Insight helps you to prevent problems in your environment and troubleshoot potential problems quickly.

Asset page data

Asset pages provide performance troubleshooting data and present summary information about a base asset (such as a virtual machine or a volume) and the related assets it uses (such as storage pools, storage nodes, and connected switch ports), with links to additional information.

Beginning with OnCommand Insight 7.3.1, all asset pages have a **Main** page and an **Additional data** page. On the Main page are a summary of the asset and different sections for charts, topology and other information. The **Additional data** page allows you to configure a customizable dashboard page for the current asset type.

A solid red circle next to a line or message on the asset page main tab indicates potential issues with the monitored environment.

Types of asset pages

Asset pages summarize the current status of an asset and contain links to additional information about the asset and its related assets.

OnCommand Insight provides asset pages for the following assets:

- Virtual machine
- Volume
- Internal volume
- Physical host
- Storage pool
- Storage
- Datastore
- Hypervisor
- Application
- Storage node
- Qtree
- Disk
- VMDK
- Port
- Switch
- Fabric
- Object storage (for example, Atmos, Centera, Amazon S3)
- Zone

Mapping and Masking information can be viewed in tables on Zone, Volume, VM, and Host/Hypervisor asset pages.



Summary information is available for object storage assets; however, you can only access this information from the Data sources detail page.

Searching your environment for specific assets

You can locate information about specific assets by using the search facility. For example, if a system user contacts the storage administrator with a complaint about a particular server, the administrator can search the server name and display an asset page summarizing the status and supplying additional linked information.

Steps

1. Open the OnCommand Insightweb UI.
2. On the toolbar, click .

The **Search assets** box is displayed.

3. Enter the name of an asset or a portion of the name.
4. Select the resource you want from the search results.

The asset page for that resource is displayed.

Advanced search techniques

Multiple search techniques can be used to search for data or objects in your monitored environment.

Wildcard search

You can perform multiple character wildcard search using the * character. For example, *applic*n* would return application.

Phrases used in search

A phrase is a group of words surrounded by double quotation marks; for example, "PAW VNX LUN 5". You can use double quotes to search for documents that contain spaces in their names or attributes.

Boolean Operators

Using Boolean operators, you can combine multiple terms to form a more complex query.

- **OR**

- The OR operator is the default conjunction operator.

If there is no Boolean operator between two terms, the OR operator is used.

- The OR operator links two terms and finds a matching document if either of the terms exists in a document.

For example, "storage OR netapp" searches for documents that contain either "storage" or "netapp".

- High scores are given to documents that match most of the terms.

- **AND**

You can use the AND operator to find documents in which both the search terms exist in a single document. For example, "aurora AND netapp" searches for documents that contain both "storage" and "netapp".

You can use the symbol && instead of the word AND.

- **NOT**

When you use the NOT operator, all the documents that contain the term after NOT are excluded from the search results. For example, "storage NOT netapp" searches for documents that contains only "storage" and not "netapp".

You can use the symbol ! instead of the word NOT.

Prefix and suffix search

- As soon as you start typing a search string, the search engine does a prefix and suffix search to find the best match.
- Exact matches are given a higher score than a prefix or suffix match. The score is calculated based on the distance of the search term from the actual search result. For example, we have three storages: “aurora”, “aurora1”, and “aurora11”. Searching for “aur” will return all three storages. However, the search result for “aurora” will have the highest score because it has the closest distance to the prefix search string.
- The search engine also searches for terms in reverse order, which allows you to perform a suffix search. For example, when you type “345” in the search box, the search engine searches for “345”.
- Search is case-insensitive.

Search using indexed terms

Searches that match more of the indexed terms result in higher scores.

The search string is split into separate search terms by space. For example, the search string “storage aurora netapp” is split into three keywords: “storage”, “aurora”, and “netapp”. The search is performed using all three terms. The documents that match most of these terms will have the highest score. The more information you provide, the better are the search results. For example, you can search for a storage by its name and mode.

The UI displays the search results across categories, with the three top results per category. If you did not find a document that you were expecting, you can include more terms in the search string to improve the search results.

The following table provides a list of indexed terms that can be added to the search string.

Category	Indexed terms
Storage	<ul style="list-style-type: none">• “storage”• name• vendor• model
StoragePool	<ul style="list-style-type: none">• “storagepool”• name• name of the storage• IP addresses of the storage• serial number of the storage• storage vendor• storage model• names for all associated internal volumes• names for all associated disks

Internal Volume	<ul style="list-style-type: none"> • “internalvolume” • name • name of the storage • IP addresses of the storage • serial number of the storage • storage vendor • storage model • name of the storage pool • names of all associated shares • names of all associated applications and business entities
Volume	<ul style="list-style-type: none"> • “volume” • name • label • names of all internal volumes • name of the storage pool • name of the storage • IP addresses of the storage • serial number of the storage • storage vendor • storage model
Storage Node	<ul style="list-style-type: none"> • “storagenode” • name • name of the storage • IP addresses of the storage • serialnumber of the storage • storage vendor • storage model
Host	<ul style="list-style-type: none"> • “host” • name • IP addresses • names of all associated applications and business entities

Datastore	<ul style="list-style-type: none"> • “datastore” • name • virtual center IP • names of all volumes • names of all internal volumes
Virtual Machines	<ul style="list-style-type: none"> • “virtualmachine” • name • DNS name • IP addresses • name of the host • IP addresses of the host • names of all datastores • names of all associated applications and business entities
Switches (regular and NPV)	<ul style="list-style-type: none"> • “switch” • IP address • wwn • name • serial number • model • domain ID • name of the fabric • wwn of the fabric
Application	<ul style="list-style-type: none"> • “application” • name • tenant • line of business • business unit • project
Tape	<ul style="list-style-type: none"> • “tape” • IP address • name • serial number • vendor

Port	<ul style="list-style-type: none"> • “port” • wwn • name
Fabric	<ul style="list-style-type: none"> • “fabric” • wwn • name


Changing the time range of displayed data

By default, an asset page displays the last 24 hours of data; however, you can change the segment of data displayed by selecting another fixed time or a custom range of time to view less or more data.

About this task

You can change the time segment of displayed data by using an option that is located on every asset page, regardless of asset type.

Steps

1. Log in to the OnCommand Insightweb UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it.
The asset page displays.
3. In the upper left corner of the page, click any of the following time icons to change the segment of data displayed:
 - **3h**

Displays the last three hours of data.
 - **24h**

Displays the last 24 hours of data.
 - **3d**


Displays the last three days of data.
 - **7d**

Displays the last seven days of data.
 - **30d**

Displays the last thirty days of data.

- **Custom**

Displays a dialog box that enables you to choose a custom range of time. You may display up to 31 days of data at a time.

4. If you chose **Custom**, do the following:
 - a. Click the date field, and select a month, day, and year for the begin date.
 - b. Click the time list, and select a begin time.
 - c. Repeat both steps a and b for the end data and time.
 - d. Click .

Determining data source acquisition status



Because data sources are the primary source of information for Insight, it is imperative that you ensure that they remain in a running state.

The ability to see the data source acquisition status is available on every asset page for all assets that are directly acquired. Either of the following acquisition scenarios can occur, in which the status is displayed in the upper right corner of the asset page:

- Acquired successfully from data source

Displays the status “Acquired xxxx”, where xxxx indicates the most recent acquisition time of the asset’s data sources.

- There is an acquisition error.

Displays the status “Acquired xxxx”, where xxxx indicates the most recent acquisition time of the asset’s one or more data sources with . When you click , a window displays each data source for the asset, the data source’s status, and the last time data was acquired. Clicking a data source displays the data source’s detail page.

If an asset is not directly acquired, no status is displayed.

Asset page sections

An asset page displays several sections containing information relevant to the asset. The sections that you see depend on the type of asset.

Summary

The Summary section on an asset page displays a summary of information about the particular asset and shows issues related to the asset, indicated by a red circle, with hyperlinks to additional information about related assets and to any performance policies assigned to the asset.

The following example shows some of the types of information available in the Summary section of an asset page for a virtual machine. Any item with a solid red circle next to it indicates potential issues with the monitored environment.


Summary

Power state:	On
Guest state:	Running
Datastore:	DS_SP1_1
CPU:	41.05%
Memory:	● 51% (1,047 / 2,048 MB)
Capacity:	10% (19.5 / 195.3 GB)
Latency:	1.93 ms (6.00 ms max)
IOPS:	1,317.33 IO/s (4,964.00 IO/s max)
Throughput:	38.79 MB/s (142.00 MB/s max)
DNS name:	VM_Cs_travBookcomp.com
IP:	10.97.133.23
OS:	Microsoft Windows Server 2008 R2(64-bit)
Processors:	4
FC Fabrics Connected:	1
Performance Policies:	VM Latency-Critical VM Latency-Warning Comp Corp.Customer Support SLA latency ● Exchange SLO

Using the Summary section

You can view the Summary section to see general information about an asset. Specifically, it is helpful to see if any metrics (for example, memory, capacity, and latency) or any performance policies are cause for concern, which OnCommand Insight indicates by displaying a red circle next to the metric or performance policy.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it. The asset page displays.



The information that displays in the Summary section depends on the type of asset page you are viewing.

3. You can click any of the asset links to view their asset pages.

For example, if you are viewing a storage node, you can click a link to view the asset page of the storage it is associated with or click to view the asset page of the HA partner.

4. You can view the metrics associated with the asset.

A red circle next to a metric indicates that you might need to diagnose and resolve potential problems.



You may notice that volume capacity might show greater than 100% on some storage assets. This is due to metadata related to the capacity of the volume being part of the consumed capacity data reported by the asset.

5. If applicable, you can click a performance policy link to view the performance policy or policies associated with the asset.


If a red circle appears next to a performance policy, this indicates an asset has crossed the performance policy's defined threshold. You should examine the performance policy to further diagnose the issue.

Topology

The Topology section, if applicable to an asset, enables you to see how a base asset is connected to its related assets.

The following shows an example of what might display in the Topology section of a virtual machine asset page.

Topology




If the topology for the asset is larger than will fit in the section, the **Click link to see the topology** hyperlink is displayed instead.

Using the Topology section

The Topology section enables you to view how the assets in your network are connected to each other and display information about related assets.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it.
The asset page displays. You can find the Topology section in the upper right-hand corner of the asset page.

If the topology for the asset is larger than will fit in the section, click the **Click link to see the topology** hyperlink.



3. To view more information about the base asset's related assets, position your cursor over a related asset in the topology and click its name, which displays its asset page.

User Data

The User Data section of an asset page displays and enables you to change any user-defined data such as applications, business entities, and annotations.

The following shows an example of what might display in the User Data section of a virtual machine asset page when an application, business entity, and annotation are assigned to the asset:




User Data

Application(s):	Concur
Business Entities:	Hybridsoft Corporation.Sales.Wes...
Birthday:	<input type="text" value="01/30/2016"/>  
+ Add	

Using the User Data section to assign or modify applications

You can assign applications running in your environment to certain assets (host, virtual machines, volumes, internal volumes, and hypervisors). The User Data section enables you to change the application assigned to an asset or assign an application or additional applications to an asset.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it. The asset page displays.
3. You can do the following:
 - To view the asset page for the application, click the application's name.
 - To change the application assigned or to assign an application or additional applications, position your cursor over the application name, if an application is assigned, or over **None**, if no application is assigned, click , type to search for an application or select one from the list, and then click .




If you choose an application that is associated with a business entity, the business entity is automatically assigned to the asset. In this case, when you place your cursor over the business entity name, the word *derived* displays. 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.

- To remove an application, click .

Using the User Data section to assign or modify business entities

You can define business entities to track and report on your environment data at a more granular level. The User Data section in an asset page enables you to change the business entity assigned to an asset or remove a business entity from an asset.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it.
The asset page displays.
3. You can do the following:
 - To change the entity assigned or to assign an entity, click  and select an entity from the list.
 - To remove a business entity, click .





You cannot remove an entity that is derived from an application that is assigned to the asset.

Using the User Data section to assign or modify annotations

When customizing OnCommand Insight to track data for your corporate requirements, you can define specialized notes, called *annotations*, and assign them to your assets. The User Data section of an asset page displays annotations assigned to an asset and also enables you to change the annotations assigned to that asset.

Steps

1. Log in to the OnCommand Insight web UI.
 2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it.
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 assigned to the asset. You can later filter assets by annotation using a query.

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.

Expert view

The Expert View section of an asset page enables you to view a performance sample for the base asset based on any number of applicable metrics in context with a chosen time period (3 hours, 24 hours, 3 days, 7 days, or a custom time period) in the performance chart and any assets related to it.

The following is an example of the Expert View section in a volume asset page:



You can select the metrics you want to view in the performance chart for the time period selected.

The Resources section shows the name of the base asset and the color representing the base asset in the performance chart. If the Top Correlated section does not contain an asset you want to view in the performance chart, you can use the **Search assets** box in the Additional resources section to locate the asset and add it to the performance chart. As you add resources, they appear in the Additional resources section.

Also shown in the Resources section, when applicable, are any assets related to the base asset in the following categories:

- Top correlated

Shows the assets that have a high correlation (percentage) with one or more performance metrics to the base asset.

- Top contributors

Shows the assets that contribute (percentage) to the base asset.

- Greedy

Shows the assets that take away system resources from the asset through sharing the same resources,

such as hosts, networks, and storage.

- Degraded

Shows the assets that are depleted of system resources due to this asset.

Expert View metric definitions

The Expert View section of an asset page displays several metrics based on the time period selected for the asset. Each metric is displayed in its own performance chart. You can add or remove metrics and related assets from the charts depending on what data you want to see.

Metric	Description
BB credit zero Rx, Tx	Number of times the receive/transmit buffer-to-buffer credit count transitioned to zero during the sampling period. This metric represents the number of times the attached port had to stop transmitting because this port was out of credits to provide.
BB credit zero duration Tx	Time in milliseconds during which the transmit BB credit was zero during the sampling interval.
Cache hit ratio (Total, Read, Write) %	Percentage of requests that result in cache hits. The higher the number of hits versus accesses to the volume, the better is the performance. This column is empty for storage arrays that do not collect cache hit information.
Cache utilization (Total) %	Total percentage of cache requests that result in cache hits
Class 3 discards	Count of Fibre Channel Class 3 data transport discards.
CPU utilization (Total) %	Amount of actively used CPU resources, as a percentage of total available (over all virtual CPUs).
CRC error	Number of frames with invalid cyclic redundancy checks (CRCs) detected by the port during the sampling period
Frame rate	Transmit frame rate in frames per second (FPS)
Frame size average (Rx, Tx)	Ratio of traffic to frame size. This metric enables you to identify whether there are any overhead frames in the fabric.

Frame size too long	Count of Fibre Channel data transmission frames that are too long.
Frame size too short	Count of Fibre Channel data transmission frames that are too short.
I/O density (Total, Read, Write)	Number of IOPS divided by used capacity (as acquired from the most recent inventory poll of the data source) for the Volume, Internal Volume or Storage element. Measured in number of I/O operations per second per TB.
IOPS (Total, Read, Write)	Number of read/write I/O service requests passing through the I/O channel or a portion of that channel per unit of time (measured in I/O per sec)
IP throughput (Total, Read, Write)	<p>Total: Aggregated rate at which IP data was transmitted and received in megabytes per second.</p> <p>Read: IP Throughput (Receive): Average rate at which IP data was received in megabytes per second.</p> <p>Write: IP Throughput (Transmit): Average rate at which IP data was transmitted in megabytes per second.</p>
Latency (Total, Read, Write)	<p>Latency (R&W): Rate at which data is read or written to the virtual machines in a fixed amount of time. The value is measured in megabytes per second.</p> <p>Latency: Average response time from the virtual machines in a data store.</p> <p>Top Latency: The highest response time from the virtual machines in a data store.</p>
Link failure	Number of link failures detected by the port during the sampling period.
Link reset Rx, Tx	Number of receive or transmit link resets during the sampling period. This metric represents the number of link resets that were issued by the attached port to this port.
Memory utilization (Total) %	Threshold for the memory used by the host.


Partial R/W (Total) %	<p>Total number of times that a read/write operation crosses a stripe boundary on any disk module in a RAID 5, RAID 1/0, or RAID 0 LUN. Generally, stripe crossings are not beneficial, because each one requires an additional I/O. A low percentage indicates an efficient stripe element size and is an indication of improper alignment of a volume (or a NetApp LUN).</p> <p>For CLARiiON, this value is the number of stripe crossings divided by the total number of IOPS.</p>
Port errors	Report of port errors over the sampling period/given time span.
Signal loss count	Number of signal loss errors. If a signal loss error occurs, there is no electrical connection, and a physical problem exists.
Swap rate (Total Rate, In rate, Out rate)	Rate at which memory is swapped in, out, or both from disk to active memory during the sampling period. This counter applies to virtual machines.
Sync loss count	Number of synchronization loss errors. If a synchronization loss error occurs, the hardware cannot make sense of the traffic or lock onto it. All the equipment might not be using the same data rate, or the optics or physical connections might be of poor quality. The port must resynchronize after each such error, which impacts system performance. Measured in KB/sec.
Throughput (Total, Read, Write)	Rate at which data is being transmitted, received, or both in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Timeout discard frames - Tx	Count of discarded transmit frames caused by timeout.
Traffic rate (Total, Read, Write)	Traffic transmitted, received, or both received during the sampling period, in mebibytes per second.
Traffic utilization (Total, Read, Write)	Ratio of traffic received/transmitted/total to receive/transmit/total capacity, during the sampling period.
Utilization (Total, Read, Write) %	Percentage of available bandwidth used for transmission (Tx) and reception (Rx).

Write pending (Total)	Number of write I/O service requests that are pending.
-----------------------	--

Using the Expert View section

The Expert view section enables you to view performance charts for an asset based on any number of applicable metrics during a chosen time period, and to add related assets to compare and contrast asset and related asset performance over different time periods.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it.
The asset page displays. By default, the performance chart shows two metrics for time period selected for the asset page. For example, for a storage, the performance chart shows latency and total IOPS by default. The Resources section displays the resource name and an Additional resources section, which enables you to search for assets. Depending on the asset, you might also see assets in the Top correlated, Top contributor, Greedy, and Degraded sections.
3. You can click **Select metrics to show**, and select a metric to add a performance chart for a metric.

A performance chart is added for the selected metric. The chart displays the data for the selected time period. You can change the time period by clicking on another time period in the top left-hand corner of the asset page.

You can perform the step again, and click to clear a metric. The performance chart for the metric is removed.


4. You can position your cursor over the chart and change the metric data that displays by clicking either of the following, depending on the asset:
 - **Read or Write**
 - **Txor Rx**
Total is the default.
5. You can drag your cursor over the data points in the chart to see how the value of the metric changes over the time period selected.
6. In the **Resources** section, you can do any of the following, if applicable, to add any related assets to the performance charts:
 - You can select a related asset in the Top correlated, Top contributors, Greedy, or Degraded sections to add data from that asset to the performance chart for each selected metric. Assets must have a minimum 15% correlation or contribution to be shown.

After you select the asset, a color block appears next to the asset to denote the color of its data points in the chart.

- For any asset shown, you can click the asset name to display its asset page, or you can click the percentage that the asset correlates or contributes to the base asset to view more information about the assets relation to the base asset.

For example, clicking the linked percentage next to a top correlated asset displays an informational message comparing the type of correlation that asset has with the base asset.

- If the Top correlated section does not contain an asset you want to display in a performance chart for comparison purposes, you can use the **Search assets** box in the Additional resources section to locate other assets.

After you select an asset, it displays in the Additional resources section. When you no longer want to view information about the asset, click .




Related Assets

If applicable, an asset page displays a Related Assets section. For example, a volume asset page might show information about assets like Storage Pools, Connected switch ports, and Compute Resources. Each section comprises a table that lists any of the related assets in that category, with links to their respective asset pages, and several performance statistics related to the asset.

Using the Related Assets section





The Related Assets section enables you to view any of the assets that are related to the base asset. Each related asset is displayed in a table along with pertinent statistics for the asset. You can export the asset information, view the asset statistics in the Expert View performance charts, or show a chart that displays statistics for only related assets.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it.
The asset page displays.
3. To control how assets display in the table:
 - Click the name of any asset to display its asset page.
 - Use the **filter** box to show only specific assets.
 - Click a page number to browse through the assets by page if there are more than five assets in the table.
 - 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.
 - Add a related asset to any performance chart in the Expert View section by placing your cursor over the related asset and clicking .
4. To export the information displayed in the table to a .CSV file:
 - a. Click .
 - b. Click **Open with** and then **OK** to open the file with Microsoft Excel and save the file to a specific location, or click **Save file** and then **OK** to save the file to your Downloads folder.

All of the object attributes for the columns currently selected for display are exported to the file. Only the attributes for the displayed columns will be exported. Note that only the first 10,000 rows of the

table are exported.

5. To display the related asset information in a chart below the table, click  and do any of the following:
 - Click **Read**, **Write**, or **Total** to change the metric data that displays. **Total** is the default.
 - Click  to select a different metric.
 - Click  to change the chart type. **Line chart** is the default.
 - Move your cursor over the data points in the chart to see how the value of the metric changes over the time period selected for each related asset.
 - Click a related asset in the chart legend to add it to or remove it from the chart.
 - Click a page number in the related asset table to view other related assets in the chart.
 - Click  to close the chart.

Violations

You can use the Violations section of an asset page to see the violations, if any, that occur in your environment as a result of a performance policy assigned to an asset. Performance policies monitor your network thresholds and enable you to 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.


The following example shows a Violations section that displays on an asset page for a hypervisor:

Violations		filter...
Time	Description	
06/05/2015 5:00:00 pm	Port balance index of 74 on esx1 exceeds the threshold of 50	
06/12/2015 8:59:54 am	2 violations for esx2 with 'Swap out rate' > 3	
06/12/2015 12:04:54 pm	esx1 violation with 'Swap out rate' > 3.00 KB/s (value of 86.85 KB/s)	
06/12/2015 12:29:54 pm	esx1 violation with 'Swap in rate' > 3.00 KB/s (value of 59.90 KB/s)	
06/12/2015 1:04:54 pm	7 violations for ds-30 with 'Latency - Total' > 50	
Showing 1 to 5 of 32 entries		< 1 2 3 4 5 >

Using the Violations section

The Violations section enables you to view and manage any of the violations that occur in your network as the result of a performance policy assigned to an asset.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it.
The asset page displays. The Violations section displays the time the violation occurred and a description of the threshold that was crossed, along with a hyperlink to the asset on which the violation occurred (for example "2 violations fir ds-30 with Latency - Total > 50").
3. You can perform any of the following optional tasks:
 - Use the **filter** box to show only specific violations.

- Click a page number to browse through the violations by page if there are more than five violations in the table.
- 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.
- Click the asset name in any description to display its asset page; a red circle indicates issues that need further investigation.

You can click the performance policy, which displays the Edit Policy dialog box, to review the performance policy and make changes to the policy if necessary.

- Click  to remove a violation from the list if you determine the issue is no longer a cause for concern.

Customizable asset page

Additional data can be displayed in customizable widgets on each asset page. Customizing the page for an asset applies the customization to the pages for all assets of that type.

You customize asset page widgets by performing the following actions:

1. Add a widget to the page
2. Create a query or expression for the widget to showcase desired data
3. Choose a filter if desired
4. Choose a rollup or grouping method
5. Save the widget
6. Repeat for all desired widgets
7. Save the asset page

You can also add variables to the custom asset page that can be used to further refine your showcased data in widgets. In addition to regular variables, each asset type can use a set of "\$this" variables to quickly identify resources directly related to the current asset, for example, all virtual machines hosted by the same hypervisor that hosts the current virtual machine.

This custom asset page is unique for each user as well as for each asset type. For example, if User A creates a custom asset page for a virtual machine, that custom page will display for any virtual machine asset page, for that user.

Users can only view, edit, or delete custom asset pages that they create.

Custom asset pages are not included in Insight's export/import functionality.

Understanding "\$this" variables

Special variables on an asset's "Additional data" customizable page allow you to easily showcase additional information that is directly related to the current asset.

About this task

To use the "\$this" variables in widgets on your asset's customizable landing page, follow the steps below. For this example, we will add a table widget.



“\$this” variables are only valid for an asset’s customizable landing page. They are not available for other Insight dashboards. The available “\$this” variables varies according to asset type.

Steps

1. Navigate to an asset page for an asset of your choosing. For this example, let’s choose a Virtual Machine (VM) asset page. Query or search for a VM and click on the link to go to that VM’s asset page.

The asset page for the VM opens.

2. Click on the **Change view:** > **Additional Virtual Machine data** drop-down to go to that asset’s customizable landing page.
3. Click on the **Widget** button and choose **Table widget**.

The Table widget opens for editing. By default, all storages are shown in the table.


4. We want to show all virtual machines. Click on the asset selector and change **Storage** to **Virtual Machine**.

All virtual machines are now shown in the table.

5. Click on the **Column selector** button  and add the **hypervisor name** field to the table.

The hypervisor name is shown for each VM in the table.

6. We only care about the hypervisor that hosts the current VM. Click on the **Filter by** field’s+button and select **hypervisor name**.
7. Click on **Any** and select the **\$this.host.name** variable. Click the check button to save the filter.



8. The table now shows all the VM’s hosted by the current VM’s hypervisor. Click **Save**.

Results

The table that you created for this virtual machine asset page will be displayed for any VM asset page you display. The use of the **\$this.host.name** variable in the widget means that only the VM’s owned by the current assets’s hypervisor will be displayed in the table.

Balancing network resources

To resolve balancing issues, use the asset pages to find the problems and identify high capacity volumes that are underused.

Steps

1. Open the Assets Dashboard in your browser.
2. In the Virtual Machines IOPS heat map, you notice the name of a VM in very large print that often reports problems.
3. Click the VM name to display the asset page.
4. Check for error messages in the summary.
5. Check the performance charts and particularly the top correlated resources to locate any volumes that might be in contention.
6. Add volumes to the performance chart to compare the patterns of activity and display more asset pages for other resources involved in the problem.
7. Scroll to the bottom of the asset page to see lists of all of the resources associated with the VM. Note any VMDKs running at high capacity. This is likely causing the contention.
8. To resolve the balancing problem, identify a resource that is under-utilized to receive the load from an over-utilized resource or remove a less demanding application from the heavily used resource.

Examining network performance

You can examine your storage environment performance and identify under-utilized and over-utilized resources and identify risks before they turn into problems.

Insight helps you to resolve or prevent performance and availability problems that are revealed through the collected storage data.

You can use Insight to perform these performance management tasks:

- Monitor performance across your environment
- Identify resources influencing the performance of other devices

The Importance of Ports

The Insight Server and Data Warehouse (DWH) server may require a number of TCP ports to be free in order to operate reliably. Some of these ports are only utilized for processes bound to the localhost adapter (127.0.0.1), but are still required for core services to operate reliably. The number of ports required is a superset of what ports are used across the network.

Insight Server Ports

Insight Servers can have software firewalls installed. The "holes" that would need to be opened would be as described below.

Inbound HTTPS 443 - assuming you have the Insight WebUI running on TCP 443, you must expose that as to allow any and all of the following consumers:

- Insight users of the WebUI

- Remote Acquisition Units seeking to connect to the Insight server
- OCI DWH servers with connectors to this Insight server.
- Any programmatic interactions with the Insight REST API

Our general recommendation for anyone looking to implement Insight server host-level firewalling is to allow HTTPS access to all corporate network IP blocks.

Inbound MySQL (TCP 3306). This port only needs to be exposed to any Insight DWH server with a connector

While Insight has dozens of data collectors, they are all poll-based - Insight will cause its Acquisition Units (AUs) to initiate outbound communication to various devices. As long as your host based firewall is "stateful" such that it allows return traffic to be allowed through the firewall, host based firewalls on the Insight Server should not impact data acquisition.

Data Warehouse Ports

For Insight DWH servers:

Inbound HTTPS 443 - assuming you have the Insight WebUI running on TCP 443, you must expose that as to allow the following consumers:

- Insight administrative users of the DWH admin portal

Inbound HTTPS (TCP 9300) - this is the Cognos reporting interface. If you will have users interacting with the Cognos reporting interface, this must be exposed remotely.

We can imagine environments where the DWH may not need to be exposed - perhaps the report authors just make RDP connections to the DWH server, and craft and schedule reports there, while having all reports scheduled to be delivered via SMTP, or written to a remote file system.

Inbound MySQL (TCP 3306). This port only needs to be exposed if your organization has any MySQL-based integrations with DWH data - are you extracting data out of the various DWH data marts for ingestion into other applications like CMDBs, chargeback systems, etc.

Analyzing slow PC performance

If you receive calls from network users complaining that their computers are running slowly, you need to analyze host performance and identify the affected resources.

Before you begin

In this example, the caller gave the host name.

Steps

1. Open Insight in your browser.
2. Enter the host name in the **Search assets** box and click the host name in the search results.

The *asset page* for the resource opens.

3. On the asset page for the host, examine the performance charts in the center of the page. You might want to show different types of data in addition to the Latency and IOPS that are usually pre-selected. Click the check boxes for other types of data, such as Throughput, Memory, CPU, or IP Throughput depending on

the device type.

4. To display a description of a point on a chart, position the mouse pointer over the point.
5. You might also want to change the time range with the selection at the top of the page to be 3 hours up to 7 days or All of the available data.
6. Examine the list of **Top correlated resources** to see if there are other resources with the same pattern of activity as the base resource.

The first resource in the list is always the base resource.

- a. Click a linked percentage beside a correlated resource to see if the correlated activity pattern is for IOPS or CPU for the base resource and another resource.
 - b. Click the check box for a correlated resource to add its data to the performance charts.
 - c. Click the linked name of the correlated resource to display its asset page.
7. For a VM, as in this example, locate the storage pool in the **Top correlated resources** and click the storage pool name.

Analyzing correlated resources


When you are researching performance problems and you open the *asset page* for a device, you should use the Top correlated resources list to refine data displayed in the performance charts. A resource with a high percentage indicates that resource has similar activity to the base resource.

About this task

You are investigating a performance problem and opened the asset page for a device.

Steps

1. In the **Top correlated resources** list, the first resource is the base resource. The correlated resources in the list are ranked by percentage of correlated activity to the first device. Click the linked percentage of correlation to see the details. In this example, the 70% correlation is in Utilization, so both the base resource and this correlated resource have equally high utilization.



- To add a correlated resource to the performance charts, select the check box in the **Top correlated resources** list for the resource you want to add. By default each resource provides the Total data available, but you can select only Read or only Write data from the menu on the check box.

Each resource in the charts has a different color so that you can compare the performance measurements for each resource. Only the appropriate type of data is plotted for the selected measurement metrics. For example, CPU data does not include Read or Write metrics, so only Total data is available.

- Click the linked name of the correlated resource to display its asset page.
- If you do not see a resource listed in the Top correlated resources that you believe should be considered in the analysis, you can use the **Search assets** box to find that resource.

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 policies 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.

Time-to-live (TTL) and downsampled data

Starting with OnCommand Insight 7.3, data retention or time-to-live (TTL) has been increased to from 7 to 90 days. Because that means much more data is processed for charts and tables and the potential for tens of thousands of datapoints, data is downsampled before being displayed.

Downsampling provides a statistical approximation of your data in charts, giving you an efficient overview of data without having to display every data point, while maintaining an accurate view of your collected data.

Why is downsampling needed?

Insight 7.3 increases the time-to-live (TTL) for data to 90 days. This means an increase in the amount of processing needed to prepare data for display in charts and graphs. In order to allow charts to display quickly and efficiently, data is downsampled in a manner that keeps the overall shape of a chart without needing to process every single data point for that chart.



No actual data is lost during downsampling. You can choose to view actual data for your chart instead of downsampled data by following the steps illustrated below.


How downsampling works

Data is downsampled under the following conditions:

- When your selected time range includes 7 days of data or less, no downsampling occurs. Charts display actual data.
- When your selected time range includes more than 7 days of data but less than 1,000 data points, no downsampling occurs. Charts display actual data.
- When your selected time range includes more than 7 days of data and more than 1,000 data points, data is downsampled. Charts display approximated data.

The following examples show downsampling in action. The first illustration shows latency and IOPS charts on a Datastore asset page for a 24-hour period, as shown by selecting **24h** on the asset page's time selector. You can also see the same data by selecting **Custom** and setting the time range to the same 24-hour period.

Since we are choosing a time range of less than 7 days and we have less than 1,000 data points to chart, the data displayed is actual data. No downsampling occurs.




However, if you are viewing data by choosing either **30d** on the asset page time selector, or by setting a custom time range of more than 7 days (or in the event that Insight has collected more than 1,000 data samples for the time period chosen), the data is downsampled before being displayed. When you zoom in on a downsampled chart, the display continues to show the approximated data.



When you zoom in on a downsampled chart, the zoom is a digital zoom. The display continues to show the approximated data.

You can see this in the following illustration, where the time range is first set to 30d, and the chart is then zoomed in to show the same 24-hour period as above.



The downsampled charts are showing the same 24-hour period as the "actual" charts above, so the lines follow the same general shape, allowing you to quickly spot interesting peaks or valleys in your performance data.



Due to the way data is approximated for downsampling, chart lines may be off slightly when comparing downsampled vs. actual data, to allow for better alignment in the graphs. However, the difference is minimal and does not affect the overall accuracy of the data displayed.

Violations on downsampled charts

When viewing downsampled charts, be aware that violations are not shown. To see violations, you can do one of two things:

- View the actual data for that time range by selecting Custom in the asset page time selector, and entering a range of time less than 7 days. Hover over each red dot. The tooltip will show the violation that occurred.
- Note the time range and find the violation(s) in the Violation Dashboard.

Pruning of inventory history

Starting with version 7.3.2, Insight keeps inventory (foundation) change history for 90 days. Previous versions of Insight kept all inventory change history from the time of installation. Following an upgrade from an older version of Insight, older inventory history is pruned down to and then kept at 90 days.

After upgrading to the current version of OnCommand Insight, history is pruned to the most recent 90 days. Insight prunes the history in 30-day chunks occurring once a day, starting with the oldest, until 90 days' worth of history remains. Then, history is pruned daily, to keep only 90 days' worth of inventory change history.

NAS path for VMs

OnCommand Insight 7.3 supports NAS paths for Virtual Machines to storage shares. These paths are similar to NAS paths for hosts to storage shares. When a VM's IP address is allowed to access a share, a NAS path is created.

NAS paths for virtual machines are displayed on the Internal Volumes landing page. This page contains a Guest Mounted Storage Resources widget which identifies the Internal Volumes that VMs have access to.

- NAS paths are created when virtual machines have access to the backend shares. There is no acknowledgment of whether the virtual machines access the shares or not.
- Correlation calculation is based on latencies and IOPs, and do not include cases where VMs have NAS paths to the backend storage.
- User can query the share by initiator IP address, but querying by path is not supported.

The Compute Resources table of the Internal Volume now also displays VM's with NAS paths. For each VM, CPU and memory, utilization and performance data is provided.

Data warehouse impact

Changes to the Data Warehouse that are present after upgrading to OnCommand Insight 7.3 include the following:

- The `dwh_inventory.nas_logical` table is removed from the Inventory data mart and replaced with a view.

Any Insight 7.2.x reports containing the NFS path table are preserved.

- The `dwh_inventory.nas_cr_logical` table is added to the Inventory data mart and includes the following:
 - Compute resource
 - Internal volume
 - Storage
 - NAS share

Capacity as Time Series

With OnCommand Insight 7.3.1, capacity information is reported and charted as time series data.

Previously, capacity information acquired from data sources has been exclusively "point-in-time" (PIT) data, meaning it could not be used in charts as time series data. Now, capacity values for assets can be used as time series data in the following ways:

- Graphed in tables, widgets, expert views, and any place where time series data is displayed
- Applied to performance thresholds with violations using existing semantics
- Used in expressions with other performance counters where appropriate

Note that if you upgrade from a previous version of Insight, previous PIT capacity values used in queries or in filters for custom dashboards will be replaced with time series capacity data. This may result in small changes in the way that capacity data is reported or filtered when compared to the equivalent data in previous Insight versions.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.