OnCommand® Insight 7.2.3

# Configuration and Administration Guide

October 2016 | 215-11672_A0
doccomments@netapp.com

NetApp®

# Contents

**OnCommand Insight overview ..... 7**
Insight architecture ..... 7
Insight Data Warehouse architecture ..... 8
Firewall-friendly architecture ..... 8
How Insight is used by administrators, managers, and planners ..... 9
**Setting up Insight ..... 10**
Accessing the web UI ..... 10
Installing your Insight licenses ..... 11
OnCommand Insight licenses ..... 12
Setting up and managing user accounts ..... 14
Insight user roles ..... 15
Configuring user definitions using LDAP ..... 16
Changing user passwords ..... 18
Editing a user definition ..... 18
Deleting a user account ..... 19
Smart Card and certificate login support ..... 19
Configuring hosts for Smart Card and certificate login ..... 19
Configuring a client to support Smart Card and certificate login ..... 22
Configuring Data Warehouse for Smart Card and certificate login ..... 22
Configuring Data Warehouse for Smart Card and certificate login for Cognos ..... 23
Importing SSL certificates ..... 25
Setting up weekly backups for your Insight database ..... 27
Configuring your email ..... 28
Email configuration prerequisites ..... 28
Configuring your email for Insight ..... 28
Configuring SNMP notifications ..... 29
Enabling the syslog facility ..... 29
Insight syslog contents ..... 30
Configuring performance and assure violation notifications ..... 31
Configuring your ASUP processing ..... 31
Creating custom dashboards ..... 32
Deleting a custom dashboard ..... 34
Defining applications ..... 34
Assigning applications to assets ..... 35
Editing applications ..... 36
Deleting applications ..... 36
How OnCommand Insight Ethernet monitoring works ..... 37
Ethernet monitoring unit network configuration ..... 37
Viewing Ethernet monitored data in the web UI ..... 40
How OnCommand Insight anomaly detection works ..... 40

Configuring applications for anomaly detection ..... 41
Using the Anomalies section ..... 42
Anomaly detection monitoring status ..... 43
Your business entities hierarchy ..... 43
Designing your business entities hierarchy ..... 44
Creating business entities ..... 44
Assigning business entities to assets ..... 45
Defining annotations ..... 45
Using annotations to monitor your environment ..... 46
Querying assets ..... 57
Creating a query ..... 57
Viewing queries ..... 59
Exporting query results to a .CSV file ..... 59
Modifying queries ..... 59
Deleting queries ..... 60
Assigning multiple applications to or removing multiple applications from assets ..... 60
Editing or removing multiple annotations from assets ..... 61
Managing performance policies ..... 62
Creating performance policies ..... 62
Performance policy evaluation precedence ..... 64
Editing performance policies ..... 65
Deleting performance policies ..... 65
Importing and Exporting user data ..... 66
**Insight data source management ..... 68**
Setting up your data sources in Insight ..... 68
Data source support information ..... 69
Adding data sources ..... 69
Cloning a data source ..... 71
Testing the data source configuration ..... 72
Vendor-specific data source reference ..... 73
3PAR InServ Firmware 2.2.2+ (SSH) data source ..... 73
Brocade Enterprise Fabric Connectivity Manager data source ..... 74
Brocade FC Switch Firmware (SSH) data source ..... 75
Brocade Sphereon/Intrepid Switch (SNMP) data source ..... 77
Cisco FC Switch Firmware (SNMP) data source ..... 78
Dell Compellent Storage Center data source ..... 80
EMC Atmos (HTTP) data source ..... 81
EMC Celerra (SSH) data source ..... 82
EMC CLARiiON (NaviCLI) data source ..... 83
EMC Data Domain data source ..... 85
EMC ECC StorageScope data source ..... 86
EMC Isilon data source ..... 86
EMC Solutions Enabler (CLI) with SMI-S Performance data source ..... 87
EMC VNX (SSH) data source ..... 90

EMC VNXe data source ..... 90
EMC VPLEX data source ..... 92
EMC XtremIO (HTTP) data source ..... 93
Fujitsu Eternus data source ..... 94
HDS HiCommand Device Manager data source ..... 95
HDS NAS (HNAS) data source ..... 97
HP CommandView AE data source ..... 98
HP EVA Storage (SSSU) data source ..... 101
IBM DS (CLI) data source ..... 102
IBM PowerVM data source ..... 103
IBM SVC (SSH) data source ..... 105
IBM Tivoli Monitoring data source ..... 106
IBM TotalStorage DS4000 data source ..... 107
IBM XIV (CLI) data source ..... 108
Infinidat InfiniBox (HTTP) data source ..... 109
Microsoft Hyper-V (WMI) data source ..... 110
NetApp Clustered Data ONTAP data source ..... 111
NetApp Clustered Data ONTAP 8.1.x for Unified Manager data source ... 112
NetApp Data ONTAP operating in 7-Mode data source ..... 112
NetApp E-Series data source ..... 114
NetApp Host and VM File Systems data source ..... 115
NetApp SolidFire data source ..... 116
Pure Storage FlashArray (HTTP) data source ..... 117
QLogic FC Switch (SNMP) data source ..... 118
Violin 6000-Series Flash Memory Array (HTTP) data source ..... 119
VMware vSphere (Web Services) data source ..... 120
Changing data source credentials ..... 121
Changes causing data collection problems ..... 122
Examining one data source in detail ..... 123
Researching a failed data source ..... 124
Controlling data source polling ..... 126
Editing data source information ..... 127
Editing information for multiple data sources ..... 127
Deleting a data source ..... 129
What data source patches are ..... 129
Applying a data source patch ..... 129
Installing a patch on one type of data source ..... 130
Managing patches ..... 130
Committing a data source patch ..... 131
Rolling back a data source patch ..... 132
**Maintaining Insight ..... 133**
Managing Insight ..... 133
Monitoring OnCommand Insight system health ..... 134
Auditing system and user activities ..... 139
Monitoring the violations in your network ..... 140

Acquisition unit status ..... 141
Restoring the Insight database ..... 142
Updating expired licenses ..... 143
Licenses no longer compliant ..... 143
Replacing licenses for older Insight versions ..... 143
Applying a service pack ..... 143
Preparing a special troubleshooting report ..... 144
Manually backing up the OnCommand Insight database ..... 144
Bundling logs for Support ..... 144
Sending information to NetApp Support ..... 145
Advanced configuration tools ..... 145
Restarting the OnCommand Insight Server ..... 146
Configuring the number of hours to ignore dynamic data ..... 147
Generating logs for Customer Support ..... 148
Displaying system information ..... 149
Listing installed OnCommand Insight components ..... 149
Calculating the number of database objects ..... 149
Moving MySQL data using the migrate option ..... 150
Forcing annotation updates ..... 151
Configuring Client settings ..... 151
Checking a tier upgrade ..... 152
Checking the status of server resources ..... 152
Scrubbing data for transfer to support ..... 152
Finding ghost data sources ..... 155
Adding a missing disk model ..... 155
Specialized administration toolkit ..... 157
**Monitoring your environment ..... 160**
Monitoring your environment ..... 160
Asset page data ..... 160
Balancing network resources ..... 179
Examining network performance ..... 181
Fibre Channel environment monitoring ..... 183
**Copyright information ..... 184**
**Trademark information ..... 185**
**How to send comments about documentation and receive update notifications ..... 186**
**Index ..... 187**

# OnCommand Insight overview

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

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

## Insight architecture

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

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

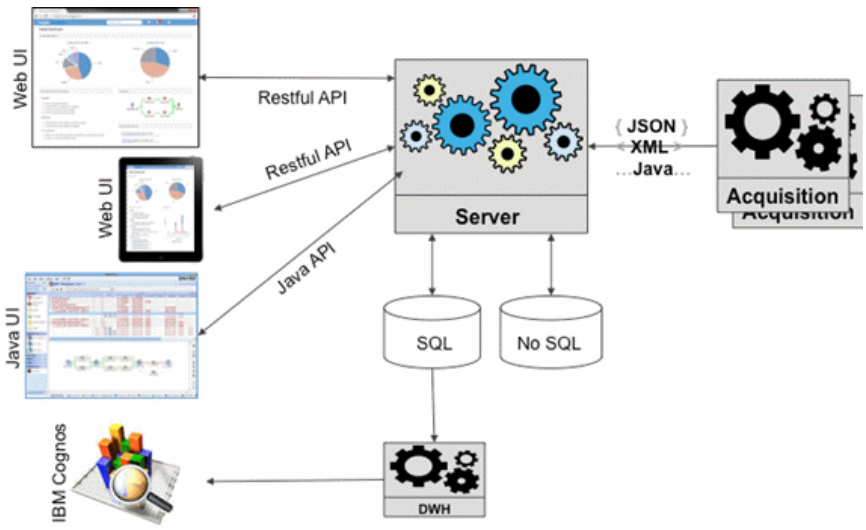


**OnCommand Insight Server**

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

**Acquisition units**

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

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

**Web UI**

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

**Java UI**

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

**Data Warehouse (DWH)**

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

**IBM Cognos**

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

## Insight Data Warehouse architecture

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

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

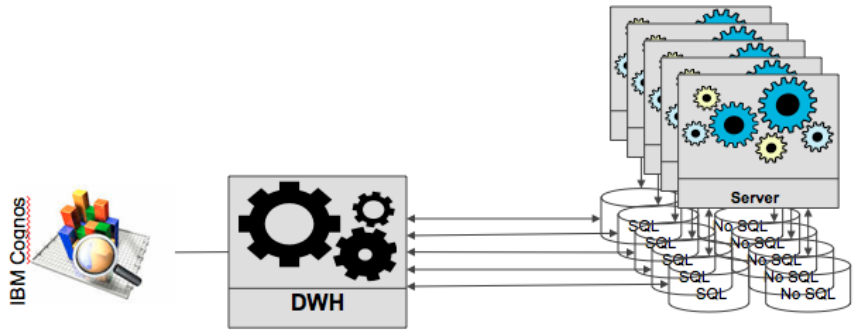


## Firewall-friendly architecture

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

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

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

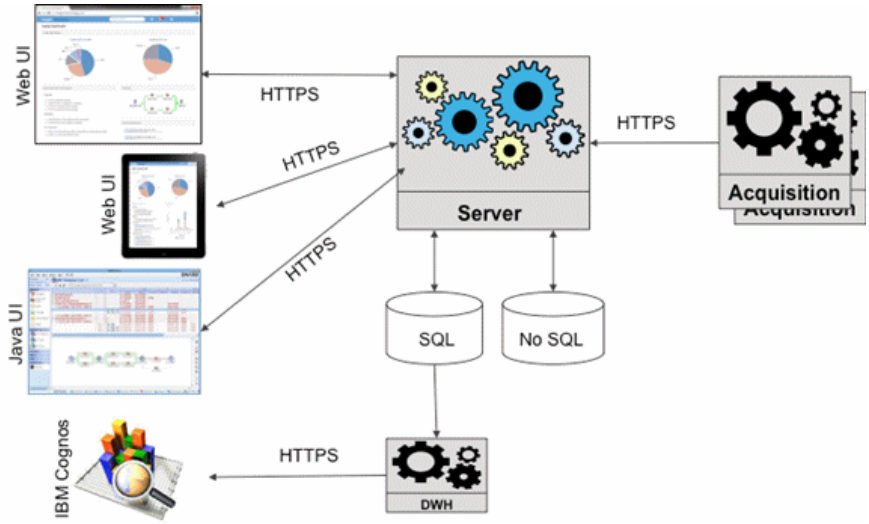


## How Insight is used by administrators, managers, and planners

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

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

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

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

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

# Setting up Insight

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

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

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

### Related concepts

*Licenses no longer compliant* on page 143

### Related tasks

*Updating expired licenses* on page 143

# Accessing the web UI

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

### Steps

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

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

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

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

   The OnCommand Insight dialog box displays:

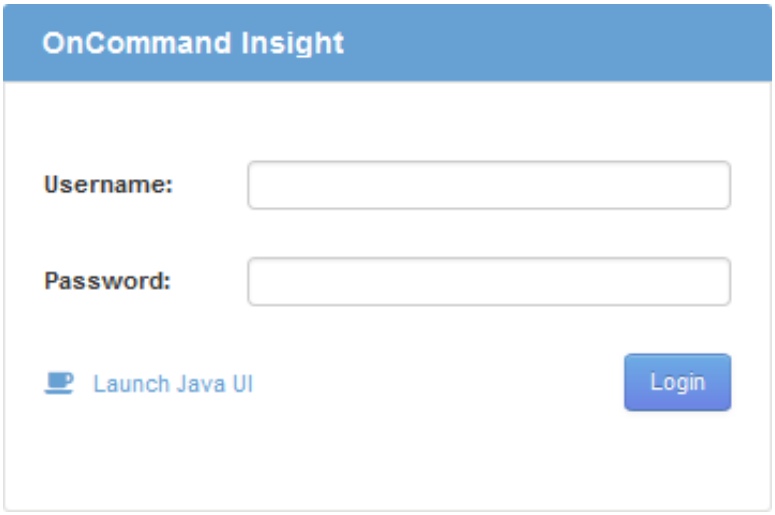


2. Enter your user name and password and click **Login**.

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

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

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

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

### Related tasks

*Updating expired licenses* on page 143

## Installing your Insight licenses

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

### About this task

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

### Steps

1. Open the license file in a text editor and copy the text.
2. Open Insight in your browser.
3. 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**.

#### After you finish

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

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

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

## OnCommand Insight licenses

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

**Discover**

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

**Assure**

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

**Perform**

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

**Plan**

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

**Host Utilization pack**

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

**Report Authoring**

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

OnCommand Insight modules are licensed for annual term or perpetual:

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

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

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

**Discover**

Acquire and manage inventory (Foundation)

Monitor changes and manage inventory policies

**Assure**

View and manage SAN path policies and violations

View and manage vulnerabilities

View and manage tasks and migrations

**Plan**

View and manage requests

View and manage pending tasks

View and manage reservation violations

View and manage port balance violations

**Perform**

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

View and manage performance policies and violations

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

| Feature (admin) | With Perform license | Without Perform license |
|---|---|---|
| Application | Yes | No performance data or charts; no anomaly detection-related widgets |
| Virtual machine | Yes | No performance data or charts |
| Hypervisor | Yes | No performance data or charts |
| Host | Yes | No performance data or charts |
| Datastore | Yes | No performance data or charts |
| VMDK | Yes | No performance data or charts |
| Internal volume | Yes | No performance data or charts |
| Volume | Yes | No performance data or charts |
| Storage pool | Yes | No performance data or charts |
| Disk | Yes | No performance data or charts |
| Storage | Yes | No performance data or charts |
| Storage node | Yes | No performance data or charts |
| Fabric | Yes | No performance data or charts |
| Switch port | Yes | No performance data or charts; “Port Errors” shows “N/A” |
| Storage port | Yes | Yes |
| NPV port | Yes | No performance data or charts |
| Switch | Yes | No performance data or charts |
| NPV switch | Yes | No performance data or charts |
| Search | Yes | Yes |
| Admin | Yes | Yes |
| Dashboard | Yes | Yes |

| Feature (admin) | With Perform license | Without Perform license |
|---|---|---|
| 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.

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

   You enter the **Name**, **Password**, **Email** address, and select one of the user **Roles** as Administrator, User, or Guest.
6. To change a user's information, select the user from the list and click the **Edit user account** symbol to the right of the user description.
7. To remove a user from the OnCommand Insight system, select the user from the list and click **Delete user account** to the right of the user description.

#### Result

When a user logs in to OnCommand Insight, the server first attempts to authenticate through LDAP, if LDAP is enabled. If OnCommand Insight cannot locate the user on the LDAP server, it searches in 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. If your user account is defined to the OnCommand Insight local user database (and not through LDAP), you can also modify your own password. This account type does not allow you to perform actions such as identifying generic devices and defining the policies in the Java UI.
- User permits all guest-level privileges, as well as access to Insight operations such as defining policy and identifying generic devices. The User account type does not allow you to perform data source operations, nor to add or edit any user accounts other than your own.
- Administrator permissions vary depending on whether you use LDAP:
  - If you authenticate and authorize users through LDAP, this level of permission allows you manage data sources.

- If you use only the local database to manage users, this level of permission allows you perform any operation, including adding new users and managing data sources.

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

## Configuring user definitions using LDAP

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

### About this task

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

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

### Steps

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

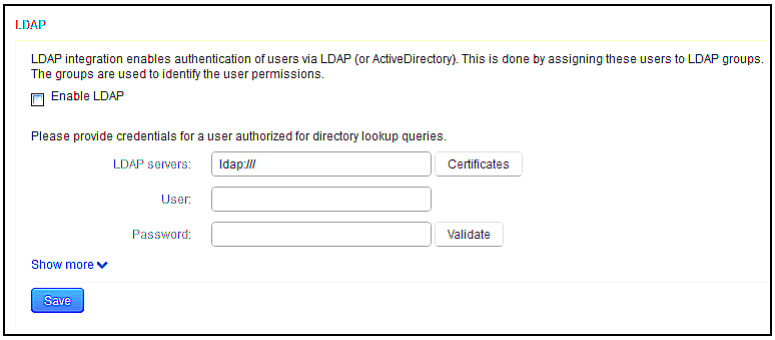


5. Click **Enable LDAP** to allow the LDAP user authentication and authorization.
6. Enter the credentials for a user authorized for directory lookup queries in these fields:
   - LDAP servers (To import the Insight certificate for this field, click **Certificates**.)
   - User name
   - Password (To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

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

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

   or use the default port:

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

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

   The default settings are valid for the Active Directory.

   **Admins group**

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

   **Users group**

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

   **Guests group**

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

   **Server admins group**

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

   **Timeout**

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

   **Domain**

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

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

   **User principal name attribute**

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

   **Role attribute**

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

   **Mail attribute**

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

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

   **Distinguished name attribute**

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

8. Click **Save**.

## Changing user passwords

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

### Before you begin

The following items must have been completed:

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

### About this task

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

### Steps

1. Log in with administrator privileges.
2. 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**.

# Smart Card and certificate login support

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

After configuring the hosts and clients to support CAC and certificates, navigating to a new session of OnCommand Insight results in the browser displaying a native dialog providing the user with a list of personal certificates to choose from. These certificates are filtered based on the set of personal certificates that have been issued by CAs trusted by the OnCommand Insight server. Most often, there is a single choice. By default, Internet Explorer skips this dialog if there is only one choice.

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

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

## Configuring hosts for Smart Card and certificate login

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

#### Before you begin

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

#### About this task

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

**Steps**

1. Use the `regedit` utility to modify registry values in `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\SANscreen Server\parameters`:
   a. Change the JVM_Option `DclientAuth=false` to `DclientAuth=want`.
   b. Change the JVM_Option `DauthClass=com.onaro.sanscreen.authenticate.CustomBasicAuthenticator` to `DauthClass=com.onaro.sanscreen.authenticate.CustomSSLAuthenticator`.
2. Contact NetApp support for a new `cacAuth.jar` file and copy that file into the `<SANScreenInstallDir>/jboss/server/onaro/lib/` folder.
3. Edit the file `<SANScreenInstallDir>/jboss/server/onaro/deploy/jbossweb.sar/web.xml` and add the following lines, replacing "X.X.X.X" with the IP address of your Insight Server:

```
    <filter>
      <filter-name>CAC Filter</filter-name>
      <filter-class>com.netapp.oci.filters.CACAuth</filter-class>
      <init-param>
          <param-name>allow</param-name>
          <param-value>127\.0\.0\.1|X\.X\.X\.X</param-value>
      </init-param>
    </filter>
    <filter-mapping>
        <filter-name>CAC Filter</filter-name>
        <url-pattern>/uiserver/webui/v1/eventBus</url-pattern>
    </filter-mapping>

    <security-constraint>
        <web-resource-collection>
            <web-resource-name>Event bus</web-resource-name>
            <description>Internal access to event bus</description>
            <url-pattern>/uiserver/webui/v1/eventBus/*</url-pattern>
            <http-method>POST</http-method>
        </web-resource-collection>
        <user-data-constraint>
            <transport-guarantee>CONFIDENTIAL</transport-guarantee>
        </user-data-constraint>
    </security-constraint>
```

   **Note:** For the IP addresses, do not forget to escape '.' by using '\.'
4. Open a command prompt specifying

   **`Run as administrator`**
5. Change directory to `C:\Program Files\SANscreen\jboss\server\onaro\cert`.
6. Delete the self-generated certificate:

   **`C:\Program Files\SANscreen\java64\bin\keytool.exe -delete -alias "ssl certificate" -keystore C:\Program Files\SANscreen\jboss\server\onaro\cert\server.keystore`**
7. Generate a new certificate:

   **`C:\Program Files\SANscreen\java64\bin\keytool.exe -genkey -alias "alias_name" -keyalg RSA -sigalg SHA1withRSA -keysize 2048 -keystore "C:\Program Files\SANscreen\jboss\server\onaro\cert\server.keystore" -dname "CN=commonName,OU=orgUnit,O=orgName,L=localityNameI,S=stateName,C=countryName"`**
8. Generate a certificate signing request:

```
C:\Program Files\SANscreen\java64\bin\keytool.exe -certreq -sigalg
SHA1withRSA -alias "alias_name" -keystore "C:\Program Files\SANscreen
\jboss\server\onaro\cert\server.keystore" -file C:\temp\server.csr"
```

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

10. Extract the certificate from the keystore:

```
C:\Program Files\SANscreen\java64\bin\keytool.exe -v -importkeystore -
srckeystore ""C:\Program Files\SANscreen\jboss\server\onaro\cert
\server.keystore"" -srcalias "alias_name" -destkeystore "C:\temp
\file.p12" -deststoretype PKCS12
```

11. Extract a private key from the p12 file:

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

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

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

13. Import the merged certificate into the keystore:

```
C:\Program Files\SANscreen\java64\bin\keytool.exe -importkeystore -
destkeystore "C:\Program Files\SANscreen\java64\bin\keytool.exe" -
srckeystore "C:\temp\servername.new.p12" -srcstoretype PKCS12 -alias
"alias_name"
```

14. Import the DoD root certificate:

```
C:\Program Files\SANscreen\java64\bin\keytool.exe -importkeysotre -
keystore "C:\Program Files\SANscreen\server\onaro\cert" -file "C:
\DoD_Root_CA_2__05__DoD_Root_CA_2.cer -trustcacerts -alias dod_root"
```

15. Import the DoD intermediate email certificate:

```
C:\Program Files\SANscreen\java64\bin\keytool.exe -importkeysotre -
keystore "C:\Program Files\SANscreen\server\onaro\cert
\server.truststore" -file "C:\DoD_Root_CA_2__45__DOD_EMAIL_CA-24.cer -
trustcacerts -alias dod_email_ca24
```

    Repeat this step for all intermediate email certificates.

16. Import the DoD intermediate certificate:

```
C:\Program Files\SANscreen\java64\bin\keytool.exe -importkeysotre -
keystore "C:\Program Files\SANscreen\server\onaro\cert
\server.truststore" -file "C:\DoD_Root_CA_2__47__DOD_CA-24.cer -
trustcacerts -alias dod_ca24
```

    Repeat this step for all intermediate certificates.

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

18. 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.

### About this task

The following are the common client configuration requirements:

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

Additional configuration steps include the following:

### Steps

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

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

## Configuring Data Warehouse for Smart Card and certificate login

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

### Before you begin

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

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

#### Steps

1. Use regedit to modify registry values in `HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\services\SANscreen Server.`

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

   b. Change the JVM_Option `DdwhauthClass=org.apache.catalina.authenticator.FormAuthenticator` to `-DdwhAuthClass=com.onaro.sanscreen.authenticate.CustomSSLAuthenticator`

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

   a. In a command window, go to `..\SANscreen\jboss\server\onaro\cert\`.

   b. Use the `keytool` utility to list the trusted CAs:

   **`C:\Program Files\SANscreen\java64\bin\keytool.exe -list -keystore server.truststore -storepass changeit`**

   The first word in each line indicates the CA alias.

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

   **`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.

## Configuring Data Warehouse for Smart Card and certificate login for Cognos

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

#### Before you begin

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

#### Steps

1. Add certificate authorities (CAs) to the Cognos trust store.

   a. In a command window, go to `..\SANscreen\cognos\c10_64\configuration\certs\`

   b. Use the `keytool` utility to list the trusted CAs:

   **`C:\Program Files\SANscreen\java64\bin\keytool.exe -list -keystore CognosTruststore.jks -storepass changeit`**

   The first word in each line indicates the CA alias.

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

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

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

**`C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore CognosTruststore.jks -alias my_alias -file 'path/to/my.pem' -v -trustcacerts`**

.

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

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

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

**2.** Modify the reporting portal registry:

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

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

**3.** Redirect the reporting portal:

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

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

**4.** Enable CAC mode:

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

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

c. Save the file.

d. Restart the server.

**5.** Start the ServletGateway:

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

b. Set the Java Home path:

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

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

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

a. Update the reporting portal in registry entry.

b. Update the reporting portal redirect.

c. Restart the ServletGateway port.

#### After you finish

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

## 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

**Note:** 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 SANscreen 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\jboss\server\onaro\cert\server.keystore
   "c:\Program Files\SANscreen\jboss\server\onaro\cert\server.keystore.old
   ```

2. List the contents of the keystore:

   ```
   C:\Program Files\SANscreen\java\bin\keytool.exe -list -v -keystore "c:
   \Program Files\SANscreen\jboss\server\onaro\cert\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:\ProgramFiles
   \SANscreen\jboss\server\onaro\cert\server.keystore
   ```

4. Generate a new key:

   ```
   C:\Program Files\SANscreen\java\bin\keytool.exe -genkey -alias
   OCI.hostname.com -keyalg RSA -keysize 2048 -keystore "c:\ProgramFiles
   \SANscreen\jboss\server\onaro\cert\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\java\bin\keytool.exe -certreq -alias
OCI.hostname.com -keystore "c:\Program Files\SANscreen\jboss\server
\onaro\cert\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:

```
C:\Program Files\SANscreen\java\bin\keytool.exe -importcert -alias
OCI.hostname.com -file c:\localhost2.DER -keystore "c:\Program Files
\SANscreen\jboss\server\onaro\cert\server.keystore"
```

a. When prompted, enter the keystore password.

The following message is displayed:

```
Certificate reply was installed in keystore
```

**8.** 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 trust stores, 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 Insight Java client or 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 trust store 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.

### Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Setup**.
3. Click the **Backup/ASUP/Proxy** tab.
4. Select **Enable weekly backup**.
5. 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.
6. Select the **Cleanup** option to keep either the last two or the last five backups.
7. Click **Save**.

### Related tasks

*Restoring the Insight database* on page 142

# Configuring your email

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

## Email configuration prerequisites

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

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

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

## Configuring your email for Insight

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

### Steps

1. 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 changes and violations. For example, SNMP notifications are sent when data source thresholds are exceeded.

### Before you begin

The following must have been completed:

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

### Steps

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

   **IP**

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

   **Port**

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

   **Community String**

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

# Enabling the syslog facility

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

### Before you begin

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

#### About this task

The syslog includes the following types of information:

- Violation changes
- Performance alerts

The following units are used in the syslog:

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

#### Steps

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

## Insight syslog contents

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

### Message types

The Insight syslog lists three types of messages:

- SAN path violations
- General violations
- Performance alerts

### Data provided

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

Performance alerts include these data:

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

# Configuring performance and assure violation notifications

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

### Before you begin

You must have 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 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

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

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

### Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Setup**.
3. Click the **Backup / ASUP / Proxy** tab.
4. In the **ASUP** section, select **Enable ASUP** to activate the AutoSupport facility.
5. If you want to change your corporate information, update the following fields:
   - **Company name**
   - **Site name**

6. Select your preferences for the following options:
   - **What to send**: Logs, inventory data, performance data
   - **How to send**: HTTP, HTTPS, FTP, email
7. Enter the size and split file limits:
   - **Size limit**: For email ASUP data transmission, enter the maximum data size to be sent.
   - **Split files if bigger than**: For email ASUP data transmission, indicate the cutoff point for file size.
     The default split file limit is 1 MB.
8. Click **Test Connection** to ensure that the connection that you specified works.
9. Click **Save**.
10. In the **Proxy** section, if you have selected the HTTP or HTTPS protocol for ASUP data transmission to connect to the Internet, enable the proxy, and define it.
11. Click **Test Connection** to ensure that the proxy that you specified works.
12. Click **Save**.

# Creating custom dashboards

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

### Before you begin

You must use an account that has Administrator privileges.

### About this task

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

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

### Steps

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

   The New Dashboard page displays.
3. From the **Customize** menu, select **Rename Dashboard**, type a unique name, and click **OK**.
4. Use one of these methods to begin to add widgets to the dashboard:

- Click the **Add a new widget** link.
- From the **Customize** menu, select **Add Widget**.

Either method opens the Widget Library dialog box, which lists all the available widgets. By default, it displays all the widgets; you can click a category to display only the widgets of a certain type.

5. To add a widget to the dashboard, click **Add Widget** under any widget.
6. To add additional widgets to the dashboard, select **Add Widgets** from the **Customize** menu, click **Add Widget**, and repeat this action until you are satisfied with the contents of the dashboard.

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

   a. Position your pointer in the over the widget title, and click .

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

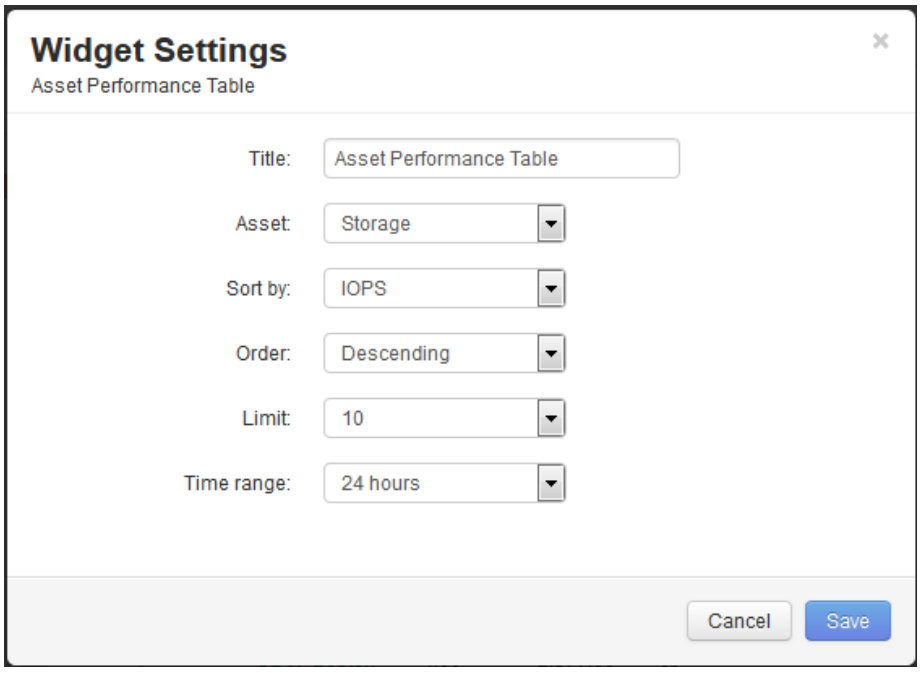


   b. Make the changes you want, and then click **Save**.
8. To delete a widget from the dashboard, perform the following actions:

   a. Click  in the upper-right corner of the widget.

   b. In the confirmation message, click **OK**.
9. To change the order in which the widgets are placed on the dashboard, perform the following actions:

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

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

#### Related references

*Importing and Exporting user data* on page 66

## Deleting a custom dashboard

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

#### Before you begin

You must use an account that has Administrator privileges.

#### Steps

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

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

# 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, storages, 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.

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

3. Click **Add**.

   The Add Application dialog box displays.

4. Enter a unique name for the application in the **Name** box.
5. Click **Priority** and select the priority (critical, high, medium, or low) for the application in your environment.

6. If you plan to use this application with a business entity, click **Business Entity** and select the entity from the list.
7. **Optional**: If you do not use volume sharing, click to clear the **Validate volume sharing** box.

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

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

#### After you finish

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

#### Related tasks

*Assigning applications to assets* on page 35
*Assigning multiple applications to or removing multiple applications from assets* on page 60

## Assigning applications to assets

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

#### About this task

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

#### Steps

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

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

   You can assign multiple applications to host, virtual machine, and internal volume; however, you can only assign one application to volume.

6. Click to assign the selected application or applications to the asset.

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

#### Related tasks

*Assigning multiple applications to or removing multiple applications from assets* on page 60

## Editing applications

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

#### Steps

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

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

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

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

## Deleting applications

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

#### Steps

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

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

# How OnCommand Insight Ethernet monitoring works

OnCommand Insight uses an Ethernet monitoring unit (EMU) to inspect Ethernet packet headers on a switch adjacent to storage, and determines the type of NFSv3 operation being performed, the destination IP address of the operation, and the volumes being addressed.

The EMU continually monitors the mirrored traffic and builds a top-k list of the "k" most heavily loaded volumes, sorted by total observed operation count. For each volume, the EMU generates a top-n list of the top "n" client IP addresses contributing to the load on that specific volume, sorted by total observed operation count. Monitored traffic is only examined to update internal counters relevant to per-volume and per-client operation accounting. The monitored traffic is not persistently stored.

Periodically, collected data is sent to the corresponding Insight server associated with the EMU. Since the Insight server is also acquiring the relevant storage details, for the same volumes and potentially the host details, the server can merge the EMU data into the standard Insight data model, correlating Volume IDs to discovered Internal Volumes and client IPs to resolved physical hosts or discovered Virtual Machines.

The information is available in OnCommand Insight, providing the following:

- Volumes with the highest number of protocol operations
- Clients driving the protocol operations on specific volumes
- Details of the protocol traffic for each volume and clients over a user defined time period

The information allows customers to pro-actively monitor NAS traffic and easily identify greedy clients that are putting the heaviest loads on specific volumes.

## Ethernet monitoring unit network configuration

The Ethernet monitoring unit (EMU) is a separate server connected to the network switches using mirror ports or Switch Port Analyzer 'SPAN' ports. These mirror ports are configured to mirror packets flowing to the ports connected to the storage links. The EMU is designed to parse every packet it receives to determine information about operations and endpoint IP addresses.

The OnCommand Insight server and the EMU server do not require access to the production VLAN traffic. The OnCommand Insight server needs connectivity to the storage array via the management VLAN. The EMU needs connectivity to the Insight server, potentially over the management VLAN. The connectivity between the EMU and the switch connected to the storage array is by way of the mirror ports only. For this configuration, no VLAN connectivity or IP addresses are required for the EMU.

The switches can be 1GbE or 10GbE capable. The amount of traffic on the switch can exceed the capability of the mirror port. Enterprise grade switches expect this to happen and will drop packets in a first-come, first-served basis, without causing network utilization issues. The recommendation is that the mirror port is 10GbE, however 1GbE ports will work, but with a higher probability of packets being dropped. When packets are dropped, the EMU receives a representative portion of all the traffic. Under these overloaded network circumstances, the greedy server creating the load, and the degraded resources continue to be visible in the correct proportions but with lower fidelity.

The illustration below is a simple representation of an Ethernet topology containing an EMU. Clients (any number of physical Hosts or Virtual Machines) are accessing NAS shares accessible through a set of Ethernet switches to a number of storage systems. Ethernet NAS traffic is flowing between the

storage systems and the clients.

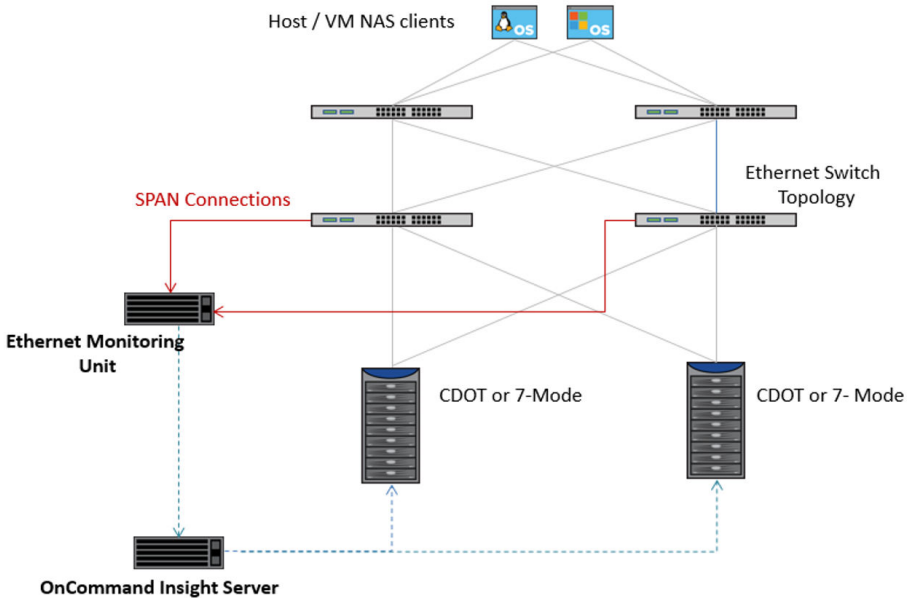


To inspect this traffic an Ethernet Monitoring Unit (EMU) is configured in the network. This is a separate server that connects to the Ethernet through SPAN ports on the switches adjacent to the storage layer. These SPAN ports are configured to mirror all packets flowing through the ports connected to the storage NICs. The EMU parses every packet to determine information about the operation type and endpoint IP addresses. For NFS this information includes:

- Client IP address
- Target (storage) IP address
- Target Volume ID
- Operation type

### Supported platforms

The following hardware platforms are supported for this release of Ethernet monitoring:

- ONTAP 9.0
- Clustered Data ONTAP
- Data ONTAP 7-Mode

### Restrictions

The following restrictions apply to this release of the EMU:

- Only Ethernet/-IPv4/NFSv3 traffic is monitored.
  If other protocols are presented to the EMU, those packets are dropped.
- If the traffic to the storage array is encrypted the EMU does not parse the packets.
- The storage arrays must be monitored by the same OnCommand Insight server that the EMU is connected to.
- If client servers are not discovered by the OnCommand Insight server, only the IP address of the initiator is displayed.

### Ethernet Monitoring Unit monitoring and management

Ethernet Monitoring Units (EMUs) are similar to Acquisition Units (AU). The web UI includes a new Ethernet Monitoring Units administration screen which allows you to perform the following operations:

- View status messages for the EMU
- Restart the EMU
- Delete an EMU (Provided it is not in "Active" status)
- Access the EMU landing page

The EMU landing page page is similar to a Data source status page. It presents the overall status of the single EMU Data source. The EMU details shows the polling status for the last 3 weeks, provides the ability to generate an error report, and the ability to install a Patch.

### Minimum server requirements

The Ethernet Monitoring Unit is a hardware server dedicated to monitoring NFS traffic in an OnCommand Insight environment. The minimum system requirements for Ethernet Monitoring Units and the supported NICs are listed below.

| Component | Required |
|---|---|
| Hardware | Rack-mounted chassis with redundant power supplies |
| Operating system | Red Hat Enterprise Linux (RHEL) 7.2 operating system |
| CPU | Two 2.2 GHz 8-core hyper-threading CPUs |
| Memory | 32 GB RAM |
| Available disk space | 240 GB hard drive |
| Network | 2 - 10GbE NICs for monitoring NFS traffic (from supported list) |
| | 2 - Short Range 10GbE Gigabit Interface Converter (GBIC) |
| | 1 - 1GbE interface for management with the OnCommand Insight server |
| Permissions | Sudo permissions are required on the Server. |

### Supported network interface cards

The Network Interface Cards (NICs) that are supported for monitoring NFS traffic are listed in the table below.

| Manufacturer | Common Linux ethernet driver | Model designation |
|---|---|---|
| Cisco | Enic | UCS Virtual Interface Card |
| Intel | e1000 | 82540, 82545, 82546 |
| Intel | e1000e | 82571,82572, 82573,82574, 82583, ICH8, ICH9, ICH10, PCH1, PCH2, I217, I218 |

| Manufacturer | Common Linux ethernet driver | Model designation |
|---|---|---|
| Intel | igb | 82575, 82576, 82580, I210, I211, I350, I354, DH89xx |
| Intel | ixgbe | 82598,82599, X540, X550 |
| Intel | i40e | X710, XL710, X722 |
| Intel | fm10k | FM10420 |
| Qlogic | bnx2x | 578xx |

## Viewing Ethernet monitored data in the web UI

Data collected by the EMU is presented on the Internal volume landing page and on the Host or Virtual Machine landing page.

On the Internal Volume landing page, the Active NFSv3 contributors table presents the top clients that have been sending NFS traffic to a specific volume, sorted by total operations.

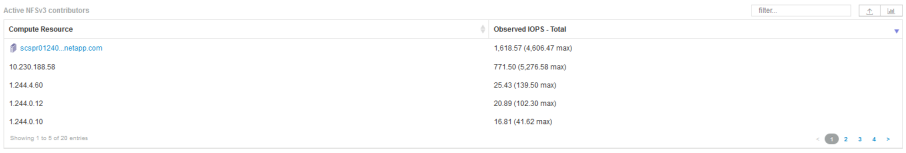


On the Host or Virtual Machine landing pages the NFSv3 activity table lists the storage system and the Internal Volumes that a Host or VM has sent traffic to.

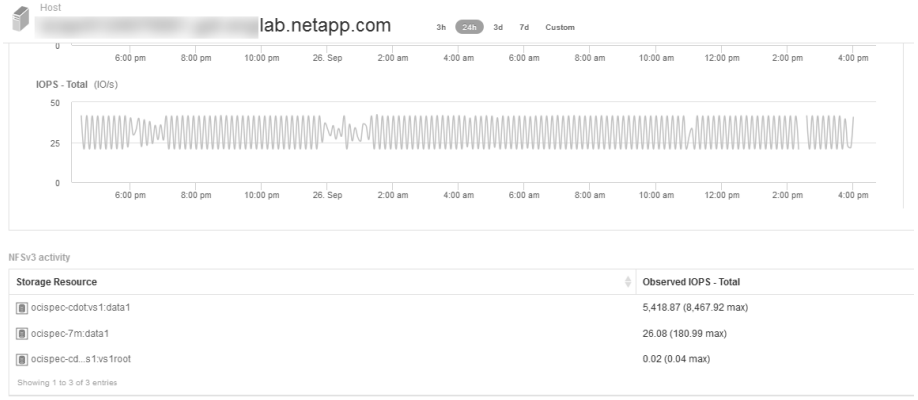


# How OnCommand Insight anomaly detection works

OnCommand Insight contains machine-learning anomaly detection analytics which provide proactive visibility into the infrastructure, supporting business critical applications and informing you of performance anomalies before they become service disruptions.

**Note:** The OnCommand Insight Anomaly detection engine is installed separately from OnCommand Insight. See the *OnCommand Insight Installation Guide* for requirements and installation instructions.

Anomaly analysis helps you identify the normal operating workload range for an application and informs when changes in performance are outside of expected levels. The application anomaly detection engine ingests performance metrics collected by Insight and identifies anomalies in the application infrastructure.

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

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

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

| Objects | Counter |
|---|---|
| VM | Latency, IOPS |
| Hypervisor | CPU utilization, IOPS |
| Edge port | BB credit zero |
| Storage Node | Latency, utilization, IOPS |
| Volume | Latency, IOPS |
| Internal volume | Latency, IOPS |
| Storage pool | IOPS, utilization |

## Configuring applications for anomaly detection

You must turn on anomaly detection in order to monitor an application's infrastructure for abnormal behavior. Anomaly detection monitoring is turned on or off for applications only.

### Before you begin

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

### Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage** and select **Applications**.
3. Locate the application assigned to the computer resource.

4. In the **Anomalies** section, ensure that the **Monitor** control box in the upper right of the section shows **ON** rather than **OFF**; if not, click anywhere in the **Monitor** control box to switch the control to **ON**.

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

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

### Related tasks

*Using the Anomalies section of the application asset page* on page 42
*Assigning applications to assets* on page 35

## Using the Anomalies section

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

### Steps

1. Turn anomaly monitoring on for the desired application and allow it to collect data.
2. While the anomaly detection engine is running, you can view the following information in the **Anomalies** section of the application asset page:
   - The anomaly detection score is shown by a blue bar, representing a score between 0 and 100.
   - Application assets view
     Located below the anomaly detection score, this shows how all the assets contribute to the current anomaly score. If an asset is not contributing significantly to the anomaly score, it is represented by . If the asset is contributing to the significance of the anomaly score, either 1, 2, or three blue bars appear, for example, . The values represented by the number of bars indicate the significance of that asset's contribution to the anomaly score as follows: 1 bar = 1-33; 2 bars = 34-67; 3 bars = 68-100.
     Assets that are assigned to an application (hosts, virtual machines, volumes, internal volumes, and hypervisors) and also assets in the infrastructure that are somehow connected to the assigned assets (for example, nodes, ports, and so on) contribute to the anomaly detection score. Both assets that contribute and those that do not contribute to the specific anomaly are displayed, because the data from the assigned application resources may flow through them.
     If there are more than 60 assets of one type, the **More** button displays, which you can click to display the remaining assets. That status of those remaining assets, for example,
     
     , appears next to the **More** button. This example indicates that of a total of 129 assets not shown, 124 of those assets did not contribute to the anomaly score; however, five of those assets did contribute some level of significance to the anomaly score.
   - Anomaly graph
     The anomaly graph displays the highest anomaly score, indicated by the orange circle, for the time range (3 hours, 24 hours, 3 days, 7 days, or a custom time range) selected for the application asset page. You can drag your pointer over the chart to see the anomaly scores, indicated by a blue circle, for other points during the time range and then click the blue circle, which will update the application assets view to display the assets that contributed to the score at that point in time.

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

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

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

   b. Click .

   The asset is added to the expert view chart.

## Anomaly detection monitoring status

Anomaly detection monitoring status is available in application queries, widgets and asset pages.

### Anomaly Detection monitoring status

For each application, you will see one of the following:

**OK**

No problems are detected.

**Learning**

The anomaly detection server is fetching up to 14 days of performance data, and loading it into the anomaly detection engine.

**Exceeds maximum assets**

The maximum number of assets being monitored for an application has been exceeded. No more than 4000 assets can be monitored for an application. Monitoring is paused until the number of assets is less than or equal to 4000. Reduce the number of assets being monitored to continue monitoring.

**Failed**

The connection to the anomaly detection server has failed.

**Not monitored**

Monitoring is turned off for the application. In queries and widgets, this field is left blank.

# 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

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

### Before you begin

You must have already created a business entity.

### About this task

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

### Steps

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

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

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

# Defining annotations

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

### Steps

1. List any industry terminology to which environment data must be associated.

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

## Using annotations to monitor your environment

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

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

- Creating or edit definitions for all annotation types.
- Displaying asset pages and associating each asset with one or more annotations.
  For example, if an asset is being leased and the lease expires within two months, you might want to apply an end-of-life annotation to the asset. This helps prevent others from using that asset for an extended time.
- Creating rules to automatically apply annotations to multiple assets of the same type.
- Using the annotation import utility to import annotations.
- Filter assets by their annotations.
  You can sort assets by annotations; however, you can only do this in the Java UI.
  See the *OnCommand Insight Inventory User Guide for the Java UI* for more information about grouping assets by annotation.
- Grouping data in reports based on annotations and generate those reports.
  See the *OnCommand Insight Reporting Guide* for more information about reports.

### Managing annotation types

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

### Default annotation types

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

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

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

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

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

| Annotation types | Description | Type |
|---|---|---|
| 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 |

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

### How annotations are assigned

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

Annotations are assigned in the following ways:

- You can assign an annotation manually to an asset.
  If an annotation is assigned directly to an asset, the annotation appears as normal text on an asset page. Annotations that are assigned manually always take precedence over annotations that are inherited or assigned by annotation rules.
- You can create an annotation rule to automatically assign annotations to assets of the same type.

If the annotation is assigned by rule, Insight displays the rule name next to the annotation name on an asset page.

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

## Associating costs with annotations

Prior to running cost-related reports, you should associate costs with the Service Level, Switch Level, and Tier system-level annotations, which enables chargeback to the storage users based on their actual usage of production and replicated capacity. For example, for the Tier level, you might have gold and silver tier values and assign a higher cost to the gold tier than to the silver tier.

### Steps

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

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

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

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

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

## Creating custom annotations

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

### Steps

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

   The Annotations page displays the list of annotations.

3. Click .

   The **Add Annotation** dialog box displays.

4. Enter a name and a description in the **Name** and **Description** fields.

   You can enter up to 255 characters in these fields.

5. Click **Type** and then select one of the following options that represents the type of data allowed in this annotation:

   - Boolean
     This creates a drop-down list with the choices of yes and no. For example, the “Direct Attached” annotation is Boolean.
   - Date
     This creates a field that holds a date. For example, if the annotation will be a date, select this.
   - List
     This can create either of the following:
     - A drop-down fixed list
       When others are assigning this annotation type on a device, they cannot add more values to the list.
     - A drop-down flexible list
       If you select the **Add new values on the fly** option when you create this list, when others are assigning this annotation type on a device, they can add more values to the list.
   - Number
     This creates a field where the user assigning the annotation can enter a number. For example, if the annotation type is “Floor”, the user could select the Value Type of “number” and enter the floor number.
   - Text
     This creates a field that allows free-form text. For example, you might enter “Language” as the annotation type, select “Text” as the value type, and enter a language as a value.

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

6. If you select **List** as the annotation type, do the following:

   a. Select **Add new values on the fly** if you want the ability to add more values to the annotation when on an asset page, which creates a flexible list.

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

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

   a. Click to add additional values.

   b. Click to remove a value.

7. Click **Save**.

   Your annotations appear in the list on the Annotations page.

### After you finish

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

### Related references

*Importing and Exporting user data* on page 66

## Manually assigning annotations to assets

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

### Before you begin

You must have created the annotation you want to assign.

### Steps

1. Log in to the OnCommand Insight web UI.
2. Locate the asset to which you want to apply the annotation by doing either of the following:
   - Click the asset in the Assets Dashboard.
   - Click on the toolbar to display the **Search assets** box, type the type of or name of the asset, and then select the asset from the list that displays.
     The asset page displays.
3. In the **User Data** section of the asset page, click .
   The Add Annotation dialog box displays.
4. Click **Annotation** and select an annotation from the list.
5. Click **Value** and do either of the following, depending on type of annotation you selected:
   - If the annotation type is list, date, or Boolean, select a value from the list.
   - If the annotation type is text, type a value.
6. Click **Save**.
   The annotation is attached to the asset. You can later filter assets by annotation using a query. If you want to group assets by annotations, you must use the Java UI. Refer to the *OnCommand Insight Inventory User Guide for the Java UI* for more information.
7. If you want to change the value of the annotation after you assign it, click and select a different value.
   If the annotation is of list type for which the **Add values dynamically upon annotation assignment** option is selected, you can type to add a new value in addition to selecting an existing value.

### Related tasks

*Editing or removing multiple annotations from assets* on page 61
*Assigning annotations to assets using annotation rules* on page 53

*Creating a query* on page 57

## Modifying annotations

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

### Steps

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

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

   The **Edit Annotation** dialog box displays.
4. You can make the following modifications to an annotation:

   a. Change the name, description, or both.

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

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

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

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

      You cannot delete an annotation value if that value is associated with an annotation contained in an annotation rule, query, or performance policy.
5. Click **Save** when you finish.

### After you finish

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

## Deleting annotations

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

### Steps

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

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

   A confirmation dialog box displays.

4. Click **OK**.

## Assigning annotations to assets using annotation rules

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

## Default storage annotation rules

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

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

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

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

## Creating annotation rules

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

### Before you begin

You must have created a query for the annotation rule.

### About this task

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

### Steps

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

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

.

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

      This name will appear in the Annotation Rules page.

b. Click **Query** and select the query that OnCommand Insight should use to apply the annotation to assets.

c. Click **Annotation** and select the annotation you want to apply.

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

For example, if you choose Birthday as the annotation, you specify a date for the value.

5. Click **Save**.

6. Click **Run all rules** if you want to run all the rules immediately; otherwise, the rules are run at a regularly scheduled interval.

#### Related tasks

*Creating a query* on page 57

## Setting annotation rule precedence

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

#### Steps

1. Log in to the Insight web UI.

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

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

3. Position your cursor over an annotation rule.

The precedence arrows appear to the right of the rule.

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

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

## Modifying annotation rules

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

#### Steps

1. Log in to the OnCommand Insight web UI.

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

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

3. Locate the rule that you want to modify:

- On the Annotation Rules page, you can filter the annotation rules by entering a value in the filter box.
- Click a page number to browse through the annotation rules by page if there are more rules than fit on a page.

4. Perform one of the following to display the **Edit Rule** dialog box:

- If you are on the Annotation Rules page, position your cursor over the annotation rule and click .
- If you are on an asset page, position your cursor over the annotation associated with the rule, position your cursor over the rule name when it displays, and then click the rule name.

5. Make the required changes and click **Save**.

## Deleting annotation rules

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

### Steps

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

   The Annotation Rules page displays the list of existing annotation rules.
3. Locate the rule that you want to delete:
   - On the Annotation Rules page, you can filter the annotation rules by entering a value in the filter box.
   - Click a page number to browse through the annotation rules by page if there are more rules than fit on a single page.
4. Point the cursor over the rule that you want to delete, and then click .

   A confirmation message is displayed, prompting whether you want to delete the rule.
5. Click **OK**.

## Importing annotation values

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

### About this task

The following rules apply:

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

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

The following object types 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 Name >*]<br>*<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 [-uusername] [-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.

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

**Note:** 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.

# Querying assets

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

## Creating a query

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

### About this task

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

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

### Steps

**1.** Log in to the OnCommand Insight web UI.

2. Click **Queries** and select **+ New Query**.
3. Click **Select Resource Type** and select a type of asset.

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

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

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

, and do either of the following:

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

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

   A column is added to the Query results list for the criteria and the results of the query in the list updates.
7. Optionally, you can click to remove an annotation or performance metric from the query results.

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

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

### Related references

*Importing and Exporting user data* on page 66

## Viewing queries

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

### Steps

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

## Exporting query results to a .CSV file

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

### Steps

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

   The Queries page is displayed.
3. Click a query.
4. Click to export query results to a `.CSV` file.
5. Do one of the following:
   - Click **Open with** and then **OK** to open the file with Microsoft Excel and save the file to a specific location.
   - Click **Save file** and then **OK** to save the file to your Downloads folder.

   All of the objects attributes are exported to the file.

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

   **Note:** 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.

## Modifying queries

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

### Steps

1. Log in to the Insight web UI.

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

   The Queries page is displayed.

3. Click the query name.

4. To remove a criterion from the query, click .

5. To add a criteria to the query, click More, and select a criteria from the list.

6. Do one of the following:

   - Click **Save** to save the query with the name that was used initially.
   - Click **Save as** to save the query with another name.
   - Click **Rename** to change the query name that you had used initially.
   - Click **Revert** to change the query name back to the one that you had used initially.

## Deleting queries

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

### Steps

1. Log in to the Insight web UI.

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

   The Queries page displays.

3. Position your cursor over the query you want to delete and click .

   A confirmation message displays, asking if you want to delete the query.

4. Click **OK**.

## Assigning multiple applications to or removing multiple applications from assets

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

### Before you begin

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

### Steps

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

   The Queries page displays.

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

   The list of assets associated with the query display.

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

The **Actions** button displays.

4. To add an application to the assets, click 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.

#### Related tasks

*Creating a query* on page 57

## Editing or removing multiple annotations from assets

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

#### Before you begin

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

#### Steps

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

   The Queries page displays.

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

   The list of assets associated with the query display.

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

   The **Actions** button displays.

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

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

#### Related tasks

*Creating a query* on page 57

# Managing performance policies

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

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

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

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

## Creating performance policies

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

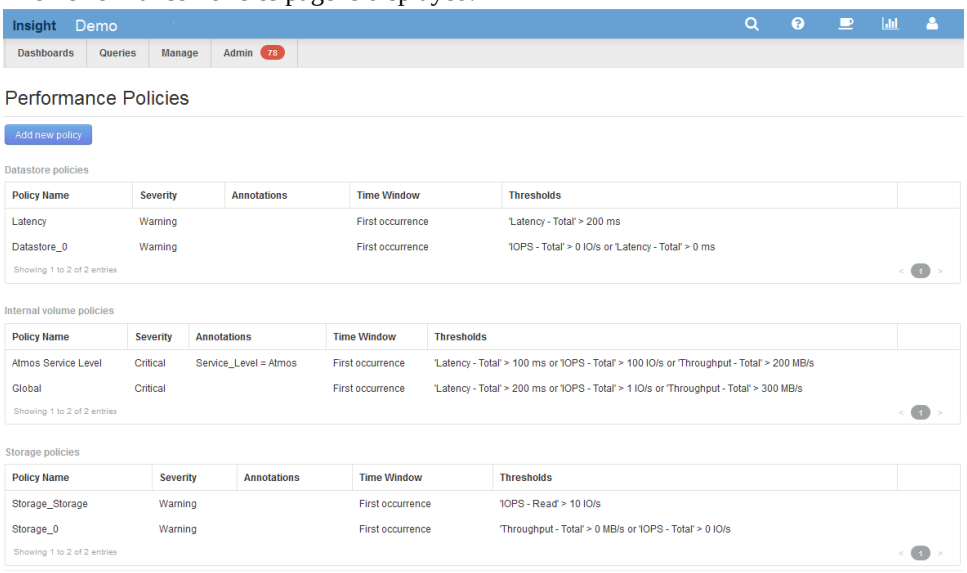


Policies are organized by object, and are evaluated in the order in which they appear in the list for that object.

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.

    #### Example

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

16. Click **Save**.

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

#### Related references

*Importing and Exporting user data* on page 66

## 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.

    **Note:** 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.

    **Note:** 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.

    **Note:** 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.

  **Note:** 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.

    **Note:** 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 data source management

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

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

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

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

#### Related concepts

*What data source patches are* on page 129
*Changes causing data collection problems* on page 122

## Setting up your data sources in Insight

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

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

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

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

#### Related concepts

*Insight data source management* on page 68

### Related tasks

*Cloning a data source* on page 71

### Related references

*Vendor-specific data source reference* on page 73

## 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 latest version of the document can be found at the NetApp Support Site: *mysupport.netapp.com/NOW/products/interoperability/*.

## 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 **Actions** button and select **Add data source**.

   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.<br>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. |

5. Click the **Configuration** link and enter the basic setup information required for the data source with your selected data collection type.
6. 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.

7. For details about configuration or advanced configuration information required or available for your specific data source, see the *Vendor-specific data source reference* on page 73.
8. Click the **Test** link to be certain that the data source is properly configured.
9. Click **Save**.

#### Related concepts

*Data source support information* on page 69

#### Related tasks

*Cloning a data source* on page 71

#### Related references

*Vendor-specific data source reference* on page 73
*Testing the data source configuration* on page 72

## 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:

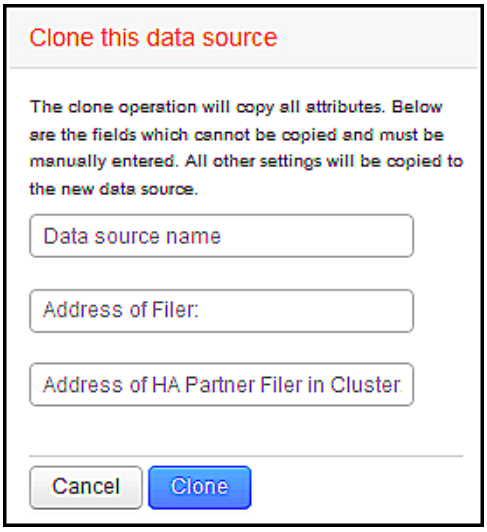


4. Enter the required information in the fields; those details cannot be copied from the existing data source.
5. Click **Clone**.

#### Result

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

#### Related concepts

*Data source support information* on page 69

#### Related references

*Vendor-specific data source reference* on page 73

## Testing the data source configuration

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

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

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

#### Related tasks

*Adding data sources* on page 69
*Editing data source information* on page 127

# Vendor-specific data source reference

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

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

### Related concepts

*Data source support information* on page 69

### Related tasks

*Adding data sources* on page 69
*Cloning a data source* on page 71

## 3PAR InServ Firmware 2.2.2+ (SSH) data source

OnCommand Insight uses the 3PAR InServ Firmware 2.2.2+ (SSH) data source to discover inventory 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 |

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

### Requirements

- For inventory, read-only user name and password to the InServ Server.
- For performance, read-write user name and password to the InServ Server.
- Validated access by using `telnet` to port 22

### 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 |

| Field | Description |
|---|---|
| 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

| | |
|---|---|
| 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 |

## 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

- 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 |

| Field | Description |
|---|---|
| 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 Firmware (SSH) data source

OnCommand Insight uses the Brocade FC Switch 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 |

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

### Requirements

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

### 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 |

| Field | Description |
|---|---|
| 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 (SNMP) data source

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

### Requirements

- There must be IP connectivity to all switches in the fabric. If you select the Discover all switches in the fabric check box, OCI identifies all the switches in the fabric; however, it needs IP connectivity to these additional switches to discover them.
- 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. |

| Field | Description |
|---|---|
| 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 |

**Note:** 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
- 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 |

| Field | Description |
|---|---|
| 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 |

| Field | Description |
|---|---|
| Primary Alias Type | Provide a first preference for resolution of the alias. Choose from the following:<br>**Device Alias**<br>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).<br>**None**<br>Do not report any alias<br>**Port Description**<br>A description to help identify the port in a list of ports<br>**Zone Alias (all)**<br>A user-friendly name for a port that can be used only for zoning configuration<br>**Zone Alias (only active)**<br>A user-friendly name for a port that can be used only for the active configuration. This is the default. |
| Secondary Alias Type | Provide a second preference for resolution of the alias |
| 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) |

## Dell Compellent Storage Center data source

The Dell Compellent Storage Center data source collects network information for analysis and validation from Dell Compellent Storage Center storage arrays using the Dell Compellent Enterprise Manager. This data source uses the Dell Java API.

### Terminology

OnCommand Insight acquires the following inventory information from the Dell Compellent 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 Folder | Disk Group |
| Storage Center | Storage |
| Controller | Storage Node |
| Storage Type | Storage Pool |

| Vendor/Model Term | Insight Term |
|---|---|
| Volume | Volume |
| Fiber Channel I/O Port | Port |

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

### Requirements

- Administrator credentials for the Dell Compellent Enterprise Manager server.

### Configuration

| Field | Description |
|---|---|
| Address of Enterprise Manager | IP address of the Compellent Enterprise Manager server |
| Port | Port used for communication with of the Compellent Enterprise Manager server |
| User Name | User name for the Compellent Enterprise Manager server |
| Password | Password used for the Compellent Enterprise Manager server |

### Advanced configuration

| Field | Description |
|---|---|
| Poll Interval (min) | Interval between inventory polls (default 60 minutes) |
| Performance Poll Interval (up to 3600 seconds) | Interval between performance polls (default 300 seconds) |

## EMC Atmos (HTTP) data source

The EMC Atmos (HTTP) data source supports the EMC cloud offering, which includes a single pool of storage.

### Requirements

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

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

### Implementation considerations

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

- Atmos data appears only in the Reporting Storage and Storage Pool Capacity data model and the Chargeback data model.
- Atmos has a physical hierarchy and a logical hierarchy that impact what you see in reports.

- Atmos resource management group (RMG) capacity is tracked in Insight reporting as storage raw capacity.
- Atmos data is not grouped by and not reported on by tier.
- Atmos data can be reported on by an Atmos service level.

### Atmos logical hierarchy and OnCommand Insight business entity divisions

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

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

### Configuration

| Field | Description |
|---|---|
| Address of Atmos | IP address or fully-qualified domain name for the Atmos node that is currently running the EMC Capacity Reporting Tool |
| User Name | User name for he EMC Capacity Reporting Tool API |
| Password | Password for the EMC Capacity Reporting Tool API |

### Advanced configuration

| Field | Description |
|---|---|
| Port | Port used for communication with Atmos |
| HTTPS Enabled | Select to use HTTPS |
| Inventory Poll Interval (min) | Interval between inventory polls (default 30 minutes) |

## EMC Celerra (SSH) data source

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

### 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 |

| Vendor/Model Term | Insight Term |
|---|---|
| Data Mover | Controller |
| File System mounted on a data Mover | File Share |
| CIFS and NFS Exports | Share |
| Disk Volume | Backend LUN |

**Note:** 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 access

### 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 |

**Note:** 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.

### 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 `navicli.exe` OR `naviseccli.exe` 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 (navseccli) |
| Scope | The secure client scope. The default is Global. |
| CLARiiON CLI Port | Port used for CLARiiON CLI |

| Field | Description |
|---|---|
| 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 |

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

### Requirements

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

| 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 r 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. |

## 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 |

**Note:** 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 |

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

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

### 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 |

**Note:** 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>stordaemon 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`

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

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

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

### Configuration

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

Definitions are supplied for Fibre Channel and Switch Performance packages.

| Field | Description |
|---|---|
| Service Name | 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 | Choose connection caching method:<br>• 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.<br>• 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.<br>• 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.<br>**Note:** This setting does not change OnCommand Insight behavior with respect to the arrays listed as REMOTE by the "symcfg list" output. OnCommand Insight gathers data only on devices shown as LOCAL by this command. |
| CLI Timeout (sec) | 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 (SSH) data source

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

### Configuration

| Field | Description |
|---|---|
| VNX IP | 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 and configuration support for EMC VNXe unified storage arrays. Insight currently supports iSCSI and NAS protocols.

The VNXe data source is CLI based; you must install the Unisphere for VNXe CLI, (uemcli.exe) onto the acquisition unit where your VNXe data source will live. 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. The EMC VNXe data source provides NAS and iSCSI support for inventory; fibre channel volumes will be discovered, but Insight will not report on FC mapping, masking, or storage ports.

### 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 |

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

### Configuration

| Field | Description |
|---|---|
| VNXe IP | IP address or fully-qualified domain name of the VNXe device |
| VNXe User Name | User name for the VNXe device |
| VNXe Password | Password for the VNXe device |
| Full path to the uemcli executable | Full path to the `uemcli.exe` executable |

### Requirements

### 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 |

**Note:** 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
- 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 |

| Field | Description |
|---|---|
| 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 (HTTP) data source

To configure the EMC XtremIO 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 |

**Note:** 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
- Administrator user name and password for the XtremIO Management Server

### 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,<br>Flexible Tier Pool,<br>Raid Group | Storage Pool |
| Standard Volume,<br>Snap Data Volume (SDV),<br>Snap Data Pool Volume (SDPV),<br>Thin Provisioning Volume (TPV) | Volume |

**Note:** 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
- Administration-level user name and password

- Access validation by using SSH to the configured IP and logging in
- Ensure that the page scroll is disabled.

### 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) |

## 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 |

**Note:** 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
- 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 - USP, USP V, VSP/R600 Storages | Device list for VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages. Each storage requires:<br>• Array's IP: IP address of the storage<br>• User Name: User name for the storage<br>• Password: Password for the storage<br>• Folder Containing Export Utility JAR Files: The folder containing the Export utility `.jar` files |

| Field | Description |
|---|---|
| SNM2Devices - WMS/SMS/AMS Storages | Device list for WMS/SMS/AMS storages. Each storage requires:<br>• Array's IP: IP address of the storage<br>• Storage Navigator CLI Path: SNM2 CLI path<br>• Account Authentication Valid: Select to choose valid account authentication<br>• User Name: User name for the storage<br>• 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 |

**Note:** 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) |

## HDS NAS (HNAS) data source

The HDS NAS (HNAS) data source is an inventory and configuration datasource 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 datasource is SSH based, so the acqusiition 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 |

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

### 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 |

| Vendor/Model Term | Insight Term |
|---|---|
| Array Group, DP Pool | Storage Pool |
| Logical Unit, LDEV | Volume |

**Note:** 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
- 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. |

| Field | Description |
|---|---|
| Devices - USP, USP V, VSP/R600 Storages | Device list for VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages. Each storage requires:<br>• Array's IP: IP address of the storage<br>• User Name: User name for the storage<br>• Password: Password for the storage<br>• Folder Containing Export Utility JAR Files: The folder containing the Export utility `.jar` files |
| SNM2Devices - WMS/SMS/AMS Storages | Device list for WMS/SMS/AMS storages. Each storage requires:<br>• Array's IP: IP address of the storage<br>• Storage Navigator CLI Path: SNM2 CLI path<br>• Account Authentication Valid: Select to choose valid account authentication<br>• User Name: User name for the storage<br>• 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 |

**Note:** 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) |

| Field | Description |
|---|---|
| Export timeout in seconds | Export utility timeout (default 300 seconds) |

## HP EVA Storage (SSSU) data source

For configuration, The EVA 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 |

**Note:** 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 `sssu.exe` 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 sssu.exe is located. To validate access, run `sssu.exe` |
| Command Timeout (sec) | `evaperf` command wait timeout (default 600 seconds) |
| Performance Exclude Devices | Comma-separated list of device names to exclude from collecting performance data |

## IBM DS (CLI) data source

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

### 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 |

| Vendor/Model Term | Insight Term |
|---|---|
| Fixed Block Volume | Volume |

**Note:** 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

### 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 `dscli.exe` 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 |

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

### Requirements

- User name and password that provide access to HMC through SSH
- 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.

### 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 (SSH) data source

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

### 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 |

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

### Inventory Requirements

- IP address of each SVC cluster
- Public and private key pair that you either generate withInsight or reuse a keypair already in use on your SVC

  If you are reusing an existing keypair, you must convert them from Putty format to OpenSSH format.
- Public key installed on the SVC cluster
- Private key needs to be identified in the data source
- Access validation: Open `ssh` session to the SVC cluster using the private key

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

### Performance Requirements

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

  **Note:** Because this access level is not required for the SVC foundation discovery package, the SVC foundation user might not work successfully.
- A private and public SSH key must be generated for this user, and the private key stored so that it is accessible from the Acquisition Unit. If the SVC foundation user has the proper permissions,

then the same user and key works. The same SSH key can be used for inventory and performance data.

- Enable data collection by connecting to the SVC cluster by SSH and running: `svctask startstats -interval 1`

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

### 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

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 (CLI) data source

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

### 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 |

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

### Requirements

- Port requirement: TCP port 5989
- 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 |

| Field | Description |
|---|---|
| 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 (HTTP) data source

OnCommand Insight used this data source 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 |

**Note:** 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 ) |

| Field | Description |
|---|---|
| Inventory Poll Interval (min) | Interval between inventory polls (default 60 minutes) |
| Connection Timeout | Connection timeout (default 60 seconds) |

## Microsoft Hyper-V (WMI) data source

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

### 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 |

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

### Requirements

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

### 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) |

| Field | Description |
|---|---|
| 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 |

**Note:** 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
- FCP license and mapped/masked volumes required for Insight discovery
- CIFS and/or NFS license and CIFS shares or NFS exports required for Insight NAS discovery
- Access validation: Telnet or SSH to the cluster management LIF from the Insight server using the username and password.

### Configuration

| Field | Description |
|---|---|
| NetApp Management IP | IP address or fully-qualified domain name of the NetApp cluster |
| User Name | name for the NetApp cluster |
| Password | Password for the NetApp cluster |

### Advanced configuration

| Field | Description |
|---|---|
| Connection Type | Choose connection type |
| Communication Port | Port used for NetApp API |
| Inventory Poll Interval (min) | Interval between inventory polls (default 20 minutes) |
| Performance Poll Interval (sec) | Interval between performance polls (default 300 seconds) |

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

This data source collects clustered Data ONTAP 8.1.x data from the Unified Manager (UM) 6.0+ database. Using this data source, Insight discovers all clusters configured and populated in UM. For efficiency, Insight does not call ZAPIs on the cluster itself. Performance is not supported in this data source.

### Configuration

| Field | Description |
|---|---|
| Unified Manager IP | IP address 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 |

| Vendor/Model Term | Insight Term |
|---|---|
| LUN | Volume |
| Volume | Internal Volume |

**Note:** 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
- User name and password for the controller and the partner
- FCP license and mapped or masked volumes required for OnCommand Insight discovery
- CIFS or NFS licenses and CIFS shares or NFS exports required for OnCommand Insight NAS discovery
- Access validation
  Use telnet to access the FAS controller from the OnCommand Insight Server using the user name and password.

### 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 |

| Vendor/Model Term | Insight Term |
|---|---|
| Volume | Volume |

**Note:** 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
- Access validation by pinging the IP address of both controllers on the array

### 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) |

## 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.

### 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<br>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.<br>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 will reside will need to be able to initiate HTTPS connections to TCP port 443 on the SolidFire cluster management IP address. the data source will need 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 |

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

### Configuration

| Field | Description |
|---|---|
| SolidFire MVIP | Management Virtual IP address of the SolidFire Cluster |
| SolidFire User Name | Name used to log into the SolidFire cluster |
| SolidFire 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) |

## Pure Storage FlashArray (HTTP) data source

The FlashArray 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 |

| Vendor/Model Term | Insight Term |
|---|---|
| Array | Storage |
| Controller | Storage Node |
| Volume | Volume |
| Port | Port |
| LUN Map (Host, Host Group, Target Port) | Volume Map, Volume Mask |

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

### 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 perfromance polls (default 300 seconds) |

## QLogic FC Switch (SNMP) data source

For configuration, this data source requires the network address for the FC Switch device, specified as an IP address, and an SNMP *read-only* community string used to access the device.

### Configuration

| Field | Description |
|---|---|
| SANSurfer Switch | 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) |

## Violin 6000-Series Flash Memory Array (HTTP) data source

This data source collects network information for analysis and validation from Violin 6000-series flash memory arrays.

### Terminology

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 |

**Note:** 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 Master Gateway | IP address or fully-qualified domain name of the Violin Memory Array Master Gateway |
| User Name | User name for the Violin Memory Array Master Gateway |
| Password | Password for the Violin Memory Array Master 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 (Web Services) data source

The VMware vSphere 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 |

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

### Requirements

- Read-only privileges on all objects within the Virtual Center.
- SDK access on the Virtual Center server
- Validate access by logging in to Virtual Center Client using your user name and password.

- Verify that the SDK is enabled by entering

  **`telnet <vc_ip> 443`**

  .
- This data source uses the VMware Infrastructure (VI) SDK 2.5.0

### Configuration

| Field | Description |
|---|---|
| Virtual Center Address | Network address for the Virtual Center or vSphere server, specified as an IP *(nnn.nnn.nnn.nnn* format) address or as a host name that can be resolved through DNS. |
| User Name | 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.

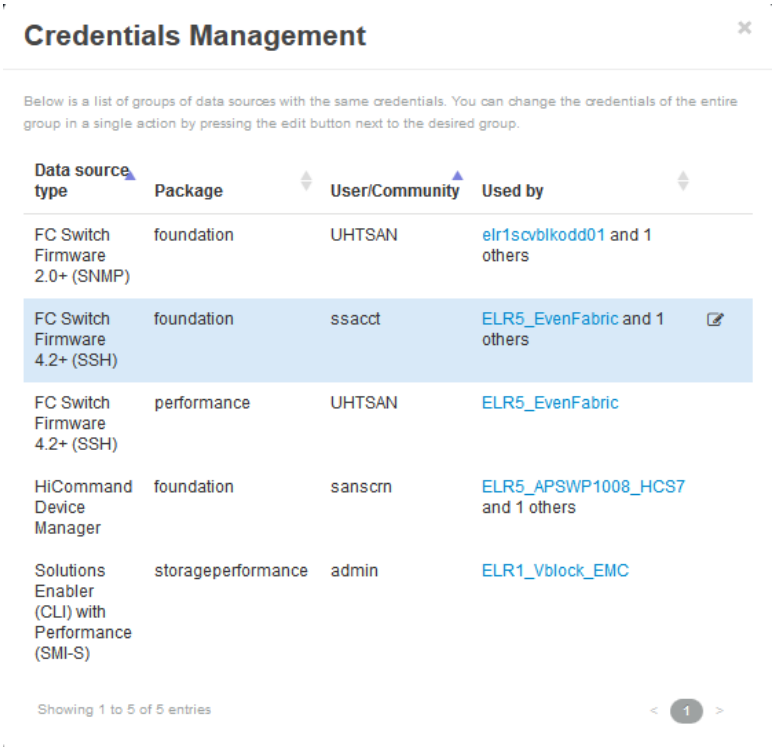


4. Click **Edit**.
5. Enter the new password and confirm it.

# 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 supportto request assistance and possibly a data source patch.
- Are all of your OnCommand Insight acquisition units using the same OnCommand Insight version? If the Remote Acquisition Units and local acquisition unit are running different OnCommand Insight versions, install the same version on all units to correct the data collection problem.
  If you need to install a new version of OnCommand Insight on all of the acquisition units, go to the support site and download the correct version.
- Have you changed any domain names or added a new domain? You must reconfigure the Auto Resolution (AR) methods in the Java UI.

### Related concepts

*What data source patches are* on page 129

### Related tasks

*Deleting a data source* on page 129

## 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.

This example shows a wide variety of information that might be provided for a failed data source.

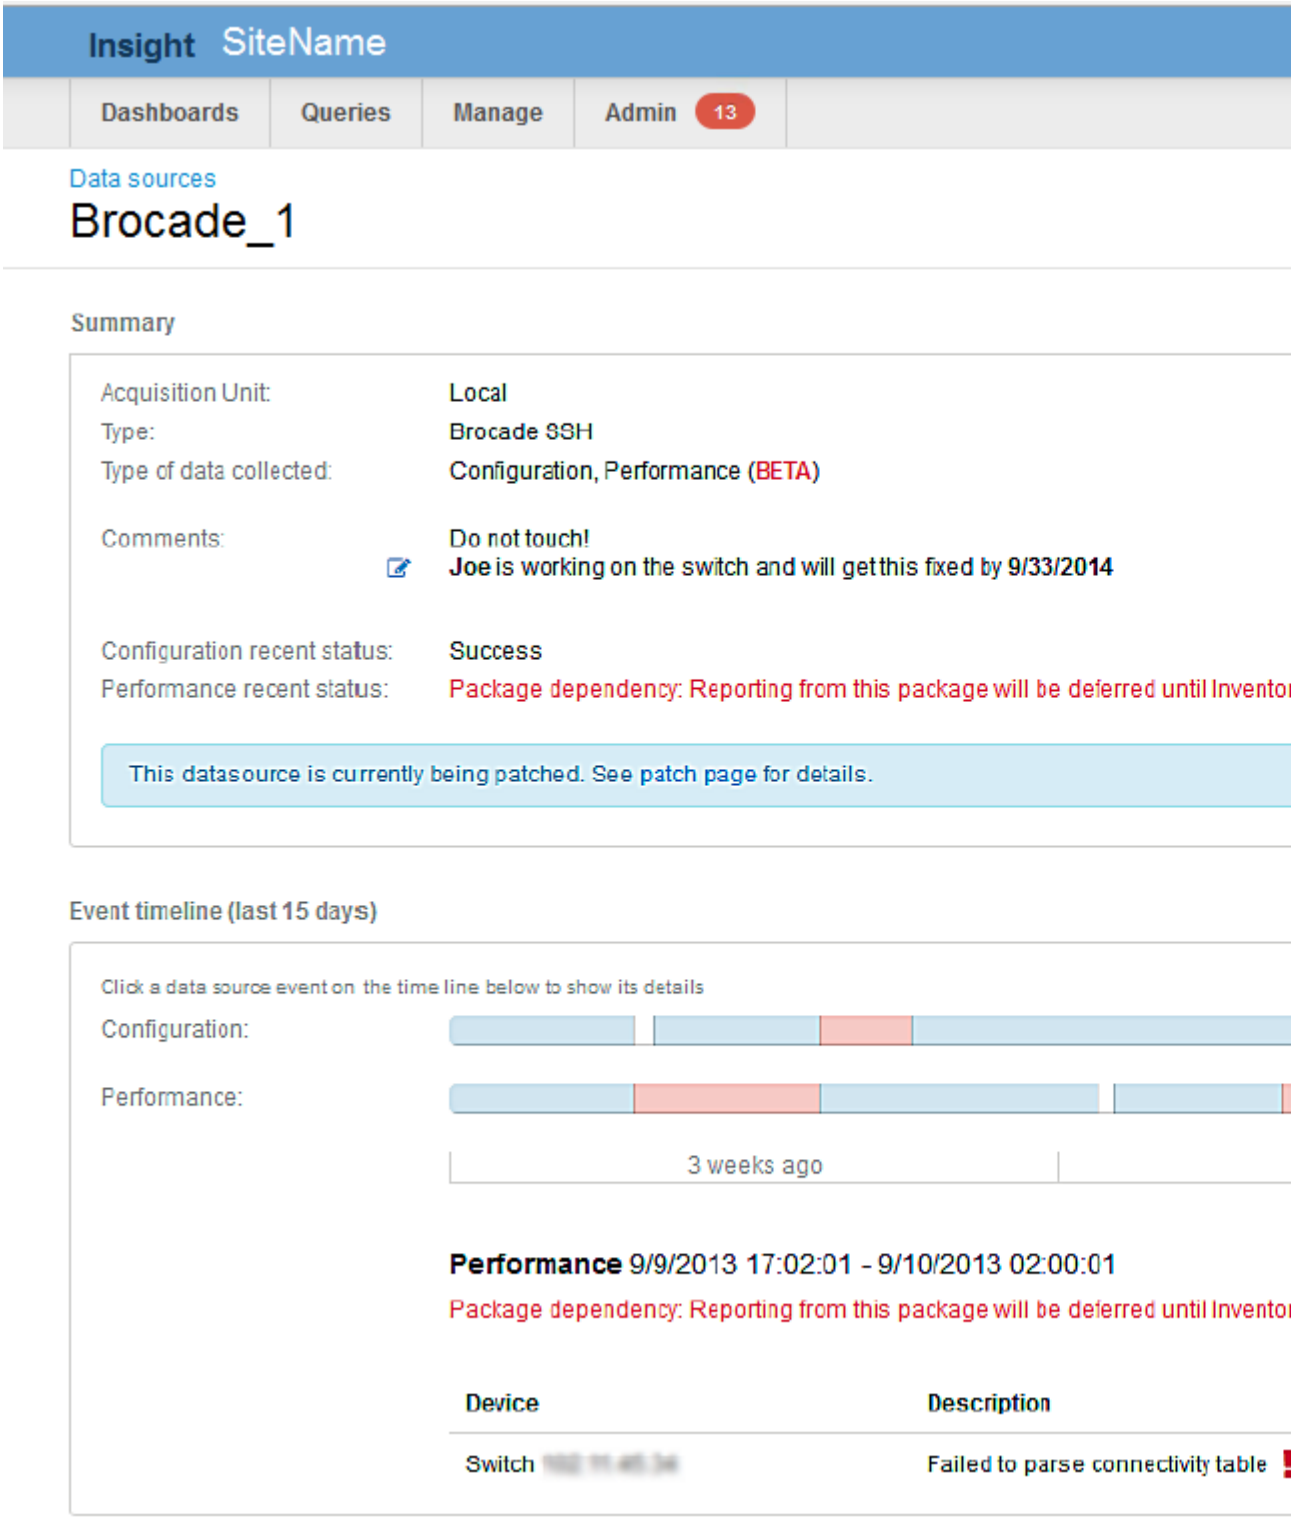


2. Check the **Comments** area to read any notes left by another engineer who might also be investigating this failure.
3. Note any performance messages.
4. If there is a patch being applied to this data source, click link to check the **patch page** to see if that has caused the problem.

5. Move your mouse pointer over the segments of the **Event timeline** graph to display additional information.
6. Select an error message for a Device and displayed below the Event timeline and click the **Error details** icon that displays to the right of the message.

   The Error details include the text of the error message, most likely causes, information in use, and suggestions of what can be tried to correct the problem.
7. In the Devices reported by this data source area, you might filter the list to display only devices of interest, and you can click the linked **Name** of a device to display the *asset page* for that device.
8. To return to previously displayed pages, use one of these techniques:
   - Click the browser back arrow.
   - Right-click the back arrow to display a list of the pages and select the page you want.
9. To display detailed information about other resources, click other linked names.
10. When you return to the data source summary page, check the **Changes** area at the bottom of the page to see if recent changes caused the problem.

### Related concepts

*Changes causing data collection problems* on page 122
*What data source patches are* on page 129
*Data source support information* on page 69

### Related tasks

*Examining one data source in detail* on page 123

## Controlling data source polling

After making a change to a data source, you might want it to poll immediately to check your changes, or you might want to postpone the data collection on a data source for one, three, or five days while you work on a problem.

### Steps

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

   The **Data sources** list opens.
2. Select the data source on which you want to control the polling.
3. Click the data source name link.
4. On the data source summary page, check the information and click one of these two polling options:
   - **Poll again** to force the data source to collect data immediately.
   - **Postpone** and select the length of the polling delay from 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** on the toolbar.

   The **Data sources** list opens.

2. Locate the data source that you want to edit.

3. Use one of these methods to begin the changes:

   - Click **Edit data source** to the right of the selected data source.
   - Click the linked name of the selected data source and click **Edit**.

   Either method opens the Edit data source dialog box.

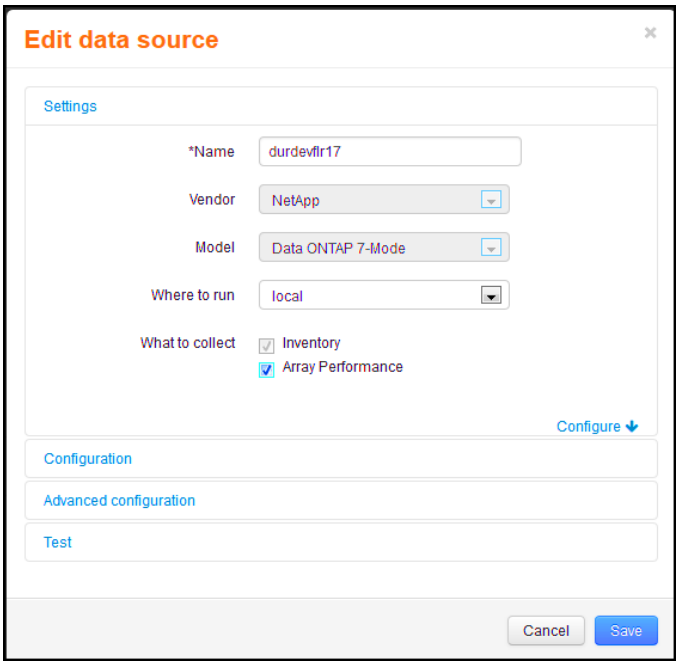


4. Select the required changes.

5. Click **Save**.

# Editing information for multiple data sources

You can edit most of the information for multiple data sources of the same vendor and model at one time. For example, if these data sources share a user name and password, you can change the password in one place and thereby update the password for all the selected data sources.

### About this task

Options that you cannot edit for the selected data sources appear dimmed or are not displayed in the Edit data source dialog box. Additionally, when an option displays a value of **Mixed**, it indicates that

the value for the option varies between the selected data sources. For example, if the **Timeout (sec)** option for two selected data sources is **Mixed**, one data source could have a timeout value of 60 and the other could have a value of 90; therefore, if you change this value to 120 and save the changes to the data sources, the timeout setting for both data sources becomes 120.

### Steps

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

   The **Data Sources** list opens.

2. Click the **Actions** button and select the **Bulk edit** option.

   The Bulk Edit Multiple Data Sources dialog box displays:

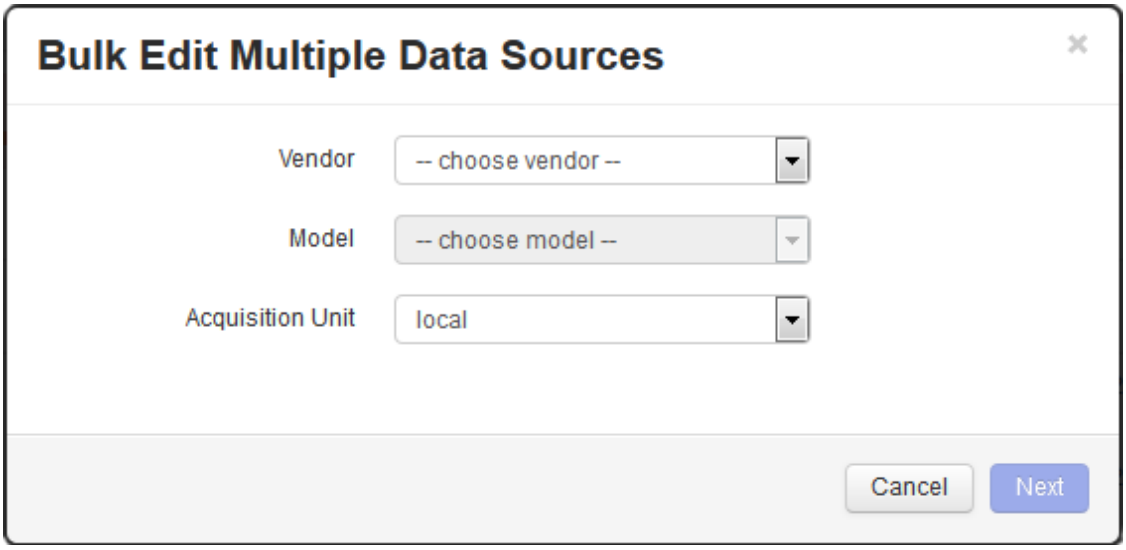


3. Select the vendor of the data sources from the Vendor list.

   If there is only one model for the vendor you selected, the dialog box lists all the data sources available.

4. Select the model of the data sources from the Model list.

   The dialog box lists all the data sources that match the vendor and model that you selected.

5. Select the location of the acquisition unit for the data sources from the Acquisition Unit list.

6. Click **Next**.

   The Edit data source dialog box displays.

7. Change any of the **Settings** options.

8. Click the **Configuration** link to change any of the basic options for the data sources.

9. Click the **Advanced Configuration** link to change any of the advanced options for the data sources.

10. Click **Save**.

# Deleting a data source

If you have removed a data source from your environment, you must also delete it from the OnCommand Insight monitoring environment.

#### Steps

1. On the Insight toolbar, click **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.

#### Related tasks

*Adding data sources* on page 69

## 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.

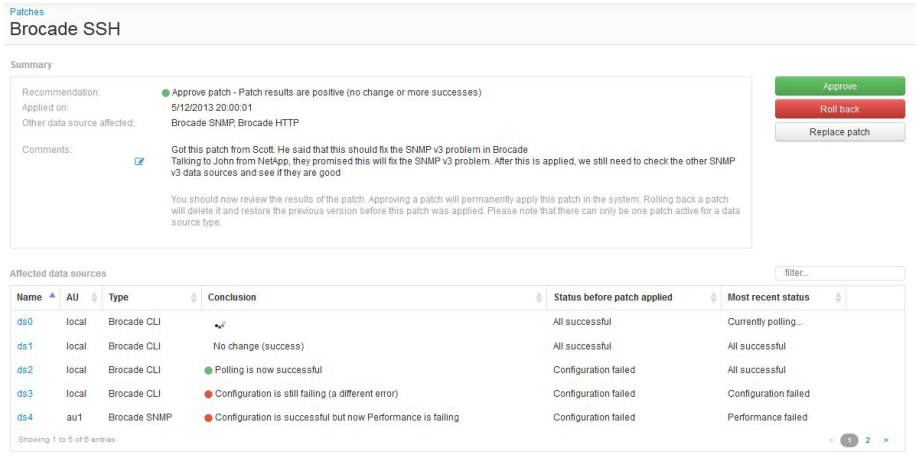


6. Check the **Data sources affected** table to see the status of each affected data source before and after the patch.

   If you are concerned that there is a problem with one of the data sources being patched, click the linked Name in the Data sources affected table.
7. If you conclude that the patch should be applied to that type of data source, click **Approve**.

   The data sources are changed and the patch is removed from Patches currently under review.

## Rolling back a data source patch

If a data source patch is not working in the manner you expected, you can roll it back. Rolling back a patch deletes it, and restores the previous version as it was before this patch was applied.

### Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Patches**.
3. In **Patches currently under review**, click the linked name of the patch that appears to be unsuccessful.
4. On the Patches page for the data source, examine this information:
   - **Summary** describes when the patch was applied, the affected data sources, and comments about the patch from you or other members of your team.
   - **Affected data sources** lists all of the data sources being patched and includes a comparison of the before and after patching status.
5. To display the details for a data source that is not successfully processing the patch, click the linked **Name**.
   a. Check the summary information.
   b. Check the **Event timeline** to see any configuration or performance data that might be affecting this data source.
6. If you conclude that the patch is not going to be successful, click the browser back arrow to return to the Patches summary page.
7. Click **Roll back** to remove that patch.

   If you know of a different patch that is more likely to be successful, click **Replace patch** and upload the new patch.

# Maintaining Insight

Whether you are new to Insight and have a new system to set up, or your system has been operating for some time, you must take steps to maintain smooth operation of Insight and your network. The key maintenance concept is that changes in your network usually need to be accommodated in Insight.

These are the most common maintenance tasks:

- Maintaining Insight backups
- Updating expired Insight licenses
- Coordinating data source patches
- Updating the Insight version on all acquisition units
- Deleting removed data sources from Insight

### Related concepts

*What data source patches are* on page 129

## Managing Insight

OnCommand Insight monitors your environment, enabling you to research potential problems before a crisis is reported. The Assets Dashboard provides summary pie charts, heat maps for IOPs, and an interactive chart of the top 10 utilized storage pools.

### Steps

1. Open the Insight **Assets Dashboard** and move your cursor over the pie charts to examine the asset distribution in these three charts:
   - Capacity by Vendor
   - Capacity by Tier
   - Switch Ports
2. View **Facts About Your Environment** to see information about your environment's used capacity, the capacity's efficiency, consumed FC resources, and virtual infrastructure statistics.
3. Position your cursor over a storage pool bar in the **Top 10 Utilized Pools** chart to view the used and unused capacity of the storage pool.
4. Click any asset name appearing in large text (which indicates that the asset has issues) in the **Storage IOP** heat map to display a page summarizing the current state of that asset.
5. In the lower right corner of the **Assets Dashboard**, click any asset name appearing in large text (which indicates the asset has issues) in the **Virtual Machine IOPS** heat map to display a page summarizing the current state of the asset.
6. On the 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 Insight web 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 | "OK" if number of concurrently-polling data sources is less than 75% of execution pool maximum (default maximum is 30).<br>"Acquisition is busy" if usage is greater than 75%, and recommends increasing polling interval or adding more remote acquisition units. |
| Anomaly Detection Engine | Engine capacity | Count of applications monitored for anomaly detection | Number and percentage of applications being monitored out of 48 possible applications. |

| Component | Test | Details | Displays |
|---|---|---|---|
| Anomaly Detection Engine | Engine status | Status of anomaly detection engine | “OK” if no errors are detected, otherwise displays information about any error found.<br>See the `prelert.log` file for more information. |
| DWH | Backup | Status of Data Warehouse scheduled backup | “OK” and the last successful DWH backup time if DWH scheduled backup is enabled.<br>Otherwise, displays information about any error found. |
| DWH | ETL | Status of Data Warehouse ETL | “OK” and the last successful DWH build time if no errors.<br>Otherwise, displays information about any error found. |
| Server | ASUP | Status of ASUP (AutoSupport Phone Home) | “ASUP Enabled” and the last successful phonehome time if available.<br>“ASUP Failed” if phonehome is enabled but encountered a problem.<br>"Invalid backup location" if backup directory is not valid.<br>Displays the last successful phonehome time as well as time of the last failed attempt if available.<br>“ASUP Disabled” if phonehome is disabled. |

| Component | Test | Details | Displays |
|---|---|---|---|
| Server | Auto resolution | Status of automatic device resolution | “OK” if no errors.<br>“Auto resolution is blocked” if identification errors prevent resolution progress.<br>“Low success rate” if less than 75% of generic devices could be identified. |
| Server | Cassandra | Status of connection to Cassandra database | “OK” if no errors.<br>“Cassandra is down” if unable to connect to Cassandra.<br>“Cassandra has file system errors” if Cassandra connection is good but file system errors are found. |
| Server | CPU | Insight CPU usage | “OK” if CPU load is less than 65%.<br>“System CPU load is high. Reduce your CPU load.” if CPU load is greater than 65%. |
| Server | Disk space | Status of disk space | Free disk space, disk space in use by Insight, and recommended disk space reserved for Insight.<br>“Low Disk Space” if disk space available for Insight is less than 100 gb. |

| Component | Test | Details | Displays |
|---|---|---|---|
| Server | EventBus | Status of EventBus | "EventBus is empty" if EventBus queue is empty.<br>"EventBus is in use" if EventBus queue length is greater than 1000.<br>"EventBus is in use. Slow down datasources" if EventBus queue length is greater than 1500.<br>"EventBus is full. Slow down datasources" if EventBus queue length is greater than 1999. |
| Server | Inventory data processing | Status of inventory data processing capability of Insight server | "OK" if Insight server is not busy.<br>"Server is busy" if the server is busy at least 75% of the time for the last hour. Recommends not adding more data sources as recommends splitting the environment to several servers. |
| Server | MySQL | Status of MySQL database | "OK" if no problems are detected.<br>"The database is having performance issues. Some queries are taking too long to run" if the number of slow queries is more than 5%.<br>"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. |

| Component | Test | Details | Displays |
|---|---|---|---|
| Server | Physical memory | Status of physical memory | "OK" if memory usage is less than 85%.<br>"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 | Violation | Status of open violations | "OK" if the number of open violations is less than 75% of the violations limit.<br>"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.<br>"Violation manager is blocked" if the number of open violations is at the violations limit.<br>Note that the violation manager cannot create new violations and recommends reviewing performance policy configuration. |
| Server | Weekly backup | Status of weekly backup | "OK" if weekly backup is enabled, otherwise displays "Weekly backup is not enabled". |

**Note:** If the anomaly detection engine displays an error, see the `prelert.log` file in the following location for more information:

- Windows: *`disk drive`*`:\`*`install directory`*`\SANscreen\jboss\server\onaro\log`
- Linux: `/var/log/netapp/oci/jboss`

# Auditing system and user activities

If you want to locate unexpected changes, you can view an audit trail of the OnCommand Insight system and its user activities.

### About this task

Insight generates audit entries for any user activities that affect the storage network or its management, including the following:

- Logging in
- Authorizing or unauthorizing a path
- Updating an authorized path
- Setting global policies or thresholds
- Adding or removing a data source
- Starting or stopping a data source
- Updating data source properties
- Adding, editing, or deleting a task
- Removing an application group
- Identifying or changing the identification for a device

### Steps

1. Open Insight in your browser.
2. Click **Admin** and select **Audit**.

   The Audit page displays the audit entries in a table.
3. You can view the following details in the table:

   **Time**

   Date and time that the changes were made

   **User**

   Name of user associated with the audit entry

   **Role**

   User account's role, which is guest, user, or administrator

   **IP**

   IP address associated with the audit entry

   **Action**

   Type of activity in the audit entry

**Details**

Details of the audit entry

If there is a user activity that affects a resource, such as a data source or an application, the details include a link to the resource's landing page.

**Note:** When a data source is deleted, the user activity details related to the data source no longer contain a link to the data source's landing page.

4. You can display audit entries by choosing a particular time period (1 hour, 3 hours, 24 hours, 3 days, and 7 days), with Insight showing a maximum number of 1000 violations for the selected time period.

   You can click a page number below the table to browse through data by page if there is more data than fits on a single page.

5. You change the sort order of the columns in a table to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header; to return to the default sort order, click any other column header.

   By default, the table displays the entries in descending order.

6. You can use the **filter** box to show only the entries you want in the table.

   **Example**

   To see only the audit entries by the user `izzyk`, type `izzyk` in the **filter** box.

# Monitoring the violations in your network

When Insight generates violations due to the thresholds set in performance policies, you can view them using the Violations Dashboard. The dashboard lists all the violations that occur in your network and enables you to locate and address issues.

**Steps**

1. Open OnCommand Insight in your browser.
2. On the Insight toolbar, click **Dashboards** and select **Violations Dashboard**.

   The Violations Dashboard displays.

3. You can use the **Violations By Policies** pie chart in the following ways:

   - You can position your cursor over any slice of a chart to display the percentage of the total violations that occurred for a particular policy or metric.
   - You can click a slice of a chart to “enlarge” it, which enables you to emphasize and study more carefully that slice by moving it away from the rest of the chart.
   - You can click the icon in the upper-right corner to display the pie chart in full screen mode, and click again to minimize the pie chart.

   A pie chart can contain a maximum of five slices; thus, if you have six policies that generate violations, Insight combines the fifth and sixth slices into an “Others” slice. Insight assigns the most violations to the first slice, the second most violations to the second slice, and so on.

4. You can use the **Violations History** chart in the following ways:

- You can position your cursor over the chart to display the total number of violations that occurred at a particular time and the number that occurred out of the total for each specified metric.
- You can click a legend label to remove the data associated with the legend from the chart. Click on the legend to display the data again.
- You can click the icon in the upper-right corner to display the chart in full screen mode, and click again to minimize the pie chart.

5. You can use the**Violations Table** in the following ways:

- You can click the icon in the upper-right corner to display the table in full screen mode, and click again to minimize the pie chart.
  If your window size is too small, then the Violations Table displays only three columns; however, when you click , additional columns (up to seven) display.
- You can display violations for a particular time period (**1h**, **3h**, **24h**, **3d**, and **7d**), with Insight showing a maximum number of 1000 violations for the selected time period.
- You can use the **filter** box to show only the violations you want.
- You can change the sort order of the columns in a table to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header; to return to the default sort order, click any other column header.
  By default, the table displays the violations in descending order.
- You can click a violation in the ID column to display the asset page for the duration of the violation.
- You can click the resource links (for example, storage pool and storage volume) in the Description column to display the asset pages associated with those resources.
- You can click the performance policy link in the Policy column to display the Edit Policy dialog box.
  You might want to adjust the thresholds for a policy if you feel it generates too few or too many violations.
- You can click a page number to browse through data by page if there is more data than fits on a single page.
- You can click to dismiss the violation.

#### Related tasks

*Editing performance policies* on page 65

## Acquisition unit status

The status of the Insight acquisition units connected to your server is displayed in the **Admin** > **Acquisition Units** table. If a unit in the list has a solid red circle in the description, that unit has a serious problem that you must investigate. You can select an entry in this table and restart that acquisition unit.

The Acquisition Units table displays this information for each unit:

- **Name**

- **IP**
- **Status** is the operating status of the acquisition unit.
- **Last reported** displays last time a data source connected to the acquisition unit reported.

If you position your pointer over one of the units in the table, you can click the **Restart Acquisition Unit** icon to restart that unit.

# Restoring the Insight database

To restore your Insight database from a verified backup file, use the Troubleshooting options. This operation completely replaces your current OnCommand Insight data.

### Before you begin

**Best practice:** Before restoring your OnCommand Insight database, use the manual backup process to create a copy of the current database. Check the backup file you plan to restore be certain that it was a successful backup containing the files you want to restore.

### Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Troubleshooting**.

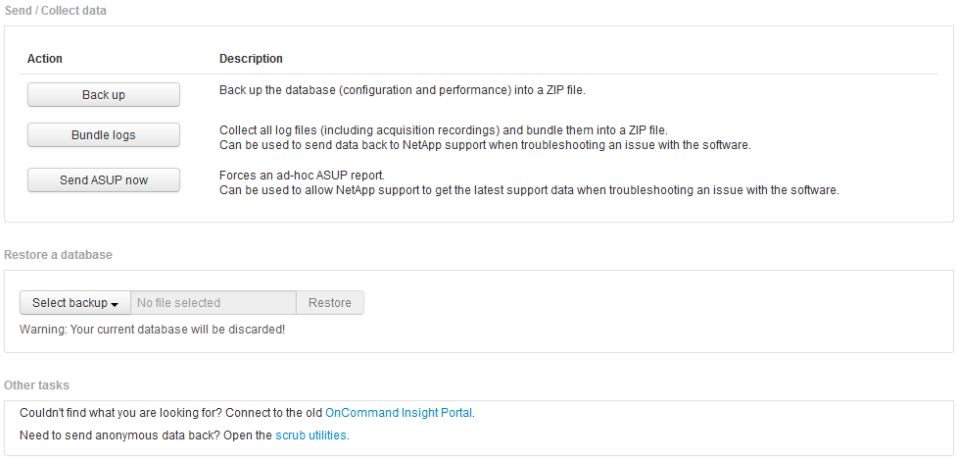


3. In the Restore a database section, select the backup file you want to restore from the **Select Backup** menu.
4. Click **Restore**.
5. On the warning that all data will be replaced, click **OK**

   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**.

## 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, *7.2*`service_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.

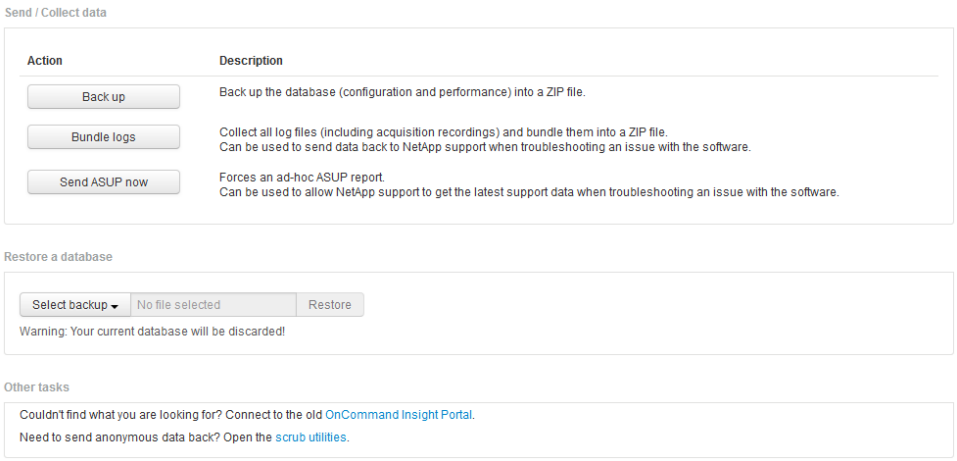


# Advanced configuration tools

To complete your OnCommand Insight configuration, you must use the advanced configuration tools. These tools run in the browser and are opened from the **Admin > Troubleshooting** page.

To open the advanced configuration tools in the browser, click the **OnCommand Insight Portal** link at the bottom of this page:

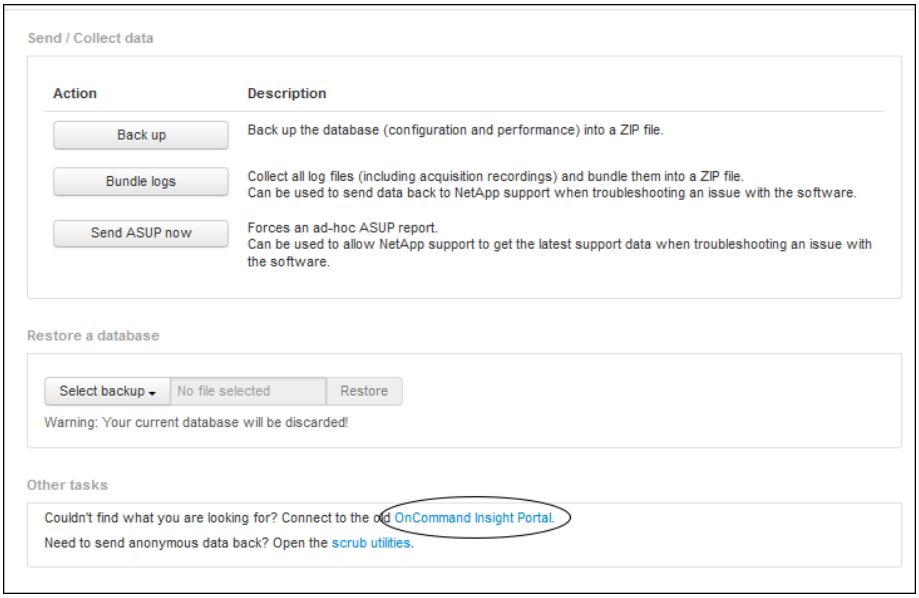


On the next page in the Advanced menu, click the **Troubleshooting** option to display these tools:

- Restart OnCommand Insight Server
- Acquisition
- Logs
- System Information
- Installed Software Packages
- Calculate Scale
- Migrate MySQL Data
- Force Update of Annotations for DWH
- Force Update of Annotations for DWH (include deleted objects)
- Configure Client Settings
- Server Resources Status
- Find Ghost Data Sources

## Restarting the OnCommand Insight Server

When you restart the OnCommand Insight Server, refresh the page and log into the OnCommand Insight Portal again.

### About this task

**Note:** Both of these options should only be used upon request by NetApp Customer Support. There is no confirmation prior to restart.

### Steps

1. 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 the **Restart OnCommand Insight Server**.

## Configuring the number of hours to ignore dynamic data

You can configure the number of hours during which OnCommand Insight ignores updating dynamic data, such as used capacity. If the default of six hours is used and no configuration changes occur, reports will not be updated with dynamic data until after the default number of hours. This option improves performance because this option defers updates when only the dynamic data changes.

### About this task

If a value is set for this option, OnCommand Insight will update dynamic data based on the following rules:

- If no configuration changes occur but capacity data changes, data will not be updated.
- Dynamic data (other than configuration changes) will be updated only after the timeout specified in this option.
- If configuration changes occur, configuration and dynamic data is updated.

Dynamic data impacted by this option includes the following:

- Capacity violation data
- File Systems Allocated Capacity and Used Capacity
- Hypervisor
  - Virtual Disk Used Capacity
  - Virtual Machine Used Capacity
- Internal Volume
  - Data Allocated Capacity
  - Data Used Capacity
  - Deduplication Savings
  - Last Known Access Time
  - Last Snapshot Time
  - Other Used Capacity
  - Snapshot Count
  - Snapshot Used Capacity
  - Total Used Capacity
- iSCSI Session Initiator IPs, Target Session ID, and Initiator Session ID
- Qtree Quota Used Capacity
- Quota Used Files and Used Capacity
- Storage Efficiency Technology, Gain/Loss, and Potential Gain/Loss
- Storage Pool
  - Data Used Capacity

  - Deduplication Savings
  - Other Used Capacity
  - Snapshot Used Capacity
  - Total Used Capacity
- Volume
  - Deduplication Savings
  - Last Known Access Time
  - Used Capacity

#### Steps

1. 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. Select the **Ignore dynamic attribute changes for x hours** option and enter the number of hours that OnCommand Insight should ignore dynamic data.
5. Click **Save Settings**.
6. (Optional) To restart the acquisition unit, click the **Restart Acquisition Unit** link.

   Restating the local acquisition unit reloads all of the OnCommand Insight data source views. This change is applied during the next poll, so you do not have to restart the Acquisition Unit.

## Generating logs for Customer Support

If requested by Customer Support, generate a server, acquisition, or remote log for troubleshooting purposes.

#### About this task

If NetApp Customer Support requests, use this option to generate the logs.

#### Steps

1. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click **OnCommand Insight Portal**.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click **Logs** and select:
   - **Server** - Server log files in {Home}\SANscreen\jboss\server\onaro\log
   - **Acq** - Acquisition log files in {Home} \SANscreen\acq\log
   - **Remote** - Remote logs from Remote Acquisition, OnCommand Insight Client, and Data Warehouse

## Displaying system information

You can display the Microsoft Windows IP configuration information about the system on which OnCommand Insight server is deployed.

### Steps

1. 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. Select **System Information**.

   The Windows IP configuration includes information such as the host name, DNS, IP address, subnet mask, OS information, memory, boot device, and connection name.

## Listing installed OnCommand Insight components

You can display a list of the installed OnCommand Insight components including, among others, inventory, capacity, dimensions, and the Data Warehouse views. Customer Support might ask you for this information, or you might want to see what software versions were installed and when they were installed.

### Steps

1. 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 **Installed Software Packages**.

   OnCommand Insight lists the package names with the software version numbers and installation date and time of each.

## Calculating the number of database objects

To determine the number of objects in the OnCommand Insight database, use the Calculate Scale feature.

### Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click **Calculate Scale**.

   The system displays the number of discovered objects in the database.

## Moving MySQL data using the migrate option

You can use migrate MySQL data directory to a different directory. You can retain the current data directory. You can use the migrate option on the Troubleshooting menu or you can use the command line. This procedure describes how to use the **Troubleshooting > Migrate MySQL data** option.

### About this task

If you retain the current data directory, it will be kept as a backup and renamed.

### Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Select **Migrate MySQL Data**.
5. Enter the path to which you want to migrate the data.
6. To retain the existing data directory, check **Keep existing data directory.**
7. Click **Migrate the data to selected directory**.

## Moving MySQL data using the command line

You can use migrate MySQL data directory to a different directory. You can retain the current data directory. You can use the migrate option on the Troubleshooting menu or alternatively, you can use the command line. This procedure describes how to use the command line.

### About this task

If you retain the current data directory, it will be kept as a backup and renamed.

You can use the Migrate MySQL Data utility or you can use a `java -jar mysqldatamigrator.jar` option in the OnCommand Insight path of `\bin\mysqldatamigrator` where the following parameters should be used:

- Mandatory parameters
  **-path**

  The new data path to which the data folder will be copied.

- Optional parameters
  **-myCnf <my .cnf file>**

  The path for the .cnf file. The default is `<install path>\mysql\my.cnf`. Use this flag only if a non-default MySQL is used.

  **-doBackup**

  If this flag is set, the current data folder will be renamed but not deleted.

### Step

1. Access the command line tool here: <*`installation path`*> \bin\mysqldatamigrator \mysqldatamigrator.jar

**Example usage**

```
java -jar mysqldatamigrator.jar -path "C:\<new path>" -doBackup
```

## Forcing annotation updates

If you have changed the annotations and want to use them in reports immediately, use one of the force annotation options.

### Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Select one of these options:
   - **Force Update of Annotations for DWH** to force the update of annotations in data warehouse to be used for reports.
   - **Force Update of Annotations for DWH (including deleted objects)** to force an annotations update (including deleted objects) in the data warehouse to be used for reports.

## Configuring Client settings

You can configure options for the OnCommand Insight Java UI Client on the server.

### Steps

1. 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. In the Advanced menu on the next page, click the **Troubleshooting** link.
4. Click **Configure Client Settings**.
5. Enter the following information:
   - **Initial Heap Size** is the size of the heap that is created when the VM (Virtual Machine) is started to run the client. This number is used to construct the `-Xms<size>` parameter to `java`. Decreasing this number means Java takes up less memory when it starts up. (However, it is then more likely that Java will need to reallocate the heap when it exceeds this number, which is an expensive operation, so you do not want to make this number too small.)
   - **Max Heap Size** is the largest that Java will allow the heap to grow in the VM that is running the client. This number is used to construct the `-Xmx<size>` parameter to `java`.
   - Log Level controls the amount of detail that is generated in the client logs. There are two possible values:
     - INFO - generates a "normal" amount of information.
     - DEBUG - generates more detailed information to the client log to help with debugging.
6. Click **Save Settings**.
7. Restart the client.

**Note:** If the client generates a "java.lang.OutOfMemoryError: Java heap space" error, you can increase the max heap size setting to see whether that fixes the problem.

## Checking a tier upgrade

After upgrading from OnCommand Insight version 6.0 to a newer version, select the **Check Tier Upgrade** option to list any volume which presently had a different tier value than the tier value it has in the new version. The newer Insight versions improve the way tiers are calculated by using disk attributes instead of volume attributes. This change increases accuracy and reduces confusion when the volume is part of a storage pool with mixed disks.

### About this task

Tier rule configuration no longer presents a disk attribute that is a combination of all the disk attributes of a storage pool, and tiers are calculated for a volume using the lowest performing disk attribute. This change might produce different values if your rules were setup to use the old volume data.

### Steps

1. 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 **Check Tier Upgrade**.
5. Locate the differences and decide if you need to adjust the tier rule.

## Checking the status of server resources

This option displays the OnCommand Insight Server's information including server memory, disk space, OS, and CPU and OnCommand Insight database information including innoDB data size and the disk free space where the database resides.

### Steps

1. 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.

## Scrubbing data for transfer to support

Customers who have secure environments need to communicate with NetApp Customer Service to troubleshoot problems that arise without compromising their database information. The OnCommand Insight Scrub utilities allow you set up a comprehensive dictionary of keywords and patterns so that you can "cleanse" sensitive data and send scrubbed files to Customer Support.

### Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.

2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click **Scrub Utilities**.

   There are three scrub utilities: Dictionary Admin, Scrub, and Lookup Dictionary.
5. Select the **Dictionary Admin** option.
6. Expand **Custom keywords**. Click **Add** to get the dialog box for entering new keywords (note that multiple entries must be line-separated
7. Enter a keyword and click **OK**. The keyword is added to the dictionary.
8. Expand **Patterns (regexp)**. Click **Add** to get the dialog box for entering a new pattern.
9. Enter a Name and a Pattern in the fields and click **OK**. The information has been added to the dictionary.

   **Note:** Patterns must be encompassed by round parentheses to identify a regular expression capturing group.
10. In the Database block, click **Build** to initiate compilation of the dictionary to the dictionary all words identified as sensitive from the OnCommand Insight database.

    On completion, you see a prompt informing you the revised dictionary is available. The Database description includes a line indicating how many keywords are in the dictionary. Check your keywords in the dictionary for accuracy. If you find problems and want to rebuild the dictionary, click **Reset** on the Database block to remove all keywords collected from the OnCommand Insight database from the dictionary. As the prompt advises, all other keywords will not be deleted. Return to the Scrub utilities and enter your Custom Keywords again.
11. Select **Scrub > Scrub log / XML files** again.
12. Browse to locate the log file and click **Scrub**.

    You can run and scrub signed MySQL queries.

    **Note:** The Scrub Utility will not accept an unsigned query. Queries can be signed only by NetApp personnel.

    While restrictions exist, the customer can scrub the query multiple times and send it in again to NetApp Customer Service if fine-tuning of the query is necessary.

### Scrub dictionary administration

The scrub utility allows you to change and enhance the dictionary manually with words and patterns that were not picked up automatically by the system. The dictionary itself is never sent to NetApp and is entirely under your control. NetApp Customer Support can address the problems without needing to view sensitive data, so the keywords and patterns they receive do not interfere with their diagnostics.

To facilitate this interaction and to enable efficient troubleshooting:

- Every troubleshooting session must use the same dictionary. Failure to do so skews the data and prevents successful troubleshooting.
- A lookup dictionary displays the confirmed entries for keywords and patterns. The scrubbing process allows both you and the engineer to adjust and edit the scrubbed files until all possible patterns and keywords have been identified and entered into the dictionary. This ensures future exchanges will be expedited because the Lookup tables will be accurate, thorough and consistent.
- The dictionary remains valid and consistent unless and until you decide to clean (reset) the dictionary itself and build it again from scratch.

### Lookup dictionary output

The Lookup dictionary is the output of the scrub utility that allows you to map back from the scrubbed fields to the real fields (for example, the name of a switch) in the original file.

The lookup utility displays the real values of what was scrubbed in the dictionary file.

**Note:** This is to be used only by the customer. This information is not available anywhere outside the customer's environment.

Information can be displayed by:

- Code assigned to a keyword.
- Keywords associated with a specific code

Therefore, you can perform a search using the {num} or by the entry that {num} replaced.

### Troubleshooting with a scrubbed file

A customer in a secure environment has an unexpected slowdown occur in one of the OnCommand Insight applications. It is not a local problem. The customer contacts NetApp Customer Support. Both Customer Support and an engineer can use the Scrub Utilities to achieve a successful outcome. The process sequence for addressing the problem would resemble the steps shown below.

**At the customer site**

1. Customer Support requests a log file from the customer. ( Types of data that may apply if you need to send NetApp include XML files, log files (for all components) and results of troubleshooting queries on the OnCommand Insight and the data warehouse database.
2. The customer has built a dictionary in the Scrub Utilities already (by clicking **Build** in the "Dictionary Admin" page). Typically the dictionary contains IP addresses, WWNs, VMs.)
3. The customer scrubs the file and checks its outcome. If there are keywords or patterns omitted, the customer edits the dictionary (using the "Dictionary Admin" and scrubs the log file again until it is correct. (These changes are made manually.)

**At NetApp**

1. NetApp Customer Support forwards the customer's scrubbed log file to an engineer.
2. The NetApp Customer Support representatives only see the results of the scrub. The engineer can identify and troubleshoot the problem without needing to know what the scrubbed material represents. Only the customer has access to the Lookup dictionary. Only the customer has the database!
3. If the engineer needs information about the scrubbed data, it can be discussed even by phone in a secure manner.

For example, if the engineer needs to know whether a device is a virtualizer, he can ask the customer this question and cite the code he sees in the log file for the address or name of the device. It could be "{123…}." When he gives the code to the customer, the customer can check using the "Lookup dictionary" by entering the code and receiving the real identity of the device. The customer is able to tell the support engineer that the device is or is not a virtualizer without ever needing to identify it to NetApp.

### Searching in the lookup dictionary

If someone at the customer site wants to know the real identity of something references in the scrubbed file, the lookup dictionary contains this information.

#### Steps

1. Log in as an administrator.
2. In the Main Menu, click the **Troubleshooting** link.
3. Click **Scrub Utilities**.
4. Select the **Lookup dictionary** option.
5. Enter a code or keyword in the field.
6. Click **Lookup**.

   If you enter an incorrect format for the code, an error message displays. It could prompt you to check the server logs or correct the format of the scrubbed value.

## Finding ghost data sources

If you have removed a device but the device data remains, you can locate any ghost data sources so that you can remove them.

#### Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click the **Find Ghost Datasources** link.

   OnCommand Insight produces a list of originators with their device information.

## Adding a missing disk model

If acquisition fails due to a missing disk model, you can add the missing disk model to the `new_disk_models.txt` file and run acquisition again.

#### Before you begin

Understand the information needed to describe this data source because inaccurate information in the `new_disk_models.txt` file produces incorrect system data.

#### About this task

Because OnCommand Insight usually reads the internal disk model daily, the new data model information is updated automatically. Although you no longer need to restart the server for the changes to take effect, you can optionally restart the server to have the changes take effect before the next update.

#### Steps

1. Open Windows Notepad or another text editor to add the missing disk model information, .
2. Open the `new_disk_models.txt` file.

This file is located in your Insight installation at `C:\Program Files\SANscreen\jboss \server\onaro\deploy\datasources.war`.

3. Back up this file to prevent it from being overwritten during an upgrade.
4. Add the required information for the new data source.
5. Save the file in the `deploy\datasources.war` directory on your server.

### Locating required information for new disk model

To locate the disk model information, identify the vendor and run an Internet search on the vendor and model number.

#### Steps

1. Vendor specifications do not always provide all of the necessary information based on the full model number. It is often useful to search for different parts of the model number string on the vendor's site to locate all of the information.
2. You need the disk vendor name, full model number, disk size and speed, and the interface type in order to define the new disk model in OnCommand Insight.

The following example shows the information used to define a Seagate disk:

```
ST3500320NS,Seagate,ST3500320NS,7200,500,SATA,,,1500,
http://h20000.www2.hp.com/bizsupport/
TechSupport/Document.jsp?
lang=en&cc=us&objectID=c00305257&prodTypeId=15351&prodSeriesId=397642
```

### Content of the new_disk_models.txt file

The new_disk_models.txt file has required and optional fields. The fields are comma separated, so do not use commas *within* the fields.

Using a text editor, add the following information in this order:

1. **key**: use the model number (required)
2. **vendor**: name (required)
3. **model number**: full number (enter the same value as in "key")
4. **rpm of the disk**: for example 10000 or 15000 (required)
5. **size**: capacity in GB (required)
6. **interface type**: ATA, SATA, FC, SCSI, SAS, or FATA (required)
7. **seek time**: in ms (optional)
8. **potential transfer rate**: in MB/sec (optional)
9. **interface transfer rate**: the rate to and from the host in MB/sec (optional).
10. Additional useful information, such as a link to the vendor's specification page for the model (used only for Insight Perform) (optional).

## Specialized administration toolkit

The Admin Tools menu provides specialized and infrequently-used options such as calculating RPO and customizing the Java UI Client Help menu.

The Admin Tools include these specialized options:

- Show OnCommand Insight logs
- Browse the SAN and list the devices in the SAN
- Calculate RPO
- Run the server
- Trigger a VM NAS job
- Customize the options in the Client Help menu

### Opening the Admin Tools

To access the administrators' specialized maintenance toolkit, open the Admin Tools menu.

#### Steps

1. To use the Insight Admin Tools, open your browser and enter this URL customized for your environment:

   ```
   http://<server_host>:<server_port>/legacy/admin/all.do
   ```

2. Log in as an administrator. The Admin Tools menu opens.

### Showing Insight logs

To research problems, you might analyze the contents of the OnCommand Insight logs. These logs contain errors and warnings and include information about the length of time that some actions took, which might also be helpful in troubleshooting.

#### Steps

1. To use the Insight Admin Tools, open your browser and enter this URL customized for your environment:

   ```
   http://<server_host>:<server_port>/legacy/admin/all.do
   ```

2. Select **Show Logs**.

   **Note:** This feature has the same function as the **Troubleshooting > Logs** option.

3. Select the server, acquisition, or remote log that you want to analyze.

   The last 100 lines from the selected log are displayed.

## Browsing the SAN to locate devices

To locate a device or port in your environment based on its WWN, you can use the SANbrowse option.

### Steps

1. To use the Insight Admin Tools, open your browser and enter this URL customized for your environment:

   ```
   http://<server_host>:<server_port>/legacy/admin/all.do
   ```

2. In the Admin Tools menu, select **SANbrowse**.
3. In the Search box, enter the WWN of the device or port you want to locate. Click **Search**.

   OnCommand Insight returns a report containing links to additional information, as shown in this example.

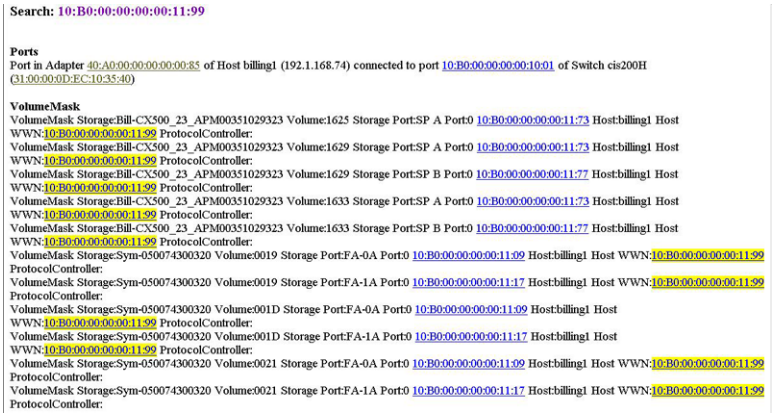


4. Click any of the links to see additional information related to a selected item. In this example, the first link in the search results was selected to produce this report with more links.

## Running a paths and violations query

If you are concerned about a recent change in the environment, you can use the **Run server** option to report changes in the paths and any new violations.

### Steps

1. To use the Insight Admin Tools, open your browser and enter this URL customized for your environment:

   ```
   http://<server_host>:<server_port>/legacy/admin/all.do
   ```

2. In the Admin Tools menu, select **Run server**.

   The results of this query list the changes found.

### Triggering VM NAS jobs

To correlate the NFS share export IP address with the VM ESX IP address, you can use the Trigger VM NAS Job option in the Admin Tools. The NFS export IP relates to the NetApp data source, and the VM ESX IP relates to the Virtual Center data source.

#### Steps

1. To use the Insight Admin Tools, open your browser and enter this URL customized for your environment:

   ```
   http://<server_host>:<server_port>/legacy/admin/all.do
   ```

2. In the Admin Tools menu, select **Trigger VM NAS Job**.

   A report listing the results displays.

# Monitoring your environment

## 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.

An asset page displays several sections containing information relevant to the asset. Depending on the type of asset, you see different sections. Any message with a solid red circle next to it on an asset page 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
- Disk
- VMDK
- Port
- Switch
- Fabric
- Object store (Atmos, Centera, and Amazon S3)

**Note:** Summary information is available for Atmos, Centera, and Amazon S3 object store 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 Insight web 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 | • “storage”<br>• name<br>• vendor<br>• model |

| **Category** | **Indexed terms** |
|---|---|
| StoragePool | • “storagepool”<br>• name<br>• name of the storage<br>• IP addresses of the storage<br>• serial number of the storage<br>• storage vendor<br>• storage model<br>• names for all associated internal volumes<br>• names for all associated disks |
| Internal Volume | • “internalvolume”<br>• name<br>• name of the storage<br>• IP addresses of the storage<br>• serial number of the storage<br>• storage vendor<br>• storage model<br>• name of the storage pool<br>• names of all associated shares<br>• names of all associated applications and business entities |
| Volume | • “volume”<br>• name<br>• label<br>• names of all internal volumes<br>• name of the storage pool<br>• name of the storage<br>• IP addresses of the storage<br>• serial number of the storage<br>• storage vendor<br>• storage model |

| Category | Indexed terms |
|---|---|
| Storage Node | • “storagenode”<br>• name<br>• name of the storage<br>• IP addresses of the storage<br>• serialnumber of the storage<br>• storage vendor<br>• storage model |
| Host | • “host”<br>• name<br>• IP addresses<br>• names of all associated applications and business entities |
| Datastore | • “datastore”<br>• name<br>• virtual center IP<br>• names of all volumes<br>• names of all internal volumes |
| Virtual Machines | • “virtualmachine”<br>• name<br>• DNS name<br>• IP addresses<br>• name of the host<br>• IP addresses of the host<br>• names of all datastores<br>• names of all associated applications and business entities |

| **Category** | **Indexed terms** |
|---|---|
| Switches (regular and NPV) | • “switch”<br>• IP address<br>• wwn<br>• name<br>• serial number<br>• model<br>• domain ID<br>• name of the fabric<br>• wwn of the fabric |
| Application | • “application”<br>• name<br>• tenant<br>• line of business<br>• business unit<br>• project |
| Tape | • “tape”<br>• IP address<br>• name<br>• serial number<br>• vendor |
| Port | • “port”<br>• wwn<br>• name |
| Fabric | • “fabric”<br>• wwn<br>• 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 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 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.
   - **7d**
     Displays the last seven days of data.
   - **Custom**
     Displays a dialog box that enables you to choose a custom range of 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.

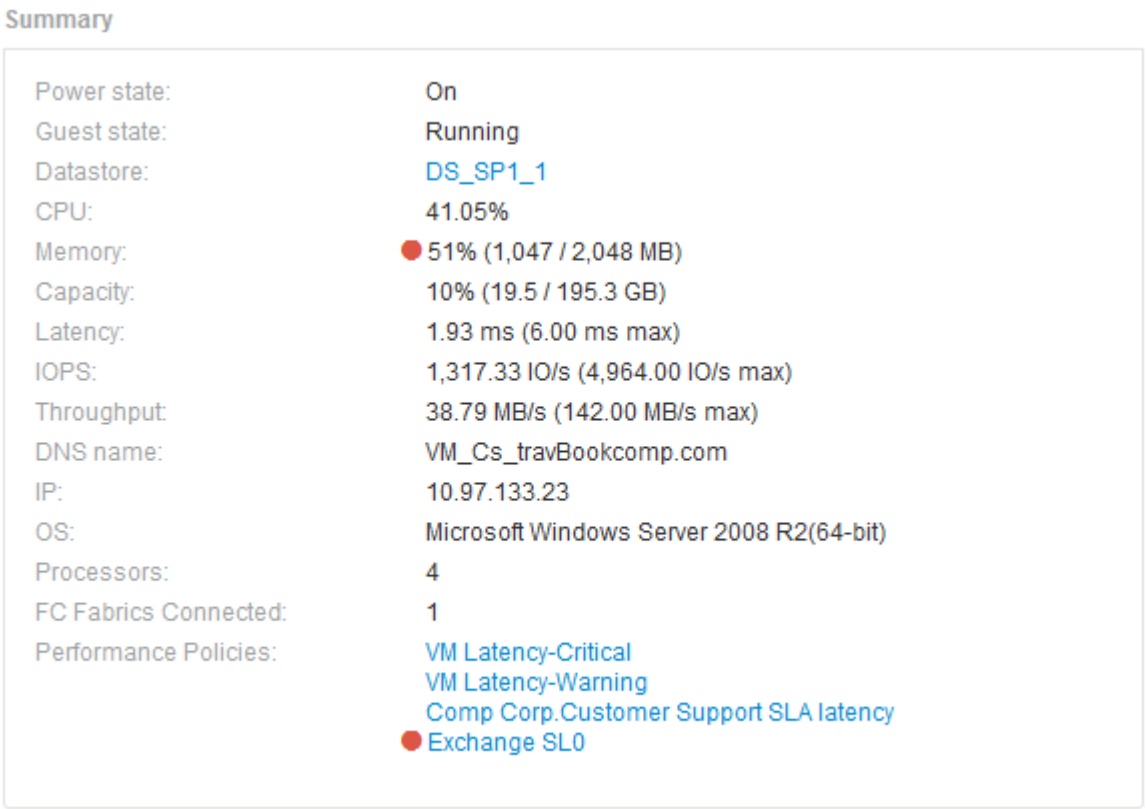


### *Using the Summary section*

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.

   **Note:** 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.

   **Note:** 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 application 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.

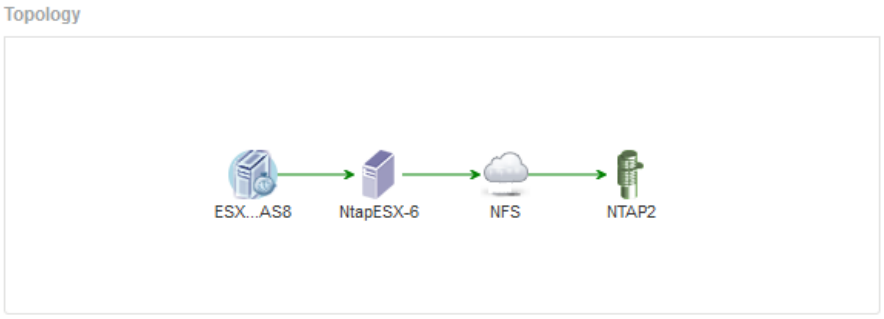


If the topology for the asset is larger than will fit in the section, the **Click link to see the topology** hyperlink is displayed instead.

### *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:

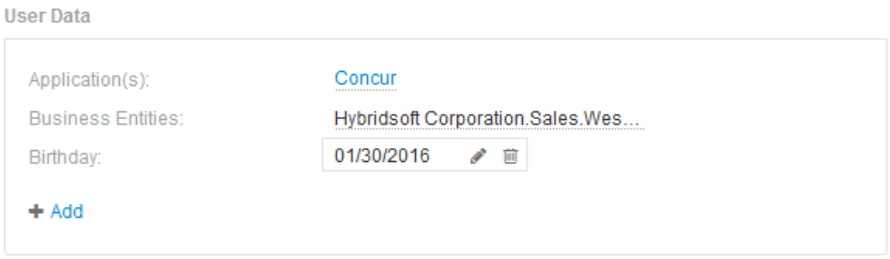


### *Using the User Data section to assign or modify applications*

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 .

     **Note:** 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. If you want to group assets by annotations, you must use the Java UI. Refer to the *OnCommand Insight Inventory User Guide for the Java UI* for more information.

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

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

## Anomalies

The Anomalies section, which displays only on an application asset page when the anomaly detection engine is installed, enables you to identify the normal operating workload range for an application and identify when changes in performance are outside of expected levels.

Anomaly analysis provides proactive visibility into the infrastructure supporting business critical applications and informs you of performance anomalies before they become service disruptions. The application anomaly detection engine ingests performance metrics collected by Insight and identifies anomalies in the application infrastructure. Anomalies are items, events or observations that do not conform to an expected pattern or other items. You use anomaly detection to:

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

### *Using the Anomalies section*

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

#### Steps

1. Turn anomaly monitoring on for the desired application and allow it to collect data.
2. While the anomaly detection engine is running, you can view the following information in the **Anomalies** section of the application asset page:
   - The anomaly detection score is shown by a blue bar, representing a score between 0 and 100.
   - Application assets view
     Located below the anomaly detection score, this shows how all the assets contribute to the current anomaly score. If an asset is not contributing significantly to the anomaly score, it is represented by . If the asset is contributing to the significance of the anomaly score, either 1, 2, or three blue bars appear, for example, . The values represented by the number of bars indicate the significance of that asset's contribution to the anomaly score as follows: 1 bar = 1-33; 2 bars = 34-67; 3 bars = 68-100.
     Assets that are assigned to an application (hosts, virtual machines, volumes, internal volumes, and hypervisors) and also assets in the infrastructure that are somehow connected to the assigned assets (for example, nodes, ports, and so on) contribute to the anomaly detection score. Both assets that contribute and those that do not contribute to the specific anomaly are displayed, because the data from the assigned application resources may flow through them.
     If there are more than 60 assets of one type, the **More** button displays, which you can click to display the remaining assets. That status of those remaining assets, for example,
     
     , appears next to the **More** button. This example indicates that of a total of 129 assets not shown, 124 of those assets did not contribute to the anomaly score; however, five of those assets did contribute some level of significance to the anomaly score.
   - Anomaly graph
     The anomaly graph displays the highest anomaly score, indicated by the orange circle, for the time range (3 hours, 24 hours, 3 days, 7 days, or a custom time range) selected for the application asset page. You can drag your pointer over the chart to see the anomaly scores, indicated by a blue circle, for other points during the time range and then click the blue circle, which will update the application assets view to display the assets that contributed to the score at that point in time.
3. To aid you in troubleshooting asset performance, you can add an asset from the application assets view to the expert view chart in the application asset page by doing the following:
   a. Click an asset icon in the applications assets view.

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

   b. Click .

   The asset is added to the expert view chart.

## 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:

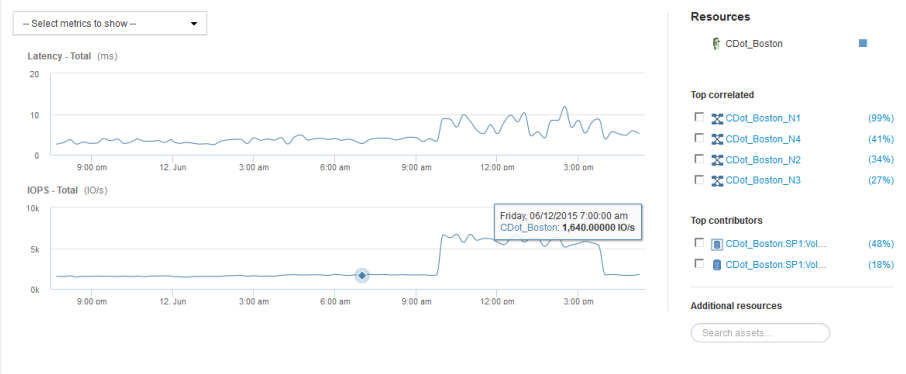


You can select the metrics you want to view in the performance chart for the time period selected.

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 |
|---|---|
| Zero BB Credit Rx | Number of times the receive buffer-to-buffer credit count transitioned to zero during the sampling period<br>This metric represents the number of times the attached port had to stop transmitting because this port was out of credits to provide. |

| Metric | Description |
|---|---|
| Zero BB Credit Tx | Number of times the transmit buffer-to-buffer credit count transitioned to zero during the sampling period<br>This metric represents the number of times this port had to stop transmitting because the attached port was out of credits to provide. |
| BB credit zero duration Tx | Time in milliseconds during which the Tx BB credit was zero during the sampling interval |
| Cache hit ratio (Total, Read, Write) % | Percentage of requests that result in cache hits<br>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, as a percentage of total available (over all virtual CPUs) |
| CRC Error | Number of frames with invalid CRCs detected by the port during the sampling period |
| Frame rate | Tx frame rate in fps |
| Frame size average - Rx and Frame size average - Tx | Ratio of traffic to frame size<br>This metric enables you to identify whether there are any overhead frames in the fabric. If this metric goes below a certain value (usually 0.2), there is an issue. |
| 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 |
| IOPS (Total, Read, Write) I/O sec | 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) This is read/write. |
| IP throughput (Total, Read, Write) MB/sec | Aggregated rate at which IP data was transmitted and received in megabytes per second<br>IP Throughput (Receive): Average rate at which IP data was received in megabytes per second<br>IP Throughput (Transmit): Average rate at which IP data was transmitted in megabytes per second |
| Latency (Total, Read, Write) milliseconds | Latency (R&W): Rate at which data is read or written to the virtual machines in a fixed amount of time<br>The value is measured in megabytes per second.<br>Latency: Average response time from the virtual machines carved from a data store<br>Top Latency: The highest response time from the virtual machines carved from a data store |
| Link failure | Number of link failures detected by the port during the sampling period |
| Link Reset Rx | Number of receive link resets during the sampling period<br>This metric represents the number of link resets that were issued by the attached port to this port. |

| Metric | Description |
|---|---|
| Link Reset Tx | Number of transmit link resets during the sampling period<br>This metric represents the number of link resets that were issued by this port to the attached port. |
| Memory utilization (Total) % | Threshold for the memory used by the host |
| Partial R/W (Total) % | Total number of times that a read/write crosses a stripe boundary on any disk module in a RAID 5, RAID 1/0, or RAID 0 LUN<br>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).<br>For CLARiiON, this value is the number of stripe crossings divided by the total number of IOPS. |
| Port errors | Indicates whether there were any port errors over the given time span |
| Signal loss count | Number of signal loss errors<br>If a signal loss error occurs, there is no electrical connection, and a physical problem exists. |
| Swap rate (Total Rate, In rate, Out rate) KB/sec | |
| Sync loss count | Number of synchronization loss errors<br>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. |
| Throughput (Total, Read, Write) MB/sec | Rate at which data is being transmitted in a fixed amount of time in response to I/O service requests (measured in MB per sec)<br>Throughput (R&W): Rate at which read/write data is being received in a fixed amount of time in response to I/O service requests (measured in MB per sec). |
| Timeout discard frames - Tx | Discarded Tx frames caused by timeout |
| Traffic Rate | |
| Traffic utilization | |
| Utilization (Total) % | Percentage of available bandwidth used for transmission (Tx) and reception (Rx) |
| 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**
   - **Tx** or **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, and Degraded sections to add data from that asset to the performance chart for each selected metric.
     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 asset's attributes are exported to the file.
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:

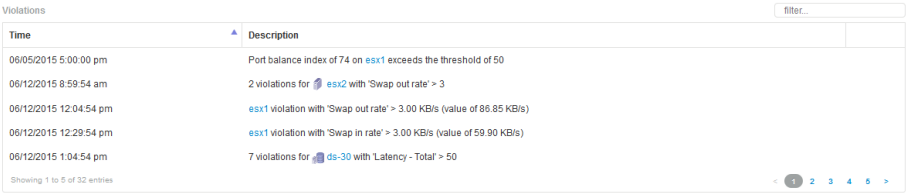


### *Using the Violations section*

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.

## 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 asset page to see lists of all of the resources associated with the VM.

   In this example, this Virtual Machine asset page shows that it is running well, but at the bottom of the page, one VMDK is at 99% capacity and probably causing the IOPs problem.

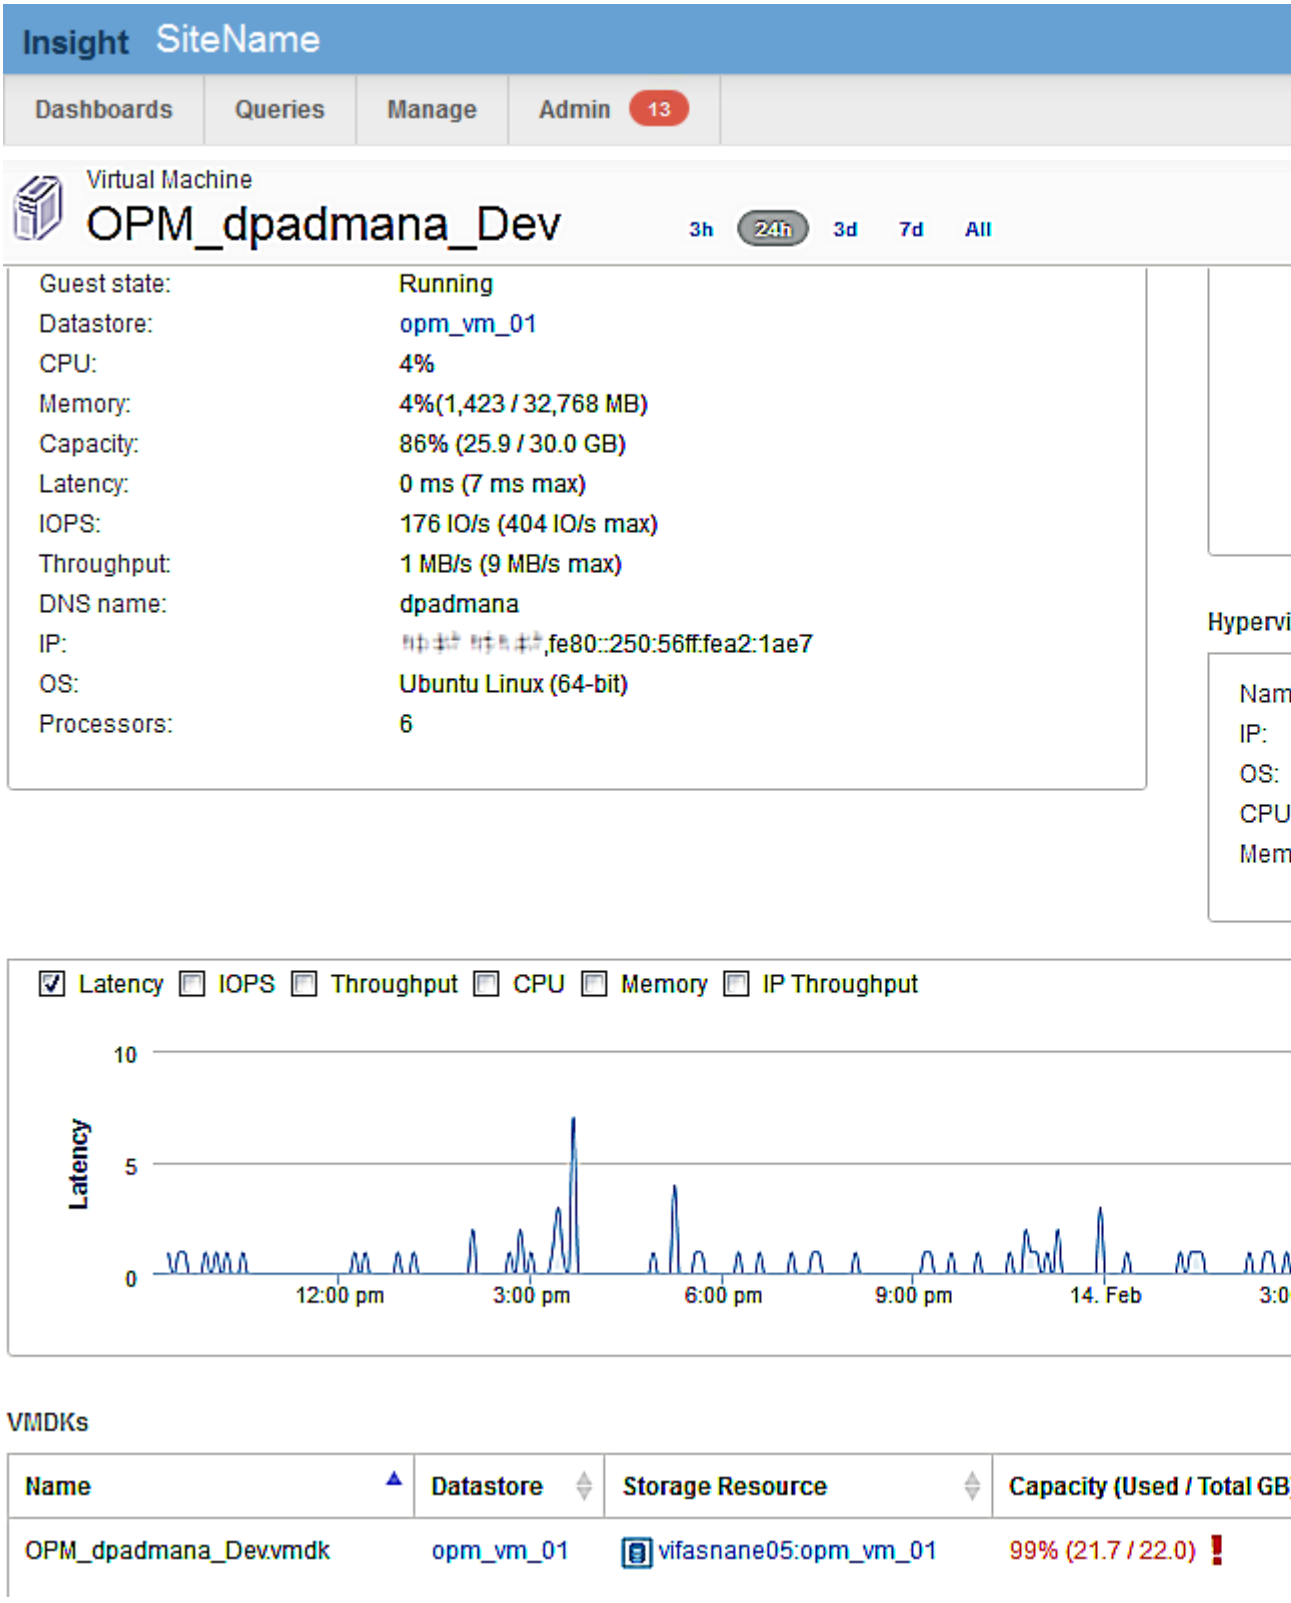


8. To resolve the balancing problem, identify a resource that is under-utilized to receive the load from an over-utilized resource or remove a less demanding application from the heavily used resource.

## 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

### 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.

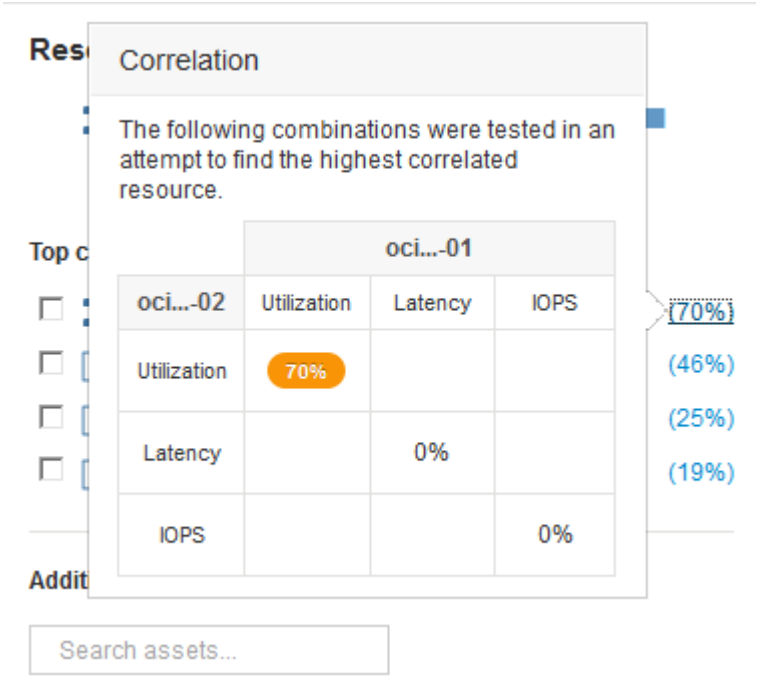


2. To add a correlated resource to the performance charts, select the check box in the **Top correlated resources** list for the resource you want to add. By default each resource provides the Total data available, but you can select only Read or only Write data from the menu on the check box.

   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.

3. Click the linked name of the correlated resource to display its asset page.

4. 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 polices for ports based on available port metrics. When a violation of a threshold occurs, Insight detects and reports it in the associated asset page by displaying a red solid circle; by email alert, if configured; and in the Violations Dashboard or any custom dashboard that reports violations.

#### Related tasks

*Creating performance policies* on page 62

# Copyright information

Copyright © 1994–2016 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

# Trademark information

Active IQ, AltaVault, Arch Design, ASUP, AutoSupport, Campaign Express, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlexArray, FlexCache, FlexClone, FlexGroup, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, Fueled by SolidFire, GetSuccessful, Helix Design, LockVault, Manage ONTAP, MetroCluster, MultiStore, NetApp, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANscreen, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, SolidFire Helix, StorageGRID, SyncMirror, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

*http://www.netapp.com/us/legal/netapptmlist.aspx*

# How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

*doccomments@netapp.com*

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

# Index


3PAR
    InServ Firmware (SSH) data source *73*
6000-Series Flash Memory Array (HTTP)
    Violin Memory data source *119*
7-Mode
    NetApp Data ONTAP data source *112*

## A

Access Gateway
    Brocade data source support for *75*
accessing
    Violations Dashboard *140*
    web UI *10*
acquisition
    recordings for Support *144*
    restarting units *141*, *147*
acquisition failure
    adding missing disk to correct failure *155*
acquisition units
    monitoring health of *134*
adding missing
    disks *155*
administrators
    changing user passwords *18*
    experience level *9*
    permissions *15*
    specialized tools *157*
annotation rules
    deleting *55*
    modifying *54*
    predefined storage tiers *53*
    setting for automatic annotation assignment *53*
    setting precedence *54*
annotations
    asset pages *170*
    assigning *170*
    assigning costs to *49*
    assignment of *48*
    configuring tier rules for automatic assignment *53*
    creating custom *49*
    default annotation rules *53*
    default types *46*
    defining *45*
    definition and purpose *46*
    deleting *52*
    importing *55*
    importing and exporting *66*
    manually assigning to assets *51*
    modifying *52*, *170*
    predefined storage tiers *53*
    updating changes for reports *151*
    using the User Data section to assign or modify *170*
    viewing on an asset page *170*
Anomalies section
    on asset pages *42*, *171*, *172*
anomaly detection
    about *40*
    enabling *41*
    monitoring status *43*
anomaly detection engine
    monitoring health of *134*
application
    enabling anomaly detection for *41*
applications
    assigning to assets *35*, *169*
    defining *34*
    deleting *36*
    editing *36*
applying
    service packs *143*
architecture
    data warehouse *8*
    firewalls *8*
    OnCommand Insight *7*
architecture components, OnCommand Insight
    acquisition units *7*
    Data Warehouse *7*
    HTML5 web-based UI *7*
    Java UI *7*
    OnCommand Insight Server *7*
asset page
    Topology section *169*
asset pages
    annotations *170*
    Anomalies section *42*, *171*, *172*
    business entities and *170*
    changing the time segment of data *166*
    correlated resources *182*
    data source acquisition status *166*
    Expert View metrics *173*
    Expert View section *173*
    opening from heat maps *133*
    opening from search results *161*
    performance analysis *181*
    performance troubleshooting data *160*
    Related Assets sections *177*
    sections of *167*
    Summary section *167*
    Topology section *168*
    types of *160*
    User Data section *169*
    Violations section *178*
asset topology
    viewing *169*
assets
    assigning applications to *35*
    manually assigning annotations to *51*
Assets Dashboard
    heat maps *179*
    monitoring system *133*
assets, related
    viewing in the Related Assets section *177*
assigning
    annotations *48*
    business entities to assets *45*
assigning multiple annotations

using queries *61*
assurance violations
configuring notifications *31*
ASUP
configuring *31*
process *145*
Atmos
EMC data source *81*
auditing
user activities *139*
authentication
LDAP *15*
user *14*
automatic storage tiering
configuring rules *53*
AutoSupport messages
recommendation to use HTTPS *31*

## B

backups
OnCommand Insight *27*
performing manually *144*
restoring *142*
balancing
network resources *179*
best practices
backup directories *27*
changing default passwords *10*
defining annotation types before annotation rules *53*
limited number of administrators *15*
restoring OCI database *142*
Brocade
Enterprise Fabric Connectivity Manager data source *74*
FC Switch Firmware (SSH) data source *75*
Sphereon/Intrepid Switch (SNMP) data source *77*
business entities
application data *34*
assigning to assets *45*
business unit *43*
creating *44*
hierarchy *44*
line of business *43*
on asset pages *170*
project *43*
tenant *43*

## C

CA
configuration for smart card *19*
CAC
browser requirements *22*
Cognos *23*
configuring for clients *22*
Data Warehouse configuration for *22*
capacity
VMDKs *179*
Celerra (SSH)
EMC data source *82*
certificate login
browser requirements *22*
client requirements *22*
configuration for Cognos *23*
configuration for Data Warehouse *22*
certificates
Cognos *23*
installing *25*
installing new *25*
replacing, after installation *25*
requirements for Smart Card and certificate login
Data Warehouse configuration for *22*
support *19*
changing the time segment of data
asset pages *166*
Cisco
FC Switch Firmware data source *78*
CLARiiON (NaviCLI)
EMC data source *83*
cloning
data sources *71*
cloud support
EMC Atmos data source *81*
Clustered Data ONTAP
NetApp data source *111*
Clustered Data ONTAP Unified Manager (UM)
NetApp data source *112*
Cognos
certificate configuration *23*
certificate login *23*
CommandViewAE
HP data source *98*
comments
how to send feedback about documentation *186*
committing
data source patches *131*
Compellent Storage Center
Dell data source *80*
components, OnCommand Insight architecture
acquisition units *7*
Data Warehouse *7*
HTML5 web-based UI *7*
Java UI *7*
OnCommand Insight Server *7*
configuration
advanced *151*
advanced tools *145*
changes *123*
testing data sources *72*
configuring
ASUP *31*
data sources *69*
email *28*
Insight *10*
Java UI Client *151*
performance policy evaluation precedence *65*
performance violation notifications *31*
SNMP notifications *29*
syslog *29*
users in LDAP *16*
cost
assigning to annotations *49*
creating
business entities *44*

custom annotations *49*
performance policies *62*
queries *57*
creating annotation rules
for automatic annotation assignment *53*
custom annotations
creating *49*
custom dashboards
deleting *34*
importing and exporting *66*
setting up *32*
Customer Support
configuring information sent to *31*
generate logs for *148*
scrubbed file example *154*
scrubbing report for *152*
sending data to *145*
customization
using annotations to monitor your environment *46*


## D


Dashboard, Violations
accessing *140*
dashboards
custom, importing and exporting *66*
setting up custom *32*
data collection
coordinating changes between your environment and Insight *122*
options *69*
Data Domain
EMC data source *85*
Data ONTAP operating in 7-Mode
NetApp data source *112*
data source acquisition status
on asset pages *166*
data sources
acquisition status on asset pages *166*
adding *69*
adding new, by patch *71*
approving patches *131*
changing passwords for a group *121*
cloning to quickly add *71*
deleting *129*
editing multiple setup information *127*
editing setup information *127*
examining detailed summaries *123*
finding serious problems *133*
importing *70*
installing a patch *129*, *130*
installing new *68*
locating ghost data *155*
managing *68*
managing patches *130*
patching groups *129*
permissions to change *15*
polling *126*
researching failures *124*
rolling back patches *132*
setting up *68*
support matrix *69*
testing configurations *72*
vendor-specific configuration information *73*
data sources, supported
3PAR InServ Firmware (SSH) *73*
Brocade Brocade Enterprise Fabric Connectivity Manager *74*
Brocade FC Switch Firmware (SSH) *75*
Brocade Sphereon/Intrepid Switch (SNMP) *77*
Cisco FC Switch Firmware (SNMP) *78*
Dell Compellent Storage Center *80*
EMC Atmos *81*
EMC Celerra (SSH) *82*
EMC CLARiiON (NaviCLI) *83*
EMC Data Domain *85*
EMC ECC StorageScope *86*
EMC Isilon *86*
EMC Solutions Enabler (CLI) with SMI-S Performance *87*
EMC VNX (SSH) *90*
EMC VNXe *90*
EMC VPLEX *92*
EMC XtremIO (HTTP) *93*
Fujitsu Eternus *94*
HDS (HiCommand Device Manager) *95*
Hitachi NAS (HNAS) *97*
HP (CommandView AE) *98*
HP EVA Storage (SSSU) *101*
IBM DS (CLI) *102*
IBM PowerVM *103*
IBM SVC (SSH) *105*
IBM Tivoli Monitoring *106*
IBM TotalStorage DS4000 *107*
IBM XIV (CLI) *108*
Infinidat InfiniBox (HTTP) *109*
Microsoft Hyper-V (WMI) *110*
NetApp Clustered Data ONTAP *111*
NetApp clustered Data ONTAP Unified Manager (UM) *112*
NetApp Data ONTAP operating in 7-Mode *112*
NetApp E-Series *114*
NetApp Host and VM File Systems *115*
NetApp SolidFire *116*
Pure Storage FlashArray (HTTP) *117*
QLogic FC Switch (SNMP) *118*
Violin Memory 6000-Series Flash Memory Array (HTTP) *119*
VMware vSphere (Web Services) *120*
Data Warehouse
architecture *8*
certificate configuration *22*
Data Warehouse server
monitoring health of *134*
databases
automatic weekly backups *27*
backing up *144*
OnCommand Insight *142*
SQL *150*
defining
applications *34*
deleting
annotation rules *55*
annotations *52*
applications *36*
custom dashboards *34*

data sources *129*
performance policies *65*
queries *60*
users *14*
Dell
Compellent Storage Center data source *80*
devices
data source setup *68*
deleting inactive *122*
examining performance *181*
disk models
locating vendor information *156*
text file format *156*
documentation
how to receive automatic notification of changes to *186*
how to send feedback about *186*
DS (CLI)
IBM data source *102*
DWH
*See* Data Warehouse


## E


E-Series
NetApp data source *114*
ECC StorageScope
EMC data source *86*
editing
annotations *52*
applications *36*
data source setups *127*
multiple data source setups *127*
performance policies *65*
user accounts *18*
editing multiple applications
using queries *60*
email
changing user *18*
configuring for receiving Insight reports *28*
prerequisites *28*
to deliver reports *28*
EMC
Atmos data source *81*
Celerra (SSH) data source *82*
CLARiiON (NaviCLI) data source *83*
Data Domain data source *85*
ECC StorageScope data source *86*
Isilon SSH data source *86*
Solutions Enabler (CLI) with SMI-S Performance data source *87*
VNX (SSH) data source *90*
VNXe data source *90*
VPLEX data source *92*
XtremIO (HTTP) data source *93*
Enterprise Fabric Connectivity Manager
Brocade data source *74*
error messages
data source *124*
non-compliant licenses *143*
ESX hosts *120*
Eternus
Fujitsu data source *94*
Ethernet
monitoring *37*
Ethernet monitoring
introduction *37*
EVA Storage (SSSU)
HP data source *101*
evaluation precedence
performance policies *64*
events
data source *123*
exclamation mark
meaning of in user interface *123*
Expert View metrics
asset pages *173*
Expert View section
asset pages *173*
populating the performance chart *176*
viewing the performance chart *176*
exporting
user data *66*
exporting to .CSV file
queries *59*


## F


FC Switch (SNMP)
Qlogic data source *118*
FC Switch Firmware
Cisco data source *78*
FC Switch Firmware (SSH)
Brocade data source *75*
feedback
how to send comments about documentation *186*
Fibre Channel environment
monitoring *183*
files
backing up *144*
license *11*
restoring OCI database *142*
firewalls
architecture *8*
FlashArray (HTTP)
Pure Storage data source *117*
Fujitsu
Eternus data source *94*


## H


HDS
HiCommand Device Manager data source *95*
heat maps
showing high IOPs *133*
HiCommand Device Manager
HDS data source *95*
Hitachi
NAS (HNAS) data source *97*
HNAS)
Hitachi data source *97*
Host and VM File Systems
NetApp data source *115*
hosts - ESX
see VMware vSphere data source *120*

HP
CommandView AE data source *98*
EVA Storage (SSSU) data source *101*
HTTP
protocol for ASUP *31*
Hyper-V (WMI)
Microsoft data source *110*

## I

IBM
DS (CLI) data source *102*
PowerVM data source *103*
SVC (SSH) data source *105*
Tivoli Monitoring data source *106*
TotalStorage DS4000 data source *107*
XIV (CLI) data source *108*
importing
annotations *55*
data sources *70*
user data *66*
InfiniBox (HTTP)
Infinidat data source *109*
Infinidat
InfiniBox (HTTP) data source *109*
information
how to send feedback about improving documentation *186*
InServ Firmware (SSH)
3PAR data source *73*
Insight
backups *27*
component list *149*
data source patches *129*
data source setup requirements *68*
data source testing *72*
license types *12*
licenses *143*
licenses, installing *11*
list of common maintenance tasks *133*
logging in *10*
maintenance *122*
number of database objects *149*
overview *7*
status *133*
Insight reports
configuring email for *28*
insight server
restart *145*
installed software packages
displaying *145*
installing
licenses *11*
Intrepid/Sphereon Switch (SNMP)
Brocade data source *77*
IOPs
heat map *179*
IP addresses
correlating export *159*
Isilon
EMC data source *86*

## L

landing page
NFS traffic
volumes *40*
NFSv3 activity *40*
LDAP
authentication *15*
rediscover user *19*
requirements for smart car and certificate login *19*
requirements for Smart Card and certificate login *23*
users *16*, *18*
licenses
for Insight *12*
installing *11*
not compliant *143*
replacing older versions *143*
types *12*
updating expired *143*
log files
accessing *145*
login
certificate *19*
Smart Card *19*
logs
audit *139*
creating a zip file for Support *144*
generate for Customer Support *148*
OnCommand Insight *157*
syslog *30*

## M

maintenance
coordinating changes between your environment and Insight *122*
specialized tools *157*
maintenance tasks
list of most common *133*
managing
performance policies *62*
Microsoft
Hyper-V (WMI) data source *110*
missing disks
adding missing disk models *155*
modifying
annotation rules *54*
annotations *52*
data source setup *127*
multiple data source setups *127*
queries *59*
monitoring
anomalies *40*
applications *40*
data sources *68*
Fibre Channel environment *183*
Insight component health *134*
violations *140*
monitoring your environment
using asset annotations *46*
moving
MySQL data *150*
MySQL data with command line *150*

MySQL
migrating data *150*

## N

NAS (HNAS)
Hitachi data source *97*
NetApp
Clustered Data ONTAP data source *111*
Clustered Data ONTAP Unified Manager (UM) data source *112*
Data ONTAP operating in 7-Mode data source *112*
E-Series data source *114*
Host and VM File Systems data source *115*
SolidFire data source *116*
NetApp Customer Support
bundling logs for troubleshooting *144*
preparing troubleshooting report *144*
sending data to *145*
network
balancing resources *179*
performance *181*
system information *149*
network maintenance
list of most common tasks *133*
notifications
performance violations *31*
notifications, SNMP
configuring *29*

## O

OCI
*See* OnCommand Insight
OnCommand Insight
managing data sources *68*
overview *7*
setup tasks *10*

## P

passwords
changing for data source groups *121*
changing user *18*
coordinating changes *122*
encrypted for ASUP *31*
LDAP *16*
new users *14*
patches
adding new data source by *71*
approving *131*
coordinating changes between your environment and Insight *122*
data source, managing *130*
groups of data sources *129*
installing data source *129*, *130*
new data source, installing *68*
rolling back *132*
performance
adding new policies *62*
alerts in syslog *29*
analyzing slow client computers *181*
correlated resources *182*
examining network *181*
monitoring *160*
performance analysis
skills required *9*
performance chart
using *176*
performance charts
correlated resources *182*
zooming in *181*
performance policies
configuring evaluation precedence *65*
creating *62*
deleting *65*
editing *65*
evaluation precedence *64*
identifying violations *167*
managing *62*
performance troubleshooting data
asset pages *160*
performance violations
configuring notifications *31*
planning
your data source support *69*
policies
importing and exporting *66*
performance *62*
polling
data sources *126*
ports
locating *158*
postpone
data source polling *126*
PowerVM
IBM data source *103*
precedence
setting for annotation rules *54*
Pure Storage
FlashArray (HTTP) data source *117*

## Q

Qlogic
FC Switch (SNMP) data source *118*
queries
assigning multiple annotations *61*
creating *57*
deleting *60*
editing multiple applications *60*
exporting to .CSV file *59*
importing and exporting *66*
modifying *59*
removing multiple annotations *61*
removing multiple applications *60*
viewing *59*

## R

red circle
meaning of in the user interface *133*
Related Assets section
viewing related assets *177*

Related Assets sections
on asset pages *177*
removing
data sources *129*
user account *19*
removing multiple annotations
using queries *61*
removing multiple applications
using queries *60*
reporting
Atmos data considerations *81*
customized with annotations *46*
resources
searching for *161*
top correlated *181*, *182*
restart
insight server *145*
restarting
acquisition units *141*
data source polling *126*
Insight servers *146*
restoring
backups *142*
resume
data source polling *126*
roles
7-Mode environments *112*
users *14*, *15*

## S

scaling
determine number of database objects *149*
search
advanced techniques *161*
boolean *161*
keywords *161*
wildcard *161*
searching
correlated resources *182*
environment assets *161*
for host *181*
sections of
asset pages *167*
security
firewall architecture *8*
user roles *15*
self-signed certificates
installing *25*
servers
backup *27*
LDAP *16*
monitoring health of *134*
restarting OnCommand Insight *146*
Service Level annotation
assigning cost to *49*
service packs
applying *143*
setting hours to ignore data updates
configuration *72*, *123*, *145*, *147*, *151*
setting up
custom dashboards *32*
Insight *10*
licenses *11*
Smart Card
support *19*
smart cards
host configuration for *19*
Smart Cards
Cognos, configuration for *23*
configuring for clients *22*
Data Warehouse configuration for *22*
SNMP
configuring notifications *29*
software
anomaly detection *40*
solid red circle
meaning *141*
SolidFire
NetApp data source *116*
Solutions Enabler (CLI) with SMI-S Performance
EMC data source *87*
Sphereon/Intrepid Switch (SNMP)
Brocade data source *77*
status
anomaly detection monitoring *43*
storage
predefined tier annotations *53*
suggestions
how to send feedback about documentation *186*
Summary section
identifying metric issues *167*
identifying performance policy violations *167*
on asset pages *167*
viewing general asset information *167*
SVC (SSH)
IBM data source *105*
Switch Level annotation
assigning cost to *49*
syslog
configuring *29*
contents *30*
system architecture
OnCommand Insight *7*

## T

terabytes
licensed *143*
testing
added data sources *69*
data source configurations *72*
Tier annotation
assigning cost to *49*
tiers
checking changes after upgrading *152*
predefined annotations *53*
Tivoli Monitoring
IBM data source *106*
Topology section
asset pages *168*
viewing an asset's topology *169*
TotalStorage DS4000
IBM data source *107*
tracking data
using asset annotations *46*

triggering VM NAS jobs *159*
troubleshooting
  data collection problems *122*
  database backup *144*
  finding changes in paths *158*
  locating a device or port *158*
  locating ghost data *155*
  locating your environment resources *161*
  logs *157*
  preparing a special report *144*
  restoring OnCommand Insight *142*
  sending report to Support *145*
  server resources status *152*
  tools *145*
  using audit log *139*
  your environment *160*
  zipping logs for Support *144*
Twitter
  how to receive automatic notification of documentation changes *186*
types of
  asset pages *160*

## U

upgrading
  check tiers after *152*
user data
  importing and exporting *66*
User Data section
  asset pages *169*
  using to assign or modify annotations *170*
users
  auditing activities *139*
  authenticated from LDAP *16*
  changing passwords *18*
  defining *14*
  deleting *19*
  editing accounts *18*
  experience level *9*
  roles *15*

## V

vendors
  finding disk model information *156*
  supported data sources *73*
vFiler unit
  connection requirements *112*
viewing
  annotations assigned to an asset *170*
  applications assigned to an asset *169*
  asset topology *169*
  business entity assigned to an asset *170*
  queries *59*
viewing an asset's Topology
  Topology section *169*
viewing related assets
  Related Assets section *177*
viewing violations
  Violations section *178*
violations
  by policy *140*
  changes in syslog *29*
  configuring
    assurance violation notifications *31*
  history *140*
  in syslog *30*
  monitoring *140*
  SNMP notifications of *29*
Violations Dashboard
  accessing *140*
Violations section
  on asset pages *178*
  viewing violations *178*
Violin Memory
  6000-Series Flash Memory Array (HTTP) data source *119*
VM NAS jobs
  triggering *159*
VMDK
  resources list *179*
VMware
  vSphere (Web Services) data source *120*
VNX (SSH)
  EMC data source *90*
VNXe
  EMC data source *90*
VPLEX
  EMC data source *92*
vSphere (Web Services)
  VMware data source *120*

## W

web UI
  access data source support matrix *69*
  accessing *10*
  advanced configuration tools *145*
WWN
  locating a device or port using *158*

## X

XIV (CLI)
  IBM data source *108*
XtremIO (HTTP) data source
  EMC data source *93*