



OnCommand Insight Documentation

OnCommand Insight

NetApp
June 10, 2024

Table of Contents

- OnCommand Insight Documentation 1
- Release Notes 2
 - Release Notes 2
- What is OnCommand Insight? 3
 - OnCommand Insight overview 3
 - Insight architecture 3
 - How Insight is used by administrators, managers, and planners 5
- Installation for Linux 6
 - Installation prerequisites 6
 - Insight installation instructions 13
 - Upgrading Insight 24
 - Uninstalling OnCommand Insight 32
- Installation for Microsoft Windows 34
 - Installation prerequisites 34
 - Insight installation instructions 42
 - Upgrading OnCommand Insight 55
 - Uninstalling the software 79
- Configuration and administration 81
 - Setting up Insight 81
 - Insight Security 171
 - Smart Card and certificate login support 184
 - Configuring Data Warehouse for Smart Card and certificate login 196
 - Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.5 through 7.3.9) 197
 - Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.10 and later) 198
 - Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.5 to 7.3.9) 200
 - Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.10 and later) 202
 - Importing SSL certificates 204
 - Your business entities hierarchy 207
 - Defining annotations 210
 - Querying assets 224
 - Insight data source management 231
 - Device resolution 330
 - Maintaining Insight 349
 - Monitoring your environment 371
- Data Warehouse administration 400
 - Welcome to OnCommand Insight Data Warehouse 400
 - Getting started with Data Warehouse 406
 - Administrative tasks you can perform using Data Warehouse 427
- Reporting 453
 - Welcome to OnCommand Insight reporting 453
 - Reporting made easy 457
 - Managing reports 466
 - Creating custom ad hoc reports 469

Reporting data model	471
FAQ	478
General Questions	478
OnCommand Insight Licensing	480
Configuration and Supported devices	481
Scale and Ease of Use	482
Performance troubleshooting	483
Managing your environment	485
Integrating Insight with other tools	485
Data ONTAP Storage IOPS	486
How-To guides	488
Getting Started with Insight	488
Creating custom dashboards	501
Creating performance policies	535
Troubleshooting Fibre Channel BB credit 0 errors	539
Analyzing your infrastructure	544
Introduction to minimizing risk in thin provisioning	549
Collecting Host and VM file system utilization data	555
Configuring your system to report chargeback data	559
Ensuring IO density reports describe only internal data volumes	565
Collecting integration data	567
Analyzing an application performance problem	576
Collecting and reporting AWS billing data	584
Integrating with ServiceNow	587
Legal notices	594
Copyright	594
Trademarks	594
Patents	594
Privacy policy	594
Notice	594

OnCommand Insight Documentation

OnCommand 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 can help you optimize your current infrastructure, allowing you to right-size operations to meet business demands. It simplifies the process of determining what and when to buy. It also reduces risk during complex technology migrations, such as moving to a hybrid cloud, by identifying which workloads are candidates for cloud migration. With Insight, you can manage the IT infrastructure as an end-to-end service by integrating the resources into the company's entire IT service delivery chain.

Release Notes

Release Notes

OnCommand Insight Release Notes are available outside the Documentation Center. You will be prompted to log in using your NetApp Support Site credentials.

[Release Notes .PDF](#) (opens in a new window)

What is OnCommand Insight?

OnCommand Insight overview

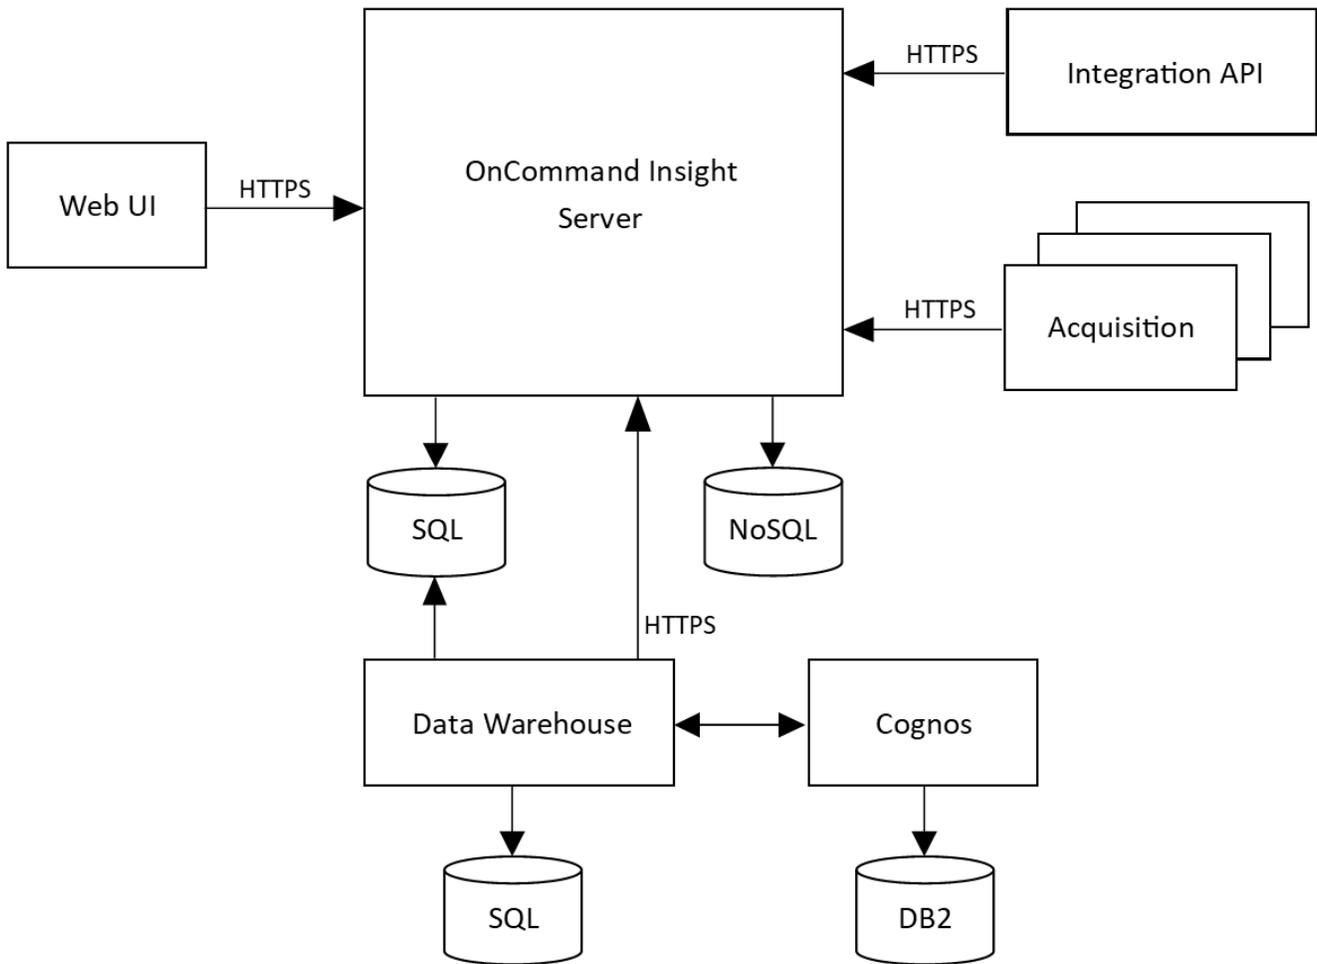
OnCommand Insight enables you to simplify operational management of complex private and hybrid cloud and virtual IT environments. Insight is a single solution to enable cross-domain, multi-vendor resource management and analysis across networks, storage, and servers in physical and virtual environments.

Insight can help you optimize your current infrastructure, allowing you to right-size operations to meet business demands. It simplifies the process of determining what and when to buy. It also reduces risk during complex technology migrations, such as moving to a hybrid cloud, by identifying which workloads are candidates for cloud migration. With Insight, you can manage the IT infrastructure as an end-to-end service by integrating the resources into the company's entire IT service delivery chain.

Insight architecture

A typical installation of OnCommand Insight includes data acquisition and data warehousing with reports, all easily accessible from a web-based UI. For more secure environments, acquisition can be done through a remote acquisition unit.

The major components of the Insight architecture are shown in the following diagram:



- **OnCommand Insight Server**

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

- **Acquisition**

The Insight collection engine is built on one or more acquisition units. Each Insight server contains a local acquisition unit and can support remote acquisition units. Each unit is a service running on the network that accesses (through modules called *data sources*) and collects data from devices in the data center. Information collected by the acquisition units is then sent to the server for analysis.

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

- **Integration API**

An API allows the collection of data from external agents. Integration data can be viewed in the web UI using queries and widgets. Dashboards can contain 'native' Insight data and integration data. You can apply filtering, roll-ups, and grouping to the data in these dashboards.

- **Web UI**

The HTML5 web-based user interface for Insight enables you to set up data sources and your monitoring

environment, including policies, thresholds and alerts. You then use the web UIAsset Dashboard and asset pages to identify and research potential problems. You can create custom dashboards with a variety of widgets, each of which provides extensive flexibility in displaying, analyzing, and charting your data.

- **Data Warehouse**

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

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

The Data Warehouse consolidates and prepares data for reporting for one or multiple installations of Insight. The data includes history, trending, inventory, chargeback, show back and data presentations to support long-term planning of the data center's infrastructure.

- **Cognos**

Cognos is the reporting engine for Insight, an IBM business intelligence tool that enables you to view pre-defined reports or create custom reports. Insight reporting generates reports from the Data Warehouse data.

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.

Installation for Linux

Installation prerequisites

Before you install OnCommand Insight, you must download the current software version, acquire the appropriate license, and set up your environment.

Before installing OnCommand Insight, ensure that you have the following:

- OnCommand Insight software files in the downloaded installation package for the current version
- A license to operate the downloaded OnCommand Insight version
- The minimum hardware and software environment

The current product might consume additional hardware resources (due to enhanced OnCommand Insight product functionality) that were not consumed with earlier versions of the OnCommand Insight product.

- A deployment plan that includes the hardware and network configurations for the OnCommand Insight Server, Data Warehouse and Reporting, and remote acquisition units.

Planning the deployment

To ensure a successful deployment, you must consider certain system elements before you install OnCommand Insight.

About this task

Planning your Insight deployment includes considering these system elements:

- Insight architecture
- Your network components to be monitored
- Insight installation prerequisites and server requirements
- Insight web browser requirements

Data source support information

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

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

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

Device identification and data source planning

As part of your deployment planning, you should collect information about the devices in

your environment.

You need the following software, connectivity, and information about each device in your environment:

- IP address or hostname resolvable by the OCI server
- Login name and password
- Type of access to the device, for example, controller and management station



Read-only access will be sufficient for most devices, but some devices require administrator permissions.

- Port connectivity to the device depending on data source port requirements
- For switches, SNMP read-only community string (user ID or password to give access to the switches)
- Any third-party software required on the device, for example, Solutions Enabler.
- See the "Vendor-specific data source reference" in the web UI Help or in the *OnCommand Insight Configuration and Administration Guide* for more information on data source permissions and requirements.

Network traffic generated by OnCommand Insight

The network traffic that OnCommand Insight generates, the amount of processed data traversing the network, and the load that OnCommand Insight places on devices differ based on many factors.

The traffic, data, and load differ across environments based on the following factors:

- The raw data
- Configuration of devices
- Deployment topology of OnCommand Insight
- Different inventory and performance data source polling intervals, which can be reduced to allow for slow devices to be discovered or bandwidth to be conserved

The raw configuration data that OnCommand Insight collects can vary significantly.

The following example illustrates how the configuration data can vary and how traffic, data, and load are affected by many configuration factors. For example, you might have two arrays each having 1,000 disks:

- Array 1: Has 1,000 SATA disks all 1 TB in size. All 1,000 disks are in one storage pool, and there are 1,000 LUNs, all presented (mapped and masked) to the same 32 nodes in an ESX cluster.
- Array 2: Has 400 2-TB data disks, 560 600-GB FC disks, and 40 SSD. There are 3 storage pools, but 320 of the FC disks are used in traditional RAID groups. The LUNs carved on the RAID groups use a traditional masking type (symmaskdb), while the thin provisioned, pool-based LUNs use a newer masking type (symaccess). There are 600 LUNs presented to 150 different hosts. There are 200 BCVs (full block replica volumes of 200 of the 600 LUNs). There are also 200 R2 volumes, remote replica volumes of volumes that exist on an array in a different site.

These arrays each have 1,000 disks and 1,000 logical volumes. They might be physically identical in the amount of rack space they consume in the data center, and they might even be running the same firmware, but the second array is much more complex in its configuration than the first array.

Uninstalling MariaDB

You must uninstall MariaDB on the Insight or Data Warehouse servers before you install OnCommand Insight or the Data Warehouse; otherwise, you can not proceed with the installation. MySQL is not compatible with MariaDB. If you attempt an installation on either server without removing MariaDB, the installation terminates with an error message instructing you to uninstall MariaDB.

Before you begin

You must have sudo privileges.

Steps

1. Log in to the Insight server.
2. Obtain a list of MariaDB components:

```
rpm -qa | grep mariadb
```

3. Type the following for each MariaDB component that is installed on the server:

```
yum remove component_name
```

Insight Server requirements

A dedicated server is recommended. Do not install Insight on a server that has any other applications installed. Both physical and virtual servers are supported, provided that the product requirements are met.

You must have sudo permissions to install the OnCommand Insight Server software.

Some Insight components may require dependent packages during installation. Ensure YUM repository is accessible prior to installing Insight.



Sizing for OnCommand Insight has multiple dependencies, such as data source type and size, number of assets in your environment, polling intervals, and more. The following sizing examples are guidelines only; they represent some of the environments in which Insight has been tested. Changing any of these or other factors in the environment can change the sizing requirements for Insight. These guidelines include disk space for up to 90 days of performance archive data.

It is recommended to contact your Sales Engineer for detailed sizing guidance before installing or upgrading Insight.

Examples:

Environment factors:	Disk space, CPUs, and Memory tested:
----------------------	--------------------------------------

80 storage arrays 4,000 Volumes 4,000 VMs 4,000 switch ports	250 GB disk space 8 cores 32 GB RAM
160 storage arrays 40,000 Volumes 8,000 VMs 8,000 switch ports	1 TB of disk space 12 cores 48 GB RAM

Requirements:

Component	Required
Operating system	<p>A computer running a licensed version of one of the following, that is running no other application-level software:</p> <ul style="list-style-type: none"> • Red Hat Enterprise Linux 7.3, 7.4, 7.5, 7.6, 7.7, 7.8, 7.9, 8.1, 8.2, 8.3, 8.4, 8.5, 8.8, 9.2 • CentOS 7.2, 7.5, 7.6, 7.7, 7.8, 7.9, CentOS 8 Stream, CentOS 9 Stream • Oracle Enterprise Linux 7.5, 7.6, 7.7, 7.8, 7.9, 8.1, 8.2, 8.3, 8.4, 8.5, 8.8 <p>A licensed version ensures that dependencies required by the installation are resolved automatically by the operating system.</p> <p>You must uninstall MariaDB before installing Insight.</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <p>Uninstalling MariaDB also removes the Postfix Mail Transport Agent.</p> </div> <p>A dedicated server is recommended.</p>
Virtual machine (VM)	<p>This component can run in a virtual environment, provided that the CPU and memory resources for your instance are reserved.</p>
Memory and CPU	<p>24 - 256 GB RAM</p> <p>8 - 32 cores</p>

<p>Available disk space</p>	<p>100 GB - 3 TB install disk space</p> <p>50 GB - 1 TB performance archive disk space</p> <p>The following partition breakdowns are recommended for an example 500 GB environment:</p> <ul style="list-style-type: none"> • /opt directory — 50 GB • /var/log directory — 100 GB • /var/lib directory — 350 GB <p>It is a best practice to mount /opt and /var on separate disks from the root file system (/).</p> <p>SSD disks are recommended for the Insight installation space.</p>
<p>Network</p>	<p>Ethernet connection and ports:</p> <ul style="list-style-type: none"> • 100 Mbps or 1 Gbps Ethernet connection with dedicated (static) IP address and IP connectivity to all components in the SAN, including FC devices and remote acquisition units. • Port requirements for the OnCommand Insight Server process are 80, 443, 1090 through 1100, 3873, 8083, 4444 through 4446, 5445, 5455, 4712 through 4714, 5500, and 5501. • Port requirements for the acquisition process are 12123 and 5679. • Port requirement for MySQL is 3306. • Port requirements for Elasticsearch are 9200 and 9310 <p>Ports 443 and 3306 require external access through any firewall that is present.</p>
<p>Permissions</p>	<p>Sudo permissions are required on the OnCommand Insight Server.</p> <p>If any of the following folders are symbolic links, ensure that the destination directories have '755' permissions.</p> <ul style="list-style-type: none"> • /opt/netapp • /var/lib/netapp • /var/log/netapp

Remote connectivity	Internet connectivity to allow WebEx access or a remote desktop connection to facilitate installation and post-installation support.
Accessibility	HTTPS access is required.
HTTP or HTTPS servers	Apache HTTP servers or other HTTPS servers should not compete for the same ports (443) as the OnCommand Insight server and should not start automatically. If they must listen to port 443, then you must configure the OnCommand Insight server to use other ports.

Data Warehouse server requirements

The Data Warehouse server must run on a computer that is compatible with established hardware and software requirements. You must ensure that Apache web server or reporting software is not already installed on this machine.



Sizing for OnCommand Insight has multiple dependencies, such as number of assets in your environment, amount of historical data retained, and more. The following data warehouse sizing examples are guidelines only; they represent some of the environments in which Insight has been tested. Changing any of these or other factors in the environment can change the sizing requirements for Insight.

It is recommended to contact your Sales Engineer for detailed sizing guidance before installing or upgrading Insight.

Examples:

Environment factors:	Disk space, CPUs, and Memory tested:
18 storage arrays 3,400 VMs 4,500 switch ports	200 GB hard disk 8 cores 32 GB RAM
110 storage arrays 11,500 VMs 14,500 switch ports	300 GB hard disk 8 cores 48 GB RAM

Requirements:

Component	Required
-----------	----------

Operating system	<p>A computer running a licensed version of one of the following, that is running no other application-level software:</p> <ul style="list-style-type: none"> • Red Hat Enterprise Linux 7.3, 7.4, 7.5, 7.6, 7.7, 7.8, 7.9, 8.1, 8.2, 8.3, 8.4, 8.5, 8.8, 9.2 • CentOS 7.2, 7.5, 7.6, 7.7, 7.8, 7.9, CentOS 8 Stream, CentOS 9 Stream • Oracle Enterprise Linux 7.5, 7.6, 7.7, 7.8, 7.9, 8.1, 8.2, 8.3, 8.4, 8.5, 8.8
Virtual machine (VM)	This component can run in a virtual environment, provided that the CPU and memory resources for your instance are reserved.
CPU	8 - 40 CPU cores
Memory	32 GB - 2 TB RAM
Available Disk Space	200 GB - 512 GB disk space There should be at least 50 GB of free disk space in the <code>/var/lib</code> partition and 25 GB of free disk space in the <code>/opt</code> and <code>/var/log</code> partitions.
Network	<ul style="list-style-type: none"> • 100 Mbps or 1 Gbps Ethernet connection • Static IP address • For the OnCommand Insight DWH server process, ports 80, 443, 1098, 1099, 3873, 8083, and 4444 through 4446 • For MySQL, port 3306

Remote Acquisition Unit server requirements

You must install a Remote Acquisition Unit (RAU) to acquire information from SAN devices that are behind a firewall, at a remote site, on a private network, or in different network segments. Before you install the RAU, you should ensure that your environment meets RAU operating system, CPU, memory, and disk space requirements.

Component	Requirement
-----------	-------------

Operating system	<p>A computer running a licensed version of one of the following, that is running no other application-level software:</p> <ul style="list-style-type: none"> • Red Hat Enterprise Linux 7.3, 7.4, 7.5, 7.6, 7.7, 7.8, 7.9, 8.1, 8.2, 8.3, 8.4, 8.5, 8.8, 9.2 • CentOS 7.2, 7.5, 7.6, 7.7, 7.8, 7.9, CentOS 8 Stream, CentOS 9 Stream • Oracle Enterprise Linux 7.5, 7.6, 7.7, 7.8, 7.9, 8.1, 8.2, 8.3, 8.4, 8.5, 8.8 <p>A dedicated server is recommended.</p>
CPU	4 CPU cores
Memory	16 GB RAM
Available disk space	40 GB
Network	100 Mbps /1 Gbps Ethernet connection, static IP address, IP connectivity to all FC devices, and a required port to the OnCommand Insight server (80 or 443).
Permissions	Sudo permissions on the RAU server

Browsers supported by OnCommand Insight

The browser-based OnCommand Insightweb UI can operate on several different browsers.

Insight supports newer, non-beta releases of the following browsers:

- Mozilla Firefox
- Google Chrome
- Microsoft Edge

For a full list of browser versions qualified for OnCommand Insight, please see the [NetApp Interoperability Matrix Tool](#).

Insight installation instructions

Installation requires installing several OnCommand Insight components, Insight Server, and Data Warehouse.

The installation includes the following major tasks:

- Downloading the OnCommand Insight installer

- Installing OnCommand Insight server
- Installing licenses
- Optionally, installing DWH and Reporting (must be installed on a separate machine or virtual machine. Reporting requires Microsoft Windows.)
- Optionally, installing a remote acquisition unit (RAU), which acquires information from your device resources that reside behind a firewall, are located at a remote site, or are on a private network

After installation, you must configure Insight to acquire information about your environment. The tasks required are described in the *OnCommand Insight Configuration and Administration Guide*.

Downloading the OnCommand Insight installer

You can download the OnCommand Insight installer from the NetApp Support Site.

Before you begin

You must have a login to the NetApp Support Site at mysupport.netapp.com.

Additionally, you must have an unzip utility with which to open the installation .ZIP files.

Steps

1. Log in to the server on which you want to install OnCommand Insight.
2. Download the installation file from the NetApp Support site.

Installing the OnCommand Insight Server

OnCommand Insight Server is installed by using the command line.

Before you begin

You must have completed all of the installation prerequisites.

Steps

1. Log in to the Insight server using an account with sudo privileges.
2. Navigate to the directory on the server where the installation files are located and type the following command:

```
unzip oci-<version>-linux-x86_64.zip
```

Ensure that you check the version number of the installation file; the version number might be different than the one shown in the command.

3. You can view syntax, command arguments, and parameter usage for `oci-install.sh`:

```
sudo ./oci-<version>-linux-x86_64/oci-install.sh --help
```

4. Run the installation script:

```
sudo ./oci-<version>-linux-x86_64/oci-install.sh
```

5. Read the License Agreement, accept it, and follow the prompts.
6. If you are using the Insight consumption licensing model, you must enable sending of usage information to NetApp. Enter `y` at this prompt.

Results

After you answer all the prompts, the installation begins and should take approximately 10 minutes, depending on the applications installed.

Installing OnCommand Insight Data Warehouse

The installation is self-contained and includes the elements required to run and operate OnCommand Insight Data Warehouse (DWH).

Before you begin

You must have completed all of the installation prerequisites.

About this task

Data Warehouse has Cognos reporting capabilities. If you install Insight on a Linux server, you can use these capabilities, however, only if you install the Data Warehouse on a Windows server. For information about installing the Data Warehouse on Windows and Cognos reporting capabilities, refer to the *OnCommand Insight Installation Guide for Microsoft Windows*.

Steps

1. Log in to the Data Warehouse server using an account with sudo privileges.
2. Navigate to the directory on the server where the installation files are located and type the following command:

```
unzip oci-dwh-<version>-linux-x86_64.zip
```

Ensure that you check the version number of the installation file; the version number might be different than the one shown in the command.

3. You can view syntax, command arguments, and parameter usage for `oci-install.sh` before you begin the installation:

```
sudo ./oci-dwh-<version>-linux-x86_64/oci-install.sh --help
```

4. Run the installation script:

```
sudo ./oci-dwh-<version>-linux-x86_64/oci-install.sh
```

5. Read the License Agreement, accept it, and follow the prompts.

Results

After you answer all the prompts, the installation begins and should take approximately 10 minutes, depending on the applications installed.

Installing a Remote Acquisition Unit

You can install one or more Remote Acquisition Units (RAUs) in your OnCommand Insight environment. Acquisition units run in the network that accesses (through modules called data *sources*) and collect data from different devices in the data center.

Before you begin

You must have completed all of the installation prerequisites.

At least one port must be open and available between the RAU server and the OnCommand Insight Server to forward change information to the server. If you are unsure about this, validate it by opening a Web browser on the RAU computer and directing it to the OnCommand Insight server:

```
https://< OnCommand Insight Server hostname >:< acquisition_port >
```

The acquisition port defaults to 443, but it might have changed during the server installation. If the connection is successful, you see a OnCommand Insight response page, indicating an open and available port between the RAU and the OnCommand Insight server.

For environments using Network Address Translation or Port Address Translation (NAT/PAT: i.e, any translation of IP addresses), Insight only supports insertion of an RAU between NAT and the Device.

- Supported: OnCommand Insight -> NAT -> RAU -> Device
- Unsupported: OnCommand Insight -> RAU -> NAT -> Device

Steps

1. Log in to the RAU server using an account with sudo privileges.
2. Navigate to the directory on the server where the installation files are located and type the following command:

```
unzip oci-rau-<version>-linux-x86_64.zip
```

3. You can view syntax, command arguments, and parameter usage for `oci-install.sh`:

```
sudo ./oci-rau-<version>-linux-x86_64/oci-install.sh --help
```

4. Run the installation script:

```
sudo ./oci-rau-<version>-linux-x86_64/oci-install.sh
```

5. Read the License Agreement, accept it, and then follow the prompts.

After you answer all the prompts, the installation begins and should take approximately 10 minutes, depending on the applications installed.

Validating the remote acquisition unit installation

To validate proper installation of the Remote Acquisition Unit, you can view the status of

the Remote Acquisition Units connected to your server.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Acquisition Units**.
3. Verify that the new Remote Acquisition Unit was registered correctly and that it has a Connected status.

If it does not have a Connected status, try restarting the service. Log into the Remote Acquisition Unit system and execute the following command:

```
oci-service.sh restart acquisition
```

If it still does not connect, contact technical support.

Checking the installation

After you complete the installation, the installation directory is located in `/opt/netapp/oci`. You can open Insight in a supported browser to check the installation. You might also want to check the Insight log files.

When you first open Insight, the license setup page opens. After you enter the license information, you must set up the data sources. See the *OnCommand Insight Configuration and Administration Guide* for information about entering data source definitions and setting up Insight users and notifications.

If you have experienced installation problems, contact technical support and provide the requested information.

Verifying that new Insight components are installed

After installation, you should verify the existence of the new components on your server.

Steps

1. To display a list of services that are currently operating on the server you are logged in to, type:

```
sudo oci-service.sh status all
```

2. Depending on the server you are logged in to, check for the following Insight services in the list and ensure they have a status of "running".
 - Insight server: wildfly, acquisition, mysql, elasticsearch
 - Data Warehouse server: wildfly, mysql
 - Remote Acquisition server: acquisition

Results

If these components are not listed, contact technical support.

Insight logs

Insight supplies many log files to assist you with research and troubleshooting. The available logs are listed in the log directory. You might want to use a log monitoring tool, such as BareTail, to display all of the logs at one time.

The log files are located in the `/var/log/netapp/oci/wildfly/` directory. Acquisition logs are located in the `/var/log/netapp/oci/acq` directory. The data files are located in `/var/lib/netapp/oci`.

Accessing the web UI

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

Steps

1. Do one of the following:

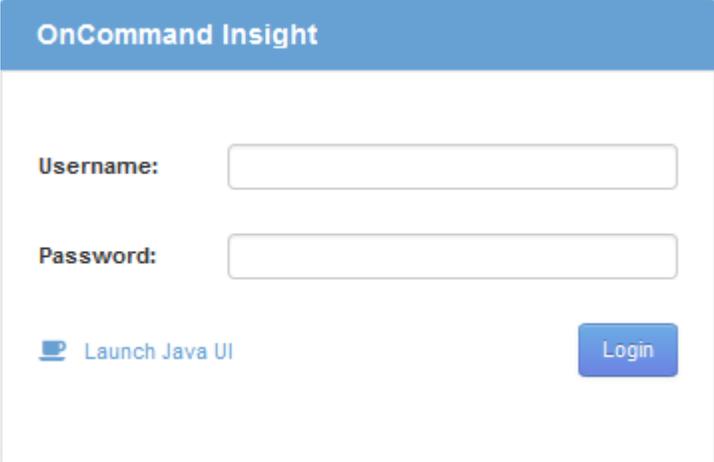
- Open Insight on the Insight server:

```
https://fqdn
```

- Open Insight from any other location:

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

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



The OnCommand Insight dialog box displays:

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

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



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

Installing your Insight licenses

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

About this task

Insight license keys are stored in a `.txt` or `.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**.
10. If you are using the Insight consumption licensing model, you must check the box to **Enable sending usage information to NetApp** in the **Send usage information** section. Proxy must be properly configured and enabled for your environment.

After you finish

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

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

OnCommand Insight licenses

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

• Discover

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

• Assure

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

• Perform

A Perform license supports performance monitoring on asset pages, dashboard widgets, queries, and so

on, as well as managing performance policies and violations.

- **Plan**

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

- **Host Utilization pack**

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

- **Report Authoring**

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

OnCommand Insight modules are licensed for annual term or perpetual:

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

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

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

- **Discover**

Acquire and manage inventory (Foundation)

Monitor changes and manage inventory policies

- **Assure**

View and manage SAN path policies and violations

View and manage vulnerabilities

View and manage tasks and migrations

- **Plan**

View and manage requests

View and manage pending tasks

View and manage reservation violations

View and manage port balance violations

- **Perform**

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

View and manage performance policies and violations

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

Feature (admin)	With Perform license	Without Perform license
Application	Yes	No performance data or charts
Virtual machine	Yes	No performance data or charts
Hypervisor	Yes	No performance data or charts
Host	Yes	No performance data or charts
Datastore	Yes	No performance data or charts
VMDK	Yes	No performance data or charts
Internal volume	Yes	No performance data or charts
Volume	Yes	No performance data or charts
Storage pool	Yes	No performance data or charts
Disk	Yes	No performance data or charts
Storage	Yes	No performance data or charts
Storage node	Yes	No performance data or charts
Fabric	Yes	No performance data or charts
Switch port	Yes	No performance data or charts; "Port Errors" shows "N/A"
Storage port	Yes	Yes
NPV port	Yes	No performance data or charts
Switch	Yes	No performance data or charts
NPV switch	Yes	No performance data or charts
Qtrees	Yes	No performance data or charts
Quota	Yes	No performance data or charts

Path	Yes	No performance data or charts
Zone	Yes	No performance data or charts
Zone member	Yes	No performance data or charts
Generic device	Yes	No performance data or charts
Tape	Yes	No performance data or charts
Masking	Yes	No performance data or charts
ISCSI sessions	Yes	No performance data or charts
ICSI network portals	Yes	No performance data or charts
Search	Yes	Yes
Admin	Yes	Yes
Dashboard	Yes	Yes
Widgets	Yes	Partially available (only asset, query, and admin widgets are available)
Violations dashboard	Yes	Hidden
Assets dashboard	Yes	Partially available (storage IOPS and VM IOPS widgets are hidden)
Manage performance policies	Yes	Hidden
Manage annotations	Yes	Yes
Manage annotation rules	Yes	Yes
Manage applications	Yes	Yes
Queries	Yes	Yes
Manage business entities	Yes	Yes

Feature	User - with Perform license	Guest - with Perform license	User - without Perform license	Guest - without Perform license
---------	-----------------------------	------------------------------	--------------------------------	---------------------------------

Assets dashboard	Yes	Yes	Partially available (storage IOPS and VM IOPS widgets are hidden)	Partially available (storage IOPS and VM IOPS widgets are hidden)
Custom dashboard	View only (no create, edit, or save options)	View only (no create, edit, or save options)	View only (no create, edit, or save options)	View only (no create, edit, or save options)
Manage performance policies	Yes	Hidden	Hidden	Hidden
Manage annotations	Yes	Hidden	Yes	Hidden
Manage applications	Yes	Hidden	Yes	Hidden
Manage business entities	Yes	Hidden	Yes	Hidden
Queries	Yes	View and edit only (no save option)	Yes	View and edit only (no save option)

Troubleshooting installations

OnCommand Insight installations are generally managed through the installation wizards. However, customers might experience problems during upgrades or with conflicts due to computer environments.

You should also be certain that you install all of the necessary OnCommand Insight licenses for installing the software.

Missing licenses

Different licenses are required for different OnCommand Insight functionality. What you see displayed in OnCommand Insight is controlled by your installed licenses. Refer to the OnCommand Insight licenses section for information on functionality controlled by each license.

Refer to the OnCommand Insight licenses section for information on functionality controlled by each license.

Submitting an online technical support request

If you have problems with the Insight installation, as a registered support customer, you can submit an online technical support request.

Before you begin

Using your corporate email address, you must register as a support customer to obtain online support services.

Registration is performed through the support site (<http://support.netapp.com>).

About this task

To assist customer support in solving the installation problem, you should gather as much information as possible, including these items:

- Insight serial number
- Description of the problem
- All Insight log files
- Screen capture of any error messages

Steps

1. Create a `.zip` file of the information you gathered to create a troubleshooting package.
2. Log in to the support site at mysupport.netapp.com and select **Technical Assistance**.
3. Click **Open a Case**.
4. Follow the instructions to your package of data.

After you finish

You can use **Check Case Status** on the Technical Assistance page to follow your request.

Upgrading Insight

When a new version of OnCommand Insight is available, you might want to upgrade to take advantage of new features and fixes to issues. You must upgrade the Insight server and Data Warehouse (DWH) separately.



You should not store any automatic or manual backups in Insight installation directories, because the entire installation folder is overwritten during the upgrade process. If you have stored backup files in any of those directories, you must move your backups to a different location before you perform any upgrade or uninstall process.

Newer versions of Insight have greater disk space, memory and CPU requirements. Before upgrading to the latest version of Insight, review the Installation requirements. It is strongly recommended to contact your Sales Engineer for detailed sizing guidance before installing or upgrading Insight.

It is Best Practice to perform a security backup and a database backup before upgrading Insight software.

Upgrading Insight to version 7.3.12 or later - Linux

Prior to upgrading from OnCommand Insight 7.3.10 - 7.3.11 to version 7.3.12 or later, you must run the OCI Data Migration Tool.

Background

OnCommand Insight versions 7.3.12 and later utilize underlying software that may be incompatible with previous versions. Insight versions 7.3.12 and later include a **Data Migration Tool** to assist with upgrading.



OnCommand Insight versions 7.3.9 and earlier are no longer supported. If you are running one of these versions, you *must* upgrade to Insight version 7.3.10 or later (7.3.11 is strongly recommended) prior to upgrading to 7.3.12 or later.

What Does The Data Migration Tool Do?

The migration tool performs an initial compatibility check and then follows one of three different upgrade paths. The path selected is based on the data compatibility of your current version.



Prior to upgrading, you must run the Data Migration Tool and follow the recommended steps.

Before you Begin

- It is strongly recommended to back up your OnCommand Insight system prior to running the Data Migration Tool.
- The Elasticsearch service on the server needs to be up and running.
- The Data Migration Tool *must* be run for the database and any performance archives before you upgrade Insight.

Running the Data Migration Tool

1. Download the latest version of the Data Migration Tool (for example, *SANScreenDataMigrationTool-x86-7.3.12-97.zip*) to your Insight server, as well as the appropriate Insight installer file. Unzip into a working folder. Downloads can be found on the [NetApp Support Site](#).
2. Open a command window and navigate to your working folder.
 - Bash shell is recommended.
3. Run the data migration tool using the following command:
 - ``sudo ./SANScreenDataMigrationTool.sh``
4. Follow the instructions as needed. The following is an example.

```
sudo ./SansscreenDataMigrationTool.sh

NetApp SANScreen Data Migration Tool 7.3.12-132

OCI 7.3.10.8.139 is installed
Elasticsearch REST port = 9200

Checking Elasticsearch service...
Elasticsearch service is up

Checking for obsolete (version 5) indexes...
Found 54 obsolete indexes. Of these,
    54 indexes may be migrated with OCI server running,
        the most recent of which is for 2021-05-13

Verifying migration component is present...
SANscreen Server service is Running

Proceed with online migration of 54 indexes (y or [n])?:
```

The Data Migration Tool will check for the presence of obsolete indexes on your system and report if any are found. If none are present the tool will exit.

Some indexes may be migrated while the SANscreen Server service is running. Others may only be migrated when the server is stopped. If there are no indexes that may be migrated the tool will exit. Otherwise follow the instructions as prompted.

After the Data Migration Tool completes it will recheck for obsolete indexes. If all indexes have been migrated, the tool will inform you that upgrade to OnCommand Insight 7.3.12 is supported. You can now proceed with upgrading Insight.

```
sudo ./SansscreenDataMigrationTool.sh

NetApp SANScreen Data Migration Tool 7.3.12-132

OCI 7.3.10.8.139 is installed
Elasticsearch REST port = 9200

Checking for obsolete (version 5) indexes...
Found 76 obsolete OCI indexes. Of these,
76 indexes may be migrated with OCI server running

SANscreen Server service is running

Proceed with online migration of 76 indexes (y or [n])? y
If you supply performance archive location, entries for any dates with
migrated
indexes will be replaced. Each original entry will be renamed and you may
delete
it after migration is completed.
When prompted enter the archive location including the site-name
directory.

Enter the location of the performance archive or blank if none:
Performance archive entries will not be updated

Running the migration application with options -u http://localhost:9200
--online -sa -

Preparing to migrate oci-timeseries-disk-2021-03-22: copied; backup;
delete old; restore new; cleanup; done.
Preparing to migrate oci-timeseries-internalvolume-2021-03-22: copied;
backup; delete old; restore new; cleanup; done.
Preparing to migrate oci-timeseries-port-2021-03-22: copied; backup;
delete old; restore new; cleanup; done.
...
Preparing to migrate oci-timeseries-disk-2021-03-27: copied; backup;
delete old; restore new; cleanup; done.
Execution time 0:08:17
Checking for obsolete (version 5) indexes...

No obsolete indexes found. Upgrade and Inline Upgrade to 7.3.12+ are
supported
```

If you were prompted to stop the SANScreen service, restart it before upgrading Insight.

Validation failures

In the event that index validation fails, the migration tool will inform you of the problem before quitting.

OnCommand Insight is not present:

```
./SanscreenDataMigrationTool.sh

NetApp SANScreen Data Migration Tool V1.0

Checking OnCommand Insight Installation...
ERROR: OnCommand Insight is not installed
```

Invalid Insight version:

```
./SanscreenDataMigrationTool.sh

NetApp SANScreen Data Migration Tool 7.3.12-105

Checking OnCommand Insight Installation...
OnCommand Insight 7.3.4 (126) is installed
ERROR: The OCI Data Migration Tool is intended to be run against OCI 7.3.5
- 7.3.11
```

Elasticsearch service is not running:

```
./SanscreenDataMigrationTool.sh
NetApp SANScreen Data Migration Tool 7.3.12-105

Checking OnCommand Insight Installation...
OnCommand Insight 7.3.11 (126) is installed

Getting installation parameters...
Elasticsearch Rest Port: 9200

Checking Elasticsearch service...
ERROR: The Elasticsearch service is not running

Please start the service and wait for initialization to complete
Then rerun OCI Data Migration Tool
```

Command-line options

The Data Migration Tool includes some optional parameters that affect its operation.

Option (Linux)	Function
-s --silent	Suppress all prompts
-a --archive	<p>If specified, existing archive entries for any date whose index(es) are migrated will be replaced. The path should point to the directory containing the archive entry zip files.</p> <p>An argument of '-' may be specified to indicate there is no performance archive to be updated.</p> <p>If this argument is present, the prompt for the archive location will be suppressed.</p>
-c --check	If present, the script will exit immediately after reporting the index counts.
-d --dryrun	If present, then the migration executable will report the actions that would be taken (to migrate data and update archive entries) but will not perform the operations.
-p --port	<p>If present, use the supplied value as Elasticsearch's REST port. If absent, obtain the value from the installation if possible; otherwise use the default value of 9200.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;">  <p>In some Linux OnCommand Insight installations, the Elasticsearch REST port might not be running on the default 9200 port. In this case use the --port option to supply the value</p> </div>
-h --help	Display usage information

Troubleshooting

If archive entries were updated, you *must* make sure that the ownership and permissions on the updated archives are correct. They should be **ocisys ocisys 644**. If they are not, navigate into the performance archive folder and run the following commands:

```
chown ocisys *
chgrp ocisys *
chmod 644 *
```

Upgrading Insight Server software

You can check for OnCommand Insight server updates after you log into the server.

Steps

1. On the Insight toolbar, click the **Help** icon.
2. Select **Check for updates**.
3. Click **OK** if the `Version is up to date` message displays.
4. If a newer version is detected, click the **download here** link in the message box.
5. In the **Download** page, click **download**. Note the download directory location.

You can also download the newer version from the NetApp support site.

6. Log in to the Insight server using an account with sudo privileges.
7. Navigate to the download directory and type the following command:

```
unzip oci-<version>-linux-x86_64.zip
```

Ensure that you have the correct the version number of the installation file.

8. You can view syntax, command arguments, and parameter usage for `oci-install.sh`:

```
sudo ./oci-<version>-linux-x86_64/oci-install.sh --help
```

9. Run the installation script:

```
sudo ./oci-<version>-linux-x86_64/oci-install.sh
```

10. Accept the License Agreement and follow the prompts.

Upgrading Data Warehouse software

After upgrading the Insight server software, you must upgrade your data warehouse software.

Steps

1. Log in to the Data Warehouse (DWH) server using an account with sudo privileges.
2. Download the Insight DWH software from the NetApp support site.
3. Navigate to the download directory and type the following command:

```
unzip oci-dwh-<version>-linux-x86_64.zip
```

Ensure that you have the correct the version number of the installation file.

4. You can view syntax, command arguments, and parameter usage for `oci-install.sh`:

```
sudo ./oci-dwh-<version>-linux-x86_64/oci-install.sh --help
```

5. Run the installation script:

```
sudo ./oci-dwh-<version>-linux-x86_64/oci-install.sh
```

6. Accept the License Agreement and follow the prompts.

Upgrading Remote Acquisition Unit software

After upgrading the Insight server software, you must upgrade your remote acquisition software.

Steps

1. Log in to the Remote Acquisition Unit (RAU) server using an account with sudo privileges.
2. Download the Insight RAU software from the NetApp support site.
3. Navigate to the download directory and type the following command:

```
unzip oci-rau-<version>-linux-x86_64.zip
```

Ensure that you have the correct the version number of the installation file.

4. You can view syntax, command arguments, and parameter usage for `oci-install.sh`:

```
sudo ./oci-rau-<version>-linux-x86_64/oci-install.sh --help
```

5. Run the installation script:

```
sudo ./oci-rau-<version>-linux-x86_64/oci-install.sh
```

6. Accept the License Agreement and follow the prompts.

Migrating from Windows to Linux

To use Insight on Linux when you have an existing Windows installation, you must perform a migration. You must perform this procedure on both the Insight server and Data Warehouse components.

Steps

1. Back up your current Insight installation on your server.

Refer to the *OnCommand Insight Configuration and Administration Guide* for information about how to back up the OCI database.

2. Install Insight for Linux.
3. Restore the database for your previous version.

Refer to the *OnCommand Insight Configuration and Administration Guide* for information about how to restore the OCI database.

4. Uninstall your previous version of Insight for Windows.

Uninstalling OnCommand Insight

You can uninstall the OnCommand Insight components if needed. You must uninstall the OnCommand Insight components separately.

Each component is uninstalled separately.

Uninstalling the OnCommand Insight Server

You can uninstall the OnCommand Insight server if needed.

Before you begin

Best practice: before uninstalling Insight, back up the OnCommand Insight database.

Steps

1. Log in to the OnCommand Insight server using an account with sudo privileges.
2. Ensure that any OnCommand Insight windows are closed.
3. You can view syntax, command arguments, and parameter usage for `oci-uninstall.sh` by entering the following command:

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

A normal uninstall does not remove the Insight license or any daily backups. To remove the entire installation, use the `--purge` option with the `oci-install.sh` command.

4. Type the following command:

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

Uninstalling Data Warehouse

You can uninstall Data Warehouse if needed.

Before you begin

Back up the current version of the OnCommand Insight Data Warehouse (DWH) database.

About this task

Uninstalling the OnCommand Insight Data Warehouse permanently deletes all previously collected data.

Steps

1. Log in to the Data Warehouse server using an account with sudo privileges.
2. Ensure that any OnCommand Insight windows are closed.
3. You can view syntax, command arguments, and parameter usage for `uninstall.sh` by entering the following command: `sudo /usr/bin/oci-uninstall.sh --help`

4. Type the following command: `sudo /usr/bin/oci-uninstall.sh`

Uninstalling a Remote Acquisition Unit

You can uninstall a Remote Acquisition Unit when you no longer need it.

Steps

1. Log in to the Remote Acquisition Unit server using an account with sudo privileges.
2. Ensure that any OnCommand Insight windows are closed.
3. You can view syntax, command arguments, and parameter usage for `uninstall.sh` by entering the following command: `sudo /usr/bin/oci-uninstall.sh --help`
4. Type the following command: `sudo /usr/bin/oci-uninstall.sh`

The uninstall script runs. Follow any prompts.

Installation for Microsoft Windows

Installation prerequisites

Before you install OnCommand Insight, you must download the current software version, acquire the appropriate license, and set up your environment.

Before installing OnCommand Insight, ensure that you have the following:

- OnCommand Insight software files in the downloaded installation package for the current version
- A license to operate the downloaded OnCommand Insight version
- The minimum hardware and software environment

The current product might consume additional hardware resources (due to enhanced OnCommand Insight product functionality) that were not consumed with earlier versions of the OnCommand Insight product.

- A deployment plan that includes the hardware and network configurations for the OnCommand Insight Server, Data Warehouse and Reporting, and remote acquisition units.
- Disabled virus scan software

During the installation of OnCommand Insight, you must completely disable all virus scanners. Following installation, the paths used by the Insight component (install, backup, and archiver paths) must be excluded from virus scanning, in addition to excluding the entire `sanscreen` directory from the scan.

Additionally, you must also exclude the IBM/Db2 folder (for example `C:\Program Files\IBM\DB2`) from anti-virus scanning following installation.

If you are performing a full installation as an upgrade or as a migration to new hardware and your existing system contains a non-default security configuration, you must back up the security configuration before you perform the installation. After the installation is complete, you must restore the security configuration before you restore the Server (which includes the local acquisition unit) or Data Warehouse database. You must restore the security configuration to all of your Insight servers before you restore the DWH Database.



For in-place upgrade (available for Insight Server only), the security configuration is properly handled and you do not need to restore it.

You use the `securityadmin` tool to create a backup of the configuration and to restore the saved configuration. For more information, search for `securityadmin` in the OnCommand Insight Documentation Center: <http://docs.netapp.com/oci-73/index.jsp>

Planning the deployment

To ensure a successful deployment, you must consider certain system elements before you install OnCommand Insight.

About this task

Planning your Insight deployment includes considering these system elements:

- Insight architecture
- Your network components to be monitored
- Insight installation prerequisites and server requirements
- Insight web browser requirements

Data source support information

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

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

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

Device identification and data source planning

As part of your deployment planning, you should collect information about the devices in your environment.

You need the following software, connectivity, and information about each device in your environment:

- IP address or hostname resolvable by the OCI server
- Login name and password
- Type of access to the device, for example, controller and management station



Read-only access will be sufficient for most devices, but some devices require administrator permissions.

- Port connectivity to the device depending on data source port requirements
- For switches, SNMP read-only community string (user ID or password to give access to the switches)
- Any third-party software required on the device, for example, Solutions Enabler.
- See the "Vendor-specific data source reference" in the web UI Help or in the *OnCommand Insight Configuration and Administration Guide* for more information on data source permissions and requirements.

Network traffic generated by OnCommand Insight

The network traffic that OnCommand Insight generates, the amount of processed data traversing the network, and the load that OnCommand Insight places on devices differ based on many factors.

The traffic, data, and load differ across environments based on the following factors:

- The raw data
- Configuration of devices

- Deployment topology of OnCommand Insight
- Different inventory and performance data source polling intervals, which can be reduced to allow for slow devices to be discovered or bandwidth to be conserved

The raw configuration data that OnCommand Insight collects can vary significantly.

The following example illustrates how the configuration data can vary and how traffic, data, and load are affected by many configuration factors. For example, you might have two arrays each having 1,000 disks:

- Array 1: Has 1,000 SATA disks all 1 TB in size. All 1,000 disks are in one storage pool, and there are 1,000 LUNs, all presented (mapped and masked) to the same 32 nodes in an ESX cluster.
- Array 2: Has 400 2-TB data disks, 560 600-GB FC disks, and 40 SSD. There are 3 storage pools, but 320 of the FC disks are used in traditional RAID groups. The LUNs carved on the RAID groups use a traditional masking type (symmaskdb), while the thin provisioned, pool-based LUNs use a newer masking type (symaccess). There are 600 LUNs presented to 150 different hosts. There are 200 BCVs (full block replica volumes of 200 of the 600 LUNs). There are also 200 R2 volumes, remote replica volumes of volumes that exist on an array in a different site.

These arrays each have 1,000 disks and 1,000 logical volumes. They might be physically identical in the amount of rack space they consume in the data center, and they might even be running the same firmware, but the second array is much more complex in its configuration than the first array.

Virus scan software disablement

If antivirus software is active on your system, OnCommand Insight installation fails. You can prevent this problem by disabling the virus scan software before installation.

To prevent an installation failure due to active virus scan software, during the installation of each OnCommand Insight component, you must completely disable all virus scanners. Following installation, the paths used by the Insight component (install, backup, and archiver paths) must be excluded from virus scanning.

Additionally, you must also exclude the IBM/Db2 folder (for example *C:\Program Files\IBM\DB2*) from anti-virus scanning following installation.

Insight Server requirements

A dedicated server is recommended. Do not install Insight on a server that has any other applications installed. Both physical and virtual servers are supported, provided that the product requirements are met.

You must have local administrator permissions to install the OnCommand Insight Server software.



Sizing for OnCommand Insight has multiple dependencies, such as data source type and size, number of assets in your environment, polling intervals, and more. The following sizing examples are guidelines only; they represent some of the environments in which Insight has been tested. Changing any of these or other factors in the environment can change the sizing requirements for Insight. These guidelines include disk space for up to 90 days of performance archive data.

It is recommended to contact your Sales Engineer for detailed sizing guidance before installing or upgrading Insight.

Examples:

Environment factors:	Disk space, CPUs, and Memory tested:
80 storage arrays 4,000 Volumes 4,000 VMs 4,000 switch ports	250 GB disk space 8 cores 32 GB RAM
160 storage arrays 40,000 Volumes 8,000 VMs 8,000 switch ports	1 TB of disk space 12 cores 48 GB RAM

Requirements:

Component	Required
Operating system	<p>A computer running 64-bit Microsoft Windows Server 2016, 2019, or 2022, with the latest service pack.</p> <p>The Resilient File System (ReFS) introduced with Windows Server 2012 is not compatible with OnCommand Insight. Windows installation of OnCommand Insight is only supported on the NTFS file system.</p> <p>A dedicated server is recommended.</p>
Virtual machine (VM)	This component can run in a virtual environment, provided that the CPU and memory resources for your instance are reserved.
Memory and CPU	<p>24 - 256 GB RAM</p> <p>8 - 32 cores</p> <p>It is strongly recommended to set the paging file size to "Windows managed". Small, fixed-size paging files may interfere with the successful storage of Insight performance data.</p>
Available disk space	<p>100 GB - 3 TB install disk space</p> <p>50 GB - 1 TB performance archive disk space</p> <p>SSD disks are recommended for the Insight installation space.</p>

<p>Network</p>	<p>Ethernet connection and ports:</p> <ul style="list-style-type: none"> • 100 Mbps or 1 Gbps Ethernet connection with dedicated (static) IP address and IP connectivity to all components in the SAN, including FC devices and remote acquisition units. • Port requirements for the OnCommand Insight Server process are 80, 443, 1090 through 1100, 3873, 8083, 4444 through 4446, 5445, 5455, 4712 through 4714, 5500, and 5501. • Port requirements for the acquisition process are 12123 and 5679. • Port requirement for MySQL is 3306. • Port requirements for Elasticsearch are 9200 and 9310 • Dynamic port requirements on Win2008/2012 are 49152 through 65535 <p>Ports 443 and 3306 require external access through any firewall that is present.</p>
<p>Permissions</p>	<p>Local administrator permissions are required on the OnCommand Insight Server.</p> <p>If any of the following folders are symbolic links, ensure that the destination directories have '755' permissions.</p> <ul style="list-style-type: none"> • /opt/netapp • /var/lib/netapp • /var/log/netapp
<p>Remote connectivity</p>	<p>Internet connectivity to allow WebEx access or a remote desktop connection to facilitate installation and post-installation support.</p>
<p>Accessibility</p>	<p>HTTPS access is required.</p>
<p>Virus scan</p>	<p>During the installation of this OnCommand Insight component, you must completely disable all virus scanners. Following installation, the paths used by the Insight component (install, backup, and archiver paths) must be excluded from virus scanning.</p> <p>Additionally, you must also exclude the IBM/Db2 folder (for example <i>C:\Program Files\IBM\DB2</i>) from anti-virus scanning following installation.</p>

HTTP or HTTPS servers	Microsoft Internet Information Services (IIS) or other HTTPS servers should not compete for the same ports (443) as the OnCommand Insight server, and should not start automatically. If they must listen to port 443, then you must configure the OnCommand Insight server to use other ports.
-----------------------	---

Data Warehouse and Reporting server requirements

You must run the Data Warehouse and the Reporting server on a computer that is compatible with established hardware and software requirements, ensuring that Apache web server or reporting software is not already installed on this machine.



Sizing for OnCommand Insight has multiple dependencies, such as number of assets in your environment, amount of historical data retained, and more. The following data warehouse sizing examples are guidelines only; they represent some of the environments in which Insight has been tested. Changing any of these or other factors in the environment can change the sizing requirements for Insight.

It is recommended to contact your Sales Engineer for detailed sizing guidance before installing or upgrading Insight.

Examples:

Environment factors:	Disk space, CPUs, and Memory tested:
18 storage arrays 3,400 VMs 4,500 switch ports	200 GB hard disk 8 cores 32 GB RAM
110 storage arrays 11,500 VMs 14,500 switch ports	300 GB hard disk 8 cores 48 GB RAM

Requirements:

Component	Required
Operating system	A computer running 64-bit Microsoft Windows Server 2016, 2019, or 2022, with the latest service pack.
Virtual machine (VM)	This component can run in a virtual environment, provided that the CPU and memory resources for your instance are reserved.
CPU	8 - 40 CPU cores

<p>Memory</p>	<p>32 GB - 2 TB RAM Best Practice: It is strongly recommended to set the paging file size to “Windows managed”. Small, fixed-size paging files may interfere with the successful storage of Insight performance data.</p>
<p>Available disk space</p>	<p>200 GB - 2 TB disk space Installation requires a minimum of 20 GB free on the C: drive.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> On Windows, Insight Data Warehouse with Reporting requires the 8dot3 name creation support be enabled on the installation drive prior to installing. The C: drive typically has this enabled by default. You can validate if 8dot3 name creation is enabled on the target installation drive by running the following command (substitute D: with target installation drive):</p> </div> <p>fsutil 8dot3name query D:</p> <p>To enable 8dot3 name creation execute the following command (substitute D: with target installation drive):</p> <p>fsutil 8dot3name set D: 0</p>
<p>Network</p>	<ul style="list-style-type: none"> • 100 Mbps or 1 Gbps Ethernet connection • Static IP address • Port 50000 must be available before installing Data Warehouse with Reporting on Windows • For the OnCommand Insight DWH server process, ports 80, 443, 1098, 1099, 3873, 8083, and 4444 through 4446 • For the reporting engine, ports 1527, 9362, 9300, and 9399 • For MySQL, port 3306 • Ensure that DNS is properly working by doing an <code>nslookup</code> against the host
<p>Virus Scan</p>	<p>During the installation of this OnCommand Insight component, you must completely disable all virus scanners. Following installation, the paths used by the Insight component (install, backup, and archiver paths) and all DWH component installation paths (SANscreen, DB2, and backup paths) must be excluded from virus scanning.</p>

Visual Studio	Visual Studio 2019 redistributables must be installed before installing Data Warehouse with Reporting on Windows.
---------------	---

Remote Acquisition Unit server requirements

You must install a Remote Acquisition Unit (RAU) to acquire information from SAN devices that are behind a firewall, at a remote site, on a private network, or in different network segments. Before you install the RAU, you should ensure that your environment meets RAU operating system, CPU, memory, and disk space requirements.

Component	Requirement
Operating system	A computer running 64-bit Microsoft Windows Server 2016, 2019, or 2022, with the latest service pack.
CPU	4 CPU cores
Memory	16 GB RAM
Available disk space	40 GB
Network	100 Mbps /1 Gbps Ethernet connection, static IP address, IP connectivity to all FC devices, and a required port to the OnCommand Insight server (80 or 443).
Permissions	Local Administrator permissions on the RAU server
Virus scan	During the installation of this OnCommand Insight component, you must completely disable all virus scanners. Following installation, the paths used by the Insight component must be excluded from virus scanning. Additionally, you must also exclude the IBM/Db2 folder (for example <i>C:\Program Files\IBM\DB2</i>) from anti-virus scanning following installation.

Browsers supported by OnCommand Insight

The browser-based OnCommand Insightweb UI can operate on several different browsers.

Insight supports newer, non-beta releases of the following browsers:

- Mozilla Firefox
- Google Chrome

- Microsoft Edge

For a full list of browser versions qualified for OnCommand Insight, please see the [NetApp Interoperability Matrix Tool](#).

Insight installation instructions

Installation requires you to install several OnCommand Insight components, including Insight Server, and Data Warehouse and Reporting.

The installation includes the following major tasks:

- Downloading the OnCommand Insight installer
- Installing OnCommand Insight server
- Installing licenses
- Optionally, installing DWH and Reporting (must be installed on a separate machine or virtual machine)
- Optionally, installing a remote acquisition unit (RAU), which acquires information from your device resources that reside behind a firewall, are located at a remote site, or are on a private network
- For upgrades, upgrading OnCommand Insight reports.

After installation, you must configure Insight to acquire information about your environment. The tasks required are described in the *OnCommand Insight Configuration and Administration Guide*.

Downloading the OnCommand Insight installer

You can download the OnCommand Insight installer from the [NetApp Support Site](#).

Before you begin

You must have a login to the NetApp Support Site at mysupport.netapp.com.

Steps

1. Log in to the server on which you want to install OnCommand Insight.
2. Download the installation file from the NetApp Support site.

Installing the OnCommand Insight Server

You can easily install the OnCommand Insight Server by using the OnCommand Insight Setup wizard.

Before you begin

You must have completed all of the installation prerequisites.

Steps

1. Log in to the Insight server using an account with administrator privileges.
2. Open Windows Explorer and navigate to the directory where the installation files are located.

3. Double-click the .MSI file that you downloaded.
4. Click **Next** to continue.
5. Read the License Agreement, select **I accept the terms in the License Agreement** check box, and then click **Next**.
6. Enter the customer name and site name in the **Customer Information** window, and click **Next**.

Best Practice: Use the customer name as a prefix for the site: for example, NetApp.

7. In the **Customer Information: Configure NetApp ASUP** window, do the following:
 - a. Select the database containing the data that you want to upload to ASUP by selecting one of the following options:
 - **No database backup:** A backup is not sent to ASUP.
 - **Backup without Performance data:** A backup is made and sent to ASUP but does not include performance data.
 - **Backup with Performance data:** A backup is made that includes performance data, but this could generate a huge *.gz file.



ASUP is delivered using HTTPS protocol.

- a. In **Logs**, select whether you want no logs, base logs, or extended logs, which contain a data source recording.
 - b. Click **Next**.
8. If you are using the Insight consumption licensing model, you must check the box to **Enable sending usage information to NetApp** in the **Send usage information** section.
9. Click **Next**
10. In the **Configure Server** window, select or set the appropriate configuration parameters to configure the OnCommand Insight Server:

Option	Description
Portal Port (HTTP)	Ports used by the OnCommand Insight Server to support user Web services, including a portal to perform administration tasks. Use the default (80); however, if the default port is in use, change this to another port.
Portal Port (HTTPS)	Port used by remote acquisition units to send SAN change information to the OnCommand Insight Server through a secure channel. Use the default (443); however, if the default port is in use, change this to another port. You specify this same port number when configuring RAUs.

Internal Database Port (SQL)	Port used internally by the PC where the OnCommand Insight Server is running, to serve as an access point to the database. Use the default (3306); however, if the default port is in use, change this to another port.
------------------------------	---

11. Click **Next**.
12. Click **Install** to proceed.

The installation should take approximately 20 minutes, depending on the applications installed.

13. Click **Finish**.

Installing OnCommand Insight Data Warehouse and Reporting

The installation is self-contained and includes the elements required to run and operate OnCommand Insight Data Warehouse (DWH) and the Reporting utilities.

Before you begin

Please note the following before installing or upgrading.

- If you are upgrading, back up DWH.
- You must have local *administrator* permissions to install OnCommand Insight Data Warehouse with Reporting.
- Make sure Windows Modules Installer service is enabled (either automatically or manually).
- If installing on non-C: drive, Short File Names must be enabled. If it is not enabled, the installer will enable it.
- For the Db2 component, the Db2 User can be either *domain* user or *local* user.
 - If the Db2 User is a *domain* user, you must have the following:
 - Db2 User must have been already created, and you must know the user name and password
 - As the user who is installing DWH with Reporting, you must be able to query the Db2 User. You can validate this using the command:


```
net user <db2 user name> /domain
```
 - If Db2 User is a *local* user, you must have the following:
 - User name and password for the user which will be used to run as Db2 User. If this user does not exist, installation will create it.
 - [NOTE]

The Db2 user name as well as the Windows login name have the following restrictions:

* Valid characters are: 'A' through 'Z'; 'a' through 'z'; '0' through '9'; '#'; '@'; ';'; '!'; ' ('; '); '{'; '}'; '-'; and '.'.

* *If using the special characters '!', ' ('; '); '{'; '}'; '-'; and '.' you must use all uppercase letters for the user name.*

* *The first character in the string must be an alphabetic character, @, #, or \$; it cannot be a number or the letter sequences _SYS, DBM, or IBM*

* It cannot exceed 128 bytes in length.

* It cannot be USERS, ADMINS, GUESTS, PUBLIC, LOCAL or any SQL reserved word.

- The Db2 user can not be the same as the user performing the installation.

Steps

1. Log in to the Data Warehouse server using an account with administrator privileges.
2. Download the Data Warehouse with Reporting .zip file and extract the files to an installation folder.
3. Navigate to the `<download location>\oci_dwh_installer\` folder and run the `install_oci_dwh.bat` script.



With OnCommand Insight 7.3.10 and later, you must run the script for proper DWH/Reporting installation. Do not run the .MSI installation executable.

4. Enter the Db2 domain, or press Enter for local domain.
5. Enter the Db2 User name. See above for user name restrictions.
6. Enter the password for the Db2 user. Re-enter the password when prompted.
7. Enter the installation path for the Db2 component, or press Enter for default.
8. You are presented with the information you entered. Verify all settings carefully. Press Enter to start installation.
9. If prompted, allow Windows to proceed with the Db2 installation.
10. Following Db2 Installation, the DWH installation wizard will run. Follow its directions to install DWH with Reporting.

Data Warehouse with Reporting Installation may take up to an hour to complete.

Locating IBM Cognos documentation

For basic information such as how to start and stop the Reporting portal software, see the IBM Cognos documentation installed with the product. You can search with a web browser for information about any of the IBM Cognos reporting products, such as Query Studio, Report Studio, Business Insight, or Business Insight Advanced on the IBM website in the Information Centers for those software products.

Steps

1. To locate the IBM Cognos documentation installed with OnCommand Insight, navigate to this directory.

```
<install_dir>\cognos\c10_64\webcontent\documentation\help_docs.html
```

2. You can also display topics describing individual IBM Cognos windows used in the OnCommand Insight Reporting Portal. Click the ? icon on the window toolbar.

Verifying the Data Warehouse and Reporting installation

After a successful OnCommand Insight Data Warehouse installation, you should ensure that all of the DWH and Reporting services are available in your Microsoft Windows services.

Steps

1. From the Windows Start menu, select **Control Panel > System and Security > Administrative Tools > Services**.
2. Ensure that the following entries appear in the list of services:

Name / State	Description
SANScreen Server / Running	The OnCommand Insight DWH server
MySQL / Running	The OnCommand Insight SQL database
IBM Cognos / Running	IBM Cognos Content Database
DB2- DB2COPY1 - DB2-0 / Running	Manage Db2 databases
DB2 Governor (DB2COPY1) / Not running	Collects statistics for applications connected to Db2 databases.
DB2 License Server (DB2COPY1) / Not running	Monitors Db2 license compliance.
DB2 Management Service (DB2COPY1) / Running	Manages Db2 registry entries for compatibility with earlier Db2 copy versions.
DB2 Remote Command Server (DB2COPY1) / Running	Supports remote Db2 command execution.
IBM Secure Shell Server for Windows / Not running	IBM Secure Shell Server for Windows

Installing a Remote Acquisition Unit (RAU)

Install one or more RAUs in your OnCommand Insight environment.

Before you begin

You must have completed all of the installation prerequisites.

At least one port needs to be open and available between the RAU server and the OnCommand Insight Server

in order to forward change information to the server. If you are unsure about this, validate it by opening a Web browser on the RAU computer and directing it to the OnCommand Insight server:

```
https://< OnCommand Insight Server hostname >:< acquisition_port >
```

The acquisition port defaults to 443, but it may have been changed during the server installation. If the connection is successful, you see a OnCommand Insight response page indicating an open and available port between the RAU and the OnCommand Insight server.

Steps

1. Log in to the RAU server using an account with administrator privileges.
2. Open Windows Explorer and navigate to the directory where the RAU installation file is located.
3. Double-click .MSI file to start the installation.
4. Click **Next** to continue to the window that shows the License Agreement. Read this and accept the terms of the License Agreement and click **Next**.
5. Select to install the RAU on a local hard drive or the entire feature on a local hard drive. (You can check the Disk Usage link to ensure you have enough space - 116MB is required.) Click **Next**.
6. In the Configure window, set these parameters that are specific to your site:
 - **OnCommand Insight Server Name or Address** - hostname or IP address to identify the OnCommand Insight Server. The RAU uses this name/IP to open a communications link with the server. If you specify a hostname, make sure it can be resolved through DNS.
 - **Acquisition Unit Name** - unique name that identifies the RAU.
 - **OnCommand Insight Secured Remote Acquisition Port (HTTPS)** - Port used by Remote Acquisition Units to send environment change information to the OnCommand Insight server. This setting should match the value entered when installing the OnCommand Insight server and must be the same on all RAUs.
7. Review your selections. Click **Back** to go back and make changes. Click **Next**.
8. Click **Install** to start the installation.

Wait for the installation to complete. This should take approximately 5 to 10 minutes.

After you finish

When the installation is complete, a final window appears. Click the **Start Remote Acquisition Service** box to start the RAU, and click **Finish** to end this operation.

Verifying the remote acquisition unit service

After a successful remote acquisition unit (RAU) installation, the OnCommand Insight RAU service should be available in the Microsoft Windows services environment.

Steps

1. To verify that the RAU was added to the Windows services, open the Windows Start menu and select the **Control Panel > Administrative Tools > Services**.

2. Locate the **OnCommand Insight Acq - OnCommand Insight's Remote Acquisition Unit (RAU)** in the list.

Validating the remote acquisition unit installation

To validate proper installation of the Remote Acquisition Unit, you can view the status of the Remote Acquisition Units connected to your server.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Acquisition Units**.
3. Verify that the new Remote Acquisition Unit was registered correctly and that it has a Connected status.

If it does not, you must contact technical support.

Checking the installation

You can open Insight in a supported browser to check the installation. You might also want to check the Insight log files.

When you first open Insight, the license setup page opens. After you enter the license information, you must set up the data sources. See the *OnCommand Insight Configuration and Administration Guide* for information about entering data source definitions and setting up Insight users and notifications.

If you have experienced installation problems, contact technical support and provide the requested information.

Verifying new Insight services

After a successful installation, you should verify that the services for the Insight components are operating on your server.

Steps

1. To display a list of services that are currently operating:
 - a. Click the **Start** button.
 - b. Click **Run**.
 - c. Type the following:

```
cmd
```
 - d. Press Enter.
 - e. Type the following in the **Command Prompt** window:

```
net start
```

2. Check for these Insight services in the list:
 - **SANscreen Server**
 - **SANscreen Acq** (the acquisition process)

- **MySql** (Insight SQL database)
 - **Elasticsearch** (Data store for Insight data)
- If these services do not display in the list, contact technical support.

Insight logs

Insight supplies many log files to assist you with research and troubleshooting. The available logs are listed in the log directory. You might want to use a log monitoring tool, such as BareTail, to display all of the logs at one time.

The log files are located in the <install directory>\SANscreen\wildfly\standalone\log directory. Acquisition logs are located in the <install directory>\SANscreen\Acq\Log directory.

Accessing the web UI

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

Steps

1. Do one of the following:

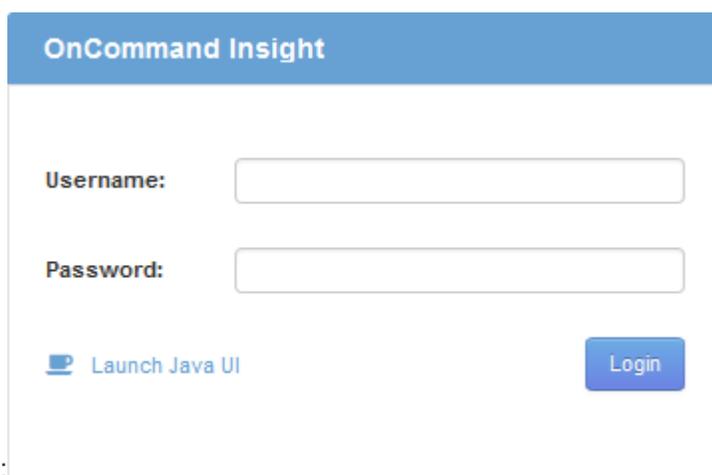
- Open Insight on the Insight server:

`https://fqdn`

- Open Insight from any other location:

`https://fqdn:port`

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



The OnCommand Insight dialog box displays:

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

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



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

Installing your Insight licenses

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

About this task

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

Steps

1. Open the license file in a text editor and copy the text.
2. Open Insight in your browser.
3. On the Insight toolbar, click **Admin**.
4. Click **Setup**.
5. Click the **Licenses** tab.
6. Click **Update License**.
7. Copy the license key text into the **License** text box.
8. Select the **Update (most common)** operation.
9. Click **Save**.
10. If you are using the Insight consumption licensing model, you must check the box to **Enable sending usage information to NetApp** in the **Send usage information** section. Proxy must be properly configured and enabled for your environment.

After you finish

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

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

OnCommand Insight licenses

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

- **Discover**

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

- **Assure**

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

- **Perform**

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

- **Plan**

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

- **Host Utilization pack**

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

- **Report Authoring**

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

OnCommand Insight modules are licensed for annual term or perpetual:

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

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

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

- **Discover**

Acquire and manage inventory (Foundation)

Monitor changes and manage inventory policies

- **Assure**

View and manage SAN path policies and violations

View and manage vulnerabilities

View and manage tasks and migrations

- **Plan**

View and manage requests

View and manage pending tasks

View and manage reservation violations

View and manage port balance violations

- **Perform**

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

View and manage performance policies and violations

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

Feature (admin)	With Perform license	Without Perform license
Application	Yes	No performance data or charts
Virtual machine	Yes	No performance data or charts
Hypervisor	Yes	No performance data or charts
Host	Yes	No performance data or charts
Datastore	Yes	No performance data or charts
VMDK	Yes	No performance data or charts
Internal volume	Yes	No performance data or charts
Volume	Yes	No performance data or charts
Storage pool	Yes	No performance data or charts
Disk	Yes	No performance data or charts
Storage	Yes	No performance data or charts
Storage node	Yes	No performance data or charts
Fabric	Yes	No performance data or charts
Switch port	Yes	No performance data or charts; “Port Errors” shows “N/A”
Storage port	Yes	Yes
NPV port	Yes	No performance data or charts
Switch	Yes	No performance data or charts

NPV switch	Yes	No performance data or charts
Qtrees	Yes	No performance data or charts
Quota	Yes	No performance data or charts
Path	Yes	No performance data or charts
Zone	Yes	No performance data or charts
Zone member	Yes	No performance data or charts
Generic device	Yes	No performance data or charts
Tape	Yes	No performance data or charts
Masking	Yes	No performance data or charts
ISCSI sessions	Yes	No performance data or charts
ICSI network portals	Yes	No performance data or charts
Search	Yes	Yes
Admin	Yes	Yes
Dashboard	Yes	Yes
Widgets	Yes	Partially available (only asset, query, and admin widgets are available)
Violations dashboard	Yes	Hidden
Assets dashboard	Yes	Partially available (storage IOPS and VM IOPS widgets are hidden)
Manage performance policies	Yes	Hidden
Manage annotations	Yes	Yes
Manage annotation rules	Yes	Yes
Manage applications	Yes	Yes

Queries	Yes	Yes
Manage business entities	Yes	Yes

Feature	User - with Perform license	Guest - with Perform license	User - without Perform license	Guest - without Perform license
Assets dashboard	Yes	Yes	Partially available (storage IOPS and VM IOPS widgets are hidden)	Partially available (storage IOPS and VM IOPS widgets are hidden)
Custom dashboard	View only (no create, edit, or save options)	View only (no create, edit, or save options)	View only (no create, edit, or save options)	View only (no create, edit, or save options)
Manage performance policies	Yes	Hidden	Hidden	Hidden
Manage annotations	Yes	Hidden	Yes	Hidden
Manage applications	Yes	Hidden	Yes	Hidden
Manage business entities	Yes	Hidden	Yes	Hidden
Queries	Yes	View and edit only (no save option)	Yes	View and edit only (no save option)

Troubleshooting installations

OnCommand Insight installations are generally managed through the installation wizards. However, customers might experience problems during upgrades or with conflicts due to computer environments.

You should also be certain that you install all of the necessary OnCommand Insight licenses for installing the software.

Missing licenses

Different licenses are required for different OnCommand Insight functionality. What you see displayed in OnCommand Insight is controlled by your installed licenses. Refer to the OnCommand Insight licenses section for information on functionality controlled by each license.

Refer to the OnCommand Insight licenses section for information on functionality controlled by each license.

Submitting an online technical support request

If you have problems with the Insight installation, as a registered support customer, you can submit an online technical support request.

Before you begin

Using your corporate email address, you must register as a support customer to obtain online support services. Registration is performed through the support site (<http://support.netapp.com>).

About this task

To assist customer support in solving the installation problem, you should gather as much information as possible, including these items:

- Insight serial number
- Description of the problem
- All Insight log files
- Screen capture of any error messages

Steps

1. Create a `.zip` file of the information you gathered to create a troubleshooting package.
2. Log in to the support site at mysupport.netapp.com and select **Technical Assistance**.
3. Click **Open a Case**.
4. Follow the instructions to your package of data.

After you finish

You can use **Check Case Status** on the Technical Assistance page to follow your request.

Upgrading OnCommand Insight

Normally, an upgrade must be performed on all of the Insight servers (Insight server, Data Warehouse server, Remote acquisition unit). You should always consult the Release Notes for the upgrade requirements for a new release of OnCommand Insight.

Unless otherwise indicated, the requirements and procedures apply to upgrading from Insight 7.x to the current version of Insight. If you are upgrading from a version prior to 7.0, contact your account representative.

Upgrading Insight to version 7.3.12 or later - Windows

Prior to upgrading from OnCommand Insight 7.3.10 - 7.3.11 to version 7.3.12 or later, you must run the OCI Data Migration Tool.

Background

OnCommand Insight versions 7.3.12 and later utilize underlying software that may be incompatible with previous versions. Insight versions 7.3.12 and later include a **Data Migration Tool** to assist with upgrading.



OnCommand Insight versions 7.3.9 and earlier are no longer supported. If you are running one of these versions, you *must* upgrade to Insight version 7.3.10 or later (7.3.11 is strongly recommended) prior to upgrading to 7.3.12 or later.

What Does The Data Migration Tool Do?

The migration tool performs an initial compatibility check and then follows one of three different upgrade paths. The path selected is based on the data compatibility of your current version.



Prior to upgrading, you must run the Data Migration Tool and follow the recommended steps.

Before you Begin

- It is strongly recommended to back up your OnCommand Insight system prior to running the Data Migration Tool.
- The Elasticsearch service on the server needs to be up and running.
- The Data Migration Tool *must* be run for the database and any performance archives before you upgrade Insight.

Running the Data Migration Tool

1. Download the latest version of the Data Migration Tool (for example, *SANScreenDataMigrationTool-x86-7.3.12-97.zip*) to your Insight server, as well as the appropriate Insight installer file. Unzip into a working folder. Downloads can be found on the [NetApp Support Site](#).
2. Open a command window and navigate to your working folder.
 - Open Powershell as Administrator.
3. Run the data migration tool using the following command:
 - `.\SANScreenDataMigrationTool.ps1``
4. Follow the instructions as needed. The following is an example.

```
.\SANSscreenDataMigrationTool.ps1

NetApp SANSscreen Data Migration Tool 7.3.12-121

Checking OnCommand Insight Installation...
OnCommand Insight 7.3.10 (139) is installed

Getting installation parameters...
Installation Directory: C:\Program Files\SANSscreen\
Elasticsearch Rest Port: 9200

Checking Elasticsearch service...
Elasticsearch service is up

Checking for obsolete (version 5) indexes...
Found 54 obsolete indexes. Of these,
    54 indexes may be migrated with OCI server running,
    the most recent of which is for 2021-05-13

Verifying migration component is present...
SANSscreen Server service is Running

Proceed with online migration of 54 indexes (y or [n])?:
```

The Data Migration Tool will check for the presence of obsolete indexes on your system and report if any are found. If none are present the tool will exit.

Some indexes may be migrated while the SANSscreen Server service is running. Others may only be migrated when the server is stopped. If there are no indexes that may be migrated the tool will exit. Otherwise follow the instructions as prompted.

After the Data Migration Tool completes it will recheck for obsolete indexes. If all indexes have been migrated, the tool will inform you that upgrade to OnCommand Insight 7.3.12 is supported. You can now proceed with upgrading Insight.

```
.\SANSscreenDataMigrationTool.ps1

NetApp SANSscreen Data Migration Tool 7.3.12-127

Checking OnCommand Insight Installation...
OnCommand Insight 7.3.10 (139) is installed

Getting installation parameters...
Installation Directory: D:\SANSscreen\
Elasticsearch Rest Port: 9200

Checking Elasticsearch service...
Elasticsearch service is up

Checking for obsolete (version 5) indexes...
Found 5 obsolete indexes. Of these,
    5 indexes need to be migrated with OCI server stopped

Verifying migration component is present...
SANSscreen Server service is Stopped

Proceed with offline migration of 5 indexes (y or [n])?: y
Preparing to perform migration...
Preparing to migrate ociint-inventory-snmp_win2012_host: copied; backup;
delete old; restore new; cleanup; done.
Preparing to migrate ociint-inventory-snmp_win2012_interface: copied;
backup; delete old; restore new; cleanup; done.
Preparing to migrate ociint-inventory-snmp_win2012_load_average: copied;
backup; delete old; restore new; cleanup; done.
Preparing to migrate ociint-inventory-snmp_win2012_storage: copied;
backup; delete old; restore new; cleanup; done.
Preparing to migrate ociint-inventory-snmp_win2012_tcp_connection: copied;
backup; delete old; restore new; cleanup; done.
Execution time 0:00:15

Checking for obsolete (version 5) indexes...
No obsolete indexes found. Upgrade to 7.3.12+ is supported.

C:\Users\root\Desktop\SANSscreenDataMigrationTool-x64-7.3.12-127>
```

If you were prompted to stop the SANSscreen service, restart it before upgrading Insight.

Validation failures

In the event that index validation fails, the migration tool will inform you of the problem before quitting.

OnCommand Insight is not present:

```
.\SANSscreenDataMigrationTool.ps1

NetApp SANSscreen Data Migration Tool V1.0

Checking OnCommand Insight Installation...
ERROR: OnCommand Insight is not installed
```

Invalid Insight version:

```
.\SANSscreenDataMigrationTool.ps1

NetApp SANSscreen Data Migration Tool 7.3.12-105

Checking OnCommand Insight Installation...
OnCommand Insight 7.3.4 (126) is installed
ERROR: The OCI Data Migration Tool is intended to be run against OCI 7.3.5
- 7.3.11
```

Elasticsearch service is not running:

```
.\SANSscreenDataMigrationTool.ps1

NetApp SANSscreen Data Migration Tool 7.3.12-105

Checking OnCommand Insight Installation...
OnCommand Insight 7.3.11 (126) is installed

Getting installation parameters...
Installation Directory: C:\Program Files\SANSscreen\
Elasticsearch Rest Port: 9200

Checking Elasticsearch service...
ERROR: The Elasticsearch service is not running

Please start the service and wait for initialization to complete
Then rerun OCI Data Migration Tool
```

Command-line options

The Data Migration Tool includes some optional parameters that affect its operation.

Option (Windows)	Function
------------------	----------

-s	Suppress all prompts
-perf_archive	<p>If specified, existing archive entries for any date whose index(es) are migrated will be replaced. The path should point to the directory containing the archive entry zip files.</p> <p>An argument of '-' may be specified to indicate there is no performance archive to be updated.</p> <p>If this argument is present, the prompt for the archive location will be suppressed.</p>
-check	If present, the script will exit immediately after reporting the index counts.
-dryrun	If present, then the migration executable will report the actions that would be taken (to migrate data and update archive entries) but will not perform the operations.

Overview of the OnCommand Insight upgrade process

Before you begin upgrading Insight, it is important to understand the upgrade process. The upgrade process is the same for most versions of Insight.

The upgrade process for Insight includes the following high-level tasks:

- Downloading the installation packages
- Backing up the Data Warehouse database

To avoid the possibility of misreporting data, you must back up the Data Warehouse database before you back up the Insight database.

- Backing up the Insight database

The Insight database is automatically backed up when you perform the in-place upgrade. It is a best practice to back up the database before the upgrade, and place the backup in a location other than on the Insight server. During the upgrade process, Insight does not collect new data. To minimize the amount of data that is not collected, you must start the database backup within an hour or two of your planned upgrade time.

- Back up the Data Warehouse and Remote Acquisition Unit security configuration if the configuration has been changed from the default configuration.

The non-default security configuration must be restored to the Data Warehouse and RAU server after the upgrade is complete and before the Data Warehouse database is restored to the system.

- Backing up any custom Data Warehouse reports

When you back up the Data Warehouse database, custom reports are included. The backup file is created

on the Data Warehouse server. It is a recommended best practice to back up the custom reports to a location other than the Data Warehouse server.

- Uninstalling the Data Warehouse and the Remote Acquisition Unit software, if applicable

The Insight server has an in-place upgrade; you do not have to uninstall the software. The in-place upgrade backs up the database, uninstalls the software, installs the new version, and then restores the database.

- Upgrading the software on the Insight server, Data Warehouse, and Remote Acquisition Unit(s)

All previously applied licenses remain in the registry; you do not have to reapply these licenses.

- Completing the post-upgrade tasks

OnCommand Insight upgrade checklist

You can use the provided checklists to record your progress as you prepare for the upgrade. These tasks are intended to help mitigate the risk for upgrade failures and to expedite recovery and restoration efforts.

Checklist for preparing for the upgrade (required)

Condition	Complete?
Ensure that you have Windows local administrator permissions, which are required to perform the upgrade process, on all Insight servers.	
If your Insight, Data Warehouse, or Remote Acquisition Unit servers reside on 32-bit platforms, you must upgrade your servers to 64-bit platforms. As of Insight 7.x, upgrades are only available for 64-bit platforms.	
<p>Ensure that you have the necessary permissions to modify or disable the antivirus software on all the servers in your environment. To prevent an upgrade failure due to active virus scan software, you must exclude the Insight installation directory (<code>disk drive:\install directory\sanscreen</code>) from access to antivirus scanning during the upgrade. After you upgrade all of the components, you can safely reactivate the antivirus software; however, ensure that you configure the scan to still exclude everything in the Insight installation directory.</p> <p>Additionally, you must also exclude the IBM/Db2 folder (for example <code>C:\Program Files\IBM\DB2</code>) from anti-virus scanning following installation.</p>	

Checklist for preparing for the upgrade (best practice)

Condition	Complete?
<p>Plan when you are going to upgrade, taking into consideration that most upgrades take a minimum of 4 to 8 hours; larger enterprises will take longer. Upgrade times might vary depending on your available resources (architecture, CPU, and memory), the size of your databases, and the number of objects monitored in your environment.</p>	
<p>Contact your account representative about your upgrade plans and provide the version of Insight you have installed and what version you would like to upgrade to.</p>	
<p>Ensure that your current resources allocated to the Insight, Data Warehouse, and Remote Acquisition Unit(s) still meet recommended specifications. See the recommend sizing guidelines for all servers. Alternatively, you can contact your account representative to discuss sizing guidelines.</p>	
<p>Ensure that you have enough disk space for the database backup and restore process. The backup and restore processes require approximately five times the disk space used by the backup file on the Insight and Data Warehouse servers. For example, a 50 GB backup requires 250 to 300 GB of free disk space.</p>	
<p>Ensure that you have access to Firefox® or the Chrome™ browser when you back up the Insight and Data Warehouse databases. Internet Explorer is not recommended, because it experiences some issues when uploading and downloading files larger than 4 GB.</p>	
<p>Delete the .tmp files on the Insight server, which you can find in the following location: <install directory>\SANscreen\wildfly\standalone\tmp.</p>	
<p>Remove duplicate data sources and decommissioned data sources from the Insight client. Removing decommissioned or duplicate data sources decreases the amount of time required to perform the upgrade and mitigates the opportunity for data corruption.</p>	

<p>If you have modified any of the default reports shipped with Insight, you should save the reports with a different name and then save them to the Customer Reports folder so that you do not lose your modified report when you upgrade or restore the system.</p>	
<p>If you have any custom or modified Data Warehouse reports created by you or professional services, create a backup of them by exporting them to XML and then moving them to the Customer Reports folder. Ensure that the backup is not located on the Data Warehouse server. If you do not move your reports to the recommended folders, these reports might not be backed up by the upgrade process. For earlier versions of Insight, failure to locate reports in the appropriate folders may result in the loss of custom and modified reports.</p>	
<p>Record all settings in the IBM Cognos Configuration utility, because these are not included in the Data Warehouse backup; you have to reconfigure these settings after the upgrade. The utility is located in the disk drive:\install directory\SANscreen\cognos\c10_64\bin64 directory on the Data Warehouse server and you run it using the cogconfigw command. Alternatively, you can perform a complete backup of Cognos and then import all of your settings. Refer to the IBM Cognos documentation for more information.</p>	

Checklist for preparing for the upgrade (if applicable)

Condition	Complete?
<p>If you have replaced the self-signed certificates that the Insight installation created due to browser security warnings with certificates signed by your internal certificate authority, back up your keystore file, which is in the following location: disk drive:\install directory\SANscreen\wildfly\standalone\configuration and restore it after the upgrade. This replaces the self-signed certificates that Insight creates with your signed certificates.</p>	

<p>If any of your data sources were modified for your environment and you are unsure if these modifications are available in the Insight version to which you are upgrading, make a copy of the following directory, which will help you troubleshoot if there are recovery issues: <code>disk drive:\install directory\SANscreen\wildfly\standalone\deployments\datasources.war</code>.</p>	
<p>Back up all custom database tables and views using the <code>mysqldump</code> command line tool. Restoring custom database tables requires privileged database access. Contact technical support for assistance with restoring these tables.</p>	
<p>Ensure that no custom integration scripts, third-party components required for Insight data sources, backups, or any other required data is stored in the <code>disk drive:\install directory\sanscreen</code> directory, because the contents of this directory is deleted by the upgrade process. Ensure that you move any of these things from the <code>\sanscreen</code> directory to another location. For example, if your environment contains custom integration scripts, ensure that you copy the following file to a directory other than the <code>\sanscreen</code> directory:</p> <pre>\install_dir\SANscreen\wildfly\standalone\deployments\datasources.war\new_disk_models.txt.</pre>	

Downloading the OnCommand Insight installation packages

You should download the installation packages for Insight, Data Warehouse, and the Remote Acquisition Unit (if applicable) prior to the day that you choose to upgrade. Download times for the packages (.msi files) vary based on your available bandwidth.

About this task

You can download the installation packages using the Insight webUI or by navigating to the appropriate OnCommand Insight link from <http://support.netapp.com/NOW/cgi-bin/software>.

To download the installation package from within the Insight server, do the following:

Steps

1. Open the Insight web UI by opening a web browser and entering one of the following:
 - On the Insight server: `https://localhost`
 - From any location: `https://IP Address:port` or `fqdn:port`

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

2. Log in to Insight.
3. Click on the Help icon and select **Check for updates**.
4. If a newer version is detected, follow the instructions in the message box.

You will be taken to the InsightDescription page for the newer version.

5. On the **Description** page, click **Continue**.
6. When the end-user license agreement (EULA) is displayed, click **Accept**.
7. Click the installation package link for each component (Insight server, Data Warehouse, Remote Acquisition Unit), etc.) and click **Save as** to save the installation package.

Before you upgrade, you should ensure that you copy the Data Warehouse and Remote Acquisition Unit installation packages to disks that are local to their respective servers.

8. Click **CHECKSUM**, and make a note of the numerical values that are associated with each installation package.
9. Verify that the installation packages are complete and without error after you download them.

Incomplete file transfers can cause issues with the upgrade process.

To generate MD5 hash values for the installation packages, you can use a third-party utility like Microsoft's [File Checksum Integrity Verifier](#) utility.

Backing up the databases

Before you upgrade, you should back up both the Data Warehouse and OnCommand Insight databases. Upgrading requires a backup of the Data Warehouse database so that you can restore the database later in the upgrade process. The in-place upgrade for Insight backs up the database; however, you should back up the database before the upgrade as a best practice.

To avoid misreporting data, you should back up the Data Warehouse database prior to backing up the Insight database. Additionally, if you have a test environment, it is recommended that you ensure you can restore the backup before you continue with the upgrade.

Backing up the Data Warehouse database

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

Steps

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

3. Click **Backup** and select your backup configuration:

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

This operation can take 30 minutes or more.

+

Data Warehouse creates a backup file and displays its name.

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

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

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

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

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

dwh_sa_staging

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

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

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

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

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



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

Backing up the OnCommand Insight database

Back up the Insight database to ensure that you have a recent backup if an issue occurs after the upgrade. During the backup and restore phase, performance data will not be collected; thus, the backup should occur as close as possible to the upgrade time.

Steps

1. Open Insight in your browser.
2. Click **Admin > Troubleshooting**.
3. On the **Troubleshooting** page, click **Backup**.

The time to back up the database might vary depending on your available resources (architecture, CPU, and memory), the size of your database, and the number of objects monitored in your environment.

When the backup is complete, you are asked if you want to download the file.

4. Download the backup file.

Backing up the security configuration

When your Insight components are using a non-default security configuration, you must back up the security configuration and then restore the configuration on all components after the new software is installed. The security configuration must be restored before the Data Warehouse database backup is restored.

About this task

You use the `securityadmin` tool to create a backup of the configuration and to restore the saved configuration. For more information, search for `securityadmin` in the OnCommand Insight Documentation Center: <http://docs.netapp.com/oci-73/index.jsp>

Backing up custom Data Warehouse reports

If you created custom reports and you do not have the `.xml` source files for them, then you should back up these reports before the upgrade. You should then copy them to a server other than the Data Warehouse server.

Steps

1. Log in to the Data Warehouse portal at `https://fqdn/dwh`.
2. On the Data Warehouse toolbar, click  to open the Reporting Portal and log in.
3. Select **File > Open**.
4. Select the folder that the report is located in, select the report, and then click **Open**.
5. Select **Tools > Copy report to clipboard**.
6. Open a text editor, paste the contents of the report, and save the file as `report_name.txt`, where `report_name` is the name of the report.
7. Store the reports on a server other than the Data Warehouse server.

Performing the software upgrade

After you complete all prerequisite tasks, you can upgrade all of the Insight components to a new release by downloading and running the applicable installation package on each server.

Upgrading Insight

After you complete all prerequisite tasks, you log in to the Insight server and run the installation package to complete the upgrade. The upgrade process uninstalls the existing software, installs the new software, and then reboots the server.

Before you begin

The Insight installation package must be located on the server.

Steps

1. Log in to the Insight server using an account that has Windows local administrator permissions.
2. Locate the Insight installation package (`SANscreenServer-x64-version_number-build_number.msi`) using Windows Explorer and double-click it.

The OnCommand InsightSetup wizard displays.

3. Move the progress window away from the center of the screen and away from the **Setup** wizard window so

that any generated errors are not hidden from view.

4. Follow the setup wizard prompts.

It is a best practice to leave all the defaults selected.

After you finish

To verify if the upgrade is successful or if errors are generated, check the upgrade log in the following location:
<install directory>\SANSscreen\wildfly\standalone\log.

Upgrading Data Warehouse

After you complete all prerequisite tasks, you can log in to the Data Warehouse server and run the installation package to complete the upgrade.

About this task

Inline upgrade is not supported by the Data Warehouse (DWH). Use the following steps to upgrade to the new version of DWH software.

When upgrading DWH, the folder containing the *securityadmin* tool vault backup is deleted. It is highly recommended to back up the vault prior to upgrading DWH. For reference, the default vault folders are as follows:



- Vault folder (vaults in use): %SANSSCREEN_HOME%\wildfly\standalone\configuration\vault
- Vault backups: %SANSSCREEN_HOME%\backup\vault

See [Managing security on the Data Warehouse](#) for more information.

Steps

1. Log in to the DWH server using an account that has Windows local administrator permissions.
2. Back up the DWH DB and Reports using the DWH portal interface.
3. Back up the security configuration if the server is using a non-default security configuration.
4. Uninstall the DWH software from the server.
5. Reboot the server to remove components from memory.
6. Install the new version of DWH on the server.

The installation takes approximately 2 hours. It is a best practice to leave all the defaults selected.

7. Restore the non-default security configuration to the DWH server.
8. Restore the DWH database to the server.

After you finish

After you upgrade, you must restore the Data Warehouse database, which can take as long or longer than the upgrade.



During an OnCommand Insight upgrade, it is not uncommon for a customer to switch to a different Insight server. If you have changed your Insight server, after you restore the data warehouse database the existing connectors will point to the previous server IP address or hostname. It is a best practice to delete the connector and create a new one, to avoid possible errors.

Preserving custom Cognos settings during a Data Warehouse upgrade

Custom Cognos settings, such as non-default SMTP email settings, are not automatically backed up as part of a Data Warehouse upgrade. You need to manually document and then restore the custom settings following an upgrade.

Prior to upgrading Data Warehouse, prepare a checklist with any custom Cognos settings that you want to preserve, and review the list prior to upgrading the system. After the upgrade is complete, you can restore the values manually to return them to the settings in the original configuration.

Backing up the security configuration

When your Insight environment is using a non-default security configuration, you must back up the security configuration and then restore the security configuration after the new software is installed. The security configuration must be restored before the Data Warehouse database backup is restored.

About this task

You use the `securityadmin` tool to create a backup of the configuration and to restore the saved configuration. For more information, search for `securityadmin` in the OnCommand Insight Documentation Center: <http://docs.netapp.com/oci-73/index.jsp>

Upgrading remote acquisition unit servers

After you complete all prerequisite tasks, you can log in to the remote acquisition unit server and run the installation package to complete the upgrade. You must perform this task on all remote acquisition servers in your environment.

Before you begin

- You must have upgraded OnCommand Insight.
- The OnCommand Insight installation package must be located on the server.

Steps

1. Log in to the remote acquisition unit server using an account that has Windows local administrator permissions.
2. Locate the Insight installation package (`RAU-x64-version_number-build_number.msi`) using Windows Explorer and double-click it.

The OnCommand Insight Setup Wizard displays.

3. Move the installation wizard progress window away from the center of the screen and away from the installation wizard window so that any generated errors are not hidden from view.

4. Follow the Setup Wizard prompts.

It is a best practice to leave all the defaults selected.

After you finish

- To verify if the upgrade is successful or if errors are generated, check the upgrade log in the following location: `<install directory>\SANscreen\bin\log`.
- Use the `securityadmin` tool to restore the saved security configuration. For more information, search for `securityadmin` in the OnCommand Insight Documentation Center: <http://docs.netapp.com/oci-73/index.jsp>
- Clear your browser's cache and history to ensure that you are receiving the latest data from the server.

Completing post-upgrade tasks

After you upgrade to the latest version of Insight, you must complete additional tasks.

Installing data source patches

If applicable, you should install the latest patches available for your data sources to take advantage of the latest features and enhancements. 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 contacted technical support and obtained the `.zip` file that contains the latest data source patches by providing them with the version you are upgrading from and the version you want to upgrade to.

Steps

1. Place the patch file on the Insight server.
2. On the Insight toolbar, click **Admin**.
3. Click **Patches**.
4. From the Actions button, select **Apply patch**.
5. In the **Apply data source patch** dialog box, click **Browse** to locate the uploaded patch file.
6. Review the **Patch name**, **Description**, and **Impacted data source types**.
7. If the selected patch is correct, click **Apply Patch**.

All data sources of the same type are updated with this patch. Insight automatically forces acquisition to restart when you add a data source. Discovery includes the detection of changes in network topology including the addition or deletion of nodes or interfaces.

8. To force the discovery process manually, click **Data Sources** and click **Poll Again** next to the data source to force it to collect data immediately.

If the data source is already in an acquisition process, Insight ignores the poll again request.

Replacing a certificate after upgrading OnCommand Insight

Opening the OnCommand Insight web UI after an upgrade results in a certification warning. The warning message is displayed because a valid self-signed certificate is not available after the upgrade. To prevent the warning message from being displayed in the future, you can install a valid self-signed certificate to replace the original certificate.

Before you begin

Your system must satisfy the minimum encryption bit level (1024 bits).

About this task

The certification warning does not impact the usability of the system. At the message prompt, you can indicate that you understand the risk, and then proceed to use Insight.

Steps

1. List the contents of the keystore: `C:\Program Files\SANscreen\java64\bin>keytool.exe -list -v -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore"`

When prompted for a password, enter `changeit`.

There should be at least one certificate in the keystore, `ssl certificate`.

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

4. Generate a certificate request file: `keytool -certreq -alias localhost -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -file c:\localhost.csr`

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

5. Submit the `c:\localhost.csr` file to your certification 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.

6. Import the approved certificate: `keytool -importcert -alias localhost -file c:\localhost2.DER -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore"`

- a. When prompted for a password, enter the keystore password.

The system displays the following message: Certificate reply was installed in keystore

7. Restart the SANscreen Server service.

Results

The web browser no longer reports certificate warnings.

Increasing Cognos memory

Before you restore the Data Warehouse database, you should increase the Java allocation for Cognos from 768 MB to 2048 MB to decrease report generation time.

Steps

1. Open a command prompt window as administrator on the Data Warehouse server.
2. Navigate to the disk drive: `\install directory\SANscreen\cognos\c10_64\bin64` directory.
3. Type the following command: `cogconfigw`

The IBM Cognos Configuration window displays.



The IBM Cognos Configuration shortcut application points to disk drive: `\Program Files\SANscreen\cognos\c10_64\bin64\cognosconfigw.bat`. If Insight is installed in the Program Files (space between) directory, which is the default, instead of ProgramFiles (no space), the `.bat` file will not work. If this occurs, right click the application shortcut and change `cognosconfigw.bat` to `cognosconfig.exe` to fix the shortcut.

4. From the navigation pane on the left, expand **Environment**, expand **IBM Cognos services**, and then click **IBM Cognos**.
5. Select **Maximum memory for Tomcat in MB** and change 768 MB to 2048 MB.
6. On the IBM Cognos Configuration toolbar, click (Save).

An informational message displays to inform you of the tasks Cognos is performing.

7. Click **Close**.
8. On the IBM Cognos Configuration toolbar, click  (Stop).
9. On the IBM Cognos Configuration toolbar, click  (Start).

Restoring the Data Warehouse database

When you back up the Data Warehouse database, Data Warehouse creates a `.zip` file that you can later use to restore that same database.

About this task

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

Steps

1. Log in to the Data Warehouse Portal at `https://fqdn/dwh`.
2. From the navigation pane on the left, click **Backup/Restore**.
3. In the **Restore Database and Reports** section, click **Browse** and locate the `.zip` file that holds the Data Warehouse backup.
4. It is a best practice to leave both of the following options selected:

- **Restore database**

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

- **Restore reports**

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

5. Click **Restore**.

Do not navigate away from the restore status. If you do, the restore status is no longer displays and you receive no indication when the restore operation is complete.

6. To check the upgrade process, view the `dwh_upgrade.log` file, which is in the following location:
`<install directory>\SANscreen\wildfly\standalone\log.`

After the restore process finishes, a message appears just below the **Restore** button. If the restore process is successful, the message indicates success. If the restore process fails, the message indicates the specific exception that occurred to cause the failure. In this case, contact technical support and provide them with `dwh_upgrade.log` file. If an exception occurs and the restore operation fails, the original database is automatically reset.



If the restore operation fails with a “Failed upgrading cognos content store” message, restore the Data Warehouse database without its reports (database only) and use your XML report backups to import your reports.

Restoring custom Data Warehouse reports

If applicable, you can manually restore any custom reports you backed up before the upgrade; however, you only need to do this if you lose reports or if they have become corrupted.

Steps

1. Open your report with a text editor, and then select and copy its contents.
2. Log in to the Reporting portal at <https://fqdn/reporting>.
3. On the Data Warehouse toolbar, click  to open the Insight Reporting portal.
4. From the Launch menu, select **Report Studio**.
5. Select any package.

Report Studio displays.

6. Click **Create new**.
7. Select **List**.
8. From the Tools menu, select **Open Report from Clipboard**.

The **Open Report from Clipboard** dialog box displays.

9. From the File menu, select **Save As** and save the report to the Custom Reports folder.
10. Open the report to verify that it was imported.

Repeat this task for each report.



You may see an “Expression parsing error” when you load a report. This means that the query contains a reference to at least one object that does not exist, which means there is no package selected in the Source window to validate the report against. In this case, right-click on a data mart dimension in the Source window, select Report Package, and then select the package associated with the report (for example, the inventory package if it is an inventory report or one of the performance packages if it’s a performance report) so Report Studio can validate it and then you can save it.

Verifying that Data Warehouse has historical data

After restoring your custom reports, you should verify that Data Warehouse is collecting historical data by viewing your custom reports.

Steps

1. Log in to the Data Warehouse portal at <https://fqdn/dwh>.
2. On the Data Warehouse toolbar, click  to open the Insight Reporting portal and log in.
3. Open the folder that contains your custom reports (for example, Custom Reports).
4. Click  to open the output format options for this report.
5. Select the options you want and click **Run** to ensure that they are populated with storage, compute, and switch historical data.

Restoring the performance archive

For systems that perform performance archiving, the upgrade process only restores seven days of archive data. You can restore the remaining archive data after the upgrade is completed.

About this task

To restore the performance archive, follow these steps.

Steps

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

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

Testing the connectors

After you upgrade, you want to test the connectors to ensure that you have a connection from the OnCommand Insight Data Warehouse to the OnCommand Insight server.

Steps

1. Log in to the Data Warehouse Portal at `https://fqdn/dwh`.
2. From the navigation pane on the left, click **Connectors**.
3. Select the first connector.

The Edit Connector page displays.

4. Click **Test**.
5. If the test is successful, click **Close**; if it fails, enter the name of the Insight server in the **Name** field and its IP address in the **Host** field and click **Test**.
6. When there is a successful connection between the Data Warehouse and the Insight server, click **Save**.

If it does not succeed, check the connection configuration and ensure the Insight server does not have any issues.

7. Click **Test**.

Data Warehouse tests the connection.

Verifying the Extract, Transform, and Load scheduling

After you upgrade, you should ensure that the Extract, Transform, and Load (ETL) process is retrieving data from the OnCommand Insight databases, transforming the data, and saving it into the data marts.

Steps

1. Log in to the Data Warehouse portal at `https://fqdn/dwh`.
2. From the navigation pane on the left, click **Schedule**.
3. Click **Edit schedule**.
4. Select **Daily** or **Weekly** from the **Type** list.

It is recommended to schedule ETL to run once a day.

5. Verify that the time selected is the time at which you want the job to run.

This ensures that the build job runs automatically.

6. Click **Save**.

Updating disk models

After upgrading, you should have any updated disk models; however, if for some reason Insight failed to discover new disk models, you can manually update them.

Before you begin

You must have obtained from technical support the `.zip` file that contains the latest data source patches.

Steps

1. Stop the SANscreen Acq service.
2. Navigate to the following directory: `<install directory>\SANscreen\wildfly\standalone\deployments\datasources.war`.
3. Move the current `diskmodels.jar` file to a different location.
4. Copy the new `diskmodels.jar` file into the `datasources.war` directory.
5. Start the SANscreen Acq service.

Verifying that business intelligence tools are running

If applicable, you should verify that your business intelligence tools are running and retrieving data after the upgrade.

Verify that business intelligence tools like BMC Atrium and ServiceNow are running and able to retrieve data. This includes the BMC connector and solutions that leverage REST.

Troubleshooting an upgrade

If you encounter issues after an OnCommand Insight upgrade, you might find it helpful to review the troubleshooting information related to some possible issues.

Unable to start Cognos from the Windows Start menu

The existence of a space before `\SANscreen\cognos` in the path name is an issue. See the following in the

NetApp Customer Success Community for more information: <https://forums.netapp.com/thread/62721>.

“Not a valid win32 application” error message

This is an issue with Microsoft Windows. To resolve this issue, you must put quotation marks around the image path in the registry. See the following documentation for more information: <https://support.microsoft.com/en-us/kb/812486/en-us>.

Annotations are not present

When a Data Warehouse ETL job queries for annotations from an Insight instance, it sometimes receives an empty response (a 0 result) in error. This error results in annotations for certain objects moving back and forth between a “present” and “not present” state in Data Warehouse. See the following for more information: <https://forums.netapp.com/docs/DOC-44167>

Differences in values displayed in reports

Prior to 7.0, reports were integer-based. They are now decimal-based; therefore, after you upgrade, you may notice a increase or decrease in how the values display.

Data does not display in reports

In 7.0.1, several model names were changed (for example, Symmetrix was changed to Symmetrix VMAX). As a result, if a report contains a filter for “Symmetrix”, you will not see any data when you run the report. To change the report, you must open the report with Query Explorer in Report Studio, search for the model name, replace it with the new model name, and save the report.

Uninstalling the software

You must uninstall the old versions of the Data Warehouse and Remote Acquisition software to install the new versions. You should do this before you attempt to upgrade any of these components. The software on the Insight server is uninstalled during the in-place upgrade.

Uninstalling the OnCommand Insight Server

You can uninstall the OnCommand Insight server if needed.

Before you begin

Best practice: before uninstalling Insight, back up the OnCommand Insight database.

Steps

1. Log in to the OnCommand Insight server using an account with administrator privileges.
2. Ensure that all of the Insight windows on the server are closed.
3. Open the **Uninstall a Program** feature from the control panel and select the OnCommand Insight application for removal.
4. Click **Uninstall** and follow the prompts.

Uninstalling the Data Warehouse software

You must uninstall the Data Warehouse software before you can upgrade.

Before you begin

If you made changes to reports you want to keep, it is critical that you create a backup before you uninstall Data Warehouse. Uninstalling Data Warehouse permanently deletes all previously collected data and removes all reports, including any newly created or edited reports.

Steps

1. Log in to the Data Warehouse server.
2. Ensure that all of the Insight windows on the server are closed.
3. To uninstall using Control Panel:
 - a. Open **Uninstall a Program** from the control panel and select the OnCommand Insight application for removal. Click **Uninstall** and follow the prompts.
 - b. Select the IBM Db2 application for removal. Click **Uninstall** and follow the prompts.
 - c. Delete the Db2 install folder (for example *C:\Program Files\IBM\DB2*) to completely remove the Db2 database.
4. To uninstall using the provided script:
 - a. Navigate to the *<download location>\oci_dwh_uninstall* folder and run the *uninstall_oci_dwh.bat* script.
5. Reboot the server.

Uninstalling the remote acquisition unit software

You must uninstall the existing version of the remote acquisition unit software before you can upgrade to a new version. You should perform this task on all remote acquisition unit servers in your environment.

Steps

1. Log in to the remote acquisition unit server.
2. Ensure that all of the OnCommand Insight windows on the server are closed.
3. Open the **Uninstall a Program** feature from the control panel and select the OnCommand Insight Remote Acquisition Unit program for removal.
4. Click **Uninstall** and follow the prompts.

Configuration and administration

Setting up Insight

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

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

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

Accessing the web UI

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

Steps

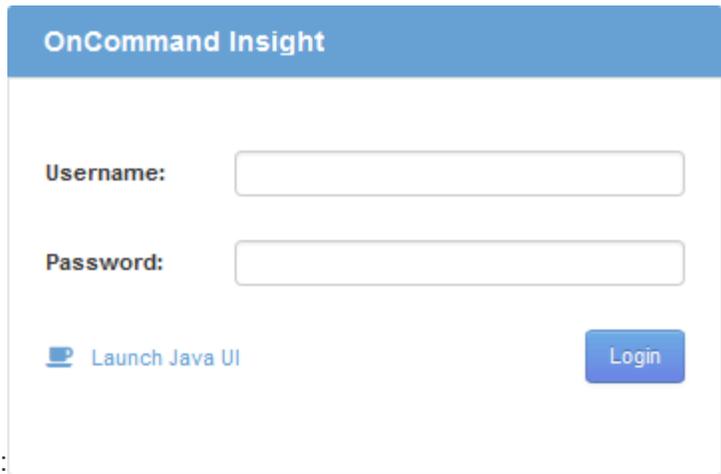
1. Do one of the following:
 - Open Insight on the Insight server:

```
https://fqdn
```

- Open Insight from any other location:

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

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



The OnCommand Insight dialog box displays:

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

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



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

Installing your Insight licenses

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

About this task

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

Steps

1. Open the license file in a text editor and copy the text.
2. Open Insight in your browser.
3. On the Insight toolbar, click **Admin**.
4. Click **Setup**.
5. Click the **Licenses** tab.
6. Click **Update License**.
7. Copy the license key text into the **License** text box.
8. Select the **Update (most common)** operation.
9. Click **Save**.
10. If you are using the Insight consumption licensing model, you must check the box to **Enable sending usage information to NetApp** in the **Send usage information** section. Proxy must be properly configured and enabled for your environment.

After you finish

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

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

OnCommand Insight licenses

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

• Discover

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

• Assure

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

• Perform

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

• Plan

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

• Host Utilization pack

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

• Report Authoring

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

OnCommand Insight modules are licensed for annual term or perpetual:

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

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

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

• Discover

Acquire and manage inventory (Foundation)

Monitor changes and manage inventory policies

- **Assure**

View and manage SAN path policies and violations

View and manage vulnerabilities

View and manage tasks and migrations

- **Plan**

View and manage requests

View and manage pending tasks

View and manage reservation violations

View and manage port balance violations

- **Perform**

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

View and manage performance policies and violations

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

Feature (admin)	With Perform license	Without Perform license
Application	Yes	No performance data or charts
Virtual machine	Yes	No performance data or charts
Hypervisor	Yes	No performance data or charts
Host	Yes	No performance data or charts
Datastore	Yes	No performance data or charts
VMDK	Yes	No performance data or charts
Internal volume	Yes	No performance data or charts
Volume	Yes	No performance data or charts
Storage pool	Yes	No performance data or charts
Disk	Yes	No performance data or charts

Storage	Yes	No performance data or charts
Storage node	Yes	No performance data or charts
Fabric	Yes	No performance data or charts
Switch port	Yes	No performance data or charts; "Port Errors" shows "N/A"
Storage port	Yes	Yes
NPV port	Yes	No performance data or charts
Switch	Yes	No performance data or charts
NPV switch	Yes	No performance data or charts
Qtrees	Yes	No performance data or charts
Quota	Yes	No performance data or charts
Path	Yes	No performance data or charts
Zone	Yes	No performance data or charts
Zone member	Yes	No performance data or charts
Generic device	Yes	No performance data or charts
Tape	Yes	No performance data or charts
Masking	Yes	No performance data or charts
ISCSI sessions	Yes	No performance data or charts
ICSI network portals	Yes	No performance data or charts
Search	Yes	Yes
Admin	Yes	Yes
Dashboard	Yes	Yes

Widgets	Yes	Partially available (only asset, query, and admin widgets are available)
Violations dashboard	Yes	Hidden
Assets dashboard	Yes	Partially available (storage IOPS and VM IOPS widgets are hidden)
Manage performance policies	Yes	Hidden
Manage annotations	Yes	Yes
Manage annotation rules	Yes	Yes
Manage applications	Yes	Yes
Queries	Yes	Yes
Manage business entities	Yes	Yes

Feature	User - with Perform license	Guest - with Perform license	User - without Perform license	Guest - without Perform license
Assets dashboard	Yes	Yes	Partially available (storage IOPS and VM IOPS widgets are hidden)	Partially available (storage IOPS and VM IOPS widgets are hidden)
Custom dashboard	View only (no create, edit, or save options)	View only (no create, edit, or save options)	View only (no create, edit, or save options)	View only (no create, edit, or save options)
Manage performance policies	Yes	Hidden	Hidden	Hidden
Manage annotations	Yes	Hidden	Yes	Hidden
Manage applications	Yes	Hidden	Yes	Hidden
Manage business entities	Yes	Hidden	Yes	Hidden
Queries	Yes	View and edit only (no save option)	Yes	View and edit only (no save option)

Setting up and managing user accounts

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

Before you begin

You must have completed the following tasks:

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



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

Steps

1. Open Insight in your browser.
2. On the Insight toolbar, click **Admin**.
3. Click **Setup**.
4. Select the **Userstab**.
5. To create a new user, click the **Actions** button and select **Add user**.

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

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

Results

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

Insight user roles

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

- Guest permits you to log into Insight and to view the various pages.

- User permits all guest-level privileges, as well as access to Insight operations such as defining policy and identifying generic devices. The User account type does not allow you to perform data source operations, nor to add or edit any user accounts other than your own.
- Administrator permits you to perform any operation, including adding new users and managing data sources.

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

Configuring Insight for LDAP(s)

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

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

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



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

User Principal Name attribute:

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

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

- **sAMAccountName**

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



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

- **mail**

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

- **referral**

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

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



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



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

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

Ports:

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

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

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

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

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

Groups:

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

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

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

Configuring user definitions using LDAP

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

Before you begin

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

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

About this task

OnCommand Insight supports LDAP and LDAPS via Microsoft Active Directory server. Additional LDAP implementations may work but have not been qualified with Insight. This procedure assumes that you are using Microsoft Active Directory Version 2 or 3 LDAP (Lightweight Directory Access Protocol).

LDAP users display along with the locally defined users in the **Admin > Setup > Users** list.

Steps

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

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

LDAP

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

Enable LDAP

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

LDAP servers:

User:

Password:

[Show more](#) ▾

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

6. Fill in the fields:

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



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

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

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

or, if using the default port:

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

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

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

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

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

- **Admins group**

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

- **Users group**

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

- **Guests group**

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

- **Server admins group**

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

- **Timeout**

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

- **Domain**

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

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

- **User principal name attribute**

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

- **Role attribute**

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

- **Mail attribute**

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



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

- **Distinguished name attribute**

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

8. Click **Save**.

Changing user passwords

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

Before you begin

The following items must have been completed:

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

About this task

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

Steps

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

Editing a user definition

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

Before you begin

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

About this task

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

Steps

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

7. Make the necessary changes.
8. Click **Save**.

Deleting a user account

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

Steps

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

Setting a Login Warning Message

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

Steps

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

The message will display upon login for all users.

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

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

Insight Security

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

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

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

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

Rekeying servers

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

Changing the Acquisition user password

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

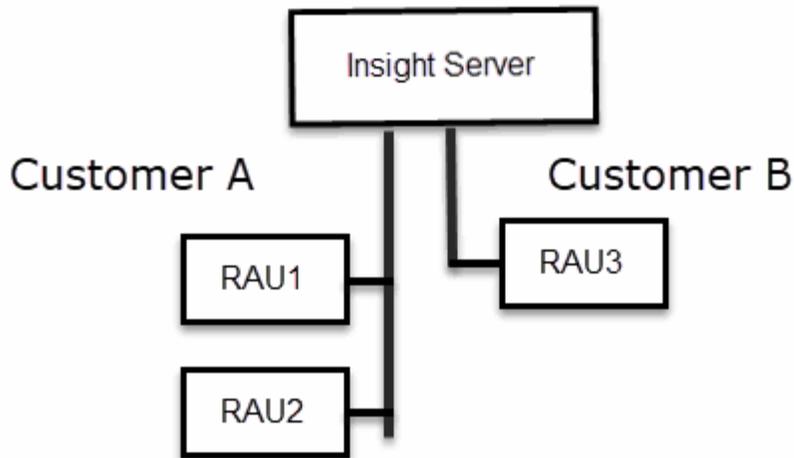
Upgrade and installation considerations

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

Managing keys in a complex service provider environment

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

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



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

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

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

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

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

Managing security on the Insight server

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

About this task

You use the `securityadmin` tool to manage security:

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

Steps

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

The system requests login credentials.

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

The following server configuration options are available:

- **Backup**

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

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

- **Restore**

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



Restore can be used to synchronize passwords and keys on multiple servers, for example:

- Change the server encryption key on one server
- Create a backup of the vault
- Restore the vault backup to the second server

- **Change Encryption Key**

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



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

- **Update Password**

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

displayed:

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



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

• **Reset to Defaults**

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

• **Exit**

Exit the `securityadmin` tool.

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

Managing security on the local acquisition unit

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

Before you begin

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

About this task

You use the `securityadmin` tool to manage security:

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

Steps

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

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

The system requests login credentials.

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

The following options are displayed:

- **Backup**

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

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

- **Restore**

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



Restore can be used to synchronize passwords and keys on multiple servers, for example:

- Change encryption keys on the LAU
- Create a backup of the vault
- Restore the vault backup to each of the RAUs

- **Change Encryption Keys**

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



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

- **Update Password**

Change password for 'acquisition' user account.



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

- **Reset to Defaults**

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

- **Exit**

Exit the `securityadmin` tool.

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

Managing security on an RAU

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

About this task

You use the `securityadmin` tool to manage security:

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

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

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

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

Steps

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

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

The system requests login credentials.

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

The system displays the menu for the RAU.

- **Backup**

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

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

- **Restore**

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



Restore can be used to synchronize passwords and keys on multiple servers, for example:

- Change encryption keys on one server
- Create a backup of the vault
- Restore the vault backup to the second server

◦ **Change Encryption Keys**

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



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

◦ **Update Password**

Change password for 'acquisition' user account.



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

◦ **Reset to Defaults**

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

◦ **Exit**

Exit the `securityadmin` tool.

Managing security on the Data Warehouse

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

About this task

You use the `securityadmin` tool to manage security:

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

Steps

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

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

The system requests login credentials.

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

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

- **Backup**

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

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

- **Restore**

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



Restore can be used to synchronize passwords and keys on multiple servers, for example:

- Change encryption keys on one server
- Create a backup of the vault
- Restore the vault backup to the second server

+

- **Change encryption keys**

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

- **Update Password**

Change password for a specific user account.

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



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

- **Reset to Defaults**

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

- **Exit**

Exit the `securityadmin` tool.

Changing OnCommand Insight internal user passwords

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

Before you begin



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

About this task

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

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

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

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

Changing passwords in the DWH Server Connection Configuration UI

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

LAU Passwords	Required changes
acquisition	Insight Server, RAU

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

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

Before you begin

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

Steps

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

The **Edit Connector** screen is displayed.

Edit Connector

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

[Advanced](#) ▾

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

The Edit Connector Advanced screen is displayed.

Edit Connector

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

[Basic ^](#)

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

7. Click save.

Changing the dwh password using the ODBC Administration tool

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

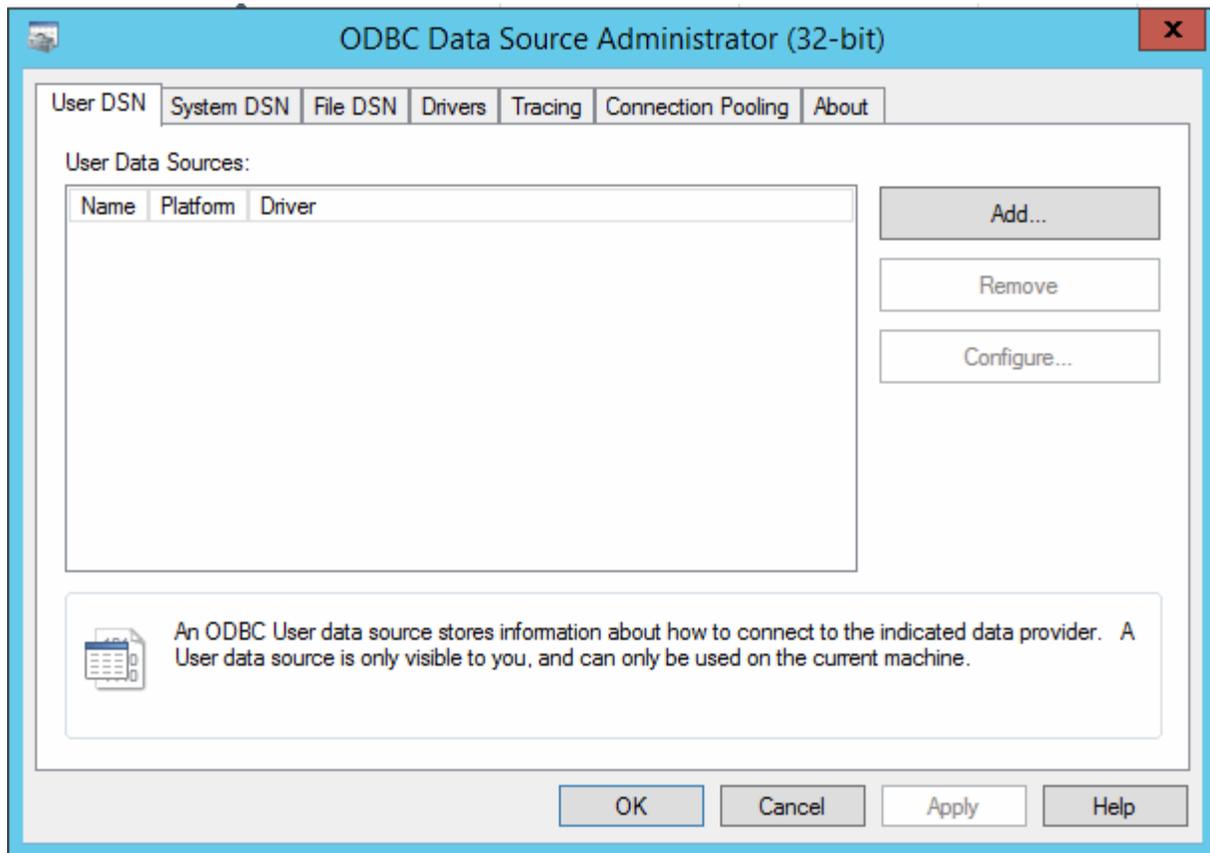
Before you begin

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

Steps

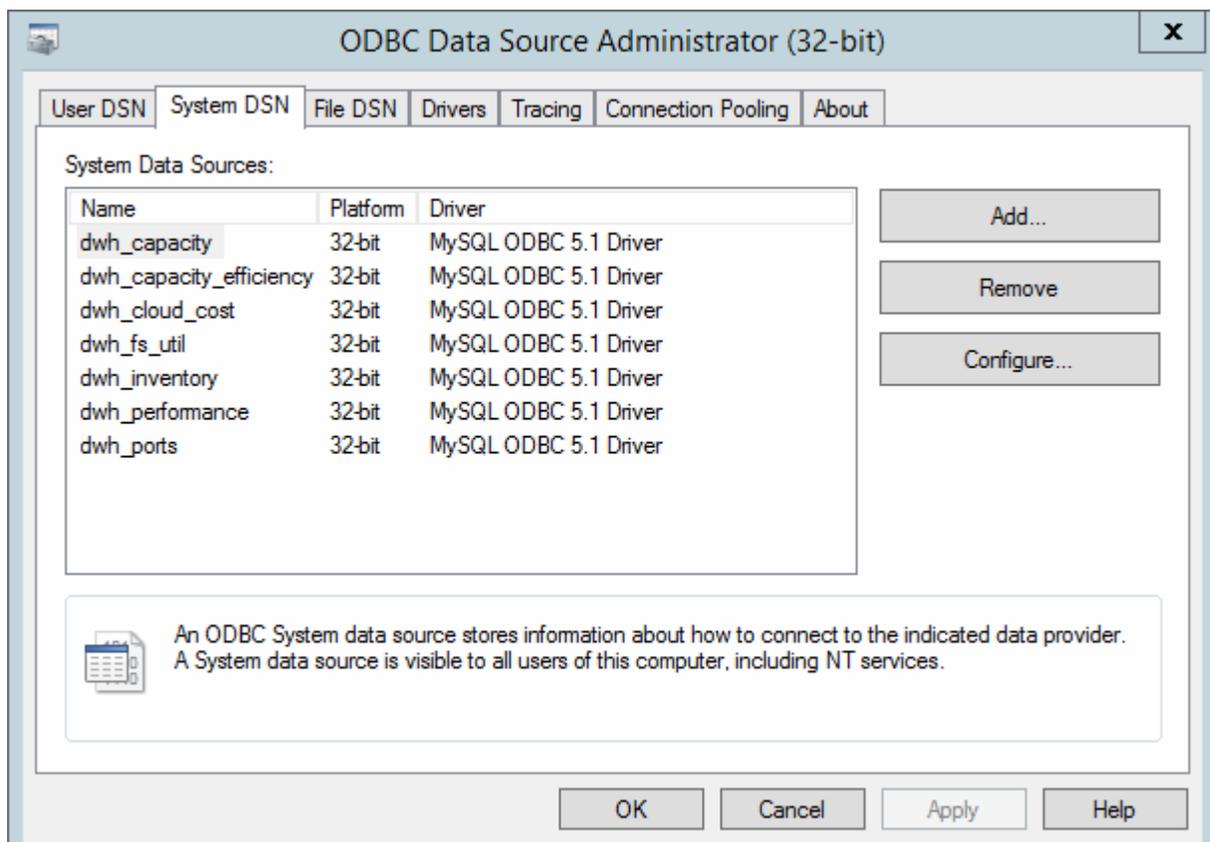
1. Perform a remote login to the server hosting that Data Warehouse.
2. Access the ODBC Administration tool at `C:\Windows\SysWOW64\odbcad32.exe`

The system displays the ODBC Data Source Administrator screen.



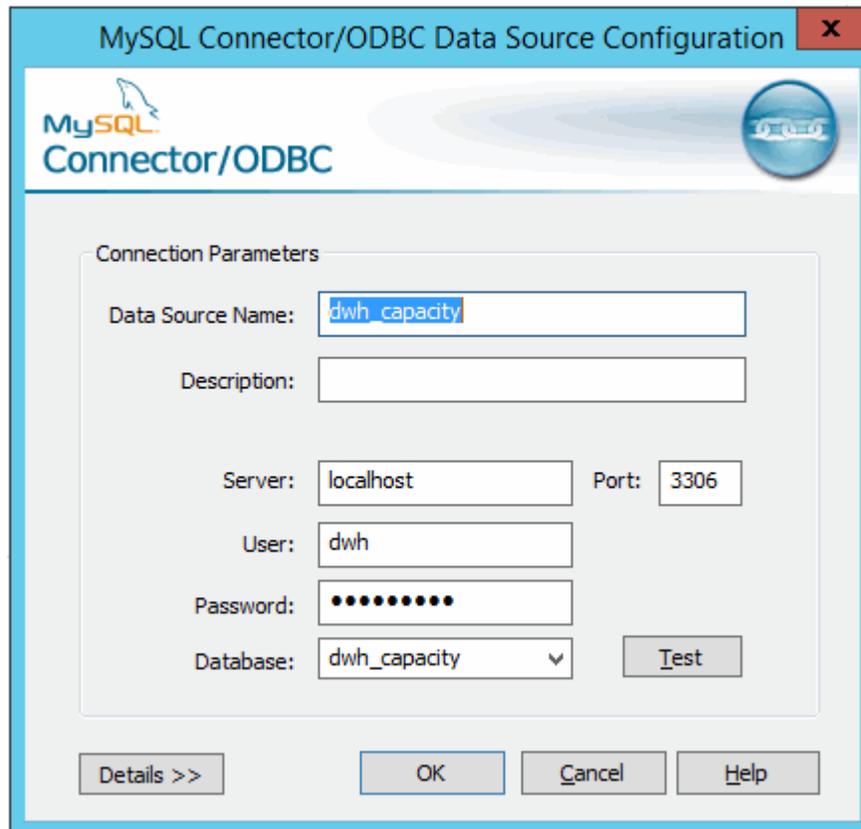
3. Click **System DSN**

The system data sources are displayed.



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

The Data Source Configuration screen is displayed.



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

Smart Card and certificate login support

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

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



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

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



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

Configuring hosts for Smart Card and certificate login

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

Before you begin

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

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



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

Steps

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

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

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

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

Repeat this step for all intermediate certificates.

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

1. Restart the server.

Configuring a client to support Smart Card and certificate login

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

Before you begin

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



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

About this task

The following are the common client configuration requirements:

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

Enabling CAC on a Linux server

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

Steps

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

Configuring Data Warehouse for Smart Card and certificate login

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

Before you begin

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

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

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

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



Steps

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

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

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

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

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

- b. Use the `keytool` utility to list the trusted CAs: `C:\Program Files\SANscreen\java64\bin\keytool.exe -list -keystore server.trustore -storepass changeit`

The first word in each line indicates the CA alias.

- c. If necessary, supply a CA certificate file, usually a `.pem` file. To include customer's CAs with Data Warehouse trusted CAs go to `..\SANscreen\wildfly\standalone\configuration` and use the `keytool import` command: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore server.trustore -alias my_alias -file 'path/to/my.pem' -v -trustcacerts`

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

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

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

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

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

- Restart the OnCommand Insight server.

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

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

Before you begin

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

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

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



Steps

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

The first word in each line indicates the CA alias.

- c. If no suitable files exist, supply a CA certificate file, usually a `.pem` file.

- d. To include customer's CAs with OnCommand Insight trusted CAs, go to `..\SANscreen\cognos\analytics\configuration\certs\`.
- e. Use the `keytool` utility to import the `.pem` file: `..\..\jre\bin\keytool.exe -importcert -keystore CAMKeystore.jks -alias my_alias -file 'path/to/my.pem' -v -trustcacerts`

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

- f. When prompted for a password, enter `NoPassWordSet`.
 - g. Answer `yes` when prompted to trust the certificate.
2. To enable CAC mode, execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`
 3. To disable CAC mode, execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`

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

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

Before you begin

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

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



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

Steps

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

The first word in each line indicates the CA alias.

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

e. Use the `keytool` utility to import the `.pem` file: `..\..\ibm-jre\jre\bin\keytool.exe -importcert -keystore CAMKeystore.jks -alias my_alias -file 'path/to/my.pem' -v -trustcacerts`

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

f. When prompted for a password, enter `NoPassWordSet`.

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

2. To enable CAC mode, do the following:

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

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

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

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

3. To disable CAC mode, do the following:

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

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

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

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

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

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

Before you begin

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

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



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

About this task

You must have admin privileges to perform this procedure.

Steps

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

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

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

This will download `fqdn.p7b` file

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

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

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

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

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

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

Before you begin

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

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



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

About this task

You must have admin privileges to perform this procedure.

Steps

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

This will download `fqdn.p7b` file

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

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


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

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

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

Configuring Data Warehouse for Smart Card and certificate login

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

Before you begin

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

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

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

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



Steps

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

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

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

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

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

- b. Use the `keytool` utility to list the trusted CAs: `C:\Program Files\SANscreen\java64\bin\keytool.exe -list -keystore server.trustore -storepass changeit`

The first word in each line indicates the CA alias.

- c. If necessary, supply a CA certificate file, usually a `.pem` file. To include customer's CAs with Data

Warehouse trusted CAs go to `..\SANscreen\wildfly\standalone\configuration` and use the `keytool import` command: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore server.trustore -alias my_alias -file 'path/to/my.pem' -v -trustcacerts`

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

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

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

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

4. Restart the OnCommand Insight server.

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

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

Before you begin

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

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



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

Steps

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

The first word in each line indicates the CA alias.
 - c. If no suitable files exist, supply a CA certificate file, usually a `.pem` file.
 - d. To include customer's CAs with OnCommand Insight trusted CAs, go to `..\SANscreen\cognos\analytics\configuration\certs\`.
 - e. Use the `keytool` utility to import the `.pem` file: `..\..\jre\bin\keytool.exe -importcert -keystore CAMKeystore.jks -alias my_alias -file 'path/to/my.pem' -v -trustcacerts`

`my_alias` is usually an alias that would easily identify the CA in the `keytool -list` operation.
 - f. When prompted for a password, enter `NoPassWordSet`.
 - g. Answer `yes` when prompted to trust the certificate.
2. To enable CAC mode, execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`
3. To disable CAC mode, execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`

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

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

Before you begin

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

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

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



Steps

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

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

The first word in each line indicates the CA alias.

- c. If no suitable files exist, supply a CA certificate file, usually a `.pem` file.
- d. To include customer's CAs with OnCommand Insight trusted CAs, go to `..\SANscreen\cognos\analytics\configuration\certs\`.
- e. Use the `keytool` utility to import the `.pem` file: `..\..\ibm-jre\jre\bin\keytool.exe -importcert -keystore CAMKeystore.jks -alias my_alias -file 'path/to/my.pem' -v -trustcacerts`

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

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

2. To enable CAC mode, do the following:

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

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

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

3. To disable CAC mode, do the following:

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

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

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

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

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

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

Before you begin

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

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



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

About this task

You must have admin privileges to perform this procedure.

Steps

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

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

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

This will download `fqdn.p7b` file

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

- c. Right-click on root CA > All Tasks > Export.
 - d. Select Base64 output.
 - e. Enter a file name identifying it as the root certificate.
 - f. Repeat steps 8a through 8c to export all of the certificates separately into .cer files.
 - g. Name the files intermediateX.cer and cognos.cer.
9. Ignore this step if you have only one CA certificate, otherwise merge both root.cer and intermediateX.cer into one file.
 - a. Open intermediate.cer with NotePad and copy the content.
 - b. Open root.cer with NotePad and save the content from 9a.
 - c. Save the file as CA.cer.
 10. Import the certificates into the Cognos keystore using the Admin CMD prompt:
 - a. cd "Program Files\sanscreen\cognos\analytics\bin"
 - b. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\CA.cer

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

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

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

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

Before you begin

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

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



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

About this task

You must have admin privileges to perform this procedure.

Steps

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

This will download `fqdn.p7b` file

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

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


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

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

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

Importing SSL certificates

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

Before you begin

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

About this task



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

Steps

1. Create a copy of the original keystore file:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

You must obtain the public certificates.

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

a. When prompted, enter the keystore password.

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

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

- a. When prompted, enter the trustore password.

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

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

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

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

10. Restart the SANscreen server service.

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

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

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

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

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

- Java and Mozilla Firefox have their own certificate stores.

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

Setting up weekly backups for your Insight database

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

About this task

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

The backup file will contain the following:

- Inventory data
- Up to 7 days of performance data

Steps

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



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

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

Results

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

What's included in the backup

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

The weekly or on-demand backup includes the following:

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

- Annotations and annotation rules

The weekly backup does not include:

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



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

Performance data archiving

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

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

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

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

Enabling performance archive

To enable performance data archiving, follow these steps.

Steps

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

You cannot specify a folder under the Insight installation folder.

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

5. Click **Save**.

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

Loading performance archive

To load the performance data archive, follow these steps.

Before you begin

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

Steps

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



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

Configuring your email

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

Email configuration prerequisites

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

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

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

Configuring your email for Insight

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

Steps

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

If you specify a hostname, ensure that the name can be resolved through DNS.

4. In the **User name** box, enter your user name.
5. In the **Password** box, enter the password for accessing the email server, which is required only if your SMTP server is password-protected. This is the same password that you use to log into the application that lets you read your email. If a password is required, you must enter it a second time for verification.
6. In the **Sender email** box, enter the sender email account that will be identified as the sender on all OnCommand Insight reports.

This account must be a valid email account within your organization.

7. In the **Email signature** box, enter the text that you want to be inserted in every email that is sent.
8. In the Recipients box, click **+**, enter an email address, and click **OK**.

To edit an email address, select the address, and click . To delete an email address, select the address, and click .

9. To send a test email to specified recipients, click .
10. Click **Save**.

Configuring SNMP notifications

OnCommand Insight supports SNMP notifications for configuration and Global Path policy changes as well as violations. For example, SNMP notifications are sent when data source thresholds are exceeded.

Before you begin

The following must have been completed:

- Identifying the IP address of the server that consolidates traps for each type of event.

You might have to consult with your system administrator to obtain this information.

- Identifying the port number through which the designated machine obtains SNMP traps, for each type of event.

The default port for SNMP traps is 162.

- Compiling the MIB at your site.

The proprietary MIB comes with the installation software to support OnCommand Insight traps. The NetApp MIB is compatible with all standard SNMP management software and can be found on the Insight server in `<install dir>\SANscreen\MIBS\sanscreen.mib`.

Steps

1. Click **Admin** and select **Notifications**.
2. Scroll down to the **SNMP** section of the page.
3. Click **Actions** and select **Add trap source**.

4. In the **Add SNMP trap recipients** dialog box, enter these values:

- **IP**

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

- **Port**

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

- **Community String**

Use “public” for SNMP trap messages.

5. Click **Save**.

Enabling the syslog facility

You can identify a location for the log of the OnCommand Insight violations and performance alerts as well as audit messages, and activate the logging process.

Before you begin

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

About this task

The syslog includes the following types of information:

- Violation messages
- Performance alerts
- Optionally, Audit log messages

The following units are used in the syslog:

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

Steps

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

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

6. In the **Facility** field, select the facility level that corresponds to the type of program that is logging the message.
7. Click **Save**.

Insight syslog contents

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

Message types

The Insight syslog lists three types of messages:

- SAN path violations
- General violations
- Performance alerts

Data provided

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

Performance alerts include these data:

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

Configuring performance and assure violation notifications

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

Before you begin

You must have configured email, syslog, and SNMP sending methods for violations.

Steps

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

4. Click **Save**.

Configuring system-level event notifications

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

Before you begin

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

Steps

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

Configuring your ASUP processing

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

Before you begin

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

About this task

ASUP data is forwarded using the HTTPS protocol.

Steps

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

What's included in the Autosupport (ASUP) package

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

The Autosupport package includes the following:

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

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

The Autosupport package does not include:

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



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

Defining applications

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

Before you begin

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

About this task

You can associate applications with the following assets: hosts, virtual machines, volumes, internal volumes, qtrees, shares, and hypervisors.

Steps

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

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

3. Click **Add**.

The Add Application dialog box displays.

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

This requires the Assure license. Set this when you want to ensure each host has access to the same volumes in a cluster. For example, hosts in high-availability clusters often need to be masked to the same volumes to allow for failover; however, hosts in unrelated applications usually have no need to access the

same physical volumes. Additionally, regulatory policies might require you to explicitly disallow unrelated applications from accessing the same physical volumes for security reasons.

8. Click **Save**.

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

After you finish

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

Assigning applications to assets

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

Steps

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

The list of available applications for the selected asset display. The applications that are currently associated with the asset are preceded by a check mark.

4. You can type in the Search box to filter the application names, or you can scroll down the list.
5. Select the applications you want to associate with the asset.

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

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

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

Editing applications

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

Steps

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

The Edit Application dialog box displays.

4. Do any of the following:
 - Click **Priority** and select a different priority.



You cannot change the application's name.

- Click **Business Entity** and select a different business entity to associate the application with or select **None** to remove the association of the application with the business entity.
- Click to clear or select **Validate volume sharing**.



This option is only available if you have the Assure license.

5. Click **Save**.

Deleting applications

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

Steps

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

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

4. Click **OK**.

Your business entities hierarchy

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

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

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

spaces.

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

Designing your business entities hierarchy

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

Steps

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

Creating business entities

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

About this task

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

Steps

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

The Business Entities page displays.

3. Click  to begin building a new entity.

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

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

Assigning business entities to assets

You can assign a business entity to an asset (host, port, storage, switch, virtual machine, qtree, share, volume, or internal volume) without having associated the business entity to an application; however, business entities are assigned automatically to an asset if that asset is associated with an application related to a business entity.

Before you begin

You must have already created a business entity.

About this task

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

Steps

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

The list of available business entities display.

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

If the business entity you choose is associated with an application, the application name is displayed. In this case, the word “derived” appears next to the business entity name. If you want to maintain the entity for only the asset and not the associated application, you can manually override the assignment of the application.

5. To override an application derived from a business entity, place your cursor over the application name and click , select another business entity, and select another application from the list.

Assigning business entities to or removing business entities from multiple assets

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

Before you begin

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

Steps

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

The **Actions** button displays.

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

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

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

Defining annotations

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

Steps

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

Using annotations to monitor your environment

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

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

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

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

- Creating rules to automatically apply annotations to multiple assets of the same type.
- Using the annotation import utility to import annotations.
- Filter assets by their annotations.
- Grouping data in reports based on annotations and generate those reports.

See the *OnCommand Insight Reporting Guide* for more information about reports.

Managing annotation types

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

Default annotation types

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

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

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

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

Annotation types	Description	Type
Alias	User-friendly name for a resource.	Text
Birthday	Date when the device was or will be brought online.	Date
Building	Physical location of host, storage, switch, and tape resources.	List
City	Municipality location of host, storage, switch, and tape resources.	List

Compute Resource Group	Group assignment used by the Host and VM Filesystems data source.	List
Continent	Geographic location of host, storage, switch, and tape resources.	List
Country	National location of host, storage, switch, and tape resources.	List
Data Center	Physical location of the resource and is available for hosts, storage arrays, switches, and tapes.	List
Direct Attached	Indicates (Yes or No) if a storage resource is connected directly to hosts.	Boolean
End of Life	Date when a device will be taken offline, for example, if the lease expired or the hardware is being retired.	Date
Fabric Alias	User-friendly name for a fabric.	Text
Floor	Location of a device on a floor of a building. Can be set for hosts, storage arrays, switches, and tapes.	List
Hot	Devices already in heavy use on a regular basis or at the threshold of capacity.	Boolean
Note	Comments that you want associated with a resource.	Text
Rack	Rack in which the resource resides.	Text
Room	Room within a building or other location of host, storage, switch, and tape resources.	List
SAN	Logical partition of the network. Available on hosts, storage arrays, tapes, switches, and applications.	List

Service Level	A set of supported service levels that you can assign to resources. Provides an ordered options list for internal volumes, qtree, and volumes. Edit service levels to set performance policies for different levels.	List
State/Province	State or province in which the resource is located.	List
Sunset	Threshold set after which no new allocations can be made to that device. Useful for planned migrations and other pending network changes.	Date
Switch Level	Includes predefined options for setting up categories for switches. Typically, these designations remain for the life of the device, although you can edit them, if needed. Available only for switches.	List
Tier	Can be used to define different levels of service within your environment. Tiers can define the type of level, such as speed needed (for example, gold or silver). This feature is available only on internal volumes, qtrees, storage arrays, storage pools, and volumes.	List
Violation Severity	Rank (for example, major) of a violation (for example, missing host ports or missing redundancy), in a hierarchy of highest to lowest importance.	List



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

How annotations are assigned

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

User Data section of the asset page.

Annotations are assigned in the following ways:

- You can assign an annotation manually to an asset.

If an annotation is assigned directly to an asset, the annotation appears as normal text on an asset page. Annotations that are assigned manually always take precedence over annotations that are inherited or assigned by annotation rules.

- You can create an annotation rule to automatically assign annotations to assets of the same type.

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

- Insight automatically associates a tier level with a storage tier model to expedite the assignment of storage annotations to your resources on acquisition of assets.

Certain storage resources are automatically associated with a predefined tier (Tier 1 and Tier 2). For example, the Symmetrix storage tier is based on the Symmetrix and VMAX family and is associated with Tier 1. You can change the default values to match your tier requirements. If the annotation is assigned by Insight (for example, Tier), you see “System-defined” when you position your cursor over the annotation’s name on an asset page.

- A few resources (children of an asset) can derive the predefined Tier annotation from their asset (parent).

For example, if you assign an annotation to a storage, the Tier annotation is derived by all the storage pools, internal volumes, volumes, qtrees, and shares belonging to the storage. If a different annotation is applied to an internal volume of the storage, the annotation is subsequently derived by all the volumes, qtrees, and shares. “Derived” appears next to the annotation name on an asset page.

Associating costs with annotations

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

Steps

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

The Annotation page displays.

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

The Edit Annotation dialog box displays.

4. Enter the values for any existing levels in the **Cost** field.

The Tier and Service Level annotations have Auto Tier and Object Storage values, respectively, which you

cannot remove.

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

Creating custom annotations

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

Steps

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

The Annotations page displays the list of annotations.

3. Click .

The **Add Annotation** dialog box displays.

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

You can enter up to 255 characters in these fields.



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

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

- Boolean

This creates a drop-down list with the choices of yes and no. For example, the "Direct Attached" annotation is Boolean.

- Date

This creates a field that holds a date. For example, if the annotation will be a date, select this.

- List

This can create either of the following:

- A drop-down fixed list

When others are assigning this annotation type on a device, they cannot add more values to the list.

- A drop-down flexible list

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

- Number

This creates a field where the user assigning the annotation can enter a number. For example, if the annotation type is “Floor”, the user could select the Value Type of “number” and enter the floor number.

- Text

This creates a field that allows free-form text. For example, you might enter “Language” as the annotation type, select “Text” as the value type, and enter a language as a value.



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

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

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

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

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

- c. Click  to add additional values.

- d. Click  to remove a value.

7. Click **Save**.

Your annotations appear in the list on the Annotations page.

Related information

[Importing and Exporting user data](#)

Manually assigning annotations to assets

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

Before you begin

You must have created the annotation you want to assign.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate the asset to which you want to apply the annotation by doing either of the following:
 - Click the asset in the Assets Dashboard.
 - Click  on the toolbar to display the **Search assets** box, type the type of or name of the asset, and then select the asset from the list that displays.

The asset page displays.

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

The Add Annotation dialog box displays.

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

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

Modifying annotations

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

Steps

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

The Annotations page displays.

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

The **Edit Annotation** dialog box displays.

4. You can make the following modifications to an annotation:
 - a. Change the name, description, or both.

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



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

- b. To add another value to an annotation of list type, click  .
- c. To remove a value from an annotation of list type, click  .

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

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

After you finish

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

Deleting annotations

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

Steps

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

The Annotations page displays.

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

A confirmation dialog box displays.

4. Click **OK**.

Assigning annotations to assets using annotation rules

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

Default storage annotation rules

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

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

- Tier 1, storage quality tier

The Tier 1 annotation is applied to the following vendors and their specified families: EMC (Symmetrix),

HDS (HDS9500V, HDS9900, HDS9900V, R600, R700, USP r, USP V), IBM (DS8000), NetApp (FAS6000 or FAS6200), and Violin (Memory).

- Tier 2, storage quality tier

The Tier 2 annotation is applied to the following vendors and their specified families: HP (3PAR StoreServ or EVA), EMC (CLARiiON), HDS (AMS or D800), IBM (XIV), and NetApp (FAS3000, FAS3100, and FAS3200).

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

Creating annotation rules

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

Before you begin

You must have created a query for the annotation rule.

About this task

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

Steps

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

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

3. Click  **Add**.

The Add Rule dialog box displays.

4. Do the following:
 - a. In the **Name** box, enter a unique name that describes the rule.

This name will appear in the Annotation Rules page.

- b. Click **Query** and select the query that OnCommand Insight should use to apply the annotation to assets.
- c. Click **Annotation** and select the annotation you want to apply.
- d. Click **Value** and select a value for the annotation.

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

5. Click **Save**.

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

Setting annotation rule precedence

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

Steps

1. Log in to the Insightweb UI.
2. Click **Manage** and select **Annotation rules**.

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

3. Position your cursor over an annotation rule.

The precedence arrows appear to the right of the rule.

4. To move a rule up or down in the list, click the up arrow or the down arrow.

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

Modifying annotation rules

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

Steps

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

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

3. Locate the rule that you want to modify:

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

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

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

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

Deleting annotation rules

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

Steps

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

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

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

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

5. Click **OK**.

Importing annotation values

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

About this task

The following rules apply:

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

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

- When storage, switches, or ports are annotated, the Application column is ignored.
- The columns of Tenant, Line_of_Business, Business_Unit, and Project make up a business entity.

Any of the values can be left empty. If an application is already related with a business entity different from the input values, the application is assigned to the new business entity.

The following object types and keys are supported in the import utility:

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

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

The CSV file should use the following format:

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

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

...

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

Steps

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

The Troubleshooting page displays.

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

```
java -jar rest-import-utility.jar [-username] [-ppassword]
[-aserver name or IP address] [-bbatch size] [-ccase
sensitive:true/false]
[-lextra logging:true/false] csv filename
```

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



There are no spaces between the options and their values.

The following keywords are reserved and prevent users from specifying them as annotation names:



- Application
- Application_Priority
- Tenant
- Line_Of_Business
- Business_Unit
- Project

Errors are generated if you attempt to import an annotation type using one of the reserved keywords. If you have created annotation names using these keywords, you must modify them so that the import utility tool can work correctly.



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

Assigning annotations to multiple assets using a query

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

Before you begin

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

About this task

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

Steps

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

The system displays the Add Annotation dialog.

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

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

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

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

Querying assets

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

Assets used in queries and dashboards

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

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

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

Creating a query

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

About this task

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

When a new data source is added which acquires assets or any annotation or application assignments are

made, you can query for those assets, annotations, or applications after the queries are indexed, which occurs at a regularly scheduled interval.

Steps

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

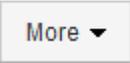
When a resource is selected for a query, a number of default columns are automatically displayed; you can remove these columns or add new ones at any time.

4. In the **Name** text box, type the name of the asset or type a portion of text to filter through the asset names.

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

- An asterisk enables you to search for everything. For example, `vol*rhel` displays all resources that start with “vol” and end with “rhel”.
- The question mark enables you to search for a specific number of characters. For example, `BOS-PRD??-S12` displays BOS-PRD12-S12, BOS-PRD13-S12, and so on.
- The OR operator enables you to specify multiple entities. For example, `FAS2240 OR CX600 OR FAS3270` finds multiple storage models.
- The NOT operator allows you to exclude text from the search results. For example, `NOT EMC*` finds everything that does not start with “EMC”. You can use `NOT *` to display fields that contain no value.

5. Click  to display the assets.

6. To add a criteria, click , and do either of the following:

- Type to search for a specific criteria and then select it.
- Scroll down the list and select a criteria.
- Enter a range of values if you choose a performance metric like IOPS - Read (IO/s). Default annotations provided by Insight are indicated by ; it is possible to have annotations with duplicate names.

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

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

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



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

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

If you have an account with an administrator role, you can create custom dashboards. A custom dashboard

can comprise any of the widgets from Widget Library, several of which, let you represent query results in a custom dashboard. For more information about custom dashboards, see the *OnCommand Insight Getting Started Guide*.

Related information

[Importing and Exporting user data](#)

Viewing queries

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

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Queries** and select **Show all queries**.
3. You can change how queries display by doing any of the following:
 - You can enter text in the **filter** box to search to display specific queries.
 - You can change the sort order of the columns in the table of queries to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header.
 - To resize a column, hover the mouse over the column header until a blue bar appears. Place the mouse over the bar and drag it right or left.
 - To move a column, click on the column header and drag it right or left.
 - When scrolling through the query results, be aware that the results may change as Insight automatically polls your data sources. This may result in some items being missing, or some items appearing out of order depending on how they are sorted.

Exporting query results to a .CSV file

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

Steps

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

The Queries page is displayed.
3. Click a query.
4. Click  to export query results to a .CSV file.
5. Do one of the following:

- Click **Open with** and then **OK** to open the file with Microsoft Excel and save the file to a specific location.
- Click **Save file** and then **OK** to save the file to your Downloads folder.
Only the attributes for the displayed columns will be exported. Some displayed columns, particularly

those that are part of complex nested relationships, are not exported.



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

+

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

+

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

+



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

+

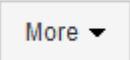
Modifying queries

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

Steps

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

The Queries page is displayed.

3. Click the query name.
4. To remove a criterion from the query, click .
5. To add a criteria to the query, click , and select a criteria from the list.

6. Do one of the following:

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

Deleting queries

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

Steps

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

The Queries page displays.

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

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

4. Click **OK**.

Assigning multiple applications to or removing multiple applications from assets

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

Before you begin

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

Steps

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

The Queries page displays.

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

The list of assets associated with the query displays.

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

The **Actions** button displays.

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

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

You can select multiple applications for hosts, internal volumes, and virtual machines; however, you

can select only one application for a volume.

b. Click **Save**.

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

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

b. Click **Delete**.

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

Editing or removing multiple annotations from assets

You can edit multiple annotations for assets or remove multiple annotations from assets by using a query instead of having to manually edit or remove them.

Before you begin

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

Steps

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

The Queries page displays.

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

The list of assets associated with the query displays.

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

The **Actions** button displays.

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

Copying table values

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

About this task

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

Steps

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

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

Managing performance policies

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

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

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

- Hypervisor

There are ESX swapping and ESX utilization policies.

- Internal volume and volume

There are two latency policies for each resource, one annotated for Tier 1 and the other annotated for Tier 2.

- Port

There is a policy for BB credit zero.

- Storage node

There is a policy for node utilization.

- Virtual machine

There are VM swapping and ESX CPU and memory policies.

- Volume

There are latency by tier and misaligned volume policies.

Creating performance policies

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

Steps

1. Open OnCommand Insight in your browser.
2. Select **Manage > Performance Policies**.

The Performance Policies page is

The screenshot shows the 'Performance Policies' page in OnCommand Insight. The page is organized into three sections, each with a table of policies. The top section is 'Database policies', the middle is 'Internal volume policies', and the bottom is 'Storage policies'. Each table has columns for Policy Name, Severity, Annotations, Time Window, and Thresholds. The 'Database policies' table shows two policies: 'Latency' and 'Database_0'. The 'Internal volume policies' table shows two policies: 'Atmos Service Level' and 'Global'. The 'Storage policies' table shows two policies: 'Storage_Storage' and 'Storage_0'. Each table has a 'Showing 1 of 2 of 2 entries' indicator at the bottom right.

Policy Name	Severity	Annotations	Time Window	Thresholds
Latency	Warning		First occurrence	'Latency - Total' > 200 ms
Database_0	Warning		First occurrence	'IOPS - Total' > 0 I/Os or 'Latency - Total' > 0 ms

Policy Name	Severity	Annotations	Time Window	Thresholds
Atmos Service Level	Critical	Service_Level = Atmos	First occurrence	'Latency - Total' > 100 ms or 'IOPS - Total' > 100 I/Os or 'Throughput - Total' > 200 MB/s
Global	Critical		First occurrence	'Latency - Total' > 200 ms or 'IOPS - Total' > 1 I/Os or 'Throughput - Total' > 300 MB/s

Policy Name	Severity	Annotations	Time Window	Thresholds
Storage_Storage	Warning		First occurrence	'IOPS - Read' > 10 I/Os
Storage_0	Warning		First occurrence	'Throughput - Total' > 0 MB/s or 'IOPS - Total' > 0 I/Os

displayed.

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

3. Click **Add new policy**.

The Add Policy dialog box is displayed.

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

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

5. From the **Apply to objects of type** list, select the type of object to which the policy applies.
6. From the **With annotation** list, select an annotation type, if applicable, and enter a value for the annotation in the **Value** box to apply the policy only to objects that have this particular annotation set.
7. If you selected **Port** as the object type, from the **Connected to** list, select what the port is connected to.
8. From the **Apply after a window of** list, select when an alert is raised to indicate a threshold violation.

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

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

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

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

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

16. Click **Save**.

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

Performance policy evaluation precedence

The Performance Policies page groups policies by object type and Insight evaluates the policies in the order in which they appear in the object's performance policy list. You can

change the order in which Insight evaluates policies in order to show the information that is most important to you in your network.

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

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

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

Changing the precedence of a performance policy

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

Steps

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

The Performance Policies page displays.

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

The precedence arrows appear to the right of the policy.

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

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

Editing performance policies

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

Steps

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

The Performance Policies page displays.

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

The Edit Policy dialog box displays.

5. Make the required changes.

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

6. Click **Save**.

Deleting performance policies

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

Steps

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

The Performance Policies page displays.

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

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

5. Click **OK**.

Importing and Exporting user data

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

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

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

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

- Annotation Types

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



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

- Annotation Rules

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



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

- Policies

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



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

+

- Queries

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



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

+

- Dashboards

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



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

+

Insight Security

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

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

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

Rekeying servers

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

Changing the Acquisition user password

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

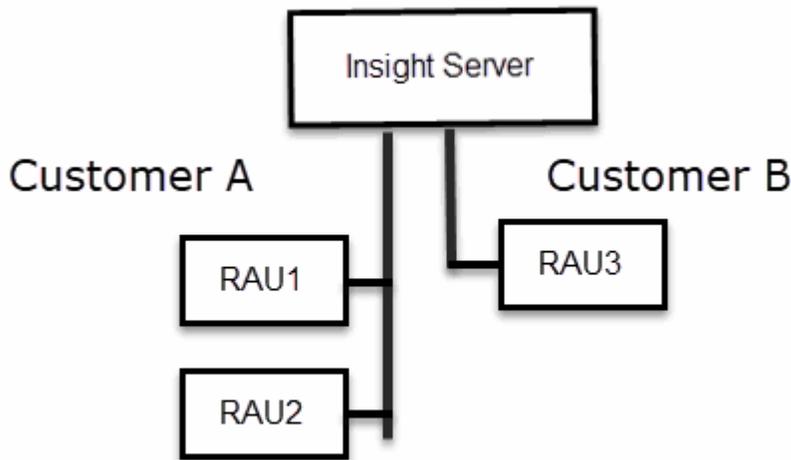
Upgrade and installation considerations

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

Managing keys in a complex service provider environment

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

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



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

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

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

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

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

Managing security on the Insight server

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

About this task

You use the `securityadmin` tool to manage security:

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

Steps

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

The system requests login credentials.

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

The following server configuration options are available:

- **Backup**

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

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

- **Restore**

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



Restore can be used to synchronize passwords and keys on multiple servers, for example:

- Change the server encryption key on one server
- Create a backup of the vault
- Restore the vault backup to the second server

- **Change Encryption Key**

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



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

- **Update Password**

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

displayed:

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



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

- **Reset to Defaults**

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

- **Exit**

Exit the `securityadmin` tool.

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

Managing security on the local acquisition unit

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

Before you begin

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

About this task

You use the `securityadmin` tool to manage security:

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

Steps

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

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

The system requests login credentials.

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

The following options are displayed:

- **Backup**

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

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

- **Restore**

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



Restore can be used to synchronize passwords and keys on multiple servers, for example:

- Change encryption keys on the LAU
- Create a backup of the vault
- Restore the vault backup to each of the RAUs

- **Change Encryption Keys**

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



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

- **Update Password**

Change password for 'acquisition' user account.



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

- **Reset to Defaults**

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

- **Exit**

Exit the `securityadmin` tool.

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

Managing security on an RAU

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

About this task

You use the `securityadmin` tool to manage security:

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

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

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

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

Steps

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

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

The system requests login credentials.

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

The system displays the menu for the RAU.

- **Backup**

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

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

- **Restore**

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



Restore can be used to synchronize passwords and keys on multiple servers, for example:

- Change encryption keys on one server
- Create a backup of the vault
- Restore the vault backup to the second server

◦ **Change Encryption Keys**

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



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

◦ **Update Password**

Change password for 'acquisition' user account.



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

◦ **Reset to Defaults**

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

◦ **Exit**

Exit the `securityadmin` tool.

Managing security on the Data Warehouse

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

About this task

You use the `securityadmin` tool to manage security:

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

Steps

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

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

The system requests login credentials.

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

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

◦ **Backup**

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

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

◦ **Restore**

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



Restore can be used to synchronize passwords and keys on multiple servers, for example:

- Change encryption keys on one server
- Create a backup of the vault
- Restore the vault backup to the second server

+

◦ **Change encryption keys**

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

◦ **Update Password**

Change password for a specific user account.

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



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

- **Reset to Defaults**

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

- **Exit**

Exit the `securityadmin` tool.

Changing OnCommand Insight internal user passwords

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

Before you begin



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

About this task

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

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

The following table lists the internal user passwords for the Data Warehouse and lists the Insight components

that have dependent passwords that need to match the new password.

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

Changing passwords in the DWH Server Connection Configuration UI

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

LAU Passwords	Required changes
acquisition	Insight Server, RAU

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

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

Before you begin

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

Steps

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

The **Edit Connector** screen is displayed.

Edit Connector

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

[Advanced](#) ▾

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

The Edit Connector Advanced screen is displayed.

Edit Connector

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

[Basic ^](#)

6. Enter the new password in the **Server password** field:
7. Click save.

Changing the dwh password using the ODBC Administration tool

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

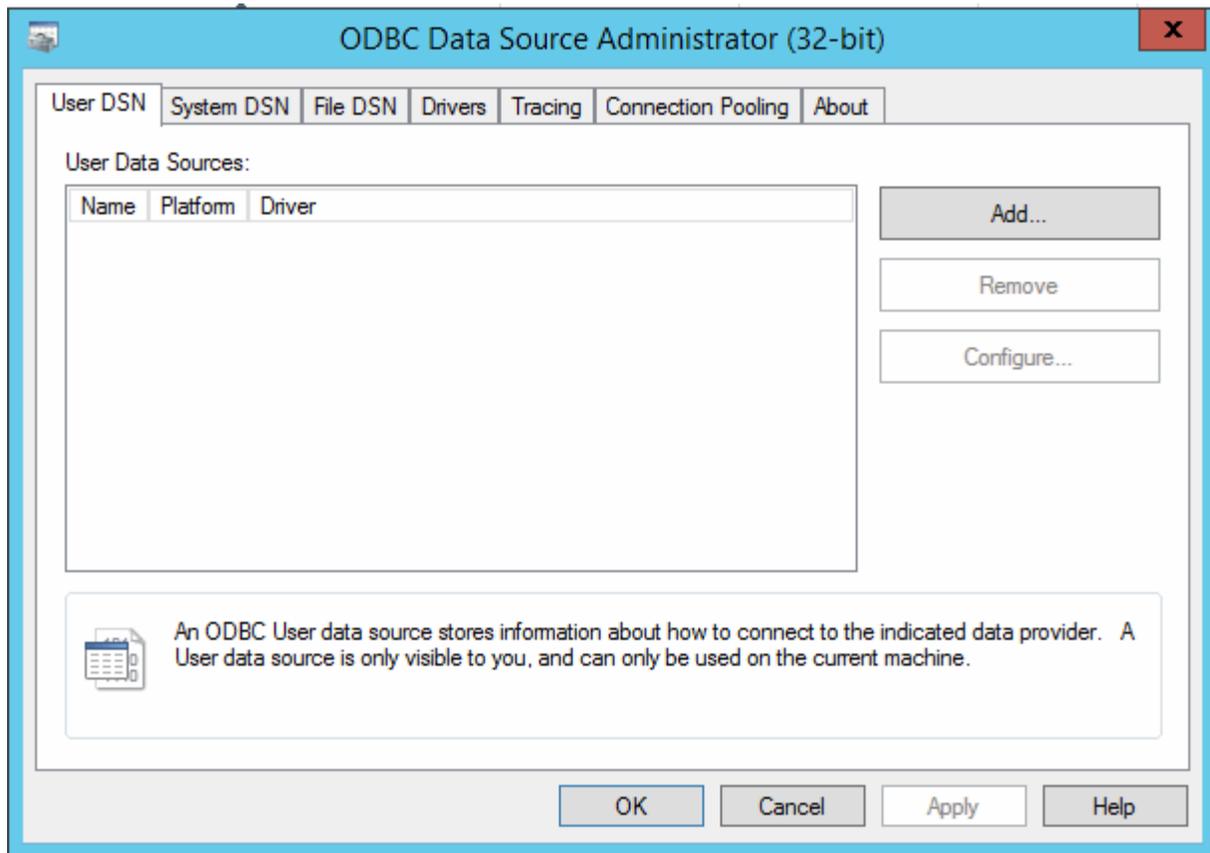
Before you begin

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

Steps

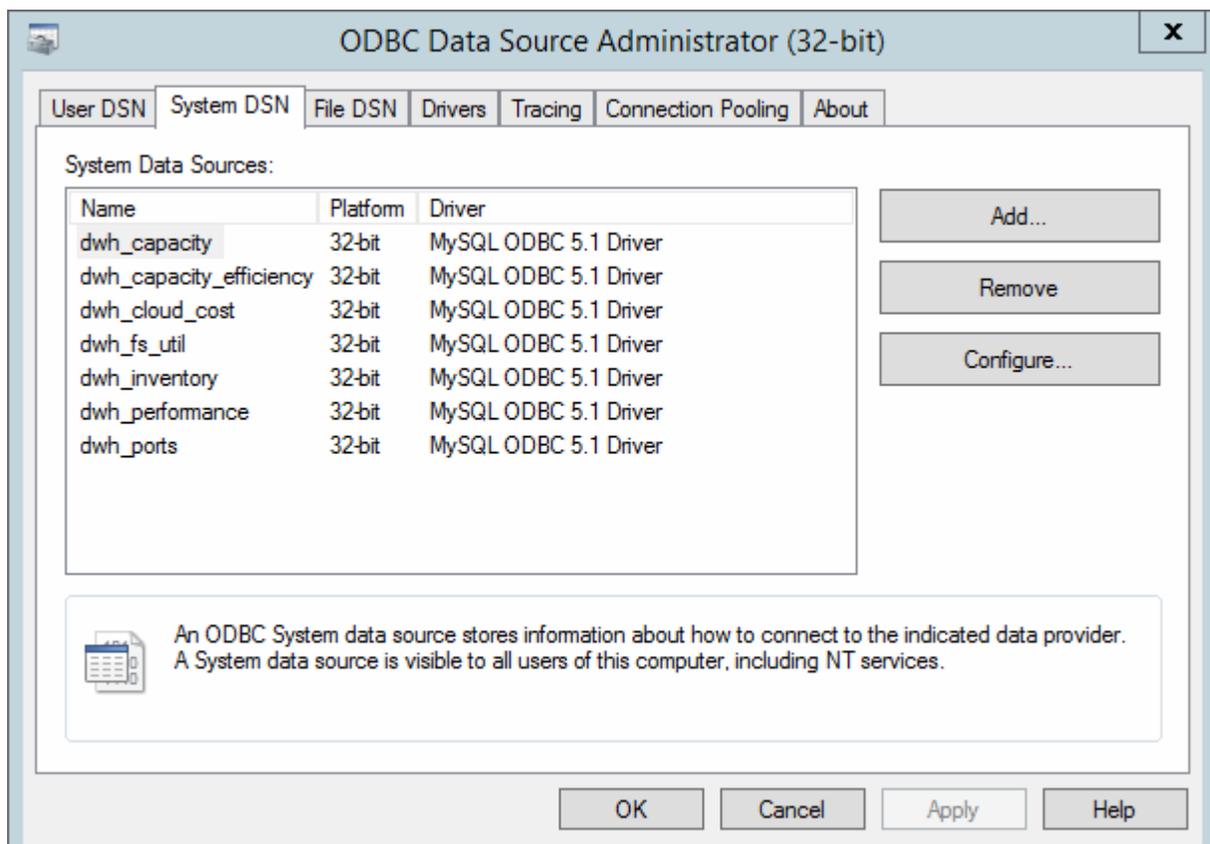
1. Perform a remote login to the server hosting that Data Warehouse.
2. Access the ODBC Administration tool at `C:\Windows\SysWOW64\odbcad32.exe`

The system displays the ODBC Data Source Administrator screen.



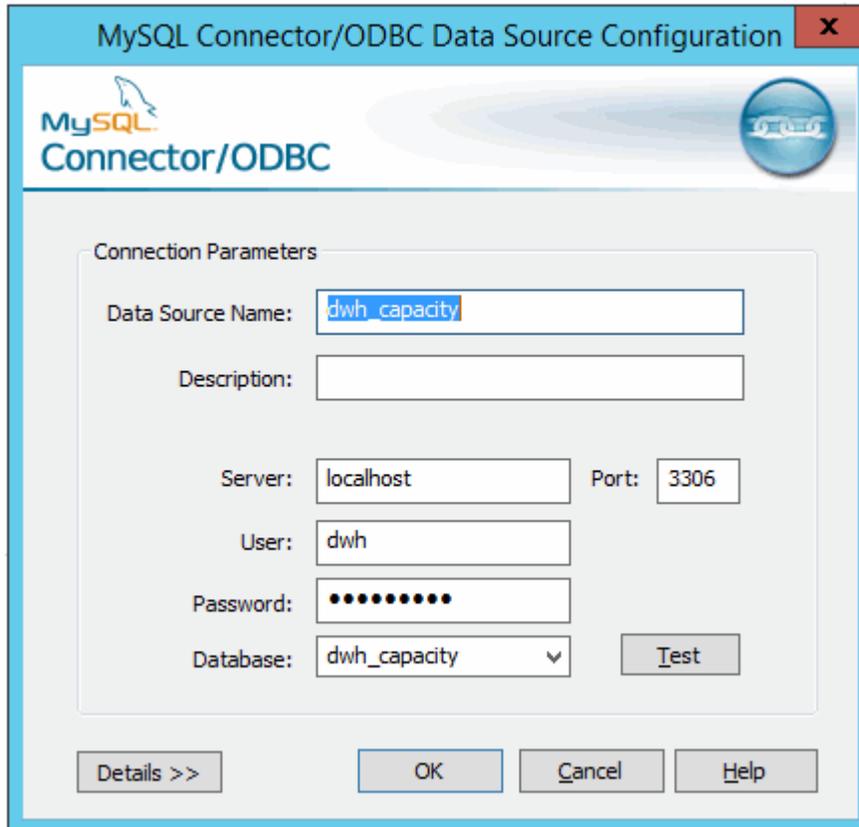
3. Click **System DSN**

The system data sources are displayed.



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

The Data Source Configuration screen is displayed.



The screenshot shows the 'MySQL Connector/ODBC Data Source Configuration' dialog box. The title bar includes the text 'MySQL Connector/ODBC Data Source Configuration' and a close button. The main window features the 'MySQL Connector/ODBC' logo and a 'Connection Parameters' section. The fields are as follows: 'Data Source Name' is 'dwh_capacity', 'Description' is empty, 'Server' is 'localhost', 'Port' is '3306', 'User' is 'dwh', 'Password' is masked with dots, and 'Database' is 'dwh_capacity'. A 'Test' button is located to the right of the 'Database' field. At the bottom of the dialog are buttons for 'Details >>', 'OK', 'Cancel', and 'Help'.

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

Smart Card and certificate login support

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

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



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

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



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

Configuring hosts for Smart Card and certificate login

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

Before you begin

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

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



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

Steps

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

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

2. Back up the keystore file: `C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore`
3. Open a command prompt specifying `Run as administrator`

4. **Delete the self-generated certificate:** `C:\Program Files\SANscreen\java64\bin\keytool.exe -delete -alias "ssl certificate" -keystore C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore`
5. **Generate a new certificate:** `C:\Program Files\SANscreen\java64\bin\keytool.exe -genkey -alias "alias_name" -keyalg RSA -sigalg SHA1withRSA -keysize 2048 -validity 365 -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -dname "CN=commonName,OU=orgUnit,O=orgName,L=localityNameI,S=stateName,C=countryName"`
6. **Generate a certificate signing request (CSR):** `C:\Program Files\SANscreen\java64\bin\keytool.exe -certreq -sigalg SHA1withRSA -alias "alias_name" -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -file C:\temp\server.csr"`
7. After the CSR is returned in step 6, import the certificate, then export the certificate in Base-64 format and place it in "C:\temp" named `servername.cer`.
8. **Extract the certificate from the keystore:** `C:\Program Files\SANscreen\java64\bin\keytool.exe -v -importkeystore -srckeystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -srcalias "alias_name" -destkeystore "C:\temp\file.p12" -deststoretype PKCS12`
9. **Extract a private key from the p12 file:** `openssl pkcs12 -in "C:\temp\file.p12" -out "C:\temp\servername.private.pem"`
10. **Merge the Base-64 certificate that you exported in step 7 with the private key:** `openssl pkcs12 -export -in "<folder>\<certificate>.cer" -inkey "C:\temp\servername.private.pem" -out "C:\temp\servername.new.p12" -name "servername.abc.123.yyy.zzz"`
11. **Import the merged certificate into the keystore:** `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -destkeystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -srckeystore "C:\temp\servername.new.p12" -srcstoretype PKCS12 -alias "alias_name"`
12. **Import the root certificate:** `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -file "C:\<root_certificate>.cer" -trustcacerts -alias "alias_name"`
13. **Import the root certificate into the server.trustore:** `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore" -file "C:\<email_certificate>.cer" -trustcacerts -alias "alias_name"`
14. **Import the intermediate certificate:** `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore "C:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore" -file "C:\<intermediate_certificate>.cer" -trustcacerts -alias "alias_name"`

Repeat this step for all intermediate certificates.

15. Specify the domain in LDAP to match this example.

1. Restart the server.

Configuring a client to support Smart Card and certificate login

Client machines require middleware and modifications to browsers to enable the use of Smart Cards and for certificate login. Customers who are already using Smart Cards should not require additional modifications to their client machines.

Before you begin

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

About this task

The following are the common client configuration requirements:

- Installing Smart Card middleware, such as ActivClient (see <http://militarycac.com/activclient.htm>)
- Modifying the IE browser (see http://militarycac.com/files/Making_AKO_work_with_Internet_Explorer_color.pdf)
- Modifying the Firefox browser (see <https://militarycac.com/firefox2.htm>)

Enabling CAC on a Linux server

Some modifications are required to enable CAC on a Linux OnCommand Insight server.

Steps

1. Navigate to `/opt/netapp/oci/conf/`
2. Edit `wildfly.properties` and change the value of `CLIENT_AUTH_ENABLED` to "True"
3. Import the "root certificate" that exists under `/opt/netapp/oci/wildfly/standalone/configuration/server.keystore`
4. Restart the server

Configuring Data Warehouse for Smart Card and certificate login

You must modify the OnCommand Insight Data Warehouse configuration to support Smart Card (CAC) and certificate logins.

Before you begin

- LDAP must be enabled on the system.
- The LDAP `User principal account name` attribute must match the LDAP field that contains a user's government ID number.

The common name (CN) stored on government-issued CACs is normally in the following format: `first.last.ID`. For some LDAP fields, such as `sAMAccountName`, this format is too long. For these fields, OnCommand Insight extracts only the ID number from the CNs.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Use `regedit` to modify registry values in `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Apache Software Foundation\Procrun2.0\SANscreen Server\Parameters\Java`

- a. Change the `JVM_Option -DclientAuth=false` to `-DclientAuth=true`.

For Linux, modify the `clientAuth` parameter in `/opt/netapp/oci/scripts/wildfly.server`

2. Add certificate authorities (CAs) to the Data Warehouse trustore:

- a. In a command window, go to `..\SANscreen\wildfly\standalone\configuration`.

- b. Use the `keytool` utility to list the trusted CAs: `C:\Program Files\SANscreen\java64\bin\keytool.exe -list -keystore server.trustore -storepass changeit`

The first word in each line indicates the CA alias.

- c. If necessary, supply a CA certificate file, usually a `.pem` file. To include customer's CAs with Data Warehouse trusted CAs go to `..\SANscreen\wildfly\standalone\configuration` and use the `keytool import` command: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore server.trustore -alias my_alias -file 'path/to/my.pem' -v -trustcacerts`

`my_alias` is usually an alias that would easily identify the CA in the `keytool -list` operation.

3. On the OnCommand Insight server, the `wildfly/standalone/configuration/standalone-full.xml` file needs to be modified by updating `verify-client` to "REQUESTED" in

/subsystem=undertow/server=default-server/https-listener=default-httpsto enable CAC. Log in to the Insight server and run the appropriate command:

OS	Script
Windows	<install dir>\SANscreen\wildfly\bin\enableCACforRemoteEJB.bat
Linux	/opt/netapp/oci/wildfly/bin/enableCACforRemoteEJB.sh

After executing the script, wait until the reload of the wildfly server is complete before proceeding to the next step.

4. Restart the OnCommand Insight server.

Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.5 through 7.3.9)

You must modify the OnCommand Insight Data Warehouse configuration to support Smart Card (CAC) and certificate logins for the Cognos server.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.5 through 7.3.9.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Add certificate authorities (CAs) to the Cognos truststore.
 - a. In a command window, go to `..\SANscreen\cognos\analytics\configuration\certs\`
 - b. Use the `keytool` utility to list the trusted CAs: `..\..\jre\bin\keytool.exe -list -keystore CAMKeystore.jks -storepass NoPassWordSet`

The first word in each line indicates the CA alias.

- c. If no suitable files exist, supply a CA certificate file, usually a .pem file.
- d. To include customer's CAs with OnCommand Insight trusted CAs, go to `..\SANscreen\cognos\analytics\configuration\certs\`.
- e. Use the `keytool` utility to import the .pem file: `..\..\jre\bin\keytool.exe -importcert -keystore CAMKeystore.jks -alias my_alias -file 'path/to/my.pem' -v -trustcacerts`

`my_alias` is usually an alias that would easily identify the CA in the `keytool -list` operation.

- f. When prompted for a password, enter `NoPassWordSet`.
 - g. Answer `yes` when prompted to trust the certificate.
2. To enable CAC mode, execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`
 3. To disable CAC mode, execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`

Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.10 and later)

You must modify the OnCommand Insight Data Warehouse configuration to support Smart Card (CAC) and certificate logins for the Cognos server.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.10 and later.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):

- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)



Steps

1. Add certificate authorities (CAs) to the Cognos truststore.
 - a. In a command window, go to `..\SANscreen\cognos\analytics\configuration\certs\`
 - b. Use the `keytool` utility to list the trusted CAs: `..\..\ibm-jre\jre\bin\keytool.exe -list -keystore CAMKeystore.jks -storepass NoPassWordSet`

The first word in each line indicates the CA alias.

- c. If no suitable files exist, supply a CA certificate file, usually a .pem file.
- d. To include customer's CAs with OnCommand Insight trusted CAs, go to
`..\SANscreen\cognos\analytics\configuration\certs\`.
- e. Use the `keytool` utility to import the .pem file: `..\..\ibm-jre\jre\bin\keytool.exe -importcert -keystore CAMKeystore.jks -alias my_alias -file 'path/to/my.pem' -v -trustcacerts`

`my_alias` is usually an alias that would easily identify the CA in the `keytool -list` operation.

- f. When prompted for a password, enter `NoPassWordSet`.
 - g. Answer `yes` when prompted to trust the certificate.
2. To enable CAC mode, do the following:
- a. Configure CAC logout page, using the following steps:
 - Logon to Cognos portal (user must be part of System Administrators group i.e. `cognos_admin`)
 - (Only for 7.3.10 and 7.3.11) Click Manage -> Configuration -> System -> Security
 - (Only for 7.3.10 and 7.3.11) Enter `cacLogout.html` against Logout Redirect URL -> Apply
 - Close browser.
 - b. Execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`
 - c. Start IBM Cognos service. Wait for Cognos service to start.
3. To disable CAC mode, do the following:
- a. Execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`
 - b. Start IBM Cognos service. Wait for Cognos service to start.
 - c. (Only for 7.3.10 and 7.3.11) Unconfigure CAC logout page, using the following steps:
 - Logon to Cognos portal (user must be part of System Administrators group i.e. `cognos_admin`)
 - Click Manage -> Configuration -> System -> Security
 - Enter `cacLogout.html` against Logout Redirect URL -> Apply
 - Close browser.

Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.5 to 7.3.9)

You can add SSL certificates to enable enhanced authentication and encryption for your Data Warehouse and Cognos environment.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.5 through 7.3.9.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

About this task

You must have admin privileges to perform this procedure.

Steps

1. Create a backup of `..\SANSscreen\cognos\analytics\configuration\cogstartup.xml`.
2. Create a backup of the “certs” and “csk” folders under `..\SANSscreen\cognos\analytics\configuration`.
3. Generate a Certificate Encryption Request from Cognos. In an Admin CMD window, run:
 - a. `cd “\Program Files\sansscreen\cognos\analytics\bin”`
 - b. `ThirdPartyCertificateTool.bat -java:local -c -e -p NoPassWordSet -a RSA -d “CN=FQDN,O=orgname,C=US” -r c:\temp\encryptRequest.csr`
4. Open the `c:\temp\encryptRequest.csr` file and copy the generated content.
5. Send the `encryptRequest.csr` to the certificate authority (CA) to obtain an SSL certificate.

Make sure to add additional attributes such as “SAN:dns=FQDN (For example, hostname.netapp.com)” to add the SubjectAltName. Google Chrome version 58 and later complains if the SubjectAltName is missing from the certificate.

6. Download the chain certificates by including root certificate by using PKCS7 format

This will download `fqdn.p7b` file

7. Get a cert in `.p7b` format from your CA. Use a name that marks it as the certificate for the Cognos Webserver.
8. `ThirdPartyCertificateTool.bat` fails to import the entire chain, so multiple steps are required to export all certificates. Split the chain by exporting them individually as follows:
 - a. Open the `.p7b` certificate in “Crypto Shell Extensions”.
 - b. Browse in the left pane to “Certificates”.
 - c. Right-click on root CA > All Tasks > Export.
 - d. Select Base64 output.

- e. Enter a file name identifying it as the root certificate.
 - f. Repeat steps 8a through 8c to export all of the certificates separately into .cer files.
 - g. Name the files intermediateX.cer and cognos.cer.
9. Ignore this step if you have only one CA certificate, otherwise merge both root.cer and intermediateX.cer into one file.
 - a. Open intermediate.cer with NotePad and copy the content.
 - b. Open root.cer with NotePad and save the content from 9a.
 - c. Save the file as CA.cer.
 10. Import the certificates into the Cognos keystore using the Admin CMD prompt:
 - a. cd "Program Files\sansscreen\cognos\analytics\bin"
 - b. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\CA.cer

This will set CA.cer as root Certificate Authority.
 - c. ThirdPartyCertificateTool.bat -java:local -i -e -r c:\temp\cognos.cer -t c:\temp\CA.cer

This will set Cognos.cer as encryption certificate which is signed by CA.cer.
 11. Open the IBM Cognos Configuration.
 - a. Select Local Configuration → Security → Cryptography → Cognos
 - b. Change "Use third party CA?" to True.
 - c. Save the configuration.
 - d. Restart Cognos
 12. Export the latest Cognos certificate into cognos.crt using the Admin CMD prompt:
 - a. "D:\Program Files\SANscreen\java\bin\keytool.exe" -exportcert -file "c:\temp\cognos.crt" -keystore "D:\Program Files\SANscreen\cognos\analytics\configuration\certs\CAMKeystore" -storetype PKCS12 -storepass NoPassWordSet -alias encryption
 13. Import the "c:\temp\cognos.crt" into dwh trustore to establish SSL communication between Cognos and DWH, using the Admin CMD prompt window.
 - a. "D:\Program Files\SANscreen\java\bin\keytool.exe" -importcert -file "c:\temp\cognos.crt" -keystore "D:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore" -storepass changeit -alias cognoscert
 14. Restart the SANscreen service.
 15. Perform a backup of DWH to make sure DWH communicates with Cognos.

Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.10 and later)

You can add SSL certificates to enable enhanced authentication and encryption for your Data Warehouse and Cognos environment.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.10 and later.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

About this task

You must have admin privileges to perform this procedure.

Steps

1. Stop Cognos using the IBM Cognos Configuration tool. Close Cognos.
2. Create backups of the `..\SANSscreen\cognos\analytics\configuration` and `..\SANSscreen\cognos\analytics\temp\cam\freshness` folders.
3. Generate a Certificate Encryption Request from Cognos. In an Admin CMD window, run:
 - a. `cd "\\Program Files\sansscreen\cognos\analytics\bin"`
 - b. `ThirdPartyCertificateTool.bat -java:local -c -e -p NoPassWordSet -a RSA -r c:\temp\encryptRequest.csr -d "CN=server.domain.com,O=NETAPP,C=US" -H "server.domain.com" -I "ipaddress"`. Note: here -H and -I are to add subjectAltNames like dns and ipaddress.
4. Open the `c:\temp\encryptRequest.csr` file and copy the generated content.
5. Input the `encryptRequest.csr` content and generate certificate using CA signing portal.
6. Download the chain certificates by including root certificate by using PKCS7 format

This will download `fqdn.p7b` file

7. Get a cert in `.p7b` format from your CA. Use a name that marks it as the certificate for the Cognos Webserver.
8. `ThirdPartyCertificateTool.bat` fails to import the entire chain, so multiple steps are required to export all certificates. Split the chain by exporting them individually as follows:
 - a. Open the `.p7b` certificate in "Crypto Shell Extensions".
 - b. Browse in the left pane to "Certificates".
 - c. Right-click on root CA > All Tasks > Export.
 - d. Select Base64 output.
 - e. Enter a file name identifying it as the root certificate.
 - f. Repeat steps 8a through 8e to export all of the certificates separately into `.cer` files.

- g. Name the files intermediateX.cer and cognos.cer.
9. Ignore this step if you have only one CA certificate, otherwise merge both root.cer and intermediateX.cer into one file.
 - a. Open root.cer with NotePad and copy the content.
 - b. Open intermediate.cer with NotePad and append the content from 9a (intermediate first and root next).
 - c. Save the file as chain.cer.
10. Import the certificates into the Cognos keystore using the Admin CMD prompt:
 - a. cd "Program Files\sanscreen\cognos\analytics\bin"
 - b. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\root.cer
 - c. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\intermediate.cer
 - d. ThirdPartyCertificateTool.bat -java:local -i -e -r c:\temp\cognos.cer -t c:\temp\chain.cer
11. Open the IBM Cognos Configuration.
 - a. Select Local Configuration → Security → Cryptography → Cognos
 - b. Change "Use third party CA?" to True.
 - c. Save the configuration.
 - d. Restart Cognos
12. Export the latest Cognos certificate into cognos.crt using the Admin CMD prompt:
 - a. cd "C:\Program Files\SANscreen"
 - b. java\bin\keytool.exe -exportcert -file c:\temp\cognos.crt -keystore cognos\analytics\configuration\certs\CAMKeystore -storetype PKCS12 -storepass NoPassWordSet -alias encryption
13. Back up the DWH server trustore


```
at. .\SANscreen\wildfly\standalone\configuration\server.trustore
```
14. Import the "c:\temp\cognos.crt" into DWH trustore to establish SSL communication between Cognos and DWH, using the Admin CMD prompt window.
 - a. cd "C:\Program Files\SANscreen"
 - b. java\bin\keytool.exe -importcert -file c:\temp\cognos.crt -keystore wildfly\standalone\configuration\server.trustore -storepass changeit -alias cognos3rdca
15. Restart the SANscreen service.
16. Perform a backup of DWH to make sure DWH communicates with Cognos.
17. The following steps should be performed even when only the "ssl certificate" is changed and the default Cognos certificates are left unchanged. Otherwise Cognos may complain about the new SANscreen certificate or be unable to create a DWH backup.
 - a. cd "%SANSSCREEN_HOME%cognos\analytics\bin\"
 - b. "%SANSSCREEN_HOME%java64\bin\keytool.exe" -exportcert -file "c:\temp\sansscreen.cer" -keystore "%SANSSCREEN_HOME%wildfly\standalone\configuration\server.keystore" -storepass changeit -alias "ssl certificate"
 - c. ThirdPartyCertificateTool.bat -java:local -i -T -r "c:\temp\sansscreen.cer"

Typically, these steps are performed as part of the Cognos certificate import process described in [How to](#)

Configuring Data Warehouse for Smart Card and certificate login

You must modify the OnCommand Insight Data Warehouse configuration to support Smart Card (CAC) and certificate logins.

Before you begin

- LDAP must be enabled on the system.
- The LDAP `User principal account name` attribute must match the LDAP field that contains a user's government ID number.

The common name (CN) stored on government-issued CACs is normally in the following format: `first.last.ID`. For some LDAP fields, such as `sAMAccountName`, this format is too long. For these fields, OnCommand Insight extracts only the ID number from the CNs.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):

- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)



Steps

1. Use `regedit` to modify registry values in `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Apache Software Foundation\Procrun2.0\SANscreen Server\Parameters\Java`
 - a. Change the `JVM_Option -DclientAuth=false` to `-DclientAuth=true`.
For Linux, modify the `clientAuth` parameter in `/opt/netapp/oci/scripts/wildfly.server`
2. Add certificate authorities (CAs) to the Data Warehouse trustore:
 - a. In a command window, go to `..\SANscreen\wildfly\standalone\configuration`.
 - b. Use the `keytool` utility to list the trusted CAs: `C:\Program Files\SANscreen\java64\bin\keytool.exe -list -keystore server.trustore -storepass changeit`

The first word in each line indicates the CA alias.

c. If necessary, supply a CA certificate file, usually a .pem file. To include customer's CAs with Data Warehouse trusted CAs go to ..\SANscreen\wildfly\standalone\configuration and use the keytool import command: C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -keystore server.trustore -alias my_alias -file 'path/to/my.pem' -v -trustcacerts

my_alias is usually an alias that would easily identify the CA in thekeytool -list operation.

3. On the OnCommand Insight server, the wildfly/standalone/configuration/standalone-full.xml file needs to be modified by updating verify-client to "REQUESTED" in /subsystem=undertow/server=default-server/https-listener=default-httpsto enable CAC. Log in to the Insight server and run the appropriate command:

OS	Script
Windows	<install dir>\SANscreen\wildfly\bin\enableCACforRemoteEJB.bat
Linux	/opt/netapp/oci/wildfly/bin/enableCACforRemoteEJB.sh

After executing the script, wait until the reload of the wildfly server is complete before proceeding to the next step.

4. Restart the OnCommand Insight server.

Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.5 through 7.3.9)

You must modify the OnCommand Insight Data Warehouse configuration to support Smart Card (CAC) and certificate logins for the Cognos server.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.5 through 7.3.9.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Add certificate authorities (CAs) to the Cognos truststore.
 - a. In a command window, go to `..\SANscreen\cognos\analytics\configuration\certs\`
 - b. Use the `keytool` utility to list the trusted CAs: `..\..\jre\bin\keytool.exe -list -keystore CAMKeystore.jks -storepass NoPassWordSet`

The first word in each line indicates the CA alias.
 - c. If no suitable files exist, supply a CA certificate file, usually a `.pem` file.
 - d. To include customer's CAs with OnCommand Insight trusted CAs, go to `..\SANscreen\cognos\analytics\configuration\certs\`.
 - e. Use the `keytool` utility to import the `.pem` file: `..\..\jre\bin\keytool.exe -importcert -keystore CAMKeystore.jks -alias my_alias -file 'path/to/my.pem' -v -trustcacerts`

`my_alias` is usually an alias that would easily identify the CA in the `keytool -list` operation.
 - f. When prompted for a password, enter `NoPassWordSet`.
 - g. Answer `yes` when prompted to trust the certificate.
2. To enable CAC mode, execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`
3. To disable CAC mode, execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`

Configuring Cognos for Smart Card and certificate login (OnCommand Insight 7.3.10 and later)

You must modify the OnCommand Insight Data Warehouse configuration to support Smart Card (CAC) and certificate logins for the Cognos server.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.10 and later.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Steps

1. Add certificate authorities (CAs) to the Cognos truststore.

- In a command window, go to `..\SANscreen\cognos\analytics\configuration\certs\`
- Use the `keytool` utility to list the trusted CAs: `..\..\ibm-jre\jre\bin\keytool.exe -list -keystore CAMKeystore.jks -storepass NoPassWordSet`

The first word in each line indicates the CA alias.

- If no suitable files exist, supply a CA certificate file, usually a `.pem` file.
- To include customer's CAs with OnCommand Insight trusted CAs, go to `..\SANscreen\cognos\analytics\configuration\certs\`.
- Use the `keytool` utility to import the `.pem` file: `..\..\ibm-jre\jre\bin\keytool.exe -importcert -keystore CAMKeystore.jks -alias my_alias -file 'path/to/my.pem' -v -trustcacerts`

`my_alias` is usually an alias that would easily identify the CA in the `keytool -list` operation.

- When prompted for a password, enter `NoPassWordSet`.
- Answer `yes` when prompted to trust the certificate.

2. To enable CAC mode, do the following:

- Configure CAC logout page, using the following steps:
 - Logon to Cognos portal (user must be part of System Administrators group i.e. `cognos_admin`)
 - (Only for 7.3.10 and 7.3.11) Click Manage -> Configuration -> System -> Security
 - (Only for 7.3.10 and 7.3.11) Enter `cacLogout.html` against Logout Redirect URL -> Apply
 - Close browser.
- Execute `..\SANscreen\bin\cognos_cac\enableCognosCAC.bat`
- Start IBM Cognos service. Wait for Cognos service to start.

3. To disable CAC mode, do the following:

- Execute `..\SANscreen\bin\cognos_cac\disableCognosCAC.bat`

- b. Start IBM Cognos service. Wait for Cognos service to start.
- c. (Only for 7.3.10 and 7.3.11) Unconfigure CAC logout page, using the following steps:
 - Logon to Cognos portal (user must be part of System Administrators group i.e. cognos_admin)
 - Click Manage -> Configuration -> System -> Security
 - Enter cacLogout.html against Logout Redirect URL -> Apply
 - Close browser.

Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.5 to 7.3.9)

You can add SSL certificates to enable enhanced authentication and encryption for your Data Warehouse and Cognos environment.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.5 through 7.3.9.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

About this task

You must have admin privileges to perform this procedure.

Steps

1. Create a backup of `..\SANSscreen\cognos\analytics\configuration\cogstartup.xml`.
2. Create a backup of the “certs” and “csk” folders under `..\SANSscreen\cognos\analytics\configuration`.
3. Generate a Certificate Encryption Request from Cognos. In an Admin CMD window, run:
 - a. `cd “\Program Files\sansscreen\cognos\analytics\bin”`
 - b. `ThirdPartyCertificateTool.bat -java:local -c -e -p NoPassWordSet -a RSA -d “CN=FQDN,O=orgname,C=US” -r c:\temp\encryptRequest.csr`
4. Open the `c:\temp\encryptRequest.csr` file and copy the generated content.

5. Send the encryptRequest.csr to the certificate authority (CA) to obtain an SSL certificate.

Make sure to add additional attributes such as "SAN:dns=FQDN (For example, hostname.netapp.com)" to add the SubjectAltName. Google Chrome version 58 and later complains if the SubjectAltName is missing from the certificate.

6. Download the chain certificates by including root certificate by using PKCS7 format

This will download fqdn.p7b file

7. Get a cert in .p7b format from your CA. Use a name that marks it as the certificate for the Cognos Webserver.
8. ThirdPartyCertificateTool.bat fails to import the entire chain, so multiple steps are required to export all certificates. Split the chain by exporting them individually as follows:
 - a. Open the .p7b certificate in "Crypto Shell Extensions".
 - b. Browse in the left pane to "Certificates".
 - c. Right-click on root CA > All Tasks > Export.
 - d. Select Base64 output.
 - e. Enter a file name identifying it as the root certificate.
 - f. Repeat steps 8a through 8c to export all of the certificates separately into .cer files.
 - g. Name the files intermediateX.cer and cognos.cer.
9. Ignore this step if you have only one CA certificate, otherwise merge both root.cer and intermediateX.cer into one file.
 - a. Open intermediate.cer with NotePad and copy the content.
 - b. Open root.cer with NotePad and save the content from 9a.
 - c. Save the file as CA.cer.
10. Import the certificates into the Cognos keystore using the Admin CMD prompt:
 - a. cd "Program Files\sansscreen\cognos\analytics\bin"
 - b. ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\CA.cer

This will set CA.cer as root Certificate Authority.
 - c. ThirdPartyCertificateTool.bat -java:local -i -e -r c:\temp\cognos.cer -t c:\temp\CA.cer

This will set Cognos.cer as encryption certificate which is signed by CA.cer.
11. Open the IBM Cognos Configuration.
 - a. Select Local Configuration → Security → Cryptography → Cognos
 - b. Change "Use third party CA?" to True.
 - c. Save the configuration.
 - d. Restart Cognos
12. Export the latest Cognos certificate into cognos.crt using the Admin CMD prompt:
 - a. "D:\Program Files\SANscreen\java\bin\keytool.exe" -exportcert -file "c:\temp\cognos.crt" -keystore "D:\Program Files\SANscreen\cognos\analytics\configuration\certs\CAMKeystore" -storetype PKCS12 -storepass NoPassWordSet -alias encryption

13. Import the "c:\temp\cognos.crt" into dwh trustore to establish SSL communication between Cognos and DWH, using the Admin CMD prompt window.
 - a. "D:\Program Files\SANscreen\java\bin\keytool.exe" -importcert -file "c:\temp\cognos.crt" -keystore "D:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore" -storepass changeit -alias cognoscert
14. Restart the SANscreen service.
15. Perform a backup of DWH to make sure DWH communicates with Cognos.

Importing CA-signed SSL certificates for Cognos and DWH (Insight 7.3.10 and later)

You can add SSL certificates to enable enhanced authentication and encryption for your Data Warehouse and Cognos environment.

Before you begin

This procedure is for systems running OnCommand Insight 7.3.10 and later.

For the most up to date CAC and Certificate instructions, see the following Knowledgebase articles (Support login required):



- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight](#)
- [How to configure Common Access Card \(CAC\) authentication for OnCommand Insight Data Warehouse](#)
- [How to create and import a Certificate Authority \(CA\) signed certificate into OnCommand Insight and OnCommand Insight Data Warehouse 7.3.x](#)
- [How to create a Self Signed Certificate within OnCommand Insight 7.3.X installed on a Windows Host](#)
- [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

About this task

You must have admin privileges to perform this procedure.

Steps

1. Stop Cognos using the IBM Cognos Configuration tool. Close Cognos.
2. Create backups of the `..\SANScreen\cognos\analytics\configuration` and `..\SANScreen\cognos\analytics\temp\cam\freshness` folders.
3. Generate a Certificate Encryption Request from Cognos. In an Admin CMD window, run:
 - a. `cd "\Program Files\sanscreen\cognos\analytics\bin"`
 - b. `ThirdPartyCertificateTool.bat -java:local -c -e -p NoPassWordSet -a RSA -r c:\temp\encryptRequest.csr -d "CN=server.domain.com,O=NETAPP,C=US" -H "server.domain.com" -I "ipaddress".` Note: here -H and -I are to add subjectAltNames like dns and ipaddress.

4. Open the `c:\temp\encryptRequest.csr` file and copy the generated content.
5. Input the `encryptRequest.csr` content and generate certificate using CA signing portal.
6. Download the chain certificates by including root certificate by using PKCS7 format

This will download `fqdn.p7b` file

7. Get a cert in `.p7b` format from your CA. Use a name that marks it as the certificate for the Cognos Webserver.
8. `ThirdPartyCertificateTool.bat` fails to import the entire chain, so multiple steps are required to export all certificates. Split the chain by exporting them individually as follows:
 - a. Open the `.p7b` certificate in “Crypto Shell Extensions”.
 - b. Browse in the left pane to “Certificates”.
 - c. Right-click on root CA > All Tasks > Export.
 - d. Select Base64 output.
 - e. Enter a file name identifying it as the root certificate.
 - f. Repeat steps 8a through 8e to export all of the certificates separately into `.cer` files.
 - g. Name the files `intermediateX.cer` and `cognos.cer`.
9. Ignore this step if you have only one CA certificate, otherwise merge both `root.cer` and `intermediateX.cer` into one file.
 - a. Open `root.cer` with NotePad and copy the content.
 - b. Open `intermediate.cer` with NotePad and append the content from 9a (intermediate first and root next).
 - c. Save the file as `chain.cer`.
10. Import the certificates into the Cognos keystore using the Admin CMD prompt:
 - a. `cd "Program Files\sansscreen\cognos\analytics\bin"`
 - b. `ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\root.cer`
 - c. `ThirdPartyCertificateTool.bat -java:local -i -T -r c:\temp\intermediate.cer`
 - d. `ThirdPartyCertificateTool.bat -java:local -i -e -r c:\temp\cognos.cer -t c:\temp\chain.cer`
11. Open the IBM Cognos Configuration.
 - a. Select Local Configuration → Security → Cryptography → Cognos
 - b. Change “Use third party CA?” to True.
 - c. Save the configuration.
 - d. Restart Cognos
12. Export the latest Cognos certificate into `cognos.crt` using the Admin CMD prompt:
 - a. `cd "C:\Program Files\SANscreen"`
 - b. `java\bin\keytool.exe -exportcert -file c:\temp\cognos.crt -keystore cognos\analytics\configuration\certs\CAMKeystore -storetype PKCS12 -storepass NoPassWordSet -alias encryption`
13. Back up the DWH server trustore


```
at .\SANscreen\wildfly\standalone\configuration\server.trustore
```
14. Import the “`c:\temp\cognos.crt`” into DWH trustore to establish SSL communication between Cognos and DWH, using the Admin CMD prompt window.

- a. `cd "C:\Program Files\SANscreen"`
 - b. `java\bin\keytool.exe -importcert -file c:\temp\cognos.crt -keystore wildfly\standalone\configuration\server.trustore -storepass changeit -alias cognos3rdca`
15. Restart the SANscreen service.
16. Perform a backup of DWH to make sure DWH communicates with Cognos.
17. The following steps should be performed even when only the “ssl certificate” is changed and the default Cognos certificates are left unchanged. Otherwise Cognos may complain about the new SANscreen certificate or be unable to create a DWH backup.
- a. `cd "%SANSSCREEN_HOME%cognos\analytics\bin\"`
 - b. `"%SANSSCREEN_HOME%java64\bin\keytool.exe" -exportcert -file "c:\temp\sansscreen.cer" -keystore "%SANSSCREEN_HOME%wildfly\standalone\configuration\server.keystore" -storepass changeit -alias "ssl certificate"`
 - c. `ThirdPartyCertificateTool.bat -java:local -i -T -r "c:\temp\sansscreen.cer"`

Typically, these steps are performed as part of the Cognos certificate import process described in [How to import a Cognos Certificate Authority \(CA\) signed certificate into OnCommand DataWarehouse 7.3.3 and later](#)

Importing SSL certificates

You can add SSL certificates to enable enhanced authentication and encryption for enhancing the security of your OnCommand Insight environment.

Before you begin

You must ensure that your system meets the minimum required bit level (1024 bits).

About this task



Before you attempt to perform this procedure, you should back up the existing `server.keystore` file, and name the backup `server.keystore.old`. Corrupting or damaging the `server.keystore` file may result in an inoperable Insight server after the Insight server is restarted. If you create a backup, you can revert to the old file if problems occur.

Steps

1. Create a copy of the original keystore file: `cp c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore.old"`
2. List the contents of the keystore: `C:\Program Files\SANscreen\java64\bin\keytool.exe -list -v -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore"`
 - a. When prompted for a password, enter `changeit`.

The system displays the contents of the keystore. There should be at least one certificate in the keystore,

"ssl certificate".

3. Delete the "ssl certificate": `keytool -delete -alias "ssl certificate" -keystore c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore`
4. Generate a new key: `C:\Program Files\SANscreen\java64\bin\keytool.exe -genkey -alias "ssl certificate" -keyalg RSA -keysize 2048 -validity 365 -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore"`
 - a. When prompted for first and last names, enter the fully qualified domain name (FQDN) that you intend to use.
 - b. Provide the following information about your organization and organizational structure:
 - Country: two-letter ISO abbreviation for your country (for example, US)
 - State or Province: name of the state or province where your organization's head office is located (for example, Massachusetts)
 - Locality: name of the city where your organization's head office is located (for example, Waltham)
 - Organizational name: name of the organization that owns the domain name (for example, NetApp)
 - Organizational unit name: name of the department or group that will use the certificate (for example, Support)
 - Domain Name/ Common Name: the FQDN that is used for DNS lookups of your server (for example, www.example.com)
The system responds with information similar to the following: Is CN=www.example.com, OU=support, O=NetApp, L=Waltham, ST=MA, C=US correct?
 - c. Enter Yes when the Common Name (CN) is equal to the FQDN.
 - d. When prompted for the key password, enter the password, or press the Enter key to use the existing keystore password.
5. Generate a certificate request file: `C:\Program Files\SANscreen\java64\bin\keytool.exe -certreq -alias "ssl certificate" -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore" -file c:\localhost.csr`

The `c:\localhost.csr` file is the certificate request file that is newly generated.

6. Submit the `c:\localhost.csr` file to your certificate authority (CA) for approval.

Once the certificate request file is approved, you want the certificate returned to you in `.der` format. The file might or might not be returned as a `.der` file. The default file format is `.cer` for Microsoft CA services.

Most organizations' CAs use a chain of trust model, including a root CA, which is often offline. It has signed the certificates for only a few child CAs, known as intermediate CAs.

You must obtain the public key (certificates) for the entire chain of trust—the certificate for the CA that signed the certificate for the OnCommand Insight server, and all the certificates between that signing CA up to and including the organizational root CA.

In some organizations, when you submit a signing request, you might receive one of the following:

- A PKCS12 file that contains your signed certificate and all the public certificates in the chain of trust
- A `.zip` file that contains individual files (including your signed certificate) and all the public certificates

in the chain of trust

- Only your signed certificate

You must obtain the public certificates.

7. Import the approved certificate for server.keystore: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -alias OCI.hostname.com -file c:\localhost2.DER -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.keystore"`

- a. When prompted, enter the keystore password.

The following message is displayed: Certificate reply was installed in keystore

8. Import the approved certificate for server.trustore: `C:\Program Files\SANscreen\java64\bin\keytool.exe -importcert -alias OCI.hostname.com -file c:\localhost2.DER -keystore "c:\Program Files\SANscreen\wildfly\standalone\configuration\server.trustore"`

- a. When prompted, enter the trustore password.

The following message is displayed: Certificate reply was installed in trustore

9. Edit the `SANscreen\wildfly\standalone\configuration\standalone-full.xml` file:

Substitute the following alias string: `alias="cbc-oci-02.muccbc.hq.netapp.com"`. For example:

```
<keystore path="server.keystore" relative-to="jboss.server.config.dir"
keystore-password="{VAULT::HttpsRealm::keystore_password::1}" alias="cbc-oci-
02.muccbc.hq.netapp.com" key-
password="{VAULT::HttpsRealm::key_password::1}"/>
```

10. Restart the SANscreen server service.

Once Insight is running, you can click the padlock icon to view the certificates that are installed on the system.

If you see a certificate containing "Issued To" information that matches "Issued By" information, you still have a self-signed certificate installed. The Insight installer-generated self-signed certificates have a 100-year expiration.

NetApp cannot guarantee that this procedure will remove digital certificate warnings. NetApp cannot control how your end user workstations are configured. Consider the following scenarios:

- Microsoft Internet Explorer and Google Chrome both utilize Microsoft's native certificate functionality on Windows.

This means that if your Active Directory administrators push your organization's CA certificates into the end user's certificate trustores, the users of these browsers will see certificate warnings disappear when the OnCommand Insight self-signed certificates have been replaced with the one signed by the internal CA infrastructure.

- Java and Mozilla Firefox have their own certificate stores.

If your system administrators do not automate ingesting the CA certificates into these applications'

trusted certificates stores, using the Firefox browser might continue to generate certificate warnings because of an untrusted certificate, even when the self-signed certificate has been replaced. Getting your organization's certificate chain installed into the trustore is an additional requirement.

Your business entities hierarchy

You can define business entities to track and report on your environment data at a more granular level.

In OnCommand Insight, the business entities hierarchy contains these levels:

- **Tenant** is primarily used by service providers to associate resources with a customer, for example, NetApp.
- **Line of Business (LOB)** is a line of business or product line within a company, for example, Data Storage.
- **Business Unit** represents a traditional business unit such as Legal or Marketing.
- **Project** is often used to identify a specific project within a business unit for which you want capacity chargeback. For example, "Patents" might be a project name for the Legal business unit and "Sales Events" might be a project name for the Marketing business unit. Note that level names may include spaces.

You are not required to use all of the levels in the design of your corporate hierarchy.

Designing your business entities hierarchy

You need to understand the elements of your corporate structure and what needs to be represented in the business entities because they become a fixed structure in your OnCommand Insight database. You can use the following information to set up your business entities. Remember you do not need to use all of the hierarchy levels to gather data in these categories.

Steps

1. Examine each level of the business entities hierarchy to determine if that level should be included in your business entity hierarchy for your company:
 - **Tenant** level is needed if your company is an ISP and you want to track customer usage of resources.
 - **Line of Business (LOB)** is needed in the hierarchy if the data for different product lines needs to be tracked.
 - **Business Unit** is required if you need to track data for different departments. This level of the hierarchy is often valuable in separating a resource that one department uses that other departments do not.
 - **Project** level can be used for specialized work within a department. This data might be useful to pinpoint, define, and monitor a separate project's technology needs compared to other projects in a company or department.
2. Create a chart showing each business entity with the names of all of the levels within the entity.
3. Check the names in the hierarchy to be certain they will be self-explanatory in OnCommand Insight views and reports.
4. Identify all applications that are associated with each business entity.

Creating business entities

After designing the business entities hierarchy for your company, you can set up applications and then associate the business entities with the applications. This process creates the business entities structure in your OnCommand Insight database.

About this task

Associating applications with business entities is optional; however, it is a best practice.

Steps

1. Log in to the Insight web UI.
2. Click **Manage** and select **Business entities**.

The Business Entities page displays.

3. Click  to begin building a new entity.

The **Add Business Entity** dialog box displays.

4. For each entity level (Tenant, Line of Business, Business Unit, and Project), you can do any of the following:
 - Click the entity level list and select a value.
 - Type a new value and press Enter.
 - Leave the entity level value as N/A if you do not want to use the entity level for the business entity.
5. Click **Save**.

Assigning business entities to assets

You can assign a business entity to an asset (host, port, storage, switch, virtual machine, qtree, share, volume, or internal volume) without having associated the business entity to an application; however, business entities are assigned automatically to an asset if that asset is associated with an application related to a business entity.

Before you begin

You must have already created a business entity.

About this task

While you can assign business entities directly to assets, it is recommended that you assign applications to assets and then assign business entities to assets.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate the asset to which you want to apply the business entity by doing either of the following:
 - Click on the asset in the Assets Dashboard.

- Click  on the toolbar to display the **Search assets** box, type the name of the asset, and then select the asset from the list.
3. In the **User Data** section of the asset page, position your cursor over **None** next to **Business Entities** and then click .

The list of available business entities display.

4. Type in the **Search** box to filter the list for a specific entity or scroll down the list; select a business entity from the list.

If the business entity you choose is associated with an application, the application name is displayed. In this case, the word “derived” appears next to the business entity name. If you want to maintain the entity for only the asset and not the associated application, you can manually override the assignment of the application.

5. To override an application derived from a business entity, place your cursor over the application name and click , select another business entity, and select another application from the list.

Assigning business entities to or removing business entities from multiple assets

You can assign business entities to or remove business entities from multiple assets by using a query instead of having to manually assign or remove them.

Before you begin

You must have already created the business entities you want to add to your desired assets.

Steps

1. Create a new query, or open an existing query.
2. If desired, filter for the assets to which you want to add business entities.
3. Select the desired assets in the list or click ▼ to select **All**.

The **Actions** button displays.

4. To add a business entity to the selected assets, click . If the selected asset type can have business entities assigned to it, you will see the menu choice to **Add Business Entity**. Select this.
5. Select the desired business entity from the list and click **Save**.

Any new business entity you assign overrides any business entities that were already assigned to the asset. Assigning applications to assets will also override the business entities assigned in the same way. Assigning business entities to an asset may also override any applications assigned to that asset.

6. To remove a business entity assigned to the assets, click  and select **Remove Business Entity**.
7. Select the desired business entity from the list and click **Delete**.

Defining annotations

When customizing OnCommand Insight to track data for your corporate requirements, you can define any specialized annotations needed to provide a complete picture of your data: for example, asset end of life, data center, building location, storage tier, or volume, and internal volume service level.

Steps

1. List any industry terminology to which environment data must be associated.
2. List corporate terminology to which environment data must be associated, which is not already being tracked using the business entities.
3. Identify any default annotation types that you might be able to use.
4. Identify which custom annotations you need to create.

Using annotations to monitor your environment

When customizing OnCommand Insight to track data for your corporate requirements, you can define specialized notes, called *annotations*, and assign them to your assets. For example, you can annotate assets with information such as asset end of life, data center, building location, storage tier, or volume service level.

Using annotations to help monitor your environment includes the following high-level tasks:

- Creating or edit definitions for all annotation types.
- Displaying asset pages and associating each asset with one or more annotations.

For example, if an asset is being leased and the lease expires within two months, you might want to apply an end-of-life annotation to the asset. This helps prevent others from using that asset for an extended time.

- Creating rules to automatically apply annotations to multiple assets of the same type.
- Using the annotation import utility to import annotations.
- Filter assets by their annotations.
- Grouping data in reports based on annotations and generate those reports.

See the *OnCommand Insight Reporting Guide* for more information about reports.

Managing annotation types

OnCommand Insight provides some default annotation types, such as asset life cycle (birthday or end of life), building or data center location, and tier, that you can customize to show in your reports. You can define values for default annotation types or create your own custom annotation types. You can later edit those values.

Default annotation types

OnCommandInsight provides some default annotation types. These annotations can be

used to filter or group data and to filter data reporting.

You can associate assets with default annotation types such as the following:

- Asset life cycle, such as birthday, sunset, or end of life
- Location information about a device, such as data center, building, or floor
- Classification of assets, such as by quality (tiers), by connected devices (switch level), or by service level
- Status, such as hot (high utilization)

The following table lists the default annotation types. You can edit any of these annotation names to suit your needs.

Annotation types	Description	Type
Alias	User-friendly name for a resource.	Text
Birthday	Date when the device was or will be brought online.	Date
Building	Physical location of host, storage, switch, and tape resources.	List
City	Municipality location of host, storage, switch, and tape resources.	List
Compute Resource Group	Group assignment used by the Host and VM Filesystems data source.	List
Continent	Geographic location of host, storage, switch, and tape resources.	List
Country	National location of host, storage, switch, and tape resources.	List
Data Center	Physical location of the resource and is available for hosts, storage arrays, switches, and tapes.	List
Direct Attached	Indicates (Yes or No) if a storage resource is connected directly to hosts.	Boolean
End of Life	Date when a device will be taken offline, for example, if the lease expired or the hardware is being retired.	Date

Fabric Alias	User-friendly name for a fabric.	Text
Floor	Location of a device on a floor of a building. Can be set for hosts, storage arrays, switches, and tapes.	List
Hot	Devices already in heavy use on a regular basis or at the threshold of capacity.	Boolean
Note	Comments that you want associated with a resource.	Text
Rack	Rack in which the resource resides.	Text
Room	Room within a building or other location of host, storage, switch, and tape resources.	List
SAN	Logical partition of the network. Available on hosts, storage arrays, tapes, switches, and applications.	List
Service Level	A set of supported service levels that you can assign to resources. Provides an ordered options list for internal volumes, qtree, and volumes. Edit service levels to set performance policies for different levels.	List
State/Province	State or province in which the resource is located.	List
Sunset	Threshold set after which no new allocations can be made to that device. Useful for planned migrations and other pending network changes.	Date
Switch Level	Includes predefined options for setting up categories for switches. Typically, these designations remain for the life of the device, although you can edit them, if needed. Available only for switches.	List

Tier	Can be used to define different levels of service within your environment. Tiers can define the type of level, such as speed needed (for example, gold or silver). This feature is available only on internal volumes, qtrees, storage arrays, storage pools, and volumes.	List
Violation Severity	Rank (for example, major) of a violation (for example, missing host ports or missing redundancy), in a hierarchy of highest to lowest importance.	List



Alias, Data Center, Hot, Service Level, Sunset, Switch Level, Service Level, Tier, and Violation Severity are system-level annotations, which you cannot delete or rename; you can change only their assigned values.

How annotations are assigned

You can assign annotations manually or automatically using annotation rules. OnCommand Insight also automatically assigns some annotations on acquisition of assets and by inheritance. Any annotations that you assign to an asset appear in the User Data section of the asset page.

Annotations are assigned in the following ways:

- You can assign an annotation manually to an asset.

If an annotation is assigned directly to an asset, the annotation appears as normal text on an asset page. Annotations that are assigned manually always take precedence over annotations that are inherited or assigned by annotation rules.

- You can create an annotation rule to automatically assign annotations to assets of the same type.

If the annotation is assigned by rule, Insight displays the rule name next to the annotation name on an asset page.

- Insight automatically associates a tier level with a storage tier model to expedite the assignment of storage annotations to your resources on acquisition of assets.

Certain storage resources are automatically associated with a predefined tier (Tier 1 and Tier 2). For example, the Symmetrix storage tier is based on the Symmetrix and VMAX family and is associated with Tier 1. You can change the default values to match your tier requirements. If the annotation is assigned by Insight (for example, Tier), you see “System-defined” when you position your cursor over the annotation’s name on an asset page.

- A few resources (children of an asset) can derive the predefined Tier annotation from their asset (parent).

For example, if you assign an annotation to a storage, the Tier annotation is derived by all the storage

pools, internal volumes, volumes, qtrees, and shares belonging to the storage. If a different annotation is applied to an internal volume of the storage, the annotation is subsequently derived by all the volumes, qtrees, and shares. “Derived” appears next to the annotation name on an asset page.

Associating costs with annotations

Prior to running cost-related reports, you should associate costs with the Service Level, Switch Level, and Tier system-level annotations, which enables chargeback to the storage users based on their actual usage of production and replicated capacity. For example, for the Tier level, you might have gold and silver tier values and assign a higher cost to the gold tier than to the silver tier.

Steps

1. Log in to the Insightweb UI.
2. Click **Manage** and select **Annotations**.

The Annotation page displays.

3. Position your cursor over the Service Level, Switch Level, or Tier annotation, and click .

The Edit Annotation dialog box displays.

4. Enter the values for any existing levels in the **Cost** field.

The Tier and Service Level annotations have Auto Tier and Object Storage values, respectively, which you cannot remove.

5. Click  to add additional levels.
6. Click **Save** when you finish.

Creating custom annotations

Using annotations, you can add custom business-specific data that matches your business needs to assets. While OnCommand Insight provides a set of default annotations, you might find that you want to view data in other ways. The data in custom annotations supplements device data already collected, such as switch manufacturer, number of ports, and performance statistics. The data you add using annotations is not discovered by Insight.

Steps

1. Log in to the Insight web UI.
2. Click **Manage** and select **Annotations**.

The Annotations page displays the list of annotations.

3. Click .

The **Add Annotation** dialog box displays.

4. Enter a name and a description in the **Name** and **Description** fields.

You can enter up to 255 characters in these fields.



Annotation names beginning or ending with a dot "." are not supported.

5. Click **Type** and then select one of the following options that represents the type of data allowed in this annotation:

◦ Boolean

This creates a drop-down list with the choices of yes and no. For example, the “Direct Attached” annotation is Boolean.

◦ Date

This creates a field that holds a date. For example, if the annotation will be a date, select this.

◦ List

This can create either of the following:

▪ A drop-down fixed list

When others are assigning this annotation type on a device, they cannot add more values to the list.

▪ A drop-down flexible list

If you select the **Add new values on the fly** option when you create this list, when others are assigning this annotation type on a device, they can add more values to the list.

◦ Number

This creates a field where the user assigning the annotation can enter a number. For example, if the annotation type is “Floor”, the user could select the Value Type of “number” and enter the floor number.

◦ Text

This creates a field that allows free-form text. For example, you might enter “Language” as the annotation type, select “Text” as the value type, and enter a language as a value.



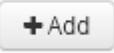
After you set the type and save your changes, you cannot change the type of the annotation. If you need to change the type, you have to delete the annotation and create a new one.

6. If you select **List** as the annotation type, do the following:

a. Select **Add new values on the fly** if you want the ability to add more values to the annotation when on an asset page, which creates a flexible list.

For example, suppose you are on an asset page and the asset has the City annotation with the values Detroit, Tampa, and Boston. If you selected the **Add new values on the fly** option, you can add additional values to City like San Francisco and Chicago directly on the asset page instead of having to go to the Annotations page to add them. If you do not choose this option, you cannot add new annotation values when applying the annotation; this creates a fixed list.

b. Enter a value and a name in **Value** and **Description** fields.

c. Click  to add additional values.

d. Click  to remove a value.

7. Click **Save**.

Your annotations appear in the list on the Annotations page.

Related information

[Importing and Exporting user data](#)

Manually assigning annotations to assets

Assigning annotations to assets helps you sort, group, and report on assets in ways that are relevant to your business. Although you can assign annotations to assets of a particular type automatically, using annotation rules, you can assign annotations to an individual asset by using its asset page.

Before you begin

You must have created the annotation you want to assign.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate the asset to which you want to apply the annotation by doing either of the following:
 - Click the asset in the Assets Dashboard.
 - Click  on the toolbar to display the **Search assets** box, type the type of or name of the asset, and then select the asset from the list that displays.

The asset page displays.

3. In the **User Data** section of the asset page, click .

The Add Annotation dialog box displays.

4. Click **Annotation** and select an annotation from the list.
5. Click **Value** and do either of the following, depending on type of annotation you selected:
 - If the annotation type is list, date, or Boolean, select a value from the list.
 - If the annotation type is text, type a value.
6. Click **Save**.
7. If you want to change the value of the annotation after you assign it, click  and select a different value.

If the annotation is of list type for which the **Add values dynamically upon annotation assignment** option is selected, you can type to add a new value in addition to selecting an existing value.

Modifying annotations

You might want to change the name, description, or values for an annotation, or delete an annotation that you no longer want to use.

Steps

1. Log in to the OnCommand Insightweb UI.
2. Click **Manage** and select **Annotations**.

The Annotations page displays.

3. Position your cursor over the annotation you want to edit and click .

The **Edit Annotation** dialog box displays.

4. You can make the following modifications to an annotation:
 - a. Change the name, description, or both.

However, note that you can enter a maximum of 255 characters for both the name and description, and you cannot change the type of any annotation. Additionally, for system-level annotations, you cannot change the name or description; however, you can add or remove values if the annotation is a list type.



If a custom annotation is published to the Data Warehouse and you rename it, you will lose historical data.

- b. To add another value to an annotation of list type, click .
- c. To remove a value from an annotation of list type, click .

You cannot delete an annotation value if that value is associated with an annotation contained in an annotation rule, query, or performance policy.

5. Click **Save** when you finish.

After you finish

If you are going to use annotations in the Data Warehouse, you need to force an update of annotations in the Data Warehouse. Refer to the *OnCommand Insight Data Warehouse Administration Guide*.

Deleting annotations

You might want to delete an annotation that you no longer want to use. You cannot delete a system-level annotation or an annotation that is used in an annotation rule, query, or performance policy.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage** and select **Annotations**.

The Annotations page displays.

3. Position your cursor over the annotation you want to delete, and click .

A confirmation dialog box displays.

4. Click **OK**.

Assigning annotations to assets using annotation rules

To automatically assign annotations to assets based on criteria that you define, you configure annotation rules. OnCommand Insight assigns the annotations to assets based on these rules. Insight also provides two default annotation rules, which you can modify to suit your needs or remove if you do not want to use them.

Default storage annotation rules

To expedite the assignment of storage annotations to your resources, OnCommand Insight includes 21 default annotation rules, which associate a tier level with a storage tier model. All of your storage resources are automatically associated with a tier upon acquisition of the assets in your environment.

The default annotation rules apply a tier annotations in the following way:

- Tier 1, storage quality tier

The Tier 1 annotation is applied to the following vendors and their specified families: EMC (Symmetrix), HDS (HDS9500V, HDS9900, HDS9900V, R600, R700, USP r, USP V), IBM (DS8000), NetApp (FAS6000 or FAS6200), and Violin (Memory).

- Tier 2, storage quality tier

The Tier 2 annotation is applied to the following vendors and their specified families: HP (3PAR StoreServ or EVA), EMC (CLARiiON), HDS (AMS or D800), IBM (XIV), and NetApp (FAS3000, FAS3100, and FAS3200).

You can edit the default settings of these rules to match your tier requirements, or you can remove them if you do not need them.

Creating annotation rules

As an alternative to manually applying annotations to individual assets, you can automatically apply annotations to multiple assets using annotation rules. Annotations set manually on an individual asset pages take precedence over rule-based annotations when Insight evaluates the annotation rules.

Before you begin

You must have created a query for the annotation rule.

About this task

Although you can edit the annotation types while you are creating the rules, you should have defined the types ahead of time.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage** and select **Annotation rules**.

The Annotation Rules page displays the list of existing annotation rules.

3. Click  **+ Add**.

The Add Rule dialog box displays.

4. Do the following:
 - a. In the **Name** box, enter a unique name that describes the rule.

This name will appear in the Annotation Rules page.

- b. Click **Query** and select the query that OnCommand Insight should use to apply the annotation to assets.
- c. Click **Annotation** and select the annotation you want to apply.
- d. Click **Value** and select a value for the annotation.

For example, if you choose Birthday as the annotation, you specify a date for the value.

5. Click **Save**.
6. Click **Run all rules** if you want to run all the rules immediately; otherwise, the rules are run at a regularly scheduled interval.

Setting annotation rule precedence

By default, OnCommand Insight evaluates annotation rules sequentially; however, you can configure the order in which OnCommand Insight evaluates annotation rules if you want Insight to evaluate rules in a specific order.

Steps

1. Log in to the Insightweb UI.
2. Click **Manage** and select **Annotation rules**.

The Annotation Rules page displays the list of existing annotation rules.

3. Position your cursor over an annotation rule.

The precedence arrows appear to the right of the rule.

4. To move a rule up or down in the list, click the up arrow or the down arrow.

By default, new rules are added sequentially to the list of rules. Annotations set manually on an individual asset pages take precedence over rule-based annotations when Insight evaluates the annotation rules.

Modifying annotation rules

You can modify an annotation rule to change the rule's name, its annotation, the annotation's value, or the query associated with the rule.

Steps

1. Log in to the OnCommand Insightweb UI.
2. Click **Manage** and select **Annotation rules**.

The Annotation Rules page displays the list of existing annotation rules.

3. Locate the rule that you want to modify:
 - On the Annotation Rules page, you can filter the annotation rules by entering a value in the filter box.
 - Click a page number to browse through the annotation rules by page if there are more rules than fit on a page.
4. Perform one of the following to display the **Edit Rule** dialog box:
 - If you are on the Annotation Rules page, position your cursor over the annotation rule and click .
 - If you are on an asset page, position your cursor over the annotation associated with the rule, position your cursor over the rule name when it displays, and then click the rule name.
5. Make the required changes and click **Save**.

Deleting annotation rules

You can delete an annotation rule when the rule is no longer required to monitor the objects in your network.

Steps

1. Log in to the OnCommand Insightweb UI.
2. Click **Manage**, and select **Annotation rules**.

The Annotation Rules page displays the list of existing annotation rules.

3. Locate the rule that you want to delete:
 - On the Annotation Rules page, you can filter the annotation rules by entering a value in the filter box.
 - Click a page number to browse through the annotation rules by page if there are more rules than fit on a single page.
4. Point the cursor over the rule that you want to delete, and then click .

A confirmation message is displayed, prompting whether you want to delete the rule.

5. Click **OK**.

Importing annotation values

If you maintain annotations on SAN objects (such as storage, hosts, and virtual machines) in a CSV file, you can import that information into OnCommand Insight. You

can import applications, business entities, or annotations such as tier and building.

About this task

The following rules apply:

- If an annotation value is empty, that annotation is removed from the object.
- When annotating volumes or internal volumes, the object name is a combination of storage name and volume name using the dash and arrow (->) separator:

```
<storage_name>-><volume_name>
```

- When storage, switches, or ports are annotated, the Application column is ignored.
- The columns of Tenant, Line_of_Business, Business_Unit, and Project make up a business entity.

Any of the values can be left empty. If an application is already related with a business entity different from the input values, the application is assigned to the new business entity.

The following object types and keys are supported in the import utility:

Type	Key
Host	id-><id> or <Name> or <IP>
VM	id-><id> or <Name>
Storage pool	id-><id> or <Storage_name>-><Storage_Pool_name>
Internal volume	id-><id> or <Storage_name>-><Internal_volume_name>
Volume	id-><id> or <Storage_name>-><Volume_name>
Storage	id-><id> or <Name> or <IP>
Switch	id-><id> or <Name> or <IP>
Port	id-><id> or <WWN>
Share	id-><id> or <Storage Name>-><Internal Volume Name>-><Share Name>-><Protocol> <Qtree> is optional if there is a default qtree.

Qtree	id-><id> or <Storage Name>-><Internal Volume Name>-><Qtree Name>
-------	--

The CSV file should use the following format:

```

, , <Annotation Type> [, <Annotation Type> ...]
[, Application] [, Tenant] [, Line_Of_Business] [,
Business_Unit] [, Project]

<Object Type Value 1>, <Object Key 1>, <Annotation Value> [,
<Annotation Value> ...] [, <Application>] [, <Tenant>] [,
<Line_Of_Business>] [, <Business_Unit>] [, <Project>]

...

<Object Type Value N>, <Object Key N>, <Annotation Value> [,
<Annotation Value> ...] [, <Application>] [, <Tenant>] [,
<Line_Of_Business>] [, <Business_Unit>] [, <Project>]

```

Steps

1. Log in to the Insight web UI.
2. Click **Admin** and select **Troubleshooting**.
The Troubleshooting page displays.
3. In the **Other tasks section** of the page, click the **OnCommand Insight Portal** link.
4. Click **Insight Connect API**.
5. Log in to the portal.
6. Click **Annotation Import Utility**.
7. Save the .zip file, unzip it, and read the `readme.txt` file for additional information and samples.
8. Place the CSV file in same folder as the .zip file.
9. In the command line window, enter the following:

```

java -jar rest-import-utility.jar [-username] [-ppassword]
[-aserver name or IP address] [-bbatch size] [-ccase
sensitive:true/false]
[-lextra logging:true/false] csv filename

```

The `-l` option, which enables extra logging, and the `-c` option, which enables case sensitivity, are set to `false` by default. Therefore, you must specify them only when you want to use the features.



There are no spaces between the options and their values.



The following keywords are reserved and prevent users from specifying them as annotation names:

- Application
- Application_Priority
- Tenant
- Line_Of_Business
- Business_Unit
- Project

Errors are generated if you attempt to import an annotation type using one of the reserved keywords. If you have created annotation names using these keywords, you must modify them so that the import utility tool can work correctly.



The Annotation Import utility requires Java 8 or Java 11. Ensure that one of those is installed prior to running the import utility. It is recommended to use the latest OpenJDK 11.

Assigning annotations to multiple assets using a query

Assigning an annotation to a group of assets helps you more easily identify or use those related assets in queries or dashboards.

Before you begin

Annotations that you wish to assign to assets must have previously been created.

About this task

You can simplify the task of assigning an annotation to multiple assets by using a query. For example, if you want to assign a custom address annotation to all of your arrays at a specific data center location.

Steps

1. Create a new query to identify the assets on which you wish to assign an annotation. Click **Queries** > **+New Query**.
2. In the **Search for...** drop-down, choose **Storage**. You can set filters to further narrow down the list of storages displayed.
3. In the list of storages displayed, select one or more by clicking on the check box beside the storage name. You may also select all the displayed storages by clicking on the main check box at the top of the list.
4. When you have selected all of the desired storages, click **Actions** > **Edit Annotation**.

The system displays the Add Annotation dialog.

5. Select the **Annotation** and **Value** you want to assign to the storages and click **Save**.

If you are displaying the column for that annotation, it will appear on all the selected storages.

6. You can now use the annotation to filter for storages in a widget or query. In a widget, you can do the following:
 - a. Create a dashboard or open an existing one. Add a **Variable** and choose the annotation you set on the storages above. The variable is added to the dashboard.
 - b. In the variable field you just added, click on **Any** and enter the appropriate Value to filter on. Click on

the check mark to save the variable value.

- c. Add a widget. In the widget's Query, click on the **Filter by+** button and select the appropriate annotation from the list.
- d. Click on **Any** and select the annotation variable you added above. Variables you have created start with "\$" and are displayed in the drop-down.
- e. Set any other filters or fields you desire, then click **Save** when the widget is customized to your liking.

The widget on the dashboard displays the data for only the storages to which you assigned the annotation.

Querying assets

Queries enable you to monitor and troubleshoot your network by searching the assets in your environment at a granular level based on user-selected criteria (annotations and performance metrics). Additionally, annotation rules, which automatically assign annotations to assets, require a query.

Assets used in queries and dashboards

Insight queries and dashboard widgets can be used with a wide range of asset types

The following asset types can be used in queries, dashboard widgets, and custom asset pages. The fields and counters available for filters, expressions, and display will vary among asset types. Not all assets can be used in all widget types.

- Application
- Datastore
- Disk
- Fabric
- Generic Device
- Host
- Internal Volume
- iSCSI Session
- iSCSI Network Portal
- Path
- Port
- Qtree
- Quota
- Share
- Storage
- Storage Node
- Storage Pool
- Switch
- Tape

- VMDK
- Virtual Machine
- Volume
- Zone
- Zone Member

Creating a query

You can create a query to enable you to search the assets in your environment at a granular level. Queries enable you to slice data by adding filters and then sorting the results to view inventory and performance data in one view.

About this task

For example, you can create a query for volumes, add a filter to find particular storages associated with the selected volume, add a filter to find a particular annotation, such as Tier 1, on the selected storages, and finally add another filter to find all storages with IOPS - Read (IO/s) greater than 25. When the results are displayed, you can then sort the columns of information associated with the query in ascending or descending order.

When a new data source is added which acquires assets or any annotation or application assignments are made, you can query for those assets, annotations, or applications after the queries are indexed, which occurs at a regularly scheduled interval.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Queries** and select **+ New Query**.
3. Click **Select Resource Type** and select a type of asset.

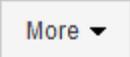
When a resource is selected for a query, a number of default columns are automatically displayed; you can remove these columns or add new ones at any time.

4. In the **Name** text box, type the name of the asset or type a portion of text to filter through the asset names.

You can use any of the following alone or combined to refine your search in any text box on the New Query page:

- An asterisk enables you to search for everything. For example, `vol*rhel` displays all resources that start with “vol” and end with “rhel”.
- The question mark enables you to search for a specific number of characters. For example, `BOS-PRD??-S12` displays BOS-PRD12-S12, BOS-PRD13-S12, and so on.
- The OR operator enables you to specify multiple entities. For example, `FAS2240 OR CX600 OR FAS3270` finds multiple storage models.
- The NOT operator allows you to exclude text from the search results. For example, `NOT EMC*` finds everything that does not start with “EMC”. You can use `NOT *` to display fields that contain no value.

5. Click  to display the assets.
- 6.

To add a criteria, click , and do either of the following:

- Type to search for a specific criteria and then select it.
- Scroll down the list and select a criteria.
- Enter a range of values if you choose a performance metric like IOPS - Read (IO/s).
Default annotations provided by Insight are indicated by ; it is possible to have annotations with duplicate names.

A column is added to the Query results list for the criteria and the results of the query in the list updates.

7. Optionally, you can click  to remove an annotation or performance metric from the query results.

For example, if your query shows maximum latency and maximum throughput for datastores and you want to show only maximum latency in the query results list, click this button, and clear the **Throughput - Max** check box. The Throughput - Max (MB/s) column is removed from the Query results list.



Depending on the number of columns displayed in the query results table, you may not be able to view additional added columns. You can remove one or more columns until your desired columns become visible.

8. Click **Save**, enter a name for the query, and click **Save** again.

If you have an account with an administrator role, you can create custom dashboards. A custom dashboard can comprise any of the widgets from Widget Library, several of which, let you represent query results in a custom dashboard. For more information about custom dashboards, see the *OnCommand Insight Getting Started Guide*.

Related information

[Importing and Exporting user data](#)

Viewing queries

You can view your queries to monitor your assets and change how your queries display the data related to your assets.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Queries** and select **Show all queries**.
3. You can change how queries display by doing any of the following:
 - You can enter text in the **filter** box to search to display specific queries.
 - You can change the sort order of the columns in the table of queries to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header.
 - To resize a column, hover the mouse over the column header until a blue bar appears. Place the mouse over the bar and drag it right or left.
 - To move a column, click on the column header and drag it right or left.
 - When scrolling through the query results, be aware that the results may change as Insight automatically polls your data sources. This may result in some items being missing, or some items

appearing out of order depending on how they are sorted.

Exporting query results to a .CSV file

You might want to export the results of a query into a .CSV file to import the data into another application.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Queries** and select **Show all queries**.

The Queries page is displayed.

3. Click a query.
4. Click  to export query results to a .CSV file.
5. Do one of the following:
 - Click **Open with** and then **OK** to open the file with Microsoft Excel and save the file to a specific location.
 - Click **Save file** and then **OK** to save the file to your Downloads folder. Only the attributes for the displayed columns will be exported. Some displayed columns, particularly those that are part of complex nested relationships, are not exported.



When a comma appears in an asset name, the export encloses the name in quotes, preserving the asset name and the proper .csv format.

+

When exporting query results, be aware that **all** rows in the results table will be exported, not just those selected or displayed on the screen, up to a maximum of 10,000 rows.

+

When opening an exported .CSV file with Excel, if you have an object name or other field that is in the format NN:NN (two digits followed by a colon followed by two more digits), Excel will sometimes interpret that name as a Time format, instead of Text format. This can result in Excel displaying incorrect values in those columns. For example, an object named "81:45" would show in Excel as "81:45:00". To work around this, import the .CSV into Excel using the following steps:

+



- Open a new sheet in Excel.
 - On the "Data" tab, choose "From Text".
 - Locate the desired .CSV file and click "Import".
 - In the Import wizard, choose "Delimited" and click Next.
 - Choose "Comma" for the delimiter and click Next.
 - Select the desired columns and choose "Text" for the column data format.
 - Click Finish.
- Your objects should show in Excel in the proper format.

+

Modifying queries

You can change the criteria that are associated with a query when you want to change the search criteria for the assets that you are querying.

Steps

1. Log in to the Insightweb UI.
2. Click **Queries** and select **Show all queries**.

The Queries page is displayed.

3. Click the query name.
4. To remove a criterion from the query, click .
5. To add a criteria to the query, click , and select a criteria from the list.
6. Do one of the following:
 - Click **Save** to save the query with the name that was used initially.
 - Click **Save as** to save the query with another name.
 - Click **Rename** to change the query name that you had used initially.
 - Click **Revert** to change the query name back to the one that you had used initially.

Deleting queries

You can delete queries when they no longer gather useful information about your assets. You cannot delete a query if it is used in an annotation rule.

Steps

1. Log in to the Insightweb UI.
2. Click **Queries** and select **Show all queries**.

The Queries page displays.

3. Position your cursor over the query you want to delete and click  .

A confirmation message displays, asking if you want to delete the query.

4. Click **OK**.

Assigning multiple applications to or removing multiple applications from assets

You can assign multiple applications to or remove multiple application from assets by using a query instead of having to manually assign or remove them.

Before you begin

You must have already created a query that finds all the assets that you to edit.

Steps

1. Click **Queries** and select **Show all queries**.

The Queries page displays.

2. Click the name of the query that finds the assets.

The list of assets associated with the query displays.

3. Select the desired assets in the list or click ▼ to select **All**.

The **Actions** button displays.

4. To add an application to the selected assets, click , and select **Edit Application**.

- a. Click **Application** and select one or more applications.

You can select multiple applications for hosts, internal volumes, and virtual machines; however, you can select only one application for a volume.

- b. Click **Save**.

5. To remove an application assigned to the assets, click and select **Remove Application**.

- a. Select the application or applications you want to remove.

- b. Click **Delete**.

Any new applications you assign override any applications on the asset that were derived from another asset. For example, volumes inherit applications from hosts, and when new applications are assigned to a volume, the new application takes precedence over the derived application.

Editing or removing multiple annotations from assets

You can edit multiple annotations for assets or remove multiple annotations from assets by using a query instead of having to manually edit or remove them.

Before you begin

You must have already created a query that finds all the assets that you want to edit.

Steps

1. Click **Queries** and select **Show all queries**.

The Queries page displays.

2. Click the name of the query that find the assets.

The list of assets associated with the query displays.

3. Select the desired assets in the list or click ▼ to select **All**.

The **Actions** button displays.

4. To add an annotation to the assets or edit the value of an annotation assigned to the assets, click , 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**.

Copying table values

You can copy values in tables for use in search boxes or other applications.

About this task

There are two methods you can use to copy values from tables or query results.

Steps

1. Method 1: Highlight the desired text with the mouse, copy it, and paste it into search fields or other applications.
2. Method 2: For single-value fields whose length exceeds the width of the table column, indicated by ellipses (...), hover over the field and click the clipboard icon. The value is copied to the clipboard for use in search fields or other applications.

Note that only values that are links to assets can be copied. Note also that only fields that include single values (i.e. non-lists) have the copy icon.

Insight data source management

Data sources are the most critical component used to maintain an OnCommand Insight environment. Because they are the primary source of information for Insight, it is imperative that you maintain data sources in a running state.

You can monitor the data sources in your network by selecting a data source to check the events related to its status and noting any changes that might have caused problems.

In addition to examining an individual data source, you can perform these operations:

- Clone a data source to create many similar data sources in Insight
- Edit data source information
- Change credentials
- Control polling
- Delete the data source
- Install data source patches
- Install a new data source from a patch
- Prepare an error report for NetApp Customer Support

Setting up your data sources in Insight

Data sources are the most critical component when trying to maintain a Insight environment. Data sources discover network information that is used for analysis and validation. You need to configure your data sources within Insight so that they can be monitored within your network.

For each data source, the specific requirements to define that data source depend on the vendor and model of the corresponding devices. Before adding the data sources, you need network addresses, account information, and passwords for all devices and possibly these additional details:

- Switches
- Device management stations
- Storage systems that have IP connectivity
- Storage management stations

- Host servers running management software for storage devices that do not have IP connectivity

For more information about your data source definitions, see the "Vendor-specific data source reference" information in this section.

Data source support information

As part of your configuration planning, you should ensure that the devices in your environment can be monitored by Insight. To do so, you can check the Data source support matrix for details about operating systems, specific devices, and protocols. Some data sources might not be available on all operating systems.

Location of the most up-to-date version of the Data Source Support Matrix

The OnCommand Insight Data Source Support Matrix is updated with each service pack release. The most current version of the document can be found at the [NetApp Support Site](#).

Adding data sources

You can add data sources quickly, using the Add data source dialog box.

Steps

1. Open OnCommand Insight in your browser and log in as a user with administrative permissions.
2. Select **Admin** and choose **Data sources**.
3. Click the **+Add** button.

The Add data source wizard opens.

4. In the **Settings** section, enter the following information:

Field	Description
Name	Enter a unique network name for this data source. NOTE: only letters, numbers and the underscore (_) character are allowed in the data source name.
Vendor	Choose the vendor of the data source from the drop-down.
Model	Choose the model of the data source from the drop-down.
Where to run	Choose Local, or you may choose a remote acquisition unit if RAU's are configured in your environment.

What to collect	For most data sources, these options will be Inventory and Performance. Inventory is always selected by default and cannot be un-selected. Note that some data sources may have different options. The collection options you select change the available fields in the Configuration and Advanced configuration sections.
-----------------	--

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](#).
8. Click the **Test** link to be certain that the data source is properly configured.
9. Click **Save**.

Importing data sources from a spreadsheet

You can import multiple data sources into OnCommand Insight from a spreadsheet. This might be helpful if you already maintain your discovery devices in a spreadsheet. This process adds new data sources, but cannot be used to update existing data sources.

About this task

OnCommand Insight includes a spreadsheet to help you create data sources. This spreadsheet has the following attributes:

- The spreadsheet can be used with Microsoft Excel 2003 or later.
- Each tab holds one data source type, for example, Brocade SSH/CLI.
- Each row represents an instance of a new data source to be created.

The spreadsheet includes a macro that creates a new data source in OnCommand Insight.

Steps

1. Locate the spreadsheet in the `<install_directory>/SANscreen/acq/bin/acqcli/SiteSurvey_DataSourceImporter_w_Macro.zip`.
2. In the spreadsheet, enter data source information in the cells with color.
3. Delete empty rows.
4. From the spreadsheet, run the `CreateDataSources` macro to create the data sources.
5. When prompted for credentials, enter the OnCommand Insight Server administration user name and password.

The results are logged in the acquisition log.

6. A prompt asks if the machine currently running the macro has OnCommand Insight installed.

Select one of the following:

- No: Select "No" if a batch file will be created that must be run on the OnCommand Insight machine. Run this batch file from the install directory.
- Yes: Select "Yes" if OnCommand Insight is already installed and no additional steps are required to generate the data source information.

7. To verify the addition of the data sources, open Insight in your browser.

8. On the Insight toolbar, click **Admin**.

9. Check the Data sources list for the data sources you imported.

Adding a new data source by patch

New data sources are released as patch files that can be loaded onto the system using the patch process. This process enables new data sources to be available between scheduled releases of OnCommand Insight.

Before you begin

You must have uploaded the patch file that you want to install.

Steps

1. On the Insight toolbar, click **Admin**.
2. Select **Patches**.
3. Select **Actions > Install service pack or patch**.
4. In the **Install Service Pack or Patch** dialog box, click **Browse** to locate and select the patch file that you uploaded.
5. Click **Next** in the **Patch Summary** dialog box.
6. Review the **Read Me** information, and click **Next** to continue.
7. In the **Install** dialog box, click **Finish**.

Cloning a data source

Using the clone facility, you can quickly add a data source that has the same credentials and attributes as another data source. Cloning allows you to easily configure multiple instances of the same device type.

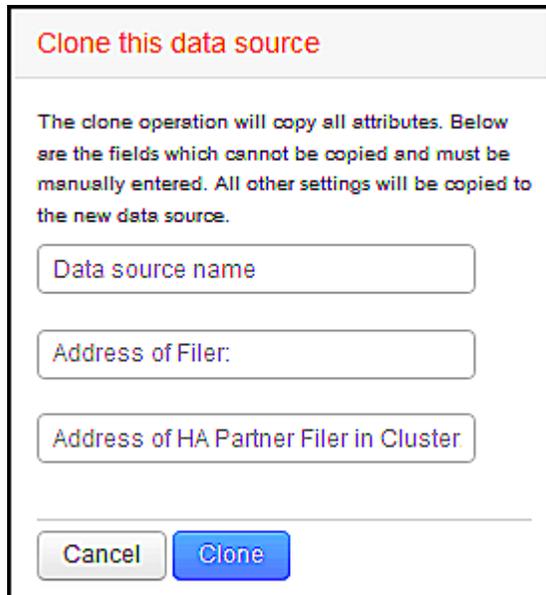
Steps

1. On the Insight toolbar, click **Admin**.

The Data sources list opens.

2. Highlight the data source that has the setup information you want to use for your new data source.
3. To the right of the highlighted data source, click the **Clone** icon.

The Clone this data source dialog box lists the information you must supply for the selected data source, as shown in this example for a NetApp data source:



Clone this data source

The clone operation will copy all attributes. Below are the fields which cannot be copied and must be manually entered. All other settings will be copied to the new data source.

Data source name

Address of Filer:

Address of HA Partner Filer in Cluster

Cancel Clone

4. Enter the required information in the fields; those details cannot be copied from the existing data source.
5. Click **Clone**.

Results

The clone operation copies all other attributes and settings to create the new data source.

Testing the data source configuration

When you are adding a data source, you can verify the correctness of configuration to communicate with the device before saving or updating that data source.

When you click the **Test** button in the data source wizard, communication with the specified device is checked. The test produces one of these results:

- **PASSED:** the data source is configured correctly.
- **WARNING:** the testing was incomplete, probably due to timing out during processing or acquisition not running.
- **FAILED:** the data source, as configured, cannot communicate with the specified device. Check your configuration settings and re-test.

Vendor-specific data source reference

The configuration details vary depending on the vendor and model of the data source being added.

If a vendor's data source requires advanced Insight configuration instructions, such as special requirements and specific commands, that information is included in this section.

3PAR InServ data source

OnCommand Insight uses the 3PAR InServ (Firmware 2.2.2+, SSH) data source to discover inventory for HP 3PAR StoreServ storage arrays.

Terminology

OnCommand Insight acquires the following inventory information from the 3PAR InServ data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Physical Disk	Disk
Storage System	Storage
Controller Node	Storage Node
Common Provisioning Group	Storage Pool
Virtual Volume	Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- IP address or FQDN of the InServ cluster
- For inventory, read-only user name and password to the InServ Server.
- For performance, read-write user name and password to the InServ Server.
- Port requirements: 22 (inventory collection), 5988 or 5989 (performance collection) [Note: 3PAR Performance is supported for InServ OS 3.x+]
- For performance collection confirm that SMI-S is enabled by logging into the 3PAR array via SSH.

Configuration

Field	Description
Cluster IP	IP address or fully-qualified domain name of the InServ cluster
User Name	User name for the InServ Server
Password	Password used for the InServ Server
SMI-S Host IP	IP address of the SMI-S Provider Host

SMI-S User Name	User name for the SMI-S Provider Host
SMI-S Password	Password used for the SMI-S Provider Host

Advanced Configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Exclude Devices	Comma-separated list of device IPs to exclude
SSH Process Wait Timeout (sec)	SSH process timeout (default 60 seconds)
Number of SSH Retries	Number of SSH retry attempts
SSH Banner Wait Timeout (sec)	SSH banner wait timeout (default 20 seconds)
SMI-S Port	Port used by SMI-S Provider Host
Protocol	Protocol used to connect to the SMI-S provider
SMI-S namespace	SMI-S namespace
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Number of SMI-S Connection Retries	Number of SMI-S connection retry attempts

Amazon AWS EC2 data source

OnCommand Insight uses this data source to discover inventory and performance for Amazon AWS EC2.

Pre-requisites:

In order to collect data from Amazon EC2 devices, you must have the following information:

- You must have the IAM Access Key ID
- You must have the Secret Access Key for your Amazon EC2 cloud account
- You must have the "list organization" privilege
- Port 433 HTTPS
- EC2 Instances can be reported as a Virtual Machine, or (less naturally) a Host. EBS Volumes can be reported as both a VirtualDisk used by the VM, as well as a DataStore providing the Capacity for the VirtualDisk.

Access keys consist of an access key ID (for example, AKIAIOSFODNN7EXAMPLE) and a secret access key (for example, wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY). You use access keys to sign programmatic requests that you make to EC@ if you use the Amazon EC2 SDKs, REST, or Query API operations. These keys are provided with your contract from Amazon.

How to configure this data source

To configure the Amazon AWS EC2 data source, you will need the AWS IAM Access Key ID and Secret Access Key for your AWS account.

Fill in the data source fields according to the tables below:

Configuration:

Field	Description
AWS Region	Choose AWS region
IAM Role	For use only when acquired on an AU in AWS. See below for more information on IAM Roles.
AWS IAM Access Key ID	Enter AWS IAM Access Key ID. Required if you do not use IAM Role.
AWS IAM Secret Access Key	Enter AWS IAM Secret Access Key. Required if you do not use IAM Role.
I understand AWS will bill me for API requests	Check this to verify your understanding that AWS bills you for API requests made by Insight polling

Advanced Configuration:

Field	Description
Include Extra Regions	Specify additional regions to include in polling.
Cross Account Role	Role for accessing resources in different AWS accounts.
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
HTTP connection and socket timeout (sec)	HTTP connection timeout (default 300 seconds)
Include AWS tags	Check this to enable support for AWS tags in Insight annotations
Performance Poll Interval (sec)	Interval between performance polls (default 1800 seconds)

Mapping AWS tags to Insight annotations

The AWS EC2 data source includes an option that allows you to populate Insight annotations with tags configured on AWS. The annotations must be named exactly as the AWS tags. Insight will always populate same-named text-type annotations, and will make a "best attempt" to populate annotations of other types (number, boolean, etc). If your annotation is of a different type and the data source fails to populate it, it may be necessary to remove the annotation and re-create it as a text type.

Note that AWS is case-sensitive, while Insight is case-insensitive. So if you create an annotation named "OWNER" in Insight, and tags named "OWNER", "Owner", and "owner" in AWS, all of the AWS variations of "owner" will map to Insight's "OWNER" annotation.

Related Information:

[Managing Access Keys for IAM Users](#)

Include Extra Regions

In the AWS Data Collector **Advanced Configuration** section, you can set the **Include extra regions** field to include additional regions, separated by comma or semi-colon. By default, this field is set to *us-.**, which collects on all US AWS regions. To collect on *all* regions, set this field to *.**.

If the **Include extra regions** field is empty, the data collector will collect on assets specified in the **AWS Region** field as specified in the **Configuration** section.

Collecting from AWS Child Accounts

Insight supports collection of child accounts for AWS within a single AWS data collector. Configuration for this collection is performed in the AWS environment:

- You must configure each child account to have an AWS Role that allows the primary account ID to access EC2 details from the children account.
- Each child account must have the role name configured as the same string
- Enter this role name string into the Insight AWS Data Collector **Advanced Configuration** section, in the **Cross Account Role** field.

Best Practice: It is highly recommended to assign the AWS predefined AmazonEC2ReadOnlyAccess policy to the ECS primary account. Also, the user configured in the data source should have at least the predefined *AWSOrganizationsReadOnlyAccess* policy assigned, in order to query AWS.

Please see the following for information on configuring your environment to allow Insight to collect from AWS child accounts:

[Tutorial: Delegate Access Across AWS Accounts Using IAM Roles](#)

[AWS Setup: Providing Access to an IAM User in Another AWS Account That You Own](#)

[Creating a Role to Delegate Permissions to an IAM User](#)

IAM Roles

When using *IAM Role* security, you must ensure that the role you create or specify has the appropriate permissions needed to access your resources.

For example, if you create an IAM role named *InstanceEc2ReadOnly*, you must set up the policy to grant EC2

read-only list access permission to all EC2 resources for this IAM role. Additionally, you must grant STS (Security Token Service) access so that this role is allowed to assume roles cross accounts.

After you create an IAM role, you can attach it when you create a new EC2 instance or any existing EC2 instance.

After you attach the IAM role *InstanceEc2ReadOnly* to an EC2 instance, you will be able to retrieve the temporary credential through instance metadata by IAM role name and use it to access AWS resources by any application running on this EC2 instance.



IAM role can be used only when the Acquisition Unit is running in an AWS instance.

Brocade Enterprise Fabric Connectivity Manager data source

OnCommand Insight uses the Brocade Enterprise Fabric Connectivity Manager (EFCM) data source to discover inventory for Brocade EFCM switches. Insight supports EFCM versions 9.5, 9.6, and 9.7.

Requirements



This data collector is not available starting with OnCommand Insight 7.3.11.

- Network address or fully-qualified domain name for the EFCM server
- EFCM version must be 9.5, 9.6, or 9.7
- IP address of the EFCM server
- Read-only username and password for the EFCM server
- Validated access to the Connectrix switch by Telnet from the Insight server, using the read-only username and password over port 51512

Configuration

Field	Description
EFC server	IP address or fully-qualified domain name of the EFC Server
User Name	User name for the switch
Password	Password used for the switch

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 15 minutes)
Fabric Name	Fabric name to be reported by the EFCM data source. Leave blank to report the fabric name as WWN.

Communication Port	Port used for communication with the switch
Enable Trapping	Select to enable acquisition upon receiving an SNMP trap from the device. If you select enable trapping, you must also activate SNMP.
Minimum Time Between Traps (sec)	Minimum time between acquisition attempts triggered by traps (default 15 seconds)
Inactive Zonesets	Comma-separated list of inactive Zonesets on which to perform acquisition, in addition to performing acquisition on the active zone sets
NIC to Use	Specify which network interface should be used on the RAU when reporting on SAN devices
Exclude Devices	Comma-separated list of unit names to include or exclude from polling
Use the EFCM switch nickname as the Insight switch name	Select to use the EFCM switch nickname as the Insight switch name
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Brocade FC Switch data source

OnCommand Insight uses the Brocade FC Switch (SSH) data source to discover inventory for Brocade or rebranded switch devices running Factored Operating System (FOS) firmware 4.2 and later. Devices in both FC switch and Access Gateway modes are supported.

Terminology

OnCommand Insight acquires the following inventory information from the Brocade FC Switch data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Switch	Switch
Port	Port
Virtual Fabric, Physical Fabric	Fabric
Zone	Zone

Logical Switch	Logical Switch
LSAN Zone	IVR Zone



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- The Acquisition Unit (local or remote) will initiate connections to TCP Port 22 on Brocade switches to collect inventory data. The AU will also initiate connections to UDP port 161 for collection of performance data.
- There must be IP connectivity to all switches in the fabric. If you select the Discover all switches in the fabric check box, OCI identifies all the switches in the fabric; however, it needs IP connectivity to these additional switches to discover them.
- The same account is needed globally across all switches in the fabric. You can use PuTTY (open source terminal emulator) to confirm access.
- If the Perform license is installed, ports 161 and 162 must be open to all switches in the fabric for SNMP performance polling.
- SNMP read-only Community String

Configuration

Field	Description
Switch IP	IP address or fully-qualified domain name of the switch
User Name	User name for the switch
Password	Password used for the switch
SNMP Version	SNMP version
SNMP Community String	SNMP read-only community string used to access the switch
SNMP User Name	SNMP version protocol user name (applies only to SNMP v3)
SNMP Password	SNMP version protocol password (applies only to SNMP v3)

Advanced configuration

Field	Description
-------	-------------

Fabric Name	Fabric name to be reported by the data source. Leave blank to report the fabric name as WWN.
Exclude Devices	Comma-separated list of device IDs to exclude from polling
Inventory Poll Interval (min)	Interval between inventory polls (default 15 minutes)
Timeout (sec)	Connection timeout (default 30 seconds)
Banner Wait Timeout (sec)	SSH banner wait timeout (default 5 seconds)
Admin Domains Active	Select if using Admin Domains
Retrieve MPR Data	Select to acquire routing data from your multiprotocol router (MPR)
Enable Trapping	Select to enable acquisition upon receiving an SNMP trap from the device. If you select enable trapping, you must also activate SNMP.
Minimum Time Between Traps (sec)	Minimum time between acquisition attempts triggered by traps (default 10 seconds)
Discover all switches in the fabric	Select to discover all switches in the fabric
Choose Favoring HBA vs. Zone Aliases	Choose whether to favor HBA or zone aliases
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
SNMP Auth Protocol	SNMP authentication protocol (SNMP v3 only)
SNMP Privacy Protocol	SNMP privacy protocol (SNMP v3 only)
SNMP Privacy Password	SNMP privacy password (SNMP v3 only)
SNMP Retries	Number of SNMP retry attempts
SNMP Timeout (ms)	SNMP timeout (default 5000 ms)

Brocade Sphereon/Intrepid Switch data source

OnCommand Insight uses the Brocade Sphereon/Intrepid Switch (SNMP) data source to discover inventory for Brocade Sphereon or Intrepid switches.

Requirements



This data collector not available starting with OnCommand Insight 7.3.11.

- There must be IP connectivity to all switches in the fabric. If you select the Discover all switches in the fabric check box, OCI identifies all the switches in the fabric; however, it needs IP connectivity to these additional switches to discover them.
- Read-only community string if using SNMP V1 or SNMP V2.
- HTTP access to the switch to obtain zoning information.
- Access validation by running the `snmpwalk` utility to the switch (see `<install_path>\bin\`).

Configuration

Field	Description
Sphereon Switch	IP address or fully-qualified domain name of the switch
SNMP Version	SNMP version
SNMP Community	SNMP read-only community string used to access the switch
User Name	SMI-S user name for the switch (SNMP v3 only)
Password	SMI-S password for the switch (SNMP v3 only)

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 15 minutes)
SNMP Auth Protocol	SNMP authentication protocol (SNMPv3 only)
SNMP Privacy Protocol	SNMP privacy protocol (SNMPv3 only)
SNMP Privacy Password	SNMP privacy password
SNMP Number of Retries	Number of SNMP retry attempts
SNMP Timeout (ms)	SNMP timeout (default 5000 ms)
Fabric Name	Fabric name to be reported by the data source. Leave blank to report the fabric name as WWN.

Enable Trapping	Select to enable acquisition upon receiving an SNMP trap from the device. If you select enable trapping, you must also activate SNMP.
Minimum Time Between Ttraps (seconds)	Minimum time between acquisition attempts triggered by traps (default 10 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Cisco FC Switch Firmware (SNMP) data source

OnCommand Insight uses the Cisco FC Switch Firmware 2.0+ (SNMP) data source to discover inventory for Cisco MDS Fibre Channel switches as well as a variety of Cisco Nexus FCoE switches on which the FC service is enabled. Additionally, you can discover many models of Cisco devices running in NPV mode with this data source.

Terminology

OnCommand Insight acquires the following inventory information from the Cisco FC Switch data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Switch	Switch
Port	Port
VSAN	Fabric
Zone	Zone
Logical Switch	Logical Switch
Name Server Entry	Name Server Entry
Inter-VSAN Routing (IVR) Zone	IVR Zone



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- An IP address of one switch in the fabric or individual switches
- Chassis discovery, to enable fabric discovery

- If using SNMP V2, read-only community string
- Port 161 is used to access the device
- Access validation using the `snmpwalk` utility to the switch (see `<install_path>\bin\`)

Configuration

Field	Description
Cisco Switch IP	IP address or fully-qualified domain name of the switch
SNMP Version	SNMP version v2 or later is required for performance acquisition
SNMP Community String	SNMP read-only community string used to access the switch (not applicable for SNMP v3)
User Name	User name for the switch (SNMP v3 only)
Password	Password used for the switch (SNMPv3 only)

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
SNMP Auth Protocol	SNMP authentication protocol (SNMPv3 only)
SNMP Privacy Protocol	SNMP privacy protocol (SNMPv3 only)
SNMP Privacy Password	SNMP privacy password
SNMP Retries	Number of SNMP retry attempts
SNMP Timeout (ms)	SNMP timeout (default 5000 ms)
Enable Trapping	Select to enable trapping. If you enable trapping, you must also activate SNMP notifications.
Minimum Time Between Traps (sec)	Minimum time between acquisition attempts triggered by traps (default 10 seconds)
Discover All Fabric Switches	Select to discover all switches in the fabric

Exclude Devices	Comma-separated list of device IPs to exclude from polling
Include Devices	Comma-separated list of device IPs to include in polling
Check Device Type	Select to accept only those devices that explicitly advertise themselves as Cisco devices
Primary Alias Type	<p>Provide a first preference for resolution of the alias. Choose from the following:</p> <ul style="list-style-type: none"> • Device Alias <p>This is a user-friendly name for a port WWN (pWWN) that can be used in all configuration commands, as required. All switches in the Cisco MDS 9000 Family support Distributed Device Alias Services (device aliases).</p> • None <p>Do not report any alias</p> • Port Description <p>A description to help identify the port in a list of ports</p> • Zone Alias (all) <p>A user-friendly name for a port that can be used only for zoning configuration</p> • Zone Alias (only active) <p>A user-friendly name for a port that can be used only for the active configuration. This is the default.</p>
Secondary Alias Type	Provide a second preference for resolution of the alias
Tertiary Alias Type	Provide a third preference for resolution of the alias
Enable SANTap Proxy Mode Support	Select if your Cisco switch is using SANTap in proxy mode. If you are using EMC RecoverPoint, then you are probably using SANTap.
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

EMC Celerra data source

The Celerra (SSH) data source collects inventory information from Celerra storage. For configuration, this data source requires the IP address of the storage processors and a *read-only* user name and password.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Celerra data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Celerra Network Server	Storage
Celerra Meta Volume / Celerra Storage Pool	Storage Pool
File System	Internal Volume
Data Mover	Controller
File System mounted on a data Mover	File Share
CIFS and NFS Exports	Share
Disk Volume	Backend LUN



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- The IP address of the storage processor
- Read-only user name and password
- SSH port 22

Configuration

Field	Description
Address of Celerra	IP address or fully-qualified domain name of the Celerra device
User Name	Name used to log in to the Celerra device
Password	Password used to log in to the Celerra device

Advanced configuration

Field	Description
Inventory Poll Interval (minutes)	Interval between inventory polls (default 20 minutes)
SSH Process Wait Timeout (sec)	SSH process timeout (default 600 seconds)
Number of Retries	Number of inventory retry attempts
SSH Banner Wait Timeout (sec)	SSH banner wait timeout (default 20 seconds)

EMC CLARiiON (NaviCLI) data source

Before configuring this data source, make sure that the EMC Navisphere CLI is installed on the target device and on the Insight server. The Navisphere CLI version must match the firmware version on the controller. For performance data collection, statistics logging must be turned on.

Navisphere Command Line Interface syntax

```
naviseccli.exe -h <IP address> -user <user> -password <password> -scope  
<scope,use 0 for global scope> -port <use 443 by default> command
```

Terminology

OnCommand Insight acquires the following inventory information from the EMC CLARiiON data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Storage	Storage
Storage Processor	Storage Node
Thin Pool, RAID Group	Storage Pool
LUN	Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- An IP address of each CLARiiON storage processor
- Read-only Navisphere username and password to the CLARiiON arrays
- NaviCLI must be installed on the Insight server/RAU
- Access validation: Run NaviCLI from the Insight server to each array using the above username and password.
- NaviCLI version should correspond with the newest FLARE code on your array
- For performance, statistics logging must be turned on.
- Port requirements: 80, 443

Configuration

Field	Description
CLARiiON storage	IP address or fully-qualified domain name of the CLARiiON Storage
User Name	Name used to log into the CLARiiON storage device.
Password	Password used to log into the CLARiiON storage device.
CLI Path to navicli.exe path or naviseccli.exe path	Full path to the <code>navicli.exe</code> OR <code>naviseccli.exe</code> executable

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Use Secure Client (naviseccli)	Select to use secure client (navseccli)
Scope	The secure client scope. The default is Global.
CLARiiON CLI Port	Port used for CLARiiON CLI
Inventory External Process Timeout (sec)	External process timeout (default 1800 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Performance External process timeout (sec)	External process timeout (default 1800 seconds)

EMC Data Domain data source

This data source collects storage and configuration information from EMC Data Domain deduplication storage systems. To add the data source, you must use specific configuration instructions and commands and be aware of data source requirements and usage recommendations.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Data Domain data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Array	Storage
Port	Port
Filesys	Internal Volume
Mtree	QTree
Quota	Quota
NFS and CIFS share	FileShare



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- IP address of the Data Domain device
- Read-only user name and password to the Data Domain storage
- SSH port 22

Configuration

Field	Description
IP address	The IP address or fully-qualified domain name of the Data Domain storage array
User name	The user name for the Data Domain storage array

Password	The password for the Data Domain storage array
----------	--

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
SSH Process Wait Timeout (sec)	SSH process timeout (default 180 seconds)
SSH Port	SSH service port

EMC ECC StorageScope data source

The EMC ECC StorageScope device has three types of data sources: 5.x, 6.0, and 6.1.

Configuration



This data collector is no longer available starting with OnCommand Insight 7.3.11.

Field	Description
ECC server	IP address or fully-qualified domain name of the ECC Server
User Name	User name for the ECC server
Password	Password for the ECC server

Advanced configuration

Field	Description
ECC Port	Port used for the ECC server
Inventory Poll Interval (min)	Interval between inventory polls (default 30 minutes)
Protocol to Connect to Database	Protocol Used to Connect to the Database
Query File System Information	Select to retrieve details for WWN Aliases and File Systems.

Dell EMC ECS data source

This data collector acquires inventory and performance data from EMC ECS storage systems. For configuration, the data collector requires an IP address of the ECS server

and an administrative level domain account..

Terminology

OnCommand Insight acquires the following inventory information from the EMC ECS data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Cluser	Storage
Tenant	Storage Pool
Bucket	Internal Volume
Disk	Disk



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- An IP address of the ECS Management Console
- Administrative level domain account for the ECS system
- Port 443 (HTTPS). Requires outbound connectivity to TCP port 443 on the ECS system.
- For performance, read-only username and password for ssh/scp access.
- For performance, port 22 is required.

Configuration

Field	Description
ECS Host	IP addresses or fully-qualified domain names of the ECS system
ECS Host Port	Port used for communication with ECS Host
ECS Vendor ID	Vendor ID for ECS
Password	Password used for ECS

Advanced configuration

Field	Description
-------	-------------

Inventory Poll Interval (minutes)	Interval between inventory polls. The default is 360 minutes.
-----------------------------------	---

EMC Isilon data source

The Isilon SSH data source collects inventory and performance from EMC Isilon scale-out NAS storage.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Isilon data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive	Disk
Cluster	Storage
Node	Storage Node
File System	Internal Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- Administrator permissions to the Isilon storage
- Validated access by using `telnet` to port 22

Configuration

Field	Description
IP address	The IP address or fully-qualified domain name of the Isilon cluster
User name	The user name for the Isilon cluster
Password	The password for the Isilon cluster

Advanced configuration

Field	Description
-------	-------------

Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
SSH Process Wait Timeout	SSH process timeout (default 60 seconds)
SSH Port	SSH service port

Running CLI Commands

Starting with OnCommand Insight version 7.3.11 and Service Pack 9, the EMC Isilon data source contains an enhancement that will result in Insight running more CLI commands. If you are using a non-root user within your data source, you will likely have configured a "sudoers" file to grant that user account the ability to run specific CLI commands via SSH.

In order for Insight to understand EMC's "Access Zones" feature, Insight will now additionally run the following new CLI commands:

- `sudo isi zone zones list --format json -verbose`
- `sudo isi zone zones list`

Insight parses the output of these commands, and runs more instances of existing commands, to obtain the logical configuration of objects like qtrees, quotas and NAS shares/exports that reside in non-default Access Zones. Insight now reports those objects for non-default Access Zones as the result of this enhancement. As Insight obtains that data by running existing commands (with different options) no sudoers file change is required in order for those to work; it is only with the introduction of the new commands above that the change is required.

Please update your sudoers file to allow your Insight service account to run those commands before upgrading to this Insight release. Failure to do so will result in your Isilon data sources failing.

"File System" statistics

Beginning with OnCommand Insight 7.3.12, the EMC Isilon data collector introduces "File System" statistics on the node object for EMC Isilon. The existing node statistics reported by OnCommand Insight are "disk" based - i.e, for IOPs and throughput of a storage node, what are the disks in this node doing in aggregate? But for workloads where reads are cached in memory and/or compression is in use, the file system workload may be substantially higher than what actually hits the disks - a data set that compresses 5:1 could therefore have a "File System Read Throughput" value 5x the storage node Read Throughput, as the latter measures the reads off of disk, which expand 5x when the node uncompresses the data to service the client's read request.

Dell EMC PowerStore data source

The Dell EMC PowerStore data collector gathers inventory information from Dell EMC PowerStore storage. For configuration, the data collector requires the IP address of the storage processors and a read-only user name and password.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Data Domain data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
host	host
host_volume_mapping	host_volume_mapping
hardware (it has Drives under "extra_details" object): Drives	Disk
Appliance	StoragePool
Cluster	Storage Array
Node	StorageNode
fc_port	Port
volume	Volume
InternalVolume	file_system
Filesys	Internal Volume
Mtree	QTree
Quota	Quota
NFS and CIFS share	FileShare



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- IP address or fully-qualified domain name of storage processor
- Read-only user name and password

Parent Serial Number explained

Traditionally Insight is capable of reporting the storage array serial number, or the individual storage node serial numbers. However, some storage array architectures do not cleanly align to this. A PowerStore cluster can be comprised of 1-4 appliances, and each appliance has 2 nodes. If the appliance itself has a serial

number, that serial number is neither the serial number for the cluster nor the nodes.

The attribute "Parent Serial Number" on the storage node object is populated appropriately for Dell/EMC PowerStore arrays when the individual nodes sit inside an intermediate appliance/enclosure that is just part of a larger cluster.

Configuration

Field	Description
PowerStore gateway(s)	IP addresses or fully-qualified domain names of PowerStore storage
User Name	User name for PowerStore
Password	Password used for PowerStore

Advanced configuration

Field	Description
HTTPS Port	Default is 443
Inventory Poll Interval (minutes)	Interval between inventory polls. The default is 60 minutes.

OnCommand Insight's PowerStore performance collection makes use of PowerStore's 5-minute granularity source data. As such, Insight polls for that data every five minutes, and this is not configurable.

EMC RecoverPoint data source

The EMC RecoverPoint data source collects inventory information from EMC recoverPoint storage. For configuration, the data source requires the IP address of the storage processors and a *read-only* user name and password.

The EMC RecoverPoint data source gathers the volume-to-volume replication relationships that RecoverPoint coordinates across other storage arrays. OnCommand Insight shows a storage array for each RecoverPoint cluster, and collects inventory data for nodes and storage ports on that cluster. No storage pool or volume data is collected.

Requirements

- IP address or fully-qualified domain name of storage processor
- Read-only user name and password
- REST API access via port 443
- SSH access via PuTTY

Configuration

Field	Description
Address of RecoverPoint	IP address or fully-qualified domain name of RecoverPoint cluster
User Name	User name for the RecoverPoint cluster
Password	Password for the RecoverPoint cluster

Advanced configuration

Field	Description
TCP Port	TCP Port used to connect to Recoverpoint cluster
Inventory Poll Interval (minutes)	Interval between inventory polls (default 20 minutes)
Excluded Clusters	Comma-separated list of cluster IDs or names to exclude when polling

EMC Solutions Enabler with SMI-S Performance data source

OnCommand Insight discovers Symmetrix storage arrays by using Solutions Enabler `symcli` commands in conjunction with an existing Solutions Enabler server in your environment. The existing Solutions Enabler server has connectivity to the Symmetrix storage array through access to gatekeeper volumes. Administrator permissions are required to access this device.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Solutions Enabler data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Disk Group	Disk Group
Storage Array	Storage
Director	Storage Node
Device Pool, Storage Resource Pool (SRP)	Storage Pool

Device, TDev	Volume
--------------	--------



These are common terminology mappings only and might not represent every case for this data source.

Requirements

Before configuring this data source, you should ensure that the OnCommand Insight server has TCP connectivity to port 2707 on the existing Solutions Enabler server. OnCommand Insight discovers all the Symmetrix arrays that are “Local” to this server, as seen in “symcfg list” output from that server.

- The EMC Solutions Enabler (CLI) with SMI-S provider application must be installed and the version must match or be earlier than the version running on the Solutions Enabler Server.
- A properly configured `{installdir}\EMC\SYMAPI\config\netcnfg` file is required. This file defines service names for Solutions Enabler servers, as well as the access method (SECURE / NOSECURE /ANY).
- If you require read/write latency at the storage node level, the SMI-S Provider must communicate with a running instance of the UNISPHERE for VMAX application.
- Administrator permissions on the Solutions Enabler (SE) Server
- Read-only user name and password to the SE software
- Solutions Enabler Server 6.5X requirements:
 - SMI-S provider 3.3.1 for SMIS-S V1.2 installed
 - After install, run `\Program Files\EMC\SYMCLI\bin>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`

Configuration



If SMI-S user authentication is not enabled, the default values in the OnCommand Insight data source are ignored.

Having symauth enabled on Symmetrix arrays might inhibit the ability of OnCommand Insight to discover them. OnCommand Insight Acquisition runs as the SYSTEM user on the OnCommand Insight / Remote Acquisition Unit server that is communicating with the Solutions Enabler server. If `hostname\SYSTEM` does not have symauth privileges, OnCommand Insight fails to discover the array.

The EMC Solutions Enabler Symmetrix CLI data source includes support for device configuration for thin provisioning and Symmetrix Remote Data Facility (SRDF).

Definitions are supplied for Fibre Channel and Switch Performance packages.

Field	Description
Service Name	Service name as specified in netcnfg file

Full path to CLI	Full path to the Symmetrix CLI
------------------	--------------------------------

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data
Inventory Exclude Devices	Comma-separated list of device IDs to include or exclude
Connection Caching	<p>Choose connection caching method:</p> <ul style="list-style-type: none"> • LOCAL means that the OnCommand Insight Acquisition service is running on the Solutions Enabler server, which has Fibre Channel connectivity to the Symmetrix arrays you seek to discover, and has access to gatekeeper volumes. This might be seen in some Remote Acquisition Unit (RAU) configurations. • REMOTE_CACHED is the default and should be used in most cases. This uses the NETCNFG file settings to connect using IP to the Solutions Enabler server, which must have Fibre Channel connectivity to the Symmetrix arrays you seek to discover, and has access to Gatekeeper volumes. • In the event that REMOTE_CACHED options make CLI commands fail, use the REMOTE option. Keep in mind that it will slow down the acquisition process (possibly to hours or even days in extreme cases). The NETCNFG file settings are still used for an IP connection to the Solutions Enabler server that has Fibre Channel connectivity to the Symmetrix arrays being discovered. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  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. </div>
CLI Timeout (sec)	CLI process timeout (default 7200 seconds)

SMI-S Host IP	IP address of the SMI-S Provider Host
SMI-S Port	Port used by SMI-S Provider Host
Protocol	Protocol used to connect to the SMI-S provider
SMI-S Namespace	Interoperability namespace that the SMI-S provider is configured to use
SMI-S User Name	User name for the SMI-S Provider Host
SMI-S Password	User name for the SMI-S Provider Host
Performance Polling Interval (sec)	Interval between performance polls (default 1000 seconds)
Performance Filter Type	Specify whether to include or exclude the array list below when collecting performance data
Performance Filter Device List	Comma-separated list of device IDs to include or exclude
RPO Polling Interval (sec)	Interval between RPO polls (default 300 seconds)

EMC VNX data source

For configuration, the EMC VNX (SSH) data source requires the IP address of the Control Station and a *read-only* username and password.

Configuration

Field	Description
VNX IP	IP address or fully-qualified domain name of the VNX Control Station
VNX User Name	User name for the VNX Control Station
VNX Password	Password for the VNX Control Station

Requirements

- An IP address of the Control Station
- Read-only username and password.
- Access validation: Verify SSH access via PuTTY.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
VNX SSH Process Wait Timeout (sec)	VNX SSH process timeout (default 600 seconds)
Celerra Command Retry Attempts	Number of Celerra command retry attempts
CLARiiON External Process Timeout for Inventory (sec)	CLARiiON external process timeout for inventory (default 1800 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
CLARiiON External Process Timeout for Performance (sec)	CLARiiON external process timeout for performance (default 1800 seconds)

EMC VNXe data source

The EMC VNXe data source provides inventory support for EMC VNXe and Unity unified storage arrays.

This data source is CLI-based and requires that you install the Unisphere for VNXe CLI (uemcli.exe) on the acquisition unit that the VNXe data source resides on. uemcli.exe uses HTTPS as the transport protocol, so the acquisition unit must be able to initiate HTTPS connections to the VNXe/Unity arrays. You must have at least a read-only user for use by the data source.

Terminology

OnCommand Insight acquires the following inventory information from the EMC VNXe data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Storage Array	Storage
Processor	Storage Node
Storage Pool	Storage Pool
General iSCSI Block info, VMWare VMFS	Volume
Shared Folder	Internal Volume

CIFS Share, NFS Share, Share from VMWare NFS datastore	Share
Replication Remote System	Synchronization
iSCSI Node	iSCSI Target Node
iSCSI Initiator	iSCSI Target Initiator



These are common terminology mappings only and might not represent every case for this data source.

Requirements

The following are requirements to configure and use this data source:

- The VNXe data collector is CLI based; you must install the Unisphere for VNXe CLI, (uemcli.exe) onto the acquisition unit where your VNXe data collector resides.
- uemcli.exe uses HTTPS as the transport protocol, so the acquisition unit will need to be able to initiate HTTPS connections to the VNXe.
- You must have at least a read-only user for use by the data source.
- IP address of the managing Solutions enabler server.
- HTTPS on Port 443 is required
- The EMC VNXe data collector provides NAS and iSCSI support for inventory; fibre channel volumes will be discovered, but Insight does not report on FC mapping, masking, or storage ports.

Configuration

Field	Description
VNXe Storage	IP address or fully-qualified domain name of the VNXe device
User Name	User name for the VNXe device
Password	Password for the VNXe device
Full path to the uemcli executable	Full path to the uemcli.exe executable

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
VNXe CLI Port	Port used for the VNXe CLI

Inventory External Process Timeout (sec)	External process timeout (default 1800 seconds)
--	---

EMC VPLEX data source

For configuration, this data source requires an IP address of the VPLEX server and an administrative level domain account.

Terminology

OnCommand Insight acquires the following inventory information from the EMC VPLEX data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Cluster	Storage
Engine	Storage Node
Device, System Extend	Backend Storage Pool
Virtual Volume	Volume
Front-End Port, Back-End Port	Port
Distributed Device	Storage Synchronization
Storage View	Volume Map, Volume Mask
Storage Volume	Backend LUN
ITLs	Backend Path



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- An IP address of the VPLEX server
- Administrative level domain account for the VPLEX server
- Port 443 (HTTPS). Requires outbound connectivity to TCP port 443 on the VPLEX management station.
- For performance, read-only username and password for ssh/scp access.
- For performance, port 22 is required.
- Validate access: Verify by using `telnet` to port 443. For a port other than the default port, with any browser use `HTTPS://<ip>:<port>`

Configuration

Field	Description
IP address of VPLEX Management Console	IP address or fully-qualified domain name of the VPLEX Management Console
User Name	User name for VPLEX CLI
Password	Password used for VPLEX CLI
Performance Remote IP Address of VPLEX Management Console	Performance Remote IP address of the VPLEX Management Console
Performance Remote User Name	Performance Remote user name of VPLEX Management Console
Performance Remote Password	Performance Remote Password of VPLEX Management Console

Advanced configuration

Field	Description
Communication Port	Port used for VPLEX CLI
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection timeout (sec)	Connection timeout (default 60 seconds)
Number of Retries	Number of inventory retry attempts
Performance Poll Interval (sec)	Interval between performance polls (default 600 seconds)
Performance SSH Process Wait Timeout (sec)	SSH process timeout (default 600 seconds)
SSH Banner Wait Timeout (sec)	SSH banner wait timeout (default 20 seconds)
Number of Retries	Number of performance retry attempts

EMC XtremIO data source

To configure the EMC XtremIO (HTTP) data source, you must have the XtremIO Management Server (XMS) Host address and an account with administrator privileges.

Terminology

OnCommand Insight acquires the following inventory information from the EMC XtremIO data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk (SSD)	Disk
Cluster	Storage
Controller	Storage Node
Volume	Volume
LUN Map	Volume Map
Initiator, Target	Volume Mask



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- An IP address of each XtremIO Management Server
- An account with Administrator privileges
- Access to port 443 (HTTPS)

Configuration

Field	Description
XMS Host	IP address or fully-qualified domain name of the XtremIO Management Server
User name	User name for the XtremIO Management Server
Password	Password for the XtremIO Management Server

Advanced configuration

Field	Description
TCP port	TCP Port used to connect to XTremIO Management Server (default 443)
Inventory poll interval (min)	Interval between inventory polls (default 60 minutes)

Connection timeout (sec)	Connection timeout (default 60 seconds)
Performance poll interval(sec)	Interval between performance polls (default 300 seconds)

Fujitsu Eternus data source

The Fujitsu Eternus data source requires the IP address of the storage. It cannot be comma delimited.

Terminology

OnCommand Insight acquires the following inventory information from the Fujitsu Eternus data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Storage	Storage
Thin Pool, Flexible Tier Pool, Raid Group	Storage Pool
Standard Volume, Snap Data Volume (SDV), Snap Data Pool Volume (SDPV), Thin Provisioning Volume (TPV)	Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- An IP address of the Eternus storage, which cannot be comma delimited
- SSH Administration-level user name and password
- Port 22
- Ensure that the page scroll is disabled. (clienv-show-more-scroll disable)

Configuration

Field	Description
IP Address of Eternus Storage	IP address of the Eternus storage

User Name	User name for Eternus storage
Password	Password used for the sternus

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
SSH Process Wait Timeout (sec)	SSH process timeout (default 600 seconds)

Hitachi Content Platform (HCP) data source

This data collector supports the Hitachi Content Platform (HCP) using the HCP Management API.

Terminology

OnCommand Insight acquires the following inventory information from the HCP data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
HCP Cluster	Storage
Tenant	Storage Pool
Namespace	Internal Volume
Node	Node



These are common terminology mappings only and might not represent every case for this data source.

Inventory Requirements

- IP address of the HCP server
- Read-only user name and password for the HCP software and peer privileges

Configuration

Field	Description
HCP Host	IP address or fully-qualified domain name of the HCP host

HCP Port	Default is 9090
HCP user ID	User name for the HCP host
HCP Password	Password used for the HCP host
HCP Authentication Type	Choose HCP_LOCAL or ACTIVE_DIRECTORY

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Performance Polling Interval (sec)	Interval between performance polls (default 900 seconds)

HDS HiCommand Device Manager data source

The HDS HiCommand and HiCommand Lite data sources support the HiCommand Device Manager server. OnCommand Insight communicates with the HiCommand Device Manager server using the standard HiCommand API.

Terminology

OnCommand Insight acquires the following inventory information from the HDS HiCommand and HiCommand Lite data sources. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
PDEV	Disk
Journal Pool	Disk Group
Storage Array	Storage
Port Controller	Storage Node
Array Group, DP Pool	Storage Pool
Logical Unit, LDEV	Volume



These are common terminology mappings only and might not represent every case for this data source.

Inventory Requirements

- IP address of the HiCommand Device Manager server
- Read-only user name and password for the HiCommand Device Manager software and peer privileges
- Port requirements: 2001 (http) or 2443 (https)
- Validate access:
 - Log in to the HiCommand Device Manager software using peer user name and password.
 - Verify access to the HiCommand Device Manager API: `telnet <HiCommand Device_Manager_server_ip\> 2001`

Performance Requirements

- HDS USP, USP V, and VSP performance
 - Performance Monitor must be licensed.
 - Monitoring switch must be enabled.
 - The Export Tool (`Export.exe`) must be copied to the OnCommand Insight Server.
 - The Export Tool version must match the microcode version of the target array.
- HDS AMS performance
 - Performance Monitor needs to be licensed.
 - The Storage Navigator Modular 2 (SNM2) CLI utility needs to be installed on the OnCommand Insight Server.
 - You must register all AMS, WMS, SMS storage arrays whose performance needs to be acquired by OnCommand Insight by using the following command:

`auunitaddauto.exe -ip<IP address of Controller0>IP address of Controller1>`
 - You must ensure that all the arrays that you registered are listed in the output of this command:
`auunitref.exe.`

Configuration

Field	Description
HiCommand Server	IP address or fully-qualified domain name of the HiCommand Device Manager server
User Name	User name for the HiCommand Device Manager server.
Password	Password used for the HiCommand Device Manager server.

Devices - VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages	<p>Device list for VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages. Each storage requires:</p> <ul style="list-style-type: none"> • Array's IP: IP address of the storage • User Name: User name for the storage • Password: Password for the storage • Folder Containing Export Utility JAR Files: The folder containing the Export utility .jar files
SNM2Devices - WMS/SMS/AMS Storages	<p>Device list for WMS/SMS/AMS storages. Each storage requires:</p> <ul style="list-style-type: none"> • Array's IP: IP address of the storage • Storage Navigator CLI Path: SNM2 CLI path • Account Authentication Valid: Select to choose valid account authentication • User Name: User name for the storage • Password: Password for the storage
Choose Tuning Manager for Performance	Choose Tuning Manager for performance and override other performance options
Tuning Manager Host	IP address or fully-qualified domain name of tuning manager
Tuning Manager Port	Port used for Tuning Manager
Tuning Manager Username	User name for Tuning Manager
Tuning Manager Password	password for Tuning Manager



In HDS USP, USP V, and VSP, any disk can belong to more than one array group.

Advanced configuration

Field	Description
HiCommand Server Port	Port used for the HiCommand Device Manager
HTTPs Enabled	Select to enable HTTPs
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)

Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data
Exclude or Include Devices	Comma-separated list of device ID's or array names to include or exclude
Query Host Manager	Select to query host manager
HTTP Timeout (sec)	HTTP connection timeout (default 60 seconds)
Performance Polling Interval (sec)	Interval between performance polls (default 300 seconds)
Export timeout in seconds	Export utility timeout (default 300 seconds)

Hitachi Ops Center data collector

This data collector uses Hitachi Ops Center's integrated suite of applications to access inventory and performance data of multiple storage devices. For inventory and capacity discovery, your Ops Center installation must include both the "Common Services" and "Administrator" components. For performance collection, you must additionally have "Analyzer" deployed.

Terminology

OnCommand Insight acquires the following inventory information from this data collector. For each asset type acquired, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

Vendor/Model Term	OnCommand Insight Term
Storage Systems	Storage
Volume	Volume
Parity Groups	Storage Pool(RAID), Disk Groups
Disk	Disk
Storage Pool	Storage Pool(Thin, SNAP)
External Parity Groups	Storage Pool(Backend), Disk Groups
Port	Storage Node → Controller Node → Port
Host Groups	Volume Mapping and Masking
Volume Pairs	Storage Synchronization

Note: These are common terminology mappings only and might not represent every case for this data collector.

Inventory Requirements

You must have the following in order to collect inventory data:

- IP address or hostname of the Ops Center server hosting the "Common Services" component
- Root/sysadmin user account and password that exist on all servers hosting Ops Center components. HDS did not implement REST API support for usage by LDAP/SSO users until Ops Center 10.8+

Performance requirements

The following requirements must be met in order to collect performance data:

- The HDS Ops Center "Analyzer" module must be installed
- Storage arrays must be feeding the Ops Center "Analyzer" module

Configuration

Field	Description
Hitachi Ops Center IP Address	IP address or fully-qualified domain name of the Ops Center server hosting the "Common Services" component
User Name	User name for the Ops Center server.
Password	Password used for the Ops Center server.

Advanced configuration

Field	Description
Connection Type	HTTPS (port 443) is the default
Override TCP Port	Specify the port to use if not the default
Inventory Poll Interval (min)	Interval between inventory polls. The default is 40.
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data.
Filter device List	Comma-separated list of device serial numbers to include or exclude
Performance Poll Interval (sec)	Interval between performance polls. The default is 300.

HDS Storage

Terms applying to objects or references that you might find on HDS storage asset landing pages.

HDS Storage Terminology

The following terms apply to objects or references that you might find on HDS storage asset landing pages. Many of these terms apply to other data collectors as well.

- Name — comes directly from HDS HiCommand Device Manager’s “name” attribute via the GetStorageArray XML API call
- Model - comes directly from HDS HiCommand Device Manager’s “arrayType” attribute via the GetStorageArray XML API call
- Vendor — HDS
- Family - comes directly from HDS HiCommand Device Manager’s “arrayFamily” attribute via the GetStorageArray XML API call
- IP — this is the management IP address of the array, not an exhaustive list of all IP addresses on the array
- Raw Capacity — a base2 value representing the sum of the total capacity of all disks in this system, regardless of disk role.

HDS Storage Pool

Terms applying to objects or references that you might find on HDS storage pool asset landing pages.

HDS Storage Pool Terminology

The following terms apply to objects or references that you might find on HDS storage pool asset landing pages. Many of these terms apply to other data collectors as well.

- Type: The value here will be one of:
 - RESERVED — if this pool is dedicated for purposes other than data volumes, i.e, journaling, snapshots
 - Thin Provisioning — if this is a HDP pool
 - Raid Group — you will not likely see these for a few reasons:

OCI takes a strong stance to avoid double counting capacity at all costs. On HDS, one typically needs to build Raid Groups from disks, create pool volumes on those Raid Groups, and construct pools (often HDP, but could be special purpose) from those pool volumes. If OCI reported both the underlying Raid Groups as is, as well as the Pools, the sum of their raw capacity would vastly exceed the sum of the disks.

Instead, OCI’s HDS HiCommand data collector arbitrarily shrinks the size of Raid Groups by the capacity of pool volumes. This may result in OCI not reporting the Raid Group at all. Additionally, any resulting Raid Groups are flagged in a way such that they are not visible in the OCI WebUI, but they do flow into the OCI Data Warehouse (DWH). The purpose of these decisions is to avoid UI clutter for things that most users do not care about — if your HDS array has Raid Groups with 50MB free, you probably cannot use that free space for any meaningful outcome.

- Node - N/A, as HDS pools are not tied to any one specific node
- Redundancy - the RAID level of the pool. Possibly multiple values for a HDP pool comprised of multiple RAID types
- Capacity % - the percent used of the pool for data usage, with the used GB and total logical GB size of the pool
- Over-committed Capacity - a derived value, stating “the logical capacity of this pool is oversubscribed by this percentage by virtue of the sum of the logical volumes exceeding the logical capacity of the pool by this percentage”
- Snapshot - shows the capacity reserved for snapshot usage on this pool

HDS Storage Node

Terms applying to objects or references that you might find on HDS storage node asset landing pages.

HDS Storage Node Terminology

The following terms apply to objects or references that you might find on HDS storage node asset landing pages. Many of these terms apply to other data collectors as well.

- Name — The name of the Front-end director (FED) or Channel Adapter on monolithic arrays, or the name of the controller on a modular array. A given HDS array will have 2 or more Storage Nodes
- Volumes — The Volume table will show any volume mapped to any port owned by this storage node

Hitachi Ops Center data collector

This data collector uses Hitachi Ops Center's integrated suite of applications to access inventory and performance data of multiple storage devices. For inventory and capacity discovery, your Ops Center installation must include both the "Common Services" and "Administrator" components. For performance collection, you must additionally have "Analyzer" deployed.

Terminology

OnCommand Insight acquires the following inventory information from this data collector. For each asset type acquired, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

Vendor/Model Term	OnCommand Insight Term
Storage Systems	Storage
Volume	Volume
Parity Groups	Storage Pool(RAID), Disk Groups
Disk	Disk
Storage Pool	Storage Pool(Thin, SNAP)
External Parity Groups	Storage Pool(Backend), Disk Groups
Port	Storage Node → Controller Node → Port
Host Groups	Volume Mapping and Masking
Volume Pairs	Storage Synchronization

Note: These are common terminology mappings only and might not represent every case for this data collector.

Inventory Requirements

You must have the following in order to collect inventory data:

- IP address or hostname of the Ops Center server hosting the "Common Services" component

- Root/sysadmin user account and password that exist on all servers hosting Ops Center components. HDS did not implement REST API support for usage by LDAP/SSO users until Ops Center 10.8+

Performance requirements

The following requirements must be met in order to collect performance data:

- The HDS Ops Center "Analyzer" module must be installed
- Storage arrays must be feeding the Ops Center "Analyzer" module

Configuration

Field	Description
Hitachi Ops Center IP Address	IP address or fully-qualified domain name of the Ops Center server hosting the "Common Services" component
User Name	User name for the Ops Center server.
Password	Password used for the Ops Center server.

Advanced configuration

Field	Description
Connection Type	HTTPS (port 443) is the default
Override TCP Port	Specify the port to use if not the default
Inventory Poll Interval (min)	Interval between inventory polls. The default is 40.
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data.
Filter device List	Comma-separated list of device serial numbers to include or exclude
Performance Poll Interval (sec)	Interval between performance polls. The default is 300.

HDS NAS (HNAS) data source

The HDS NAS (HNAS) data source is an inventory and configuration data source to support discovery of HDS NAS clusters. Insight supports discovering NFS and CIFS shares, file systems (Insight Internal Volumes), and spans (Insight Storage Pools).

This data source is SSH based, so the acquisition unit that will host it needs to be able to initiate SSH sessions to TCP 22 on the HNAS itself, or the Systems Management Unit (SMU) that the cluster is connected to.

Terminology

OnCommand Insight acquires the following inventory information from the HNAS data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Tier	Disk Group
Cluster	Storage
Node	Storage Node
Span	Storage Pool
File System	Internal Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

The following are requirements to configure and use this data source:

- Device IP address
- Port 22, SSH protocol
- Username and password - privilege level: Supervisor
- NOTE: This data collector is SSH based, so the AU that hosts it must be able to initiate SSH sessions to TCP 22 on the HNAS itself, or the Systems Management Unit (SMU) that the cluster is connected to.



This data collector is SSH based, so the AU that hosts it must be able to initiate SSH sessions to TCP 22 on the HNAS itself, or the Systems Management Unit (SMU) that the cluster is connected to.

Configuration

Field	Description
HNAS Host	IP address or fully-qualified domain name of HNAS Management Host
User Name	User name for HNAS CLI
Password	Password used for HNAS CLI

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 30 minutes)
SSH Banner Wait Timeout (sec)	SSH banner wait timeout (default 15 seconds)

SSH Command Timeout (sec)	SSH command timeout (default 30 seconds)
---------------------------	--

HP CommandView AE data source

The HP CommandView Advanced Edition (AE) and CommandView AE CLI/SMI (AE Lite) data sources support inventory and performance from a CommandView (also referred to as HiCommand) Device Manager server.

Terminology

OnCommand Insight acquires the following inventory information from the HP CommandView AE and AE Lite data sources. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
PDEV	Disk
Journal Pool	Disk Group
Storage Array	Storage
Port Controller	Storage Node
Array Group, DP Pool	Storage Pool
Logical Unit, LDEV	Volume



These are common terminology mappings only and might not represent every case for this data source.

Inventory Requirements

- IP address of the HiCommand Device Manager server
- Read-only user name and password for the CommandView AE software and peer privileges
- The CommandView AE Lite version of the device manager has only the CLI licensed
- Port requirement: 2001

Performance Requirements

- HDS USP, USP V, and VSP performance
 - Performance Monitor must be licensed.
 - Monitoring switch must be enabled.
 - The Export Tool (`Export.exe`) must be copied to the OnCommand Insight Server.
 - The Export Tool version must match the microcode version of the target array.

- HDS AMS performance
 - Performance Monitor needs to be licensed.
 - The Storage Navigator Modular 2 (SNM2) CLI utility needs to be installed on the OnCommand Insight Server.
 - You must register all AMS, WMS, SMS storage arrays whose performance needs to be acquired by OnCommand Insight by using the following command:

```
auunitaddauto.exe -ip<IP address of Controller0>IP address of Controller1>
```

- You must ensure that all the arrays that you registered are listed in the output of this command:
auunitref.exe.

Configuration

Field	Description
HiCommand Server	IP address or fully-qualified domain name of the HiCommand Device Manager server
User Name	User name for the HiCommand Device Manager server.
Password	Password used for the HiCommand Device Manager server.
Devices - USP, USP V, VSP/R600 Storages	Device list for VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages. Each storage requires: <ul style="list-style-type: none"> • Array's IP: IP address of the storage • User Name: User name for the storage • Password: Password for the storage • Folder Containing Export Utility JAR Files: The folder containing the Export utility .jar files
SNM2Devices - WMS/SMS/AMS Storages	Device list for WMS/SMS/AMS storages. Each storage requires: <ul style="list-style-type: none"> • Array's IP: IP address of the storage • Storage Navigator CLI Path: SNM2 CLI path • Account Authentication Valid: Select to choose valid account authentication • User Name: User name for the storage • Password: Password for the storage
Choose Tuning Manager for Performance	Choose Tuning Manager for performance and override other performance options

Tuning Manager Host	IP address or fully-qualified domain name of tuning manager
Tuning Manager Port	Port used for Tuning Manager
Tuning Manager Username	User name for Tuning Manager
Tuning Manager Password	password for Tuning Manager



In HDS USP, USP V, and VSP, any disk can belong to more than one array group.

Advanced configuration

Field	Description
HiCommand Server Port	Port used for the HiCommand Device Manager
HTTPs Enabled	Select to enable HTTPs
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data
Exclude or Include Devices	Comma-separated list of device ID's or array names to include or exclude
Query Host Manager	Select to query host manager
HTTP Timeout (sec)	HTTP connection timeout (default 60 seconds)
Performance Polling Interval (sec)	Interval between performance polls (default 300 seconds)
Export timeout in seconds	Export utility timeout (default 300 seconds)

HP EVA Storage data source

For configuration, The EVA Storage (SSSU) data source requires the IP address of the Command View (CV) server and a *read-only* username and password to the CV software. The user must be defined in the CV software.

Terminology

OnCommand Insight acquires the following inventory information from the HP EVA data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Disk Group	Disk Group (not modeled)
Storage Cell	Storage
Virtual Disk	Storage Pool
Virtual Disk	Volume



These are common terminology mappings only and might not represent every case for this data source.

Inventory Requirements

- IP address of the CV server
- Read-only username and password to the CV software. The user must be defined in the CV software.
- Third-party software installed on the OnCommand Insight Server/RAU: `sssu.exe`. The `sssu.exe` version should correspond to the CV version.
- Access validation: Run `sssu.exe` commands using username and password.

Performance Requirements

The HP StorageWorks Command View EVA software suite must be installed on the OnCommand Insight Server. Alternatively, you can install a Remote Acquisition Unit (RAU) on the EVA server:

1. Install HP StorageWorks Command View EVA Software Suite on the OnCommand Insight Server, or install a Remote Acquisition Unit on the Command View EVA server.
2. Locate the `evaperf.exe` command. For example, `c:\Program Files\Hewlett-Packard\EVA Performance Monitor\`
3. Using the IP of the Command View server, perform these steps:
 - a. Run this command where 860 is the default port `Evaperf.exe server <Command View Server IP> 860 <username>`
 - b. Enter the Command View server password at the password prompt.

This should return a command line prompt and nothing else.

4. Verify the setup by running `evaperf.exe ls`.

You should see a list of arrays or controllers managed by the Command View server. Each line shows a controller on an EVA array.

Configuration

Field	Description
-------	-------------

CommandView Server	IP address or fully-qualified domain name of the EVA Storage Manager
User Name	User name for the Command View manager. The name must be defined in Command View.
Password	Password used for the Command View manager.
Performance User Name	For performance, the user name for the Command View manager. The name must be defined in Command View.
Performance Password	For performance, the password used for the Command View manager.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
CLI Home	Full pathname to the CLI home directory where <code>sssu.exe</code> is located
Inventory Exclude Devices	Comma-separated list of device names to include
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Performance CLI Home	For Array Performance, full pathname to the CLI home directory where <code>sssu.exe</code> is located. To validate access, run <code>sssu.exe</code>
Command Timeout (sec)	<code>evaperf</code> command wait timeout (default 600 seconds)
Performance Exclude Devices	Comma-separated list of device names to exclude from collecting performance data

HPE Nimble data source

The HPE Nimble data collector supports inventory and performance data for HPE Nimble storage arrays.

Terminology

OnCommand Insight acquires the following inventory information from the HPE Nimble data source. For each

asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Array	Storage
Disk	Disk
Pool	Storage Pool
Volume	Volume
Initiator	Storage Host Alias
Controller	Storage Node
Fibre Channel Interface	Controller



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- The array must be installed and configured, and reachable from the client through its fully qualified domain name (FQDN) or array management IP address.
- The array must be running NimbleOS 2.3.x or later.
- You must have a valid user name and password to the array.
- Port 5392 must be open on the array.

Configuration

Field	Description
Array Management IP Address	Fully qualified domain name (FQDN) or array management IP address.
User Name	User name for the Nimble array
Password	Password for the Nimble array

Advanced configuration

Field	Description
Port	Port used by Nimble REST API. The default is 5392.

Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
-------------------------------	---

Note: The default performance poll interval is 300 seconds and can not be changed. This is the only interval supported by Nimble.

Huawei OceanStor data source

OnCommand Insight uses the Huawei OceanStor (REST/HTTPS) data source to discover inventory for Huawei OceanStor storage.

Terminology

OnCommand Insight acquires the following inventory and performance information from the Huawei OceanStor. For each asset type acquired by OnCommand Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

Vendor/Model Term	OnCommand Insight Term
Storage Pool	Storage Pool
File System	Internal Volume
Controller	Storage Node
FC Port (Mapped)	Volume Map
Host FC Initiator (Mapped)	Volume Mask
NFS/CIFS Share	Share
Share	iSCSI Target Node
iSCSI Link Initiator	iSCSI Initiator Node
Disk	Disk
LUN	Volume

Requirements

The following are requirements to configure and use this data collector:

- Device IP
- Credentials to access OceanStor device manager
- Port 8088 must be available

Configuration

Field	Description
OceanStor Host IP Address	IP address or fully-qualified domain name of the OceanStor Device Manager
User Name	Name used to log into the OceanStor Device Manager
Password	Password used to log into the OceanStor Device Manager

Advanced configuration

Field	Description
TCP Port	TCP Port used to connect to OceanStor Device Manager (default 8088)
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Connection Timeout (sec)	Connection timeout (default 60 seconds)

IBM Cleversafe data source

This data source collects inventory and performance data for IBM Cleversafe.

Requirements

The following are requirements for configuring this data source:

- Manager IP Address or Host Name
- A username and password for same
- Port 9440

Configuration

Field	Description
Cleversafe manager Host Name or IP Address	Host IP address of the CleverSafe device
User Name	Name used to log into the Cleversafe
Password	Password used to log into the Cleversafe

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Default is 60 minutes
HTTP Connection Timeout)	Default is 60 seconds

IBM DS data source

The IBM DS (CLI) data source supports DS6xxx and DS8xxx devices only. DS3xxx, DS4xxx, and DS5xxx devices are supported by the NetApp E-Series data source. You should refer to the Insight data source support matrix for supported models and firmware versions.

Terminology

OnCommand Insight acquires the following inventory information from the IBM DS data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk Drive Module	Disk
Storage Image	Storage
Extent Pool	Storage Pool
Fixed Block Volume	Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- IP address of each DS array
- Storage Display Name is optional and cosmetic only
- Read-only username and password on each DS array
- Third-party software installed on the Insight server: IBM dscli
- Access validation: Run `dscli` commands using the username and password
- Port requirements: 80, 443, & 1750

Configuration

Field	Description
-------	-------------

DS storage	IP address or fully-qualified domain name of the DS Storage Host
User Name	Name used for the DS CLI
Password	Password used for the DS CLI
Executable dscli.exe Path	Full path to the <code>dscli.exe</code> utility.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Storage Display Name	Name of the IBM DS storage array
Inventory Exclude Devices	Comma-separated list of device serial numbers to exclude from inventory collection
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Performance Filter Type	Include: Data collected only from devices on list. Exclude: No data from these devices is collected
Performance Filter Device List	Comma-separated list of device IDs to include or exclude from performance collection

IBM PowerVM data source

The IBM PowerVM (SSH) data source collects information about virtual partitions running on IBM POWER hardware instances managed by a hardware management console (HMC). For configuration, this data source requires the user name to log in to the HMC through SSH, and the view-level permission on HMC configurations.

Terminology

OnCommand Insight acquires the following inventory information from the IBM PowerVM data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
hdisk	Virtual Disk
Managed System	Host

LPAR, VIO Server	Virtual Machine
Volume Group	Data Store
Physical Volume	LUN



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- IP address of the Hardware Management Console (HMC)
- User name and password that provide access to HMC through SSH
- Port requirement SSH-22
- View permission on all management systems and logical partition security domains

The user must also have View permission on HMC configurations and the ability to collect VPD information for the HMC console security grouping. The user must also be allowed Virtual IO Server Command access under the Logical Partition security grouping. It is a best practice to start from a role of an operator and then remove all roles. Read-only users on the HMC do not have privileges to run proxied commands on AIX hosts.

- IBM best practice is to have the devices monitored by two or more HMCs. Be aware that this may cause OnCommand Insight to report duplicated devices, therefore it is highly recommended to add redundant devices to the "Exclude Devices" list in the Advanced Configuration for this data collector.

Configuration

Field	Description
Hardware Management Console (HMC) Address	IP address or fully-qualified domain name of the PowerVM Hardware Management Console
HMC User	User name for the Hardware Management Console
Password	Password used for the Hardware Management Console

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
SSH Port	Port used for SSH to the PowerVM
SSH Process Wait Timeout (sec)	SSH process timeout (default 600 seconds)

Number of Retries	Number of inventory retry attempts
Exclude Devices	Comma-separated list of device IDs or display names to exclude

IBM SVC data source

The IBM SVC data source collects inventory and performance data using SSH, supporting a variety of devices that run the SVC operating system. The list of supported devices includes models such as the SVC, the v7000, the v5000, and the v3700. Refer to the Insight data source support matrix for supported models and firmware versions.

Terminology

OnCommand Insight acquires the following inventory information from the IBM SVC data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive	Disk
Cluster	Storage
Node	Storage Node
Mdisk Group	Storage Pool
Vdisk	Volume
Mdisk	Backend LUN



These are common terminology mappings only and might not represent every case for this data source.

Inventory Requirements

- IP address of each SVC cluster
- Port 22 available
- Public and private key pair that you either generate with Insight or reuse a keypair already in use on your SVC

If you are reusing an existing keypair, you must convert them from Putty format to OpenSSH format.

- Public key installed on the SVC cluster
- Private key needs to be identified in the data source

- Access validation: Open `ssh` session to the SVC cluster using the private key



No third-party software needs to be installed.

Performance Requirements

- SVC Console, which is mandatory for any SVC cluster and required for the SVC discovery foundation package.
- Administrative access level required only for copying performance data files from cluster nodes to the config node.



Because this access level is not required for the SVC foundation discovery package, the SVC foundation user might not work successfully.

- Port 22 required
- A private and public SSH key must be generated for this user, and the private key stored so that it is accessible from the Acquisition Unit. If the SVC foundation user has the proper permissions, then the same user and key works. The same SSH key can be used for inventory and performance data.
- Enable data collection by connecting to the SVC cluster by SSH and running: `svctask startstats -interval 1`



Alternatively, enable data collection using the SVC management user interface.

Parent Serial Number explained

Traditionally Insight is capable of reporting the storage array serial number, or the individual storage node serial numbers. However, some storage array architectures do not cleanly align to this. An SVC cluster can be comprised of 1-4 appliances, and each appliance has 2 nodes. If the appliance itself has a serial number, that serial number is neither the serial number for the cluster nor the nodes.

The attribute "Parent Serial Number" on the storage node object is populated appropriately for IBM SVC arrays when the individual nodes sit inside an intermediate appliance/enclosure that is just part of a larger cluster.

Configuration

Field	Description
Cluster/s IP	IP address of fully-qualified domain name for the SVC storage
Choose 'Password' or 'OpenSSH Key File' to specify credential type	The credential type used to connect to the device via SSH
Inventory User Name	User name for the SVC CLI
Inventory Password	Password for the SVC CLI
Full Path to Inventory Private Key	Full path to the Inventory private key file

Performance User Name	User name for the SVC CLI for performance collection
Performance Password	Password for the SVC CLI for performance collection
Full Path to Performance Private Key	Full path to the Performance private key file

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Exclude Devices	Comma-separated list of device IDs to exclude from inventory collection
SSH Process Wait Timeout (sec)	SSH process timeout (default 200 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Performance Exclude Devices	Comma-separated list of device IDs to exclude from performance collection
Performance SSH Process Wait Timeout (sec)	SSH process timeout (default 200 seconds)
To clean up dumped stats files	Select to clean up dumped stats files

IBM Tivoli Monitoring data source

This data source is used solely for File System Utilization. It communicates directly with the Tivoli Monitoring Database, also known as the Tivoli Monitoring Data Warehouse. Oracle and DB2 databases are supported.

Oracle error message



This data collector is no longer available starting with OnCommand Insight 7.3.11.

If the specified SID results in the error message containing "ORA-12154" on attempting to connect, double-check your Oracle DB network service configuration. If the access configuration specifies a fully qualified hostname (for example, "NAMES.DEFAULT_DOMAIN"), try inserting the fully qualified service name in the SID field. A simple example would be that the connection to SID `testdb` is failing and your Oracle configuration specifies a domain of `company.com`. The following string can be used instead of the base SID to try to connect: `testdb.company.com`.

Configuration

Field	Description
-------	-------------

Tivoli Monitoring Database IP	IP address or fully-qualified domain name of the Tivoli Monitoring server
User Name	User name for the Tivoli Monitoring server
Password	Password for the Tivoli Monitoring server

Advanced configuration

Field	Description
Tivoli Monitoring Database Port	Port used for Tivoli monitoring database
Oracle SID or DB2 Database Name	Oracle listener service ID or DB2 database name
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Database Driver to Use	Choose Database Driver to use
Protocol Used to Connect to the Database	Protocol Used to Connect to the Database
Database Schema	Enter Database Schema

IBM TotalStorage DS4000 data source

This data source collects inventory and performance information. There are two possible configurations (firmware 6.x and 7.x+), and they both have the same values. The API collects the volume data statistics.

Configuration

Field	Description
Comma Separated List of Array SANtricity Controller IPs	IP addresses or fully-qualified domain names of controllers, separated by commas

Requirements

- IP address of each DS5 or FAStT array
- Access validation: Ping the IP address of both controllers on each array.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 30 minutes)

Performance Poll Interval (up to 3600 seconds)	Interval between performance polls (default 300 seconds)
--	--

IBM XIV data source

IBM XIV (CLI) data source inventory is performed by using the XIV command-line interface. XIV performance is accomplished by making SMI-S calls to the XIV array, which runs a SMI-S provider on port 5989.

Terminology

OnCommand Insight acquires the following inventory information from the IBM XIV data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Storage System	Storage
Storage Pool	Storage Pool
Volume	Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- Port requirement: TCP port 7778
- IP address of the XIV management interface
- Read-only user name and password
- The XIV CLI must be installed on the Insight server or RAU
- Access validation: Log in to the XIV user interface from the Insight server using the user name and password.

Configuration

Field	Description
IP Address	IP address or fully-qualified domain name for the XIV storage
User Name	User name for the XIV storage

Password	Password for the XIV storage
Full path to XIV CLI directory	Full path to the XIV CLI directory

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
CLI Process Wait Timeout (ms)	CLI process timeout (default 7200000 ms)
SMI-S Host IP	IP address of the SMI-S Provider Host
SMI-S Port	Port used by SMI-S Provider Host
SMI-S Protocol	Protocol used to connect to the SMI-S provider
SMI-S Namespace	SMI-S namespace
Username	User name for the SMI-S Provider Host
Password	Password for the SMI-S Provider Host
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Number of SMI-S Connection Retries	Number of SMI-S connection retry attempts

Infinidat InfiniBox data source

The Infinidat InfiniBox (HTTP) data source is used to collect information from the Infinidat InfiniBox storage. You must have access to the InfiniBox Management Node.

Terminology

OnCommand Insight acquires the following inventory information from the InfiniBox data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive	Disk
InfiniBox	Storage

Node	Storage Node
Pool	Storage Pool
Volume	Volume
FC Port	Port
Filesystem	Internal Volume
Filesystem	FileShare
Filesystem Exports	Share



These are common terminology mappings only and might not represent every case for this data source.

Configuration

Field	Description
InfiniBox Host	IP address or fully-qualified domain name of the InfiniBox Management Node
User Name	User name for InfiniBox Management Node
Password	Password for the InfiniBox Management Node

Advanced configuration

Field	Description
TCP Port	TCP Port used to connect to InfiniBox Server (default 443)
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Connection Timeout	Connection timeout (default 60 seconds)

Microsoft Azure compute data source

OnCommand Insights uses the Azure compute data collector to acquire inventory and performance data from Azure compute instances.

Requirements

You need the following information to configure this data collector:

- Port requirement: 443 HTTPS
- Azure Management Rest IP (management.azure.com)
- Azure Service Principal Application (Client) ID (user account)
- Azure Service Principal Authentication key (user password)

You need to set up an Azure account for Insight discovery. Once the account is properly configured and you register the application in Azure, you will have the credentials required to discover the Azure instance with Insight. The following link describes how to set up the account for discovery:<https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-create-service-principal-portal>

Configuration

Enter data into the data source fields according to the table below:

Field	Description
Azure Service Principal Application (Client) ID (Reader role required)	Sign-in ID to Azure. Requires Reader Role access.
Azure tenant ID	Microsoft tenant ID
Azure Service Principal Authentication Key	Login authentication key
I understand Microsoft bills me for API requests	Check this to verify your understanding that Microsoft bills you for API requests made by Insight polling.

Advanced Configuration

Enter data into the data source fields according to the table below:

Field	Description
Inventory Poll Interval (min)	The default is 60
Choose 'Exclude' or 'Include' to Apply to Filter VMs by Tags	Specify whether to include or exclude VM's by Tags when collecting data. If 'Include' is selected, the Tag Key field can not be empty.
Tag Keys and Values on which to Filter VMs	Click + Filter Tag to choose which VMs (and associated disks) to include/exclude by filtering for keys and values that match keys and values of tags on the VM. Tag Key is required, Tag Value is optional. When Tag Value is empty, the VM is filtered as long as it matches the Tag Key.

Performance Poll Interval (sec)	The default is 300
---------------------------------	--------------------

Azure NetApp Files data source

This data source acquires inventory and performance data for Azure NetApp Files (ANF).

Requirements

The following are requirements for configuring this data source:

- Port requirement: 443 HTTPS
- Azure Management Rest IP (management.azure.com)
- Azure Service Principal Application (Client) ID (user account)
- Azure Service Principal authentication key (user password)
- You need to set up an Azure account for Cloud Insights discovery.

Once the account is properly configured and you register the application in Azure, you will have the credentials required to discover the Azure instance with Cloud Insights. The following link describes how to set up the account for discovery:

<https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-create-service-principal-portal>

Configuration

Field	Description
Azure Service Principal Application (Client) ID	Sign-in ID to Azure
Azure Tenant ID	Azure Tenant ID
Azure Service Principal Authentication Key	Login authentication key
I understand Microsoft bills me for API requests	Check this to verify your understanding that Microsoft bills you for API requests made by Insight polling.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Default is 60 minutes

Microsoft Hyper-V data source

For configuration, the Microsoft Hyper-V data source requires the IP address or the resolvable DNS name for the physical host (hypervisor). This data source uses Powershell (previously used WMI).

Terminology

OnCommand Insight acquires the following inventory information from the Hyper-V data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Virtual hard Disk	Virtual Disk
Host	Host
Virtual Machine	Virtual Machine
Cluster Shared Volumes (CSV), Partition Volume	Data Store
Internet SCSI Device, Multi Path SCSI LUN	LUN
Fiber Channel Port	Port



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- The Hyper-V requires port 5985 opened for data collection and remote access/management.
- IP address of Clustering group node
- Local Administrator user & password on the hypervisor
- Administrative-level user account
- Port requirements: Port 135 and Dynamic TCP ports assigned 1024-65535 for Windows 2003 and older and 49152-65535 for Windows 2008.
- DNS resolution must succeed, even if the data collector is pointed at only an IP address.
- Each Hyper-V hypervisor must have “Resource Metering” turned on for every VM, on every host. This allows each hypervisor to have more data available for Cloud Insights on each guest. If this is not set, fewer performance metrics are acquired for each guest. More information on Resource metering can be found in the microsoft documentation:

[Hyper-V Resource Metering Overview](#)

[Enable-VMResourceMetering](#)

Configuration

Field	Description
Physical Host IP Address	The IP address or fully-qualified domain name for the physical host (hypervisor)

User Name	Administrator user name for the hypervisor
Password	Password for the hypervisor
NT Domain	The DNS name used by the nodes in the cluster

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection Timeout (ms)	Connection timeout (default 60000 ms)

NetApp Clustered Data ONTAP data source

This data source should be used for storage systems using Clustered Data ONTAP, and requires an administrator account used for read-only API calls.

Terminology

OnCommand Insight acquires the following inventory information from the Clustered Data ONTAP data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Raid Group	Disk Group
Cluster	Storage
Node	Storage Node
Aggregate	Storage Pool
LUN	Volume
Volume	Internal Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- Administrator account used for read-only API calls
- Target IP is the cluster management LIF
- Username (with read-only role name to ontapi application to the default Vserver) and password to log into NetApp cluster
- Port requirements: 80 or 443
- License requirements: FCP license and mapped/masked volumes required for discovery

Configuration

Field	Description
NetApp Management IP	IP address or fully-qualified domain name of the NetApp cluster
User Name	User name for the NetApp cluster
Password	Password for the NetApp cluster

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Clustered Data ONTAP Storage

Terms applying to objects or references that you might find on NetApp Clustered Data ONTAP storage asset landing pages.

Clustered Data ONTAP Storage Terminology

The following terms apply to objects or references that you might find on NetApp Clustered Data ONTAP storage asset landing pages. Many of these terms apply to other data collectors as well.

- Model — A comma delimited list of the unique, discrete node model names within this cluster. If all the nodes in the clusters are the same model type, just one model name will appear.
- Vendor — same Vendor name you would see if you were configuring a new data source.
- Serial number — The array serial number. On cluster architecture storage systems like NetApp Clustered Data Ontap, this serial number may be less useful than the individual “Storage Nodes” serial numbers.
- IP — generally will be the IP(s) or hostname(s) as configured in the data source.
- Microcode version — firmware.
- Raw Capacity — base 2 summation of all the physical disks in the system, regardless of their role.

- Latency — a representation of what the host facing workloads are experiencing, across both reads and writes. Ideally, OCI is sourcing this value directly, but this is often not the case. In lieu of the array offering this up, OCI is generally performing an IOPs-weighted calculation derived from the individual internal volumes' statistics.
- Throughput — aggregated from internal volumes.
- Management — this may contain a hyperlink for the management interface of the device. Created programmatically by the Insight data source as part of inventory reporting.

Clustered Data ONTAP Storage Pool

Terms applying to objects or references that you might find on NetApp Clustered Data ONTAP storage pool asset landing pages.

Clustered Data ONTAP Storage Pool Terminology

The following terms apply to objects or references that you might find on NetApp Clustered Data ONTAP storage pool asset landing pages. Many of these terms apply to other data collectors as well.

- Storage — what storage array this pool lives on. Mandatory.
- Type — a descriptive value from a list of an enumerated list of possibilities. Most commonly will be "Aggregate" or "RAID Group".
- Node — if this storage array's architecture is such that pools belong to a specific storage node, its name will be seen here as a hyperlink to its own landing page.
- Uses Flash Pool — Yes/No value — does this SATA/SAS based pool have SSDs used for caching acceleration?
- Redundancy — RAID level or protection scheme. RAID_DP is dual parity, RAID_TP is triple parity.
- Capacity — the values here are the logical used, usable capacity and the logical total capacity, and the percentage used across these.
- Over-committed capacity — If by using efficiency technologies you have allocated a sum total of volume or internal volume capacities larger than the logical capacity of the storage pool, the percentage value here will be greater than 0%.
- Snapshot — snapshot capacities used and total, if your storage pool architecture dedicates part of its capacity to segments areas exclusively for snapshots. Ontap in MetroCluster configurations are likely to exhibit this, while other Ontap configurations are less so.
- Utilization — a percentage value showing the highest disk busy percentage of any disk contributing capacity to this storage pool. Disk utilization does not necessarily have a strong correlation with array performance — utilization may be high due to disk rebuilds, deduplication activities, etc in the absence of host driven workloads. Also, many arrays' replication implementations may drive disk utilization while not showing as internal volume or volume workload.
- IOPS — the sum IOPs of all the disks contributing capacity to this storage pool.
- Throughput — the sum throughput of all the disks contributing capacity to this storage pool.

Clustered Data ONTAP Storage Node

Terms applying to objects or references that you might find on NetApp Clustered Data ONTAPs storage node asset landing pages.

Clustered Data ONTAP Storage Node Terminology

The following terms apply to objects or references that you might find on NetApp Clustered Data ONTAP storage pool asset landing pages. Many of these terms apply to other data collectors as well.

- Storage — what storage array this node is part of. Mandatory.
- HA Partner — on platforms where a node will fail over to one and only one other node, it will generally be seen here.
- State — health of the node. Only available when the array is healthy enough to be inventoried by a data source.
- Model — model name of the node.
- Version — version name of the device.
- Serial number — The node serial number.
- Memory — base 2 memory if available.
- Utilization — On Ontap, this is a controller stress index from a proprietary algorithm. With every performance poll, a number between 0 and 100% will be reported that is the higher of either WAFL disk contention, or average CPU utilization. If you observe sustained values > 50%, that is indicative of undersizing — potentially a controller/node not large enough or not enough spinning disks to absorb the write workload.
- IOPS — Derived directly from Ontap ZAPI calls on the node object.
- Latency — Derived directly from Ontap ZAPI calls on the node object.
- Throughput — Derived directly from Ontap ZAPI calls on the node object.
- Processors — CPU count.

NetApp Clustered Data ONTAP for Unified Manager data source

This data source collects ONTAP 8.1.x data from the Unified Manager (UM) 6.0+ database. Using this data source, Insight discovers all clusters configured and populated in UM. For efficiency, Insight does not call ZAPIs on the cluster itself. Performance is not supported in this data source.

Configuration



This data collector is no longer available starting with OnCommand Insight 7.3.11.

Field	Description
Unified Manager IP	IP address or fully-qualified domain name of the Unified Manager
User Name	User name for the Unified Manager
Password	Password for the Unified Manager
Port	Port used for communication with the Unified Manager (default 3306)

Advanced configuration

Field	Description
Inventory Poll Interval (min) Interval	Interval between inventory polls (default 15 minutes)
Exclude Clusters	Comma-separated list of cluster IPs to exclude

NetApp Data ONTAP operating in 7-Mode data source

For storage systems using Data ONTAP software operating in 7-Mode, you should use the ONTAPI data source, which uses the CLI to obtain capacity numbers.

Terminology

OnCommand Insight acquires the following inventory information from the NetApp Data ONTAP 7-Mode data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Raid Group	Disk Group
Filer	Storage
Filer	Storage Node
Aggregate	Storage Pool
LUN	Volume
Volume	Internal Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- IP address of the FAS storage controller and partner
- Port 443
- User name and password for the controller and the partner
- A custom admin level username and password for controller and partner controller with the following role capabilities for 7-Mode:
 - "api-*": Use this to allow OnCommand Insight to execute all NetApp storage API commands.
 - "login-http-admin": Use this to allow OnCommand Insight to connect to the NetApp storage via HTTP.

- "security-api-vfiler": Use this to allow OnCommand Insight to execute NetApp storage API commands to retrieve vFiler unit information.
- "cli-options": Use this to read storage system options.
- "cli-lun": Access these commands for managing LUNs. Displays the status (LUN path, size, online/offline state, and shared state) of the given LUN or class of LUNs.
- "cli-df": Use this to display free disk space.
- "cli-ifconfig": Use this to display interfaces and IP addresses.

Configuration

Field	Description
Address of Filer	IP address or fully-qualified domain name for the NetApp Filer
User Name	User name for the NetApp Filer
Password	password for the NetApp Filer
Address of HA Partner Filer in Cluster	IP address or fully-qualified domain name for the HA Partner Filer
User Name of HA Partner Filer in Cluster	User name for the NetApp HA Partner Filer
Password of HA Partner Filer in Cluster	password for the NetApp HA Partner Filer

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection Type	Choose connection type
Connection Port	Port used for NetApp API
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Storage systems connection

As an alternative to using the default administrative user for this data source, you can configure a user with administrative rights directly on the NetApp storage systems so that this data source can acquire data from NetApp storage systems.

Connecting to NetApp storage systems requires that the user, who is specified when acquiring the main pfiler (on which the storage system exist), meet the following conditions:

- The user must be on vfiler0 (root filer/pfiler).

Storage systems are acquired when acquiring the main pfiler.

- The following commands define the user role capabilities:
 - "api-*": Use this to allow OnCommand Insight to execute all NetApp storage API commands. This command is required to use the ZAPI.
 - "login-http-admin": Use this to allow OnCommand Insight to connect to the NetApp storage via HTTP. This command is required to use the ZAPI.
 - "security-api-vfiler": Use this to allow OnCommand Insight to execute NetApp storage API commands to retrieve vFiler unit information.
 - "cli-options": For "options" command and used for partner IP and enabled licenses.
 - "cli-lun": Access these command for managing LUNs. Displays the status (LUN path, size, online/offline state, and shared state) of the given LUN or class of LUNs.
 - "cli-df": For "df -s", "df -r", "df -A -r" commands and used to display free space.
 - "cli-ifconfig": For "ifconfig -a" command and used for getting filer IP address.
 - "cli-rdfile": For "rdfile /etc/netgroup" command and used for getting netgroups.
 - "cli-date": For "date" command and used to get full date for getting Snapshot copies.
 - "cli-snap": For "snap list" command and used for getting Snapshot copies.

If cli-date or cli-snap permissions are not provided, acquisition can finish, but Snapshot copies are not reported.

To acquire a 7-Mode data source successfully and generate no warnings on the storage system, you should use one of the following command strings to define your user roles. The second string listed here is a streamlined version of the first:

```
login-http-admin,api-*,security-api-vfile,cli-rdfile,cli-options,cli-
df,cli-lun,cli-ifconfig,cli-date,cli-snap,
or
login-http-admin,api-*,security-api-vfile,cli-*
```

NetApp E-Series data source

The NetApp E-Series data source collects inventory and performance information. There are two possible configurations (firmware 6.x and firmware 7.x+), and they both have the same values.

Terminology

OnCommand Insight acquires the following inventory information from the NetApp E-Series data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
-------------------	--------------

Drive	Disk
Volume Group	Disk Group
Storage Array	Storage
Controller	Storage Node
Volume Group	Storage Pool
Volume	Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- The IP address of each controller on the array
- Port requirement 2463

Configuration

Field	Description
Comma-separated list of Array SANtricity Controller IPs	IP addresses and/or fully-qualified domain names for the array controllers

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 30 minutes)
Performance Poll Interval (up to 3600 seconds)	Interval between performance polls (default 300 seconds)

E-Series Storage

Terms applying to objects or references that you might find on NetApp E-Series storage asset landing pages.

E-Series Storage Terminology

The following terms apply to objects or references that you might find on NetApp E-Series storage asset landing pages. Many of these terms apply to other data collectors as well.

- Model — model name of the device.

- Vendor — same Vendor name you would see if you were configuring a new data source.
- Serial number — The array serial number. On cluster architecture storage systems like NetApp Clustered Data Ontap, this serial number may be less useful than the individual “Storage Nodes” serial numbers.
- IP — generally will be the IP(s) or hostname(s) as configured in the data source.
- Microcode version — firmware.
- Raw Capacity — base 2 summation of all the physical disks in the system, regardless of their role.
- Latency — a representation of what the host facing workloads are experiencing, across both reads and writes. Insight calculates an IOPs-weighted average derived from the volumes in the storage.
- Throughput — the array’s total host facing throughput. Insight sums the volumes’ throughput to derive this value.
- Management — this may contain a hyperlink for the management interface of the device. Created programmatically by the Insight data source as part of inventory reporting.

E-Series Storage Pool

Terms applying to objects or references that you might find on NetApp E-Series storage pool asset landing pages.

E-Series Storage Pool Terminology

The following terms apply to objects or references that you might find on NetApp E-Series storage pool asset landing pages. Many of these terms apply to other data collectors as well.

- Storage — what storage array this pool lives on. Mandatory.
- Type — a descriptive value from a list of an enumerated list of possibilities. Most commonly will be “Thin Provisioning” or “RAID Group”.
- Node — if this storage array’s architecture is such that pools belong to a specific storage node, its name will be seen here as a hyperlink to its own landing page.
- Uses Flash Pool — Yes/No value.
- Redundancy — RAID level or protection scheme. E-Series reports “RAID 7” for DDP pools.
- Capacity — the values here are the logical used, usable capacity and the logical total capacity, and the percentage used across these. These value both include E-Series “preservation” capacity, resulting both in numbers and the percentage being higher than what the E-Series own user interface may show.
- Over-committed capacity — If by using efficiency technologies you have allocated a sum total of volume capacities larger than the logical capacity of the storage pool, the percentage value here will be greater than 0%.
- Snapshot — snapshot capacities used and total, if your storage pool architecture dedicates part of its capacity to segments areas exclusively for snapshots.
- Utilization — a percentage value showing the highest disk-busy percentage of any disk contributing capacity to this storage pool. Disk utilization does not necessarily have a strong correlation with array performance — utilization may be high due to disk rebuilds, deduplication activities, etc in the absence of host-driven workloads. Also, many arrays’ replication implementations may drive disk utilization while not showing as volume workload.
- IOPS — the sum IOPs of all the disks contributing capacity to this storage pool.
- Throughput — the sum throughput of all the disks contributing capacity to this storage pool.

E-Series Storage Node

Terms applying to objects or references that you might find on NetApp E-Series storage node asset landing pages.

E-Series Storage Node Terminology

The following terms apply to objects or references that you might find on NetApp E-Series storage pool asset landing pages. Many of these terms apply to other data collectors as well.

- Storage — what storage array this node is part of. Mandatory.
- HA Partner — on platforms where a node will fail over to one and only one other node, it will generally be seen here.
- State — health of the node. Only available when the array is healthy enough to be inventoried by a data source.
- Model — model name of the node.
- Version — version name of the device.
- Serial number — The node serial number.
- Memory — base 2 memory if available.
- Utilization — Utilization is not currently available for NetApp E-Series.
- IOPS — Calculated by summing all the IOPs for volumes that belong exclusively to this node.
- Latency — a number representing the typical host latency or response time on this controller. Insights calculates an IOPs weighted average from volumes that belong exclusively to this node.
- Throughput — a number representing the host driven throughput on this controller. Calculated by summing all the throughput for volumes that belong exclusively to this node.
- Processors — CPU count.

NetApp Host and VM File Systems data source

You can use the NetApp Host and VM File Systems data source to retrieve file system details and storage resource mappings for all Microsoft Windows host and VM (virtual machine) file systems and for all supported Linux VMs (those that are virtually mapped only) existing in the Insight server that are annotated with the configured Compute Resource Group (CRG).

General Requirements

- This feature must be purchased separately.

You can contact your Insight representative for assistance.

- You should check the Insight support matrix to verify that your host or virtual machine operating system is supported.

To verify that links from file systems to storage resources are created, check that the relevant storage or virtualization vendor type and version report the volume or virtual disk identification data required.

Microsoft Windows Requirements

- This data source uses Window Management Instrumentation (WMI) data structures to retrieve data.

This service must be operational and available remotely. In particular, port 135 must be accessible and must be opened if behind a firewall.

- Windows domain users must have the appropriate permissions to access WMI structures.
- Administrator permissions are required.
- Dynamic TCP ports assigned 1024-65535 for Windows 2003 and older
- Ports 49152—65535 for Windows 2008



As a general rule, when trying to use a firewall between Insight, an AU, and this data source, you should consult with your Microsoft team to identify the ports they believe will be required.

Linux Requirements

- This data source uses a Secure Shell (SSH) connection to execute commands on Linux VMs.

The SSH service must be operational and available remotely. In particular, port 22 must be accessible and must be opened if behind a firewall.

- SSH users must have sudo permissions to execute read-only commands on Linux VMs.

You must use the same password to log in to SSH and to answer any sudo password challenge.

Usage Recommendations

- You should annotate a group of hosts and virtual machines that have common operating system credentials using the same Compute Resource Group annotation.

Each group has an instance of this data source discovering file system details from those hosts and virtual machines.

- If you have an instance of this data source for which the success rate is low (for example, OnCommand Insight is discovering file system details for only 50 of 1000 hosts and virtual machines in a group), you should move the hosts and virtual machines for which discovery is successful into a separate Compute Resource Group.

Configuration

Field	Description
User Name	Operating system user with appropriate rights to retrieve file system data For Windows operating system users, this must include the domain prefix.
Password	Password for the operating system user

Compute Resource Group	Annotation value used to flag host and virtual machines for the data source discovers file systems. A blank value indicates that the data source discovers file systems for all hosts and virtual machines not currently annotated with any Compute Resource Group.
------------------------	---

Advanced configuration

Field	Description
Inventory poll interval (min)	Interval between inventory polls (default 360 minutes)

NetApp SolidFire data source

The NetApp SolidFire data source supports both iSCSI and Fibre Channel SolidFire configurations, for both inventory and performance collection.

The SolidFire data source utilizes the SolidFire REST API. The acquisition unit where the data source resides needs to be able to initiate HTTPS connections to TCP port 443 on the SolidFire cluster management IP address. The data source needs credentials capable of making REST API queries on the SolidFire cluster.

Terminology

OnCommand Insight acquires the following inventory information from the NetApp SolidFire data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive	Disk
Cluster	Storage
Node	Storage Node
Volume	Volume
Fibre Channel Port	Port
Volume Access Group, LUN Assignment	Volume Map
iSCSI Session	Volume Mask



These are common terminology mappings only and might not represent every case for this data source.

Requirements

The following are requirements for configuring this data source:

- Management Virtual IP Address
- Port 443

Configuration

Field	Description
Management Virtual IP Address (MVIP)	Management Virtual IP address of the SolidFire Cluster
User Name	Name used to log into the SolidFire cluster
Password	Password used to log into the SolidFire cluster

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
TCP Port	TCP Port used to connect to SolidFire Server (default 443)
Connection Timeout (sec)	Connection timeout (default 60 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Troubleshooting

When SolidFire reports an error it is displayed in OnCommand Insight as follows:

An error message was received from a SolidFire device while trying to retrieve data. The call was <method> (<parameterString>). The error message from the device was (check the device manual): <message>

Where:

- The <method> is an HTTP method, such as GET or PUT.
- The <parameterString> is a comma separated list of parameters that were included in the REST call.
- The <message> is whatever the device returned as the error message.

NetApp StorageGRID data source

This data source collects inventory and performance data for StorageGRID.

Requirements

The following are requirements for configuring this data source:

- StorageGRID Host IP Address
- A username and password for a user that has had the Metric Query and Tenant Access roles assigned
- Port 443

Configuration

Field	Description
StorageGRID Host IP Address (MVIP)	Host IP address of the StorageGRID
User Name	Name used to log into the StorageGRID
Password	Password used to log into the StorageGRID

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Performance Poll Interval (sec)	Interval between performance polls (default 900 seconds)

OpenStack data source

The OpenStack (REST API / KVM) data source collects information about OpenStack hardware instances. This data source collects inventory data for all OpenStack instances, and optionally, VM performance data.

Requirements

The following are requirements for configuring the OpenStack data source.

- IP address of the OpenStack controller
- OpenStack admin role credentials and sudo access to the Linux KVM hypervisor are recommended.



If you are not using an admin account or admin equivalent privileges, you can still acquire data from the data source. You will need to modify the policy configuration file (i.e. `etc/nova/policy.json`) to allow users with non-admin role to call the API:

- `"os_compute_api:os-availability-zone:detail": ""`
- `"os_compute_api:os-hypervisors": ""`
- `os_compute_api:servers:detail:get_all_tenants": ""`
- For performance collection the OpenStack Ceilometer module must be installed and configured.

Configuring the Ceilometer is done by editing the `nova.conf` file for each hypervisor and then restart the Nova Compute service on each hypervisor. The option name changes for different releases of OpenStack:

- Icehouse
- Juno
- Kilo
- Liberty
- Mitaka
- Newton
- Ocata
- For CPU stats, "compute_monitors=ComputeDriverCPUMonitor" needs to be turned on in `/etc/nova/nova.conf` on compute nodes.
- Port requirements:
 - 5000 for http and 13000 for https, for the Keystone service
 - 22 for KVM SSH
 - 8774 for Nova Compute Service
 - 8776 for Cinder Block Service
 - 8777 for Ceilometer Performance Service
 - 9292 for Glance Image Service



The port binds to the specific service, and the service may run on the controller or another host in larger environments.

Configuration

Field	Description
OpenStack Controller IP Address	IP address or fully-qualified domain name of the OpenStack Controller
OpenStack Administrator	User name for an OpenStack Admin
OpenStack Password	Password used for the OpenStack Admin
OpenStack Administrator Tenant	OpenStack Administrator Tenant
KVM Sudo User	KVM Sudo User name
Choose 'Password' or 'OpenSSH Key File' to specify credential type	The credential type used to connect to the device via SSH
Full Path to Inventory Private Key	Full Path to Inventory Private Key
KVM Sudo Password	KVM Sudo Password

Advanced configuration

Field	Description
Enable hypervisor inventory discovery through SSH	Check this to enable hypervisor inventory discovery through SSH
OpenStack Admin URL port	OpenStack Admin URL port
Use HTTPS	Check to use secure HTTP
HTTP Connection Timeout (sec)	Timeout for HTTP connection (default 300 seconds)
SSH Port	Port used for SSH
SSH Process Wait Timeout (sec)	SSH process timeout (default 30 seconds)
SSH Process Retries	Number of inventory retry attempts
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)

Oracle ZFS data source

The Oracle ZFS data source supports inventory and performance collection.

Terminology

OnCommand Insight acquires the following inventory information from this data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk (SDD)	Disk
Cluster	Storage
Controller	Storage Node
LUN	Volume
LUN Map	Volume Map
Initiator, Target	Volume mask
Share	Internal Volume



These are common terminology mappings only and might not represent every case for this data source.

Requirements

The following are requirements for configuring this data source:

- Host names for the ZFS Controller-1 and the ZFS Controller-2
- Administrator user name and credentials
- Port requirement: 215 HTTP/HTTPS

Configuration

ZFS Controller-1 Hostname	Host name for storage controller 1
ZFS Controller-2 Hostname	Host name for storage controller 2
User name	User name for the storage system administrator user account
Password	Password for the administrator user account

Advanced configuration

Field	Description
TCP port	TCP Port used to connect to ZFS (default 215)
Connection Type	HTTP or HTTPS
Inventory poll interval	Inventory poll interval (default 60 minutes)
Connection Timeout	Default is 60 seconds
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Troubleshooting

Some things to try if you encounter problems with this data collector:

Problem:	Try This:
"Invalid login credentials"	validate Zfs user account and password

"Configuration error" with error message "REST Service is disabled"	Verify REST service is enabled on this device.
"Configuration error " with error message "User unauthorized for command"	<p>Likely due to certain roles (for example, 'advanced_analytics') are not included for the configured user <userName>.Possible Solution:</p> <ul style="list-style-type: none"> • Correct the Analytics (statistic) scope for the user \${user} with the read only role:- From the Configuration → Users screen, put your mouse over the role and double click to allow editing • Select "Analytics" from the Scope drop down menu. A list of the possible properties appears. • Click the top most check box and it will select all three properties.- Click the Add button on the right side. • Click the Apply button at the top right of the pop-up window. The pop-up window will close.

Pure Storage FlashArray data source

The Pure Storage FlashArray (HTTP) data source is used to collect information from the Pure Storage Flash Array. Insight supports both inventory and performance collection.

Terminology

OnCommand Insight acquires the following inventory information from the Pure Storage FlashArray data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive (SSD)	Disk
Array	Storage
Controller	Storage Node
Volume	Volume
Port	Port
LUN Map (Host, Host Group, Target Port)	Volume Map, Volume Mask



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- Storage system IP address
- User name and password for the Administrator account of the Pure storage system.
- Port requirement: HTTP/HTTPS 80/443

Configuration

Field	Description
FlashArray Host	IP address or fully-qualified domain name of FlashArray Management Server
User Name	User name for the FlashArray Management Server
Password	Password for the FlashArray Management Server

Advanced configuration

Field	Description
Connection Type	Management Server
TCP Port	TCP Port used to connect to FlashArray Server (default 443)
Connection Timeout (sec)	Connection timeout (default 60 seconds)
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

QLogic FC Switch data source

For configuration, the QLogic FC Switch (SNMP) data source requires the network address for the FC Switch device, specified as an IP address, and an SNMP *read-only* community string used to access the device.

Configuration

Field	Description
SANSurfer Switch	IP address or fully-qualified domain name for the SANSurfer switch
SNMP version	SNMP version

SNMP community	SNMP Community String
User Name	User name for the SANSurfer switch
Password	Password for the SANSurfer switch

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 15 minutes)
SNMP Auth Protocol	SNMP authentication protocol (SNMPv3 only)
SNMP Retries	Number of SNMP retry attempts
SNMP Timeout (ms)	SNMP timeout (default 5000 ms)
Enable Trapping	Select to enable trapping
Minimum Time Between Traps (sec)	Minimum time between acquisition attempts triggered by traps (default 10 seconds)
Fabric Name	Fabric name to be reported by the data source. Leave blank to report the fabric name as WWN.
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Red Hat (RHEV) data source

The Red Hat Enterprise Virtualization (REST) data source collects information about RHEV instances via HTTPS.

Requirements

- IP address of the RHEV server over port 443 via REST API
- Read-only username and password
- RHEV Version 3.0+

Configuration

Field	Description
RHEV Server IP Address	IP address or fully-qualified domain name of the RHEV server

User Name	User name for the RHEV server
Password	Password used for the RHEV server

Advanced configuration

Field	Description
HTTPS Communication Port	Port used for HTTPS communication to RHEV
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection timeout (sec)	Connection timeout (default 60 seconds)

Violin Flash Memory Array data source

The Violin 6000-Series Flash Memory Array (HTTP) data source collects network information for analysis and validation from Violin 6000-series flash memory arrays.

Terminology



This data collector is no longer available starting with OnCommand Insight 7.3.11.

OnCommand Insight acquires the following inventory information from the Violin 6000-Series Flash Memory Array data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Violin Intelligent Memory Module (VIMM)	Disk
Container	Storage
Memory Gateway	Storage Node
LUN	Volume
Initiator, Initiator Group, Target	Volume Map, Volume Mask



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- You need a read-only user name and password to the storage.
- Validate access with a web browser using the storage IP address.

Configuration

Field	Description
IP address or FQDN of Violin Memory Array Main Gateway	IP address or fully-qualified domain name of the Violin Memory Array Main Gateway
User Name	User name for the Violin Memory Array Main Gateway
Password	Password for the Violin Memory Array Main Gateway

Advanced configuration

Field	Description
Communication Port	Port used for communication with Violin array
HTTPS Enabled	Select to use HTTPS
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection timeout (sec)	Connection timeout (default 60 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

VMware vSphere data source

The VMware vSphere (Web Services) data source collects ESX Host information and requires *read-only* privileges on all objects within the Virtual Center.

Terminology

OnCommand Insight acquires the following inventory information from the VMware vSphere data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Virtual Disk	Disk
Host	Host
Virtual Machine	Virtual Machine
Data Store	Data Store
LUN	LUN

Fiber Channel Port	Port
--------------------	------



These are common terminology mappings only and might not represent every case for this data source.

Requirements

- IP address of the Virtual Center server
- Read-only username and password in Virtual Center
- Read-only privileges on all objects within the Virtual Center.
- SDK access on the Virtual Center server
- Port requirements: http-80 https-443
- Validate access by logging in to Virtual Center Client using your user name and password and verifying that the SDK is enabled by entering `telnet <vc_ip> 443`.

Configuration

Field
Description
Virtual Center Address
Network address for the Virtual Center or vSphere server, specified as an IP (<i>nnn.nnn.nnn.nnn</i> format) address or as a host name that can be resolved through DNS.
User Name
User name for the VMware server.
Password
Password for the VMware server.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection Timeout (ms)	Connection timeout (default 60000 ms)
Filter VMs by	Choose how to filter VMs
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the VM list below when collecting data

List of VMs to filter (Comma Separated, or Semicolon Separated If Comma Is Used in the Value)	Comma-separated or semicolon-separated list of VMs to include or exclude from polling
Number of Retries for Requests to vCenter	Number of vCenter Request retry attempts
Communication Port	Port used for VMware server
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Changing data source credentials

If multiple data sources of the same type are sharing a username and password, you can change the password for all devices in the group at the same time.

Steps

1. On the Insight toolbar, click **Admin**.

The **Data sources** list opens.

2. Click the **Actions** button and select the **Change credentials** option.
3. In the Credentials Management dialog box, select one of the data source groups from the list.

The Edit icon, a pen on a sheet of paper, becomes active to the right.

Credentials Management

Below is a list of groups of data sources with the same credentials. You can change the credentials of the entire group in a single action by pressing the edit button next to the desired group.

Data source type	Package	User/Community	Used by
FC Switch Firmware 2.0+ (SNMP)	foundation	UHTSAN	elr1scvbkodd01 and 1 others
FC Switch Firmware 4.2+ (SSH)	foundation	ssacct	ELR5_EvenFabric and 1 others
FC Switch Firmware 4.2+ (SSH)	performance	UHTSAN	ELR5_EvenFabric
HiCommand Device Manager	foundation	sanscm	ELR5_APSWP1008_HCS7 and 1 others
Solutions Enabler (CLI) with Performance (SMI-S)	storageperformance	admin	ELR1_Vblock EMC

Showing 1 to 5 of 5 entries

4. Click **Edit**.
5. Enter the new password and confirm it.

Changes causing data collection problems

If you are experiencing data collection problems in OnCommand Insight, changes in your environment are a likely cause. As a general maintenance rule, you should accommodate any changes in your environment in Insight as well.

You can use this checklist to identify changes to your network that might be causing problems:

- Have you changed any passwords? Were those passwords changed in Insight?
- Did you remove a device from your network? You must also remove the device from OnCommand Insight to prevent it from being rediscovered and reintroduced.
- Did you upgrade infrastructure software (such as HP CommandView EVA or EMC Solutions Enabler)?

Ensure that the appropriate versions of the client tools are installed on the acquisition unit. If data source failures persist, you need to contact technical support to request assistance and possibly a data source patch.

- Are all of your OnCommand Insight acquisition units using the same OnCommand Insight version? If the Remote Acquisition Units and local acquisition unit are running different OnCommand Insight versions, install the same version on all units to correct the data collection problem.

If you need to install a new version of OnCommand Insight on all of the acquisition units, go to the support site and download the correct version.

- Have you changed any domain names or added a new domain? You must update your Device Resolution (formerly Auto Resolution) methods.

Examining one data source in detail

If you see that a data source has failed or slowed, you might want to examine a detailed summary of information for that data source to determine the cause of the problem. Data sources with conditions requiring your attention are marked with a solid red circle.

Steps

1. On the Insight toolbar, click **Admin**.

The **Data sources** list opens. Any listed data sources with potential problems are marked with a solid red circle. The most serious problems are at the top of the list.

2. Select the data source that is causing concern.
3. Click the data source name link.
4. On the data source summary page, check the information in any of these sections:

- **Event timeline**

Lists events tied to the current status shown in the Data sources list. Events in this summary are displayed per device. Errors are shown in red. You can position your mouse pointer on timeline items to

display additional information.

- **Devices reported by this data source**

Lists the types of devices, their IP addresses, and links to more detailed information for each device.

- **Changes reported by this data source (last 3 weeks)**

Lists any devices that were added or removed or had a change to the configuration.

5. After examining the data source information, you might want to perform one of these operations using the buttons at the top of the page:

- **Edit** the description of the data source to correct the problem.
- **Poll again** forces polling to reveal if the problem was persistent or intermittent.
- **Postpone** data source polling for 3, 7, or 30 days to give you time to research the problem and stop the warning messages.
- **Install a patch** on the data source to correct the problem.
- Prepare an **Error report** for technical support.
- **Delete** the data source from your Insight monitoring environment.

Researching a failed data source

If a data source has the "**Inventory failed !**" or "**Performance failed !**" message and a High or Medium Impact, you need to research this problem using the data source summary page with its linked information.

Steps

1. Click the linked **Name** of the data source to open the Summary page.
2. On the Summary page, check the **Comments** area to read any notes left by another engineer who might also be investigating this failure.
3. Note any performance messages.
4. If there is a patch being applied to this data source, click link to check the **patch page** to see if that has caused the problem.
5. Move your mouse pointer over the segments of the **Event timeline** graph to display additional information.
6. Select an error message for a Device and displayed below the Event timeline and click the **Error details** icon that displays to the right of the message.

The Error details include the text of the error message, most likely causes, information in use, and suggestions of what can be tried to correct the problem.

7. In the Devices reported by this data source area, you might filter the list to display only devices of interest, and you can click the linked **Name** of a device to display the *asset page* for that device.
8. To return to previously displayed pages, use one of these techniques:
 - Click the browser back arrow.
 - Right-click the back arrow to display a list of the pages and select the page you want.
9. To display detailed information about other resources, click other linked names.

10. When you return to the data source summary page, check the **Changes** area at the bottom of the page to see if recent changes caused the problem.

Controlling data source polling

After making a change to a data source, you might want it to poll immediately to check your changes, or you might want to postpone the data collection on a data source for one, three, or five days while you work on a problem.

Steps

1. Click **Admin** and navigate to the data source list view
2. Select the data source for which you want to control the polling.
3. Click the data source name link.
4. On the data source summary page, check the information and click one of these two polling options:
 - **Poll again** to force the data source to collect data immediately.
 - **Postpone** and select the length of the polling delay from 3, 7, or 30 days.

After you finish

If you postponed the data collection on a data source and want to restart collection, click **Resume** on the summary page.

Editing data source information

You can quickly edit data source setup information.

Steps

1. Click **Admin** and navigate to the data source list view
2. Locate the data source that you want to edit.
3. Use one of these methods to begin the changes:
 - Click **Edit data source** to the right of the selected data source.
 - Click the linked name of the selected data source and click **Edit**. Either method opens the Edit data source dialog box.
4. Make the desired changes and Click **Save**.

Editing information for multiple data sources

You can edit most of the information for multiple data sources of the same vendor and model at one time. For example, if these data sources share a user name and password, you can change the password in one place and thereby update the password for all the selected data sources.

About this task

Options that you cannot edit for the selected data sources appear dimmed or are not displayed in the Edit data source dialog box. Additionally, when an option displays a value of **Mixed**, it indicates that the value for the option varies between the selected data sources. For example, if the **Timeout (sec)** option for two selected data sources is **Mixed**, one data source could have a timeout value of 60 and the other could have a value of 90; therefore, if you change this value to 120 and save the changes to the data sources, the timeout setting for both data sources becomes 120.

Steps

1. Click **Admin** and navigate to the data source list view
2. Select the data sources you want to modify. Selected data sources must belong to same vendor, model and acquisition unit.
3. Click the **Actions** button and select the **Edit** option.
4. In the edit dialog, change any of the **Settings** as needed.
5. Click the **Configuration** link to change any of the basic options for the data sources.
6. Click the **Advanced Configuration** link to change any of the advanced options for the data sources.
7. Click **Save**.

Mapping data source tags to annotations

When a data source is configured to poll tag data, Insight automatically sets annotation values for an existing Insight annotation with the same name as a tag.

When the Insight annotation exists before the tags are enabled in the data source, the data source tag data is automatically added to the Insight annotation.

When you create an annotation after the tag is enabled, initial polling of the data source does not automatically update the annotation. There is a delay in the time it takes to replace or populate the Insight annotation. To avoid the delay, you can force the tag to annotation update by postponing and then resuming the data source.

Deleting a data source

If you have removed a data source from your environment, you must also delete it from the OnCommand Insight monitoring environment.

Steps

1. On the Insight toolbar, click **Admin**.

The Data sources list opens.
2. Select the data source that you want to delete.
3. Click the linked data source name.
4. Check the information for the selected data source on the summary page to be certain that it is the one you want to delete.
5. Click **Delete**.
6. Click **OK** to confirm the operation.

What data source patches are

Data source patches fix issues with existing patches and also enable you to easily add new data source types (vendors and models). For each data source type in your network, you can upload data source patches. You can also install, test, and manage the patching process. However, only one patch can be active for a data source type at a time.

For each patch, you can perform these tasks:

- Check the before and after comparison of each data source receiving the patch.
- Write comments to explain decisions or summarize research.
- Make changes to a data source that is not responding well to the patch.
- Approve the patch to be committed to your Insight server.
- Roll back a patch that is not operating as you intended.
- Replace a failing patch with a different one.

Applying a data source patch

Data source patches are periodically available and enable you to fix issues with an existing data source, add a data source for a new vendor, or add a new model for a vendor.

Before you begin

You must have obtained the `.zip` file that contains the latest data source `.patch` files from technical support.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Patches**.
3. From the Actions button, select **Apply patch**.
4. In the **Apply data source patch** dialog box, click **Browse** to locate the `.patch` file.
5. Inspect the **Patch name**, **Description**, and **Impacted data source types**.
6. If the selected patch is correct, click **Apply Patch**.

If you are applying a patch that fixes issues with a data source, all data sources of the same type are updated with the patch and you must approve the patch. Patches that do not affect any configured data sources are automatically approved.

After you finish

If you are applying a patch that adds a data source for a new vendor or a new model, you must add the data source after applying the patch.

Installing a patch on one type of data source

After uploading a data source patch, you can install it on all of the data sources of the

same type.

Before you begin

You must have uploaded a patch file that you want to install on one type of data source.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Patches**.
3. From the Actions button, select **Apply patch**.
4. In the **Apply data source patch** dialog box, click **Browse** to locate the uploaded patch file.
5. Check the **Patch name**, **Description**, and **Impacted data source types**.
6. If the selected patch is correct, click **Apply Patch**.

All data sources of the same type are updated with this patch.

Managing patches

You can review the current status of all of the data source patches being applied to your network. If you want to perform an action on a patch, you can click the linked name in the Patches currently under review table.

Before you begin

You must have already uploaded and be installing at least one patch.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Patches**.

If no patches are being installed, the table of Patches currently under review is empty.

3. In **Patches currently under review**, check the status of the data source patches currently being applied.
4. To examine the details associated with a specific patch, click the linked name of the patch.
5. For the selected patch, you might click any of these options to perform the next action on the patch:
 - **Approve patch** commits the patch to the data sources.
 - **Rollback** removes the patch.
 - **Replace patch** enables you to select a different patch for those data sources.

Committing a data source patch

You use the information in the Patches summary to decide if the patch is performing as expected and then commit the patch to your network.

Before you begin

You have installed a patch and need to decide if the patch is successful and should be approved.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Patches**.

If no patches are being installed, the Patches currently under review is empty.

3. In **Patches currently under review**, check the status of the data source patches currently being applied.
4. To examine the details associated with a specific patch, click the linked name of the patch.
5. In the Patches summary information, shown in this example, check the **Recommendation** and **Comments** to assess the progress on the patch.

Patches
Brocade SSH

Summary

Recommendation: **Approve patch - Patch results are positive (no change or more successes)**

Applied on: 5/12/2013 20:00:01

Other data source affected: Brocade SNMP, Brocade HTTP

Comments: Got this patch from Scott. He said that this should fix the SNMP v3 problem in Brocade. Talking to John from NetApp, they promised this will fix the SNMP v3 problem. After this is applied, we still need to check the other SNMP v3 data sources and see if they are good.

You should now review the results of the patch. Approving a patch will permanently apply this patch to the system. Rolling back a patch will delete it and restore the previous version before this patch was applied. Please note that there can only be one patch active for a data source type.

Buttons: Approve, Roll back, Replace patch

affected data sources

Name	Ali	Type	Conclusion	Status before patch applied	Most recent status
ds0	local	Brocade CLI	No change (success)	All successful	Currently polling...
ds1	local	Brocade CLI	No change (success)	All successful	All successful
ds2	local	Brocade CLI	Polling is now successful	Configuration failed	All successful
ds3	local	Brocade CLI	Configuration is still failing (a different error)	Configuration failed	Configuration failed
ds4	au1	Brocade SNMP	Configuration is successful but now Performance is failing	Configuration failed	Performance failed

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.

Device resolution

You need to discover all of the devices you want to monitor with OnCommand Insight. Discovery is necessary in order to accurately track performance and inventory in your environment. Typically the majority of devices in your environment are discovered through automatic device resolution.



If you are performing an upgrade and have inactive Auto Resolution rules in the system you are upgrading from, these rules will be deleted during the upgrade. To preserve inactive Auto Resolution rules, activate the rules (check the box) before the upgrade is performed.

After you install and configure data sources, devices in your environment, including switches, storage arrays and your virtual infrastructure of hypervisors and VMs are identified. However, this does not normally identify 100% of the devices in your environment.

After data source type devices have been configured, best practice is to leverage device resolution rules to help identify the remaining unknown devices in your environment. Device resolution can help you resolve unknown devices as the following device types:

- physical hosts
- storage arrays
- tapes
- switches

Devices remaining as “unknown” after device resolution are considered generic devices, which you can also

show in queries and on dashboards.

The rules created in turn will automatically identify new devices with similar attributes as they are added to your environment. In some cases, Device resolution also allows for manual identification bypassing the device resolution rules for undiscovered devices within Insight.

Incomplete identification of devices can result in issues including:

- Incomplete paths
- Unidentified multipath connections
- The inability to group applications
- Inaccurate topology views
- Inaccurate data in the Data warehouse and reporting

The Device resolution feature (**Manage > Device resolution**) includes the following tabs, each of which plays a role in device resolution planning and viewing results:

- “FC identify” contains a list WWNs and port information of Fibre Channel devices that were not resolved through automatic device resolution. The tab also identifies the percentage of devices that have been identified.
- “IP identify” contains a list of devices accessing CIFS shares and NFS shares that were not identified through automatic device resolution. The tab also identifies the percentage of devices that have been identified.
- “Auto resolution rules” contains the list of rules that are run when performing Fibre channel device resolution. These are rules you create to resolve unidentified Fibre channel devices.
- “Preferences” provides configuration options that you use to customize device resolution for your environment.

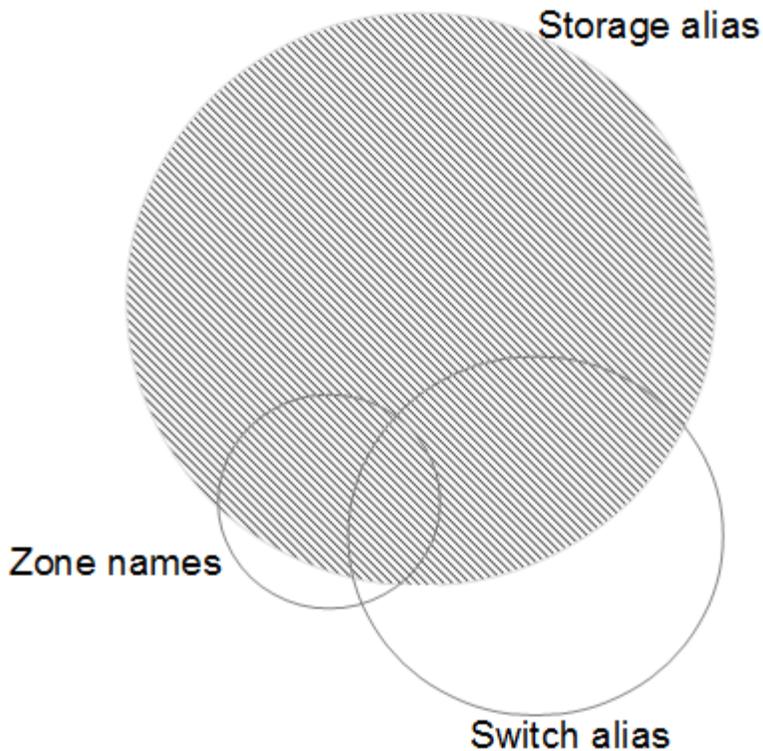
Before you begin

You need to know how your environment is configured before you define the rules for identifying devices. The more you know about your environment the easier it will be to identify devices.

You need to answer questions similar to the following to help you create accurate rules:

- Does your environment have naming standards for zones or hosts and what percentage of these are accurate?
- Does your environment use a switch alias or storage alias and do they match the host name?
- Does your environment use an SRM tool and can you use it to identify host names? What coverage does the SRM provide?
- How often do naming schemes change in your environment?
- Have there been any acquisitions or mergers that introduced different naming schemes?

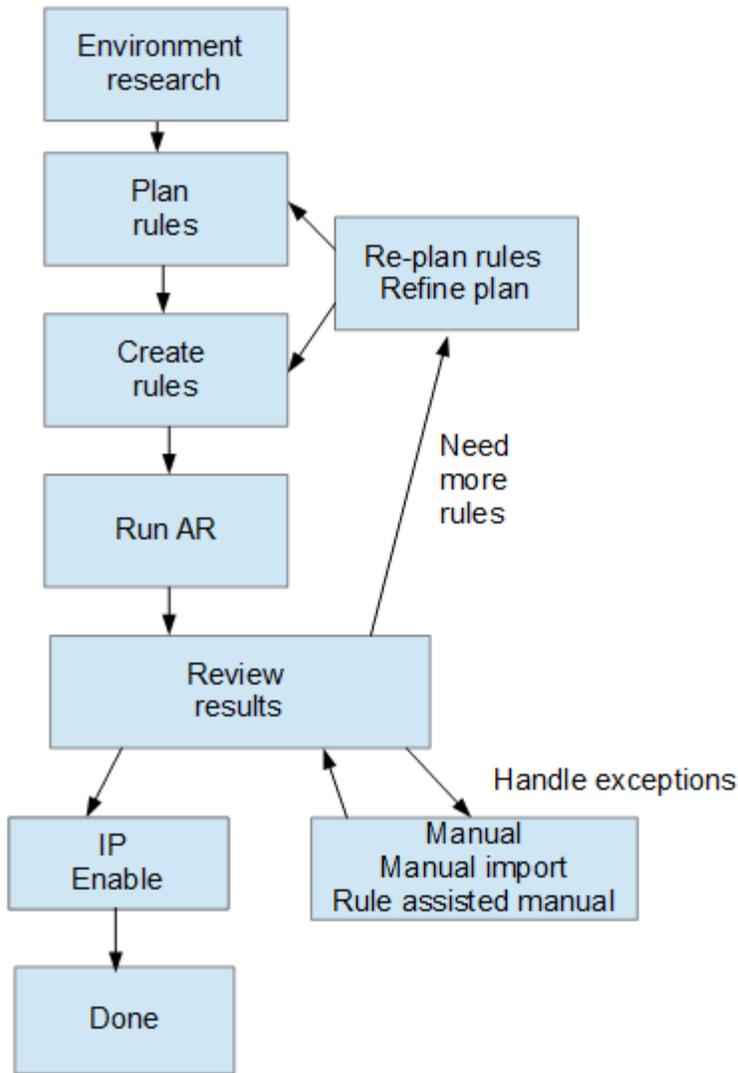
After analyzing your environment, you should be able to identify what naming standards exist that you can expect to reliably encounter. The information you gathered might be represented graphically in a figure similar to the following:



In this example the largest number of devices are reliably represented by storage aliases. Rules that identify hosts using storage aliases should be written first, rules using switch aliases should be written next, and the last rules created should use zone aliases. Due to the overlap of the use of zone aliases and switch aliases, some storage alias rules might identify additional devices, leaving less rules required for zone aliases and switch aliases.

Steps to defining devices in your environment

Typically, you would use workflow similar to the following to identify devices in your environment. Identification is an iterative process and might require multiple steps of planning and refining rules.



If you have unidentified devices (otherwise known as “unknown” or generic devices) in your environment and you subsequently configure a data source that identifies those devices upon polling, they will no longer be displayed or counted as generic devices.

Planning device resolution rules for your environment

Using rules to identify devices in your environment is typically an iterative process that requires a thorough analysis of your environment and the creation of multiple rules to identify as many devices as possible. The best case scenario is to set a goal to identify 100% of the devices in your environment.

The most efficient order for rules is to place the most restrictive rules first, resulting in most entries not pattern matching, with the process proceeding to less restrictive rules. This allows Insight to apply more patterns to each entry increasing the possibility of patterns matching and of positive host identification.

When you create rules, your objective should be to create rules that address the largest number of unidentified devices possible. For example, creating rules that follow a pattern of coverage similar to the following is far more efficient than creating 30 rules with lower percentages of coverage:

Rule	Percentage of coverage
Rule 1	60%
Rule 2	25%
Rule 3	8%
Rule 4	4%
Rule 5	1%

Creating device resolution rules

You create device resolution rules to identify hosts, storage, and tapes that are not automatically identified currently by OnCommand Insight. The rules that you create identify devices currently in your environment and also identify similar devices as they are added to your environment.

About this task

When you create rules you start by identifying the source of information that the rule runs against, the method used to extract information, and whether DNS lookup is applied to the results of the rule.

Source that is used to identify the device
<ul style="list-style-type: none"> • SRM aliases for hosts • Storage alias containing an embedded host or tape name • Switch alias containing an embedded host or tape name • Zone names containing an embedded host name
Method that is used to extract the device name from the source
<ul style="list-style-type: none"> • As is (extract a name from an SRM) • Delimiters • Regular expressions
DNS lookup
Specifies if you use DNS to verify the host name.

You create rules in the Auto Resolution Rules tab. The following steps describe the rule creation process.

Steps

1. Click **Manage > Device resolution**
2. In the **Auto resolution rules** tab, click **+Add**

The New Rule screen is displayed.



The New Rule screen includes a ? icon, that provides help and examples for creating regular expressions.

3. In the **Type** list select the device you want to identify.

You can select Host or Tape.

4. In the **Source** list, select the source you want to use to identify the host.

Depending on the source you chose, Insight displays the following response:

- Zones lists the zones and WWN that need to be identified by Insight.
- SRM lists the unidentified aliases that need to be identified by Insight
- Storage alias lists storage aliases and WWN that need to be identified by Insight
- Switch alias lists the switch aliases that need to be identified by Insight

5. In the **Method** list select the method you want to employ to identify the host.

Source	Method
SRM	“As is”, “Delimiters”, “Regular expressions”
Storage alias	“Delimiters”, or “Regular expressions”
Switch alias	“Delimiters”, or “Regular expressions”
Zones	“Delimiters”, or “Regular expressions”

- Rules using “Delimiters” require the delimiters and the minimum length of the host name.

The minimum length of the host name is number of characters that Insight should use to identify a host. Insight performs DNS lookups only for host names that are this long or longer.

For rules using Delimiters, the input string is tokenized by the delimiter and a list of host name candidates is created by making several combinations of the adjacent token. The list is then sorted, largest to smallest. For example, for vipsnq03_hba3_emc3_12ep0 the list would result in the following:

- vipsnq03_hba3_emc3_12ep0
- vipsnq03_hba3_emc3
- hba3_emc3_12ep0
- vipsnq03_hba3
- emc3_12ep0

- hba3_emc3
 - vipsnq03
 - 12ep0
 - emc3
 - hba3
- Rules using “Regular expression” require a regular expression, the format, and cases sensitivity selection.

6.

Click  to run all rules, or click the down-arrow in the button to run the rule you created (and any other rules that have been created since the last full run of AR.)

Results

The results of the rule run are displayed in the FC identify tab.

Starting an automatic device resolution update

A device resolution update commits manual changes that have been added since the last full automatic device resolution run. Running an update can be used to commit and run only the new manual entries made to the device resolution configuration. No full device resolution run is performed.

Steps

1. Log into the Insight web UI.
2. Click **Manage > Device Resolution**
3. In the **Device resolution** screen, click the down-arrow in the **Run AR** button.
4. Click **Update** to start the update.

Rule assisted manual identification

This feature is used for special cases where you want to run a specific rule or a list of rules (with or without a one-time reordering) to resolve unknown hosts, storage, and tape devices or group of them.

Before you begin

You have a number of devices that have not been identified and you also have multiple rules that successfully identified other devices.

About this task



If your source only contains part of a host or device name, use a regular expression rule and format it to add the missing text.

Steps

1. Log into the OnCommand Insight web UI.

2. Click **Manage > Device resolution**

3. Click the **FC Identify** tab.

The system displays the identified and unidentified devices.

4. Select multiple unidentified devices.

5. Click **Identify > Set host resolution** or **> Set tape resolution**

The system displays the Identify screen which contains a list of all of the rules that successfully identified devices.

6. Change the order of the rules to an order that meets your needs.

The order of the rules are changed in the Identify screen, but are not changed globally.

7. Select the method that that meets your needs.

OnCommand Insight executes the host resolution process in the order in which the methods appear, beginning with those at the top.

When rules that apply are encountered, rule names are shown in the rules column and identified as manual.

Fibre Channel device resolution

The FC Identify screen displays the WWN and WWPN of Fibre Channel devices whose hosts have not been identified by automatic device resolution. The screen also displays any devices that have been resolved by manual device resolution.

Devices that have been resolved by manual resolution contain a status of "OK" and identify the rule used to identify the device. Missing devices have a status of "Unidentified". The total coverage for identification of devices is listed on this page.

FC Identify (10)

<input type="checkbox"/>	WWN	Port WWN	IP	Name	Type	Status	Rule
<input type="checkbox"/>	30:E0:00:00:00:00:00:00	10:B0:00:00:00:00:28:20	1.1.1.1	ResolvedHost1	Host	OK	Hosts by zone
<input type="checkbox"/>	30:E0:00:00:00:00:00:02	10:B0:00:00:00:00:28:22	2.2.2.2	ResolvedHost2	Host	OK	Rule deleted
<input type="checkbox"/>	30:E0:00:00:00:00:00:03	10:B0:00:00:00:00:28:23			Unknown	Unidentified	
<input type="checkbox"/>	30:E0:00:00:00:00:00:04	10:B0:00:00:00:00:28:24			Unknown	Unidentified	
<input type="checkbox"/>	30:E0:00:00:00:00:00:05	10:B0:00:00:00:00:28:25			Unknown	Unidentified	

Showing 1 to 5 of 10 entries

Total coverage: 30% (3/10)

You perform bulk actions by selecting multiple devices on the left-hand side of the FC identify screen. Actions can be performed on a single device by hovering over a device and selecting the identify or unidentify buttons on the far right of the list.

The Total coverage link displays a list of the "number of devices identified/number of devices available" for your configuration:

- SRM alias

- Storage alias
- Switch alias
- Zones
- User defined

Adding a Fibre Channel device manually

You can manually add a Fibre Channel device to OnCommand Insight using the manual add feature available in the Device resolution FC Identify tab. This process might be used for pre-identification of a device that is expected to be discovered in the future.

Before you begin

To successfully add a device identification to the system you need to know the WWN or IP address and the device name.

About this task

You can add a Host, Storage, Tape or Unknown Fibre Channel device manually.

Steps

1. Log in to the Insight web UI
2. Click **Manage > Device resolution**
3. Click the **FC Identify** tab.
4. Click the add button.

The Add Device dialog is displayed

5. Enter the WWN or IP address, the device name, and select the device type.

Results

The device you enter is added to the list of devices in the FC Identify tab. The “Rule” is identified as Manual.

Importing Fibre Channel device identification from a CSV file

You can manually import Fibre Channel device identification into OnCommand Insight Device Resolution feature using a list of devices in a CSV file.

Before you begin

You must have a correctly formatted CSV file in order to import device identifications directly into the Device Resolution feature. The CSV file for Fibre Channel devices requires the following information:

WWN
IP

Name
Type



As a best practice, it is recommended to first export the FC Identify information to a CSV file, make your desired changes in that file, and then import the file back into FC Identify. This ensures that the expected columns are present and in the proper order.

To import FC Identify information:

Steps

1. Log into the Insight web UI.
2. Click **Manage > Device Resolution**
3. Select the **FC identify** tab.
4. Click **Identify > Identify from file**
5. a. Navigate to the folder containing your CSV files for import and select the desired file.

The devices you enter are added to the list of devices in the FC Identify tab. The “Rule” is identified as “Manual”.

Exporting Fibre Channel device identifications to a CSV file

You can export existing Fibre Channel device identifications to a CSV file from the OnCommand Insight device resolution feature. You might want to export a device identification so that you can modify it and then import it back into Insight where it is then used to identify devices that are similar to those originally matching the exported identification.

About this task

This scenario might be used when devices have similar attributes that can be easily edited in the CSV file and then imported back into the system.

When you export a Fibre Channel device identification to a CSV file, the file contains the following information in the order shown:

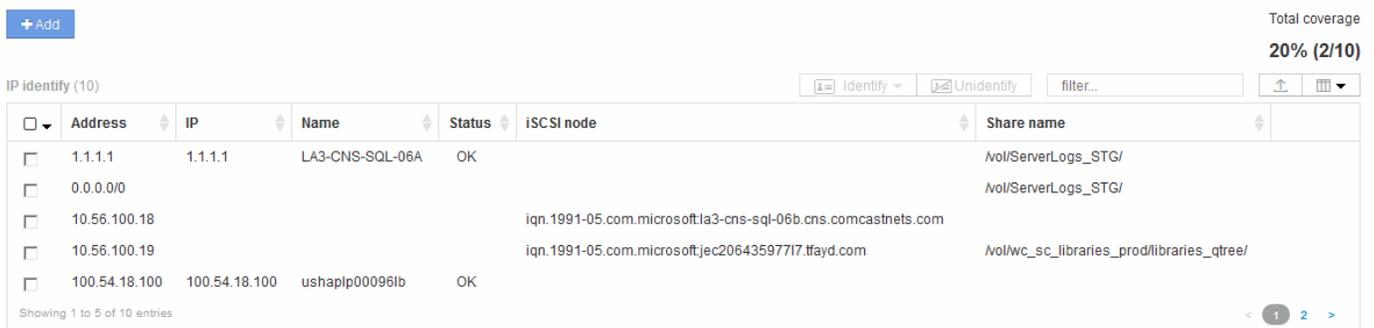
WWN
IP
Name
Type

Steps

1. Log into the Insight web UI.
2. Click **Manage > Device Resolution**
3. Select the **FC identify** tab.
4. Select the Fibre Channel device or devices whose identification you want to export.
5. Click the export  icon.
6. Chose if you want to open the CSV file or save the file.

IP device resolution

The IP Identify screen displays any iSCSI and CIFS or NFS shares that have been identified by automatic device resolution or by manual device resolution. Unidentified devices are also shown. The screen includes the IP address, Name, Status, iSCSI node, and share name for devices. The percentage of devices that have been successfully identified is also displayed.



IP Identify (10)

<input type="checkbox"/>	Address	IP	Name	Status	iSCSI node	Share name
<input type="checkbox"/>	1.1.1.1	1.1.1.1	LA3-CNS-SQL-06A	OK		/vol/ServerLogs_STG/
<input type="checkbox"/>	0.0.0.0/0					/vol/ServerLogs_STG/
<input type="checkbox"/>	10.56.100.18				iqn.1991-05.com.microsoft:la3-cns-sql-06b.cns.comcastnets.com	
<input type="checkbox"/>	10.56.100.19				iqn.1991-05.com.microsoft:jec20643597717.tfayd.com	/vol/wc_sc_libraries_prod/libraries_qtree/
<input type="checkbox"/>	100.54.18.100	100.54.18.100	ushapip000961b	OK		

Showing 1 to 5 of 10 entries

Adding IP devices manually

You can manually add an IP device to OnCommand Insight using the manual add feature available in the IP Identify screen.

Steps

1. Log in to the Insight web UI.
2. Click **Manage > Device resolution**
3. Click the **IP Identify** tab.
4. Click the add button.

The Add Device dialog is displayed

5. Enter the address, IP address, and a unique device name.

Results

The device you enter is added to the list of devices in the IP Identify tab.

Importing IP device identification from a CSV file

You can manually import IP device identifications into the Device Resolution feature using a list of device identifications in a CSV file.

Before you begin

You must have a correctly formatted CSV file in order to import device identifications. The CSV file for IP devices requires the following information:

Address
IP
Name



As a best practice, it is recommended to first export the IP Identify information to a CSV file, make your desired changes in that file, and then import the file back into IP Identify. This ensures that the expected columns are present and in the proper order.

To import IP Identify information:

Steps

1. Log into the Insight web UI.
2. Click **Manage > Device Resolution**
3. Select the **IP identify** tab.
4. Click **Identify > Identify from file**
5. a. Navigate to the folder containing your CSV files for import and select the desired file.

The devices you enter are added to the list of devices in the IP Identify tab.

Exporting IP device identification to a CSV file

You can export existing IP device identifications from Insight using the Device Resolution feature. You might want to export a device identification so that you can modify it and then import it back into Insight so that it can be used to identify devices that are similar to those in the exported identification.

About this task

When you export an IP device identification to a CSV file, the file contains the following information in the order shown:

Address
IP

Name

Steps

1. Log into the Insight web UI.
2. Click **Manage > Device Resolution**
3. Select the **IP Identify** tab.
4. Select the IP device or devices whose identification you want to export.
5. Click the export  icon.
6. Chose if you want to open the CSV file or save the file.

Setting options in the Preferences tab

The device resolution preferences tab lets you create an auto resolution schedule, specify storage and tape venders to include or exclude from identification, and set DNS lookup options.

Auto resolution schedule

An auto resolution schedule can specify when automatic device resolution is run:

Option	Description
Every	Use this option to run automatic device resolution on intervals of days, hours, or minutes.
Every day	Use this option to run automatic device resolution daily at a specific time.
Manually	Use this option to only run automatic device resolution manually.
On every environment change	Use this option to run automatic device resolution whenever there is a change in the environment.

If you specify manually, nightly automatic device resolution is disabled.

DNS processing options

DNS processing options allow you to select the following features:

- When DNS lookup result processing is enabled, you can add a list of DNS names to append to resolved devices.
- You can select "Auto resolution of IPs:" to enables automatic host resolution for iSCSI initiators and hosts accessing NFS shares by using DNS lookup. If this is not specified, only FC-based resolution is performed.
- You can choose to allow underscores in host names and to use a "connected to" alias instead of the standard port alias in results.

Including or excluding specific storage and tape vendors

You can include or exclude specific storage and tape vendors for automatic resolution. You might want to exclude specific vendors if you know, for example, that a specific host will become a legacy host and should be excluded from your new environment. You can also re-add vendors that you earlier excluded but no longer want excluded.



Device resolution rules for tape only work for WWNs where the Vendor for that WWN is set to **Included as Tape only** in the Vendors preferences.

Regular expression examples

If you have selected the regular expression approach as your source naming strategy, you can use the regular expression examples as guides for your own expressions used in the OnCommand Insight automatic resolution methods.

Formatting regular expressions

When creating regular expressions for OnCommand Insight automatic resolution, you can configure output format by entering values in a field named `FORMAT`.

The default setting is `\1`, which means that a zone name that matches the regular expression is replaced by the contents of the first variable created by the regular expression. In a regular expression, variable values are created by parenthetical statements. If multiple parenthetical statements occur, the variables are referenced numerically, from left to right. The variables can be used in the output format in any order. Constant text can also be inserted in the output, by adding it to the `FORMAT` field.

For example, you might have the following zone names for this zone naming convention:

```
[Zone number]_[data center]_[hostname]_[device type]_[interface number]
```

- S123_Miami_hostname1_filer_FC1
- S14_Tampa_hostname2_switch_FC4
- S3991_Boston_hostname3_windows2K_FC0
- S44_Raleigh_hostname4_solaris_FC1

And you might want the output to be in the following format:

```
[hostname]-[data center]-[device type]
```

To do this, you need to capture the host name, data center, and device type fields in variables, and use them in the output. The following regular expression would do this:

```
. *? _ ( [a-zA-Z0-9]+ ) _ ( [a-zA-Z0-9]+ ) _ ( [a-zA-Z0-9]+ ) _ . *
```

Because there are three sets of parentheses, the variables `\1`, `\2` and `\3` would be populated.

You could then use the following format to receive output in your preferred format:

```
\2-\1-\3
```

Your output would be as follows:

```
hostname1-Miami-filer  
hostname2-Tampa-switch  
hostname3-Boston-windows2K  
hostname4-Raleigh-solaris
```

The hyphens between the variables provide an example of constant text that is inserted in the formatted output.

Example 1 showing zone names

In this example, you use the regular expression to extract a host name from the zone name. You could create a regular expression if you have something similar to the following zone names:

- S0032_myComputer1Name-HBA0
- S0434_myComputer1Name-HBA1
- S0432_myComputer1Name-HBA3

The regular expression that you could use to capture the host name would be:

```
S[0-9]+_([a-zA-Z0-9]*)[_-]HBA[0-9]
```

The outcome is a match of all zones beginning with S that are followed by any combination of digits, followed by an underscore, the alphanumeric hostname (myComputer1Name), an underscore or hyphen, the capital letters HBA, and a single digit (0-9). The hostname alone is stored in the `\1` variable.

The regular expression can be broken into its components:

- "S" represents the zone name and begins the expression. This matches only an "S" at the beginning of the zone name.
- The characters [0-9] in brackets indicate that what follows "S" must be a digit between 0 and 9, inclusive.
- The + sign indicates that the occurrence of the information in the preceding brackets has to exist 1 or more times.
- The _ (underscore) means that the digits after S must be followed immediately by only an underscore character in the zone name. In this example, the zone naming convention uses the underscore to separate the zone name from the host name.
- After the required underscore, the parentheses indicate that the pattern contained within will be stored in the `\1` variable.
- The bracketed characters [a-zA-Z0-9] indicate that the characters being matched are all letters (regardless of case) and numbers.

- The * (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.
- The bracketed characters [_-] (underscore and dash) indicate that the alphanumeric pattern must be followed by an underscore or a dash.
- The letters HBA in the regular expression indicate that this exact sequence of characters must occur in the zone name.
- The final set of bracketed characters [0-9] match a single digit from 0 through 9, inclusive.

Example 2

In this example, skip up to the first underscore "", then match E and everything after that up to the second "", and then skip everything after that.

Zone: Z_E2FHDBS01_E1NETAPP

Hostname: E2FHDBS01

RegExp: .? (E.?) .*?

Example 3

The parentheses "(" around the last section in the Regular Expression (below) identifies which part is the hostname. If you wanted VSAN3 to be the host name, it would be: _([a-zA-Z0-9]).*

Zone: A_VSAN3_SR48KENT_A_CX2578_SPA0

Hostname: SR48KENT

RegExp: _[a-zA-Z0-9]+_([a-zA-Z0-9]).*

Example 4 showing a more complicated naming pattern

You could create a regular expression if you have something similar to the following zone names:

- myComputerName123-HBA1_Symm1_FA3
- myComputerName123-HBA2_Symm1_FA5
- myComputerName123-HBA3_Symm1_FA7

The regular expression that you could use to capture these would be:

```
([a-zA-Z0-9]*)_.*
```

The \1 variable would contain only myComputerName123 after being evaluated by this expression.

The regular expression can be broken into its components:

- The parentheses indicate that the pattern contained within will be stored in the \1 variable.
- The bracketed characters [a-zA-Z0-9] mean that any letter (regardless of case) or digit will match.
- The * (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.

- The `_` (underscore) character in the regular expression means that the zone name must have an underscore immediately following the alphanumeric string matched by the preceding brackets.
- The `.` (period) matches any character (a wildcard).
- The `*` (asterisk) indicates that the preceding period wildcard may occur 0 or more times.

In other words, the combination `.*` indicates any character, any number of times.

Example 5 showing zone names without a pattern

You could create a regular expression if you have something similar to the following zone names:

- `myComputerName_HBA1_Symm1_FA1`
- `myComputerName123_HBA1_Symm1_FA1`

The regular expression that you could use to capture these would be:

```
(.*?)_.*
```

The `\1` variable would contain *myComputerName* (in the first zone name example) or *myComputerName123* (in the second zone name example). This regular expression would thus match everything prior to the first underscore.

The regular expression can be broken into its components:

- The parentheses indicate that the pattern contained within will be stored in the `\1` variable.
- The `.*` (period asterisk) match any character, any number of times.
- The `*` (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.
- The `?` character makes the match non-greedy. This forces it to stop matching at the first underscore, rather than the last.
- The characters `_*` match the first underscore found and all characters that follow it.

Example 6 showing computer names with a pattern

You could create a regular expression if you have something similar to the following zone names:

- `Storage1_Switch1_myComputerName123A_A1_FC1`
- `Storage2_Switch2_myComputerName123B_A2_FC2`
- `Storage3_Switch3_myComputerName123T_A3_FC3`

The regular expression that you could use to capture these would be:

```
.*?_.*?_([a-zA-Z0-9]*[ABT])_.*
```

Because the zone naming convention has more of a pattern, we could use the above expression, which will match all instances of a hostname (*myComputerName* in the example) that ends with either an A, a B, or a T, placing that hostname in the `\1` variable.

The regular expression can be broken into its components:

- The .* (period asterisk) match any character, any number of times.
- The ? character makes the match non-greedy. This forces it to stop matching at the first underscore, rather than the last.
- The underscore character matches the first underscore in the zone name.
- Thus, the first .*?_ combination matches the characters *Storage1_* in the first zone name example.
- The second .*?_ combination behaves like the first, but matches *Switch1_* in the first zone name example.
- The parentheses indicate that the pattern contained within will be stored in the \1 variable.
- The bracketed characters [a-zA-Z0-9] mean that any letter (regardless of case) or digit will match.
- The * (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.
- The bracketed characters in the regular expression [ABT] match a single character in the zone name which must be A, B, or T.
- The _ (underscore) following the parentheses indicates that the [ABT] character match must be followed up an underscore.
- The .* (period asterisk) match any character, any number of times.

The result of this would therefore cause the \1 variable to contain any alphanumeric string which:

- was preceded by some number of alphanumeric characters and two underscores
- was followed by an underscore (and then any number of alphanumeric characters)
- had a final character of A, B or T, prior to the third underscore.

Example 7

Zone: myComputerName123_HBA1_Symm1_FA1

Hostname: myComputerName123

RegExp: ([a-zA-Z0-9]+)_.*

Example 8

This example finds everything before the first _.

Zone: MyComputerName_HBA1_Symm1_FA1

MyComputerName123_HBA1_Symm1_FA1

Hostname: MyComputerName

RegExp: (.*?)_.

Example 9

This example finds everything after the 1st _ and up to the second _.

Zone: Z_MyComputerName_StorageName

Hostname: MyComputerName

RegExp: .? (.?) .*?

Example 10

This example extracts "MyComputerName123" from the zone examples.

Zone: Storage1_Switch1_MyComputerName123A_A1_FC1

Storage2_Switch2_MyComputerName123B_A2_FC2

Storage3_Switch3_MyComputerName123T_A3_FC3

Hostname: MyComputerName123

RegExp: .? .? ([a-zA-Z0-9]+) [ABT]_.

Example 11

Zone: Storage1_Switch1_MyComputerName123A_A1_FC1

Hostname: MyComputerName123A

RegExp: .? .? ([a-zA-Z0-9]+) .*?

Example 12

The ^ (circumflex or caret) **inside square brackets** negates the expression, for example, [^Ff] means anything except uppercase or lowercase F, and [^a-z] means everything except lowercase a to z, and in the case above, anything except the _. The format statement adds in the "-" to the output host name.

Zone: mhs_apps44_d_A_10a0_0429

Hostname: mhs-apps44-d

RegExp: `([^_])_([AB]).*`Format in OnCommand Insight: `1-2 ()_`

([^_])_() . *Format in OnCommand Insight: 1-2-3

Example 13

In this example, the storage alias is delimited by "\" and the expression needs to use "\\" to define that there are actually "\" being used in the string, and that those are not part of the expression itself.

Storage Alias: \Hosts\E2DOC01C1\E2DOC01N1

Hostname: E2DOC01N1

RegExp: \\ .? \\ .? \\ (. *?)

Example 14

This example extracts "PD-RV-W-AD-2" from the zone examples.

Zone: PD_D-PD-RV-W-AD-2_01

Hostname: PD-RV-W-AD-2

RegExp: [^_]- (.-\d+) .+

Example 15

The format setting in this case adds the "US-BV-" to the hostname.

Zone: SRV_USBVM11_F1

Hostname: US-BV-M11

RegExp: SRV_USBV([A-Za-z0-9]+)_F[12]

Format: US-BV-\1

Maintaining Insight

Whether you are new to Insight and have a new system to set up, or your system has been operating for some time, you must take steps to maintain smooth operation of Insight and your network. The key maintenance concept is that changes in your network usually need to be accommodated in Insight.

These are the most common maintenance tasks:

- Maintaining Insight backups
- Updating expired Insight licenses
- Coordinating data source patches
- Updating the Insight version on all acquisition units
- Deleting removed data sources from Insight

Managing Insight

OnCommand Insight monitors your environment, enabling you to research potential problems before a crisis is reported. The Assets Dashboard provides summary pie charts, heat maps for IOPS, and an interactive chart of the top 10 utilized storage pools.

Steps

1. Open the Insight **Assets Dashboard** and move your cursor over the pie charts to examine the asset distribution in these three charts:
 - Capacity by Vendor shows the total raw capacity for storage by each vendor.
 - Capacity by Tier shows the total useable capacity for each storage tier.

- Switch Ports pie chart shows the manufacturers of ports and shows the percentage of ports used.
2. View **Facts About Your Environment** to see information about your environment's used capacity, the capacity's efficiency, consumed FC resources, and virtual infrastructure statistics.
 3. Position your cursor over a storage pool bar in the **Top 10 Utilized Pools** chart to view the used and unused capacity of the storage pool.
 4. Click any asset name appearing in large text (which indicates that the asset has issues) in the **Storage IOP** heat map to display a page summarizing the current state of that asset.
 5. In the lower right corner of the **Assets Dashboard**, click any asset name appearing in large text (which indicates the asset has issues) in the **Virtual Machine IOPS** heat map to display a page summarizing the current state of the asset.
 6. On the Insight toolbar, click **Admin**.
 7. Note any areas showing solid red circles.

In the OnCommand Insightweb UI, potential problems are marked with a solid red circle.

8. Click **Data Sources** to examine a list of all monitored data sources.

Examine any data source with a **Status** column containing a message with a solid red circle and with an **Impact** listed as High or Medium. These are at the top of the table. The problems with those data sources affect a significant portion of your network, which you need to address.

9. Click **Acquisition Units** to note the status for each IP address running Insight and to restart an acquisition unit, if necessary
10. Click **Health** to see high-level instance monitoring of the Insight servers.

Monitoring OnCommand Insight system health

You should periodically check the current status of your Insight system components by viewing the health page, which shows the status of each component and alerts you when there is an issue.

Steps

1. Log in to the Insightweb UI.
2. Click **Admin** and select **Health**.

The Health page is displayed.

3. View the summary of the current status of the components paying particular attention to any attention status in the **Details** column that is preceded by a red circle, which indicates an issue that requires your immediate attention.

The Health page displays information about any or all of the following Insight components based on your system configuration:

Component	Test	Details	Displays
-----------	------	---------	----------

Acquisition	Inventory data processing	Status of local acquisition unit	<p>“OK” if number of concurrently-polling data sources is less than 75% of execution pool maximum (default maximum is 30).</p> <p>“Acquisition is busy” if usage is greater than 75%, and recommends increasing polling interval or adding more remote acquisition units.</p>
DWH	Backup	Status of Data Warehouse scheduled backup	<p>“OK” and the last successful DWH backup time if DWH scheduled backup is enabled.</p> <p>Otherwise, displays information about any error found.</p>
DWH	ETL	Status of Data Warehouse ETL	<p>“OK” and the last successful DWH build time if no errors.</p> <p>Otherwise, displays information about any error found.</p>
Server	ASUP	Status of ASUP	<p>“ASUP Enabled” and the last successful phonehome time if available. “ASUP Failed” if phonehome is enabled but encountered a problem.</p> <p>+ "Invalid backup location" if backup directory is not valid.</p> <p>+ Displays the last successful phonehome time as well as time of the last failed attempt if available.</p> <p>+ “ASUP Disabled” if phonehome is disabled.</p>

Server	Auto resolution	Status of automatic device resolution	<p>“OK” if no errors. “Auto resolution is blocked” if identification errors prevent resolution progress.</p> <p>+ “Low success rate” if less than 75% of generic devices could be identified.</p>
Server	Elasticsearch	Status of elastic search data store	<p>“OK” if no errors. “Service unavailable” if unable to connect to elastic search service.</p> <p>+ "Cluster mode detected" if more than one node is detected.</p> <p>+ "High memory utilization" if heap space used is more than 85%.</p> <p>+ "Status: RED" indicates an error reported by elastic search. Displays information about the error and recommends contacting customer support.</p>
Server	CPU	Insight CPU usage	<p>“OK” if CPU load is less than 65%. “System CPU load is high. Reduce your CPU load.” if CPU load is greater than 65%.</p>
Server	Disk space	Status of disk space	<p>Free disk space, disk space in use by Insight, and recommended disk space reserved for Insight. “Low Disk Space” if disk utilization is more than 80%.</p>

Server	EventBus	Status of EventBus	“EventBus is empty” if EventBus queue is empty, otherwise displays status of EventBus queue.
Server	Inventory data processing	Status of inventory data processing capability of Insight server	“OK” if Insight server is not busy. “Server is busy” if the server is busy at least 75% of the time for the last hour. Recommends not adding more data sources and recommends splitting the environment to several servers.
Server	MySQL	Status of MySQL database	“OK” if no problems are detected. “The database is having performance issues. Some queries are taking too long to run” if the number of slow queries is more than 5%. + “The database log file grew more than <size> in the past hour. Check MySQL log file” if the error log grows to more than 20 KB.
Server	Performance archive	Status of performance archive	“Performance archive is enabled” or “Performance archive is not enabled”.
Server	Physical memory	Status of physical memory	“OK” if memory usage is less than 85%. “Memory usage is high. Reduce your overall memory footprint for system stability” if memory usage is greater than 85%.

Server	Service pack	Service pack availability	Displays whether a service pack is available for Insight. If a service pack is available, displays instructions.
Server	Usage information	Status of sending of usage information	<p>Displays whether sending of usage information to NetApp is enabled or disabled. Recommends enabling if disabled. Displays last attempted or last successful send time.</p> <p>+ Displays information on any problems encountered.</p>
Server	Violation	Status of open violations	<p>“OK” if the number of open violations is less than 75% of the violations limit. “Maximum number of open violations allowed is <number>” if the number of open violations is greater than 75% of the violations limit. Recommends reviewing performance policy configuration.</p> <p>+ “Violation manager is blocked” if the number of open violations is at the violations limit.</p> <p>+ Note that the violation manager cannot create new violations and recommends reviewing performance policy configuration.</p>
Server	Weekly backup	Status of weekly backup	“OK” if weekly backup is enabled, otherwise displays “Weekly backup is not enabled”.

Deleting inactive devices

Deleting devices that are inactive helps keep your data cleaner and easier to navigate.

About this task

To delete inactive devices from Insight, do the following:

Steps

1. Create a new query or open an existing query.
2. Choose either the *generic device*, *host*, *storage*, *switch*, or *tape* asset type.
3. Add a filter for **Is active**, and set the filter to **No**.

The results table displays only assets that are not active.

4. Select the devices that you want to delete.
5. Click the **Actions** button and select **Delete Inactive Devices**.

Your inactive devices are deleted and will no longer be displayed in Insight.

Auditing system and user activities

If you want to locate unexpected changes, you can view an audit trail of the OnCommand Insight system and its user activities. Audit log messages can optionally be sent to syslog in addition to being displayed on the Audit page.

About this task

Insight generates audit entries for any user activities that affect the storage network or its management, including the following:

- Logging in
- Authorizing or unauthorizing a path
- Updating an authorized path
- Setting global policies or thresholds
- Adding or removing a data source
- Starting or stopping a data source
- Updating data source properties
- Adding, editing, or deleting a task
- Removing an application group
- Identifying or changing the identification for a device
- Create a user
- Delete a user
- User role change

- Modify a user (Guest à Admin)
- Logout of a user (either forced logout or manual logout)
- Deleting an acquisition unit
- Update License
- Enabling backup
- Disabling Backup
- Enabling ASUP (Enabling Proxy on same page is reported in audit log)
- Disabling ASUP (Disabling Proxy on same page is reported in audit log)
- Security - re-key, change system passwords.
- Removing/adding annotations on assets
- CAC user logon / logoff
- CAC user session timeout

Steps

1. Open Insight in your browser.
2. Click **Admin** and select **Audit**.

The Audit page displays the audit entries in a table.

3. You can view the following details in the table:

- **Time**

Date and time that the changes were made

- **User**

Name of user associated with the audit entry

- **Role**

User account's role, which is guest, user, or administrator

- **IP**

IP address associated with the audit entry

- **Action**

Type of activity in the audit entry

- **Details**

Details of the audit entry

If there is a user activity that affects a resource, such as a data source or an application, the details include a link to the resource's landing page.



When a data source is deleted, the user activity details related to the data source no longer contain a link to the data source's landing page.

4. You can display audit entries by choosing a particular time period (1 hour, 3 hours, 24 hours, 3 days, and 7 days), with Insight showing a maximum number of 1000 violations for the selected time period.

You can click a page number below the table to browse through data by page if there is more data than fits on a single page.

5. You change the sort order of the columns in a table to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header; to return to the default sort order, click any other column header.

By default, the table displays the entries in descending order.

6. You can use the **filter** box to show only the entries you want in the table.

To see only the audit entries by the user `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, 7d, and 30d**), with Insight showing a maximum number of 1000 violations for the selected time period.
- You can use the **filter** box to show only the violations you want.
- You can change the sort order of the columns in a table to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header; to return to the default sort order, click any other column header.

By default, the table displays the violations in descending order.

- You can click a violation in the ID column to display the asset page for the duration of the violation.
- You can click the resource links (for example, storage pool and storage volume) in the Description column to display the asset pages associated with those resources.
- You can click the performance policy link in the Policy column to display the Edit Policy dialog box.

You might want to adjust the thresholds for a policy if you feel it generates too few or too many violations.

- You can click a page number to browse through data by page if there is more data than fits on a single page.
- You can click  to dismiss the violation.

Acquisition unit status

The Acquisition Unit screen provides a view of all your acquisition units, including status and any errors present.

The status of the Insight acquisition units connected to your server is displayed in the **Admin > Acquisition Units** table. This table displays the following information for each acquisition unit:

- **Name**
- **IP**
- **Status** is the operating status of the acquisition unit.
- **Last reported** displays the last time a data source connected to the acquisition unit reported.
- **Note** displays a user-entered note related to the AU.

If an acquisition unit in the list has a problem, the Status field will show a red circle with brief information about the problem. You should investigate any acquisition unit problems, as they likely affect data collection.

To restart an acquisition unit, hover over the unit and click on the *Restart Acquisition Unit* button that appears..

To add a text note, hover over an acquisition unit and click the *Add Note* button that appears. Only the most recently entered note is displayed.

Restoring the Insight database

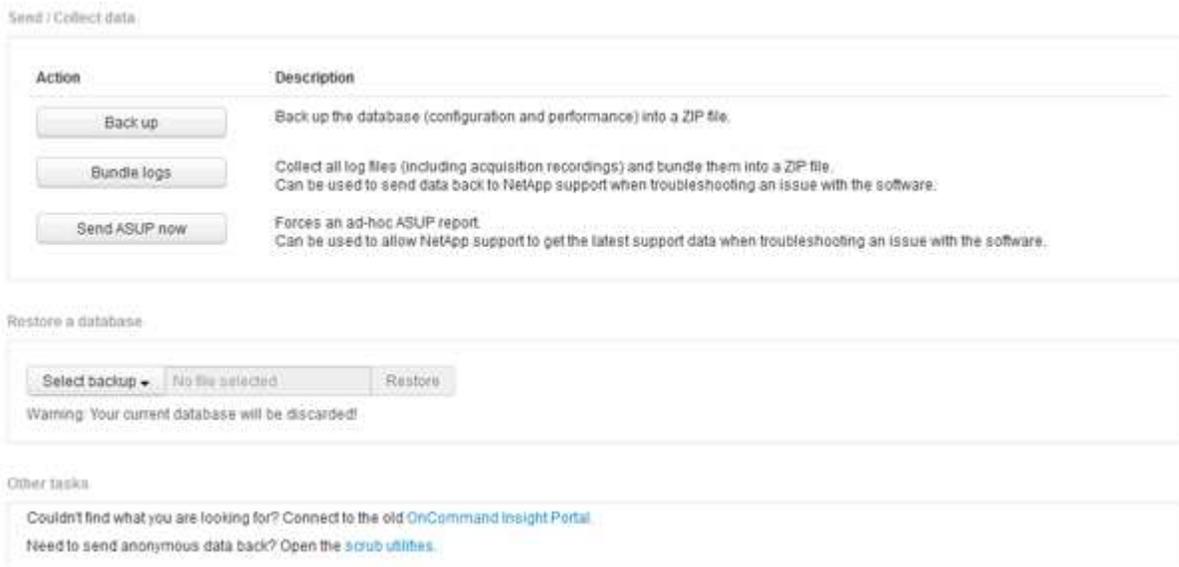
To restore your Insight database from a verified backup file, use the Troubleshooting options. This operation completely replaces your current OnCommand Insight data.

Before you begin

Best practice: Before restoring your OnCommand Insight database, use the manual backup process to create a copy of the current database. Check the backup file you plan to restore be certain that it was a successful backup containing the files you want to restore.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Troubleshooting**.



3. In the Restore a database section, select the backup file you want to restore from the **Select Backup** menu.
4. Click **Restore**.
5. On the warning that all data will be replaced, click **OK**

The status of the restore activity is displayed on the restore page.

Updating expired licenses

If one or more of your Insight licenses expired, you can update the licenses quickly using the same procedure as you did to install the licenses originally.

Steps

1. In a text editor, such as Notepad, open the new license file you received from NetApp Support and copy the license key text to your Windows Clipboard.
2. Open OnCommand Insight in your browser.
3. Click on **Admin** on the toolbar.
4. Click **Setup**.
5. Click the **Licenses** tab.
6. Click **Update License**.
7. Copy the license key text into the **License** text box.
8. Select the **Update (most common)** operation.

This operation adds your new licenses to any currently active Insight licenses.

9. Click **Save**.
10. If you are using the Insight consumption licensing model, you must check the box to **Enable sending usage information to NetApp** in the usage section. Proxy must be properly configured and enabled for your environment.

Licenses no longer compliant

If you notice the "Not Compliant" message on your Insight Licenses page, Insight is managing more terabytes than your company licensed.

The "Not Compliant" message means your company paid for fewer terabytes than Insight is currently managing. The difference between the managed terabytes and the licensed number of terabytes is shown beside the non-compliance message.

The operation of your Insight system is not affected, but you should contact your NetApp representative to increase your license coverage and update the appropriate license.

Replacing licenses for older Insight versions

If you have purchased a new Insight version that is not backward compatible with your older version of the product, you must replace the older licenses with the new licenses.

When you are installing the new licenses, you must select the **Replace** operation before you save the license key text.

Applying a service pack

Periodically, service packs are available, which you can apply to take advantage of fixes and enhancements to OnCommand Insight.

Before you begin

- You must have downloaded the service pack file (for example, `7.2service_pack_1.patch`) from the NOW site.
- You must have approved all patches.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Patches**.
3. From the Actions button, select **Apply patch**.
4. In the **Apply data source patch** dialog box, click **Browse** to locate the service pack file.
5. Inspect the **Patch name**, **Description**, **Impacted data source types**, which shows if any data sources are affected, and **Details**, which describes the enhancements that the service pack contains.
6. If the selected service pack is correct, click **Apply Patch**.

Service packs are approved automatically; no further action is required.

Preparing a special troubleshooting report

Insight sends information to NetApp Customer Support automatically through the ASUP system you set up after installing the software. However, you might want to create a troubleshooting report and open a case with the Support team for a specific problem.

You can use tools in Insight to perform a manual Insight backup, bundle the logs, and send that information to NetApp Customer Support.

Manually backing up the OnCommand Insight database

If you enabled weekly backups for the OnCommand Insight database, you are automatically generating copies that you can use to restore the database, if necessary. If you need to create a backup before a restore operation, or to send to NetApp technical support for assistance, you can create a backup .zip file manually.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Troubleshooting**.
3. In the Send/Collect data section, click **Backup**.
4. Click **Save File**.
5. Click **OK**.

Bundling logs for Support

When troubleshooting a problem with Insight software, you can quickly generate a zip file (using the "gz" format) of the logs and acquisition recordings to send to NetApp Customer Support.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Troubleshooting**.

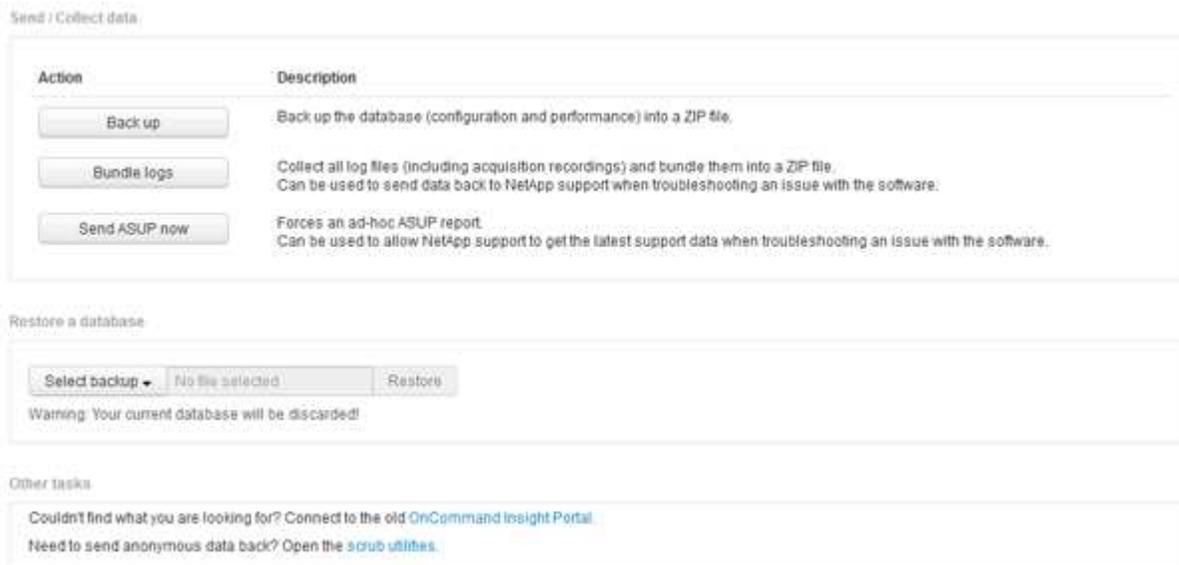
3. In the Send / Collect data section, click **Bundle logs**.
4. Click **Save File**.
5. Click **OK**.

Sending information to NetApp Support

The NetApp automated support (ASUP) facility sends troubleshooting information directly to the NetApp Customer Support team. You can force a special report to be sent.

Steps

1. On the Insight toolbar, click **Admin**.
2. Click **Setup**.
3. Click the **Backup/ASUP** tab.
4. In the Send/Collect data area, click **Send ASUP now** to submit your logs, recordings, and backup to NetApp Support.



Scrubbing data for transfer to support

Customers who have secure environments need to communicate with NetApp Customer Service to troubleshoot problems that arise without compromising their database information. The OnCommand Insight Scrub utilities allow you set up a comprehensive dictionary of keywords and patterns so that you can "cleanse" sensitive data and send scrubbed files to Customer Support.

Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **Scrub utilities** link.

There are several scrub sections: Lookup in Dictionary, Scrub data, and Build dictionary, Custom keywords, and Regular expressions.

3. a. In the **Lookup in dictionary** section, Enter a code to display the value it replaces, or enter a value to see the code that replaces it. Note: before you can do a lookup, you must **Build** the dictionary to identify values to scrub from the support data.
4. To add your own keywords to scrub from the support data, in the **Custom keywords** section, click **Actions** › **Add custom keyword**. Enter a keyword and click **Save**. The keyword is added to the dictionary.
5. Expand **Patterns (regex)**. Click **Add** to get the dialog box for entering a new pattern.
6. To use a regular expression to identify words or phrases to scrub, enter a pattern or patterns in the **Regular expressions** section. Click **Actions** › **Add regular expression**, enter a Name for the pattern and the Regular expression in the fields and click **Save**. The information has been added to the dictionary.



Patterns must be encompassed by round parentheses to identify a regular expression capturing group.

7. In the **Build dictionary** section, click **Build** to initiate compilation of the dictionary of all words identified as sensitive from the OnCommand Insight database.

On completion, you see a prompt informing you the revised dictionary is available. The Database description includes a line indicating how many keywords are in the dictionary. Check your keywords in the dictionary for accuracy. If you find problems and want to rebuild the dictionary, click **Reset** on the Database block to remove all keywords collected from the OnCommand Insight database from the dictionary. As the prompt advises, no other keywords will be deleted. Return to the Scrub utilities and enter your Custom Keywords again.

8. After you create a Scrub dictionary, you can use it to scrub a log, XML, or other text file to make the data anonymous.
9. To scrub a log, XML, or other text file, in the **Scrub data** section, Browse to locate the file and click **Scrub file**.

Advanced troubleshooting

To complete your OnCommand Insight configuration, you must use the advanced troubleshooting tools. These tools run in the browser and are opened from the **Admin** › **Troubleshooting** page.

To open the advanced troubleshooting tools in the browser, click the **Advanced Troubleshooting** link at the bottom of the page.

The advanced troubleshooting tools allow you to view various reports, system information, installed packages, and logs, as well as perform numerous actions such as restarting the server or acquisition units, update DWH annotations, and import annotations.

See the Advanced Troubleshooting page for all available options.

Configuring the number of hours to ignore dynamic data

You can configure the number of hours during which OnCommand Insight ignores updating dynamic data, such as used capacity. If the default of six hours is used and no configuration changes occur, reports will not be updated with dynamic data until after the default number of hours. This option improves performance because this option defers updates when only the dynamic data changes.

About this task

If a value is set for this option, OnCommand Insight will update dynamic data based on the following rules:

- If no configuration changes occur but capacity data changes, data will not be updated.
- Dynamic data (other than configuration changes) will be updated only after the timeout specified in this option.
- If configuration changes occur, configuration and dynamic data is updated.

Dynamic data impacted by this option includes the following:

- Capacity violation data
- File Systems Allocated Capacity and Used Capacity
- Hypervisor
 - Virtual Disk Used Capacity
 - Virtual Machine Used Capacity
- Internal Volume
 - Data Allocated Capacity
 - Data Used Capacity
 - Deduplication Savings
 - Last Known Access Time
 - Last Snapshot Time
 - Other Used Capacity
 - Snapshot Count
 - Snapshot Used Capacity
 - Total Used Capacity
- iSCSI Session Initiator IPs, Target Session ID, and Initiator Session ID
- Qtree Quota Used Capacity
- Quota Used Files and Used Capacity
- Storage Efficiency Technology, Gain/Loss, and Potential Gain/Loss
- Storage Pool
 - Data Used Capacity
 - Deduplication Savings
 - Other Used Capacity
 - Snapshot Used Capacity
 - Total Used Capacity
- Volume
 - Deduplication Savings
 - Last Known Access Time
 - Used Capacity

Steps

1. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **Advanced Troubleshooting** link.
3. Click the **Advanced settings** tab, in the Acquisition Dynamic Attributes section enter the number of hours that OnCommand Insight should ignore dynamic data for Acquisition Dynamic Attributes.
4. Click **Save**.
5. (Optional) To restart the acquisition unit, click the **Restart Acquisition Unit** link.

Restarting the local acquisition unit reloads all of the OnCommand Insight data source views. This change is applied during the next poll, so you do not have to restart the Acquisition Unit.

Generating logs for Customer Support

If requested by Customer Support, generate a server, acquisition, or remote log for troubleshooting purposes.

About this task

If NetApp Customer Support requests, use this option to generate the logs.

Steps

1. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click **Advanced Troubleshooting**.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click the **Logs** tab and select the log file to download.

A dialog box opens allowing you to open the log or save the log locally.

Displaying system information

You can display the Microsoft Windows IP configuration information about the system on which OnCommand Insight server is deployed.

Steps

1. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **Advanced Troubleshooting** link.
3. On the Advanced Troubleshooting page, click the **Reports** tab.
4. Click **System Information**.

The Windows IP configuration includes information such as the host name, DNS, IP address, subnet mask, OS information, memory, boot device, and connection name.

Listing installed OnCommand Insight components

You can display a list of the installed OnCommand Insight components including, among

others, inventory, capacity, dimensions, and the Data Warehouse views. Customer Support might ask you for this information, or you might want to see what software versions were installed and when they were installed.

Steps

1. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **Advanced Troubleshooting** link.
3. On the Advanced Troubleshooting page, click the **Reports** tab.
4. Click **Installed Software Packages**.

Calculating the number of database objects

To determine the number of objects in the OnCommand Insight database, use the Calculate Scale feature.

Steps

1. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **Advanced Troubleshooting** link.
3. On the Advanced Troubleshooting page, click the **Reports** tab.
4. Click **Calculated Scale**.

Restarting the OnCommand Insight Server

When you restart the OnCommand Insight Server, refresh the page and log into the OnCommand Insight Portal again.

About this task



Both of these options should only be used upon request by NetApp Customer Support. There is no confirmation prior to restart.

Steps

1. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **Advanced Troubleshooting** link.
3. On the next page in the Advanced menu, click the **Actions** tab.
4. Click **Restart Server**.

Moving MySQL data using the migrate option

You can use migrate MySQL data directory to a different directory. You can retain the current data directory. You can use the migrate option on the Troubleshooting menu or you can use the command line. This procedure describes how to use the **Troubleshooting > Migrate MySQL data** option.

About this task

If you retain the current data directory, it will be kept as a backup and renamed.

Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.
2. Click **Advanced Troubleshooting**.
3. Select the **Actions** tab
4. Select **Migrate MySQL Data**.
5. Enter the path to which you want to migrate the data.
6. To retain the existing data directory, check **Keep existing data directory**.
7. Click **Migrate**.

Moving MySQL data using the command line

You can use migrate MySQL data directory to a different directory. You can retain the current data directory. You can use the migrate option on the Troubleshooting menu or alternatively, you can use the command line. This procedure describes how to use the command line.

About this task

If you retain the current data directory, it will be kept as a backup and renamed.

You can use the Migrate MySQL Data utility or you can use a `java -jar mysqldatamigrator.jar` option in the OnCommand Insight path of `\bin\mysqldatamigrator` where the following parameters should be used:

- Mandatory parameters

- **-path**

The new data path to which the data folder will be copied.

- Optional parameters

- **-myCnf <my .cnf file>**

The path for the .cnf file. The default is `<install path>\mysql\my.cnf`. Use this flag only if a non-default MySQL is used.

- **-doBackup**

If this flag is set, the current data folder will be renamed but not deleted.

Steps

1. Access the command line tool here: `<installation path>\bin\mysqldatamigrator\mysqldatamigrator.jar```

Example usage

```
java -jar mysqldatamigrator.jar -path "C:\<new path>" -doBackup
```

Forcing annotation updates

If you have changed the annotations and want to use them in reports immediately, use one of the force annotation options.

Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.
2. On the bottom of the page, click the **Advanced Troubleshooting** link.
3. Click the **Actions** tab.
4. Select one of these options:
 - **Update DWH Annotations** to force the update of annotations in data warehouse to be used for reports.
 - **Update DWH Annotations (incl. deleted)** to force an annotations update (including deleted objects) in the data warehouse to be used for reports.

Checking the status of server resources

This option displays the OnCommand Insight Server's information including server memory, disk space, OS, and CPU and OnCommand Insight database information including innoDB data size and the disk free space where the database resides.

Steps

1. On the Insight toolbar, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **OnCommand Insight Portal** link.
3. On the next page in the Advanced menu, click the **Troubleshooting** link.
4. Click **Server Resources Status**.

For advanced OnCommand Insight users: The administrator can run some SQL tests to check the database and server's response time from the button at the end of the information summary. This option displays a warning if server resource is low.

Finding ghost data sources

If you have removed a device but the device data remains, you can locate any ghost data sources so that you can remove them.

Steps

1. In the web UI, click **Admin** and select **Troubleshooting**.
2. At the bottom of the page in the Other tasks area, click the **Advanced Troubleshooting** link.

3. On the **Reports** tab, click the **Ghost Data Sources** link.

OnCommand Insight produces a list of originators with their device information.

Adding a missing disk model

If acquisition fails due to an unknown disk model, you can add the missing disk model to the `new_disk_models.txt` file and run acquisition again.

About this task

As part of a poll of a storage device by OnCommand Insight acquisition, the disk models on the storage device are read. If a vendor has added new disk models to their array that Insight doesn't know about, or if there is a mismatch between the model number Insight looks for and the one returned by the storage device, acquisition of that data source will fail with an error. In order to prevent these errors, it is necessary to update the disk model information known to Insight. New disk models are added to Insight with updates, patches and maintenance releases. However, you may decide to update this information manually instead of waiting for a patch or update.

Because OnCommand Insight reads the disk model file every five minutes, any new data model information you enter is updated automatically. You do not need to restart the server for the changes to take effect, but you can opt to restart the server and any remote acquisition units (RAUs) to have the changes take effect before the next update.

Disk model updates are added to the `new_disk_models.txt` file located in the `<SANSscreenInstallDir>\wildfly\standalone\deployments\datasources.war` directory. Understand the information needed to describe your new disk model before updating the `new_disk_models.txt` file. Inaccurate information in the file produces incorrect system data and could result in failed acquisition.

Follow these instructions to manually update Insight disk models:

Steps

1. Locate the proper information for your disk model.
2. Using a text editor, open the `new_disk_models.txt` file.
3. Add the required information for the new data source.
4. Save the file in the `<SANSscreenInstallDir>\wildfly\standalone\deployments\datasources.war` directory on your server.
5. Back up the `new_disk_models.txt` file to a safe location. During any subsequent OnCommand Insight upgrade, this file will be overwritten. If your disk model information is not present in the upgraded file, you will need to re-enter it.

Locating required information for new disk model

To locate the disk model information, identify the vendor and model number and run an Internet search.

About this task

Locating disk model information is as simple as running an internet search. Be sure to note the vendor name and disk model number before searching.

Steps

1. It is recommended to use an advanced internet search for the vendor, model, and document type “PDF” to find the vendor’s data sheet and/or installation guide for the drive. These data sheets are usually the best source for vendor disk information.
2. Vendor specifications do not always provide all of the necessary information based on the full model number. It is often useful to search for different parts of the model number string on the vendor’s site to locate all of the information.
3. Locate the disk vendor name, full model number, disk size and speed, and the interface type in order to define the new disk model in OnCommand Insight You can use the following table as a guide to help note this information as you find it:

For this field:	Which is:	Enter this:
Model number (aka Key)	Required	
Vendor	Required	
Disk speed (RPM)	Required	
Size (in GB)	Required	
Interface Type (select one)	Required	ATA, SATA, SATA2, SATA3, FC, SAS, FATA, SSD, OTHER
Seek time in ms	Optional	
Maximum transfer rate in MB/sec	Optional	
Interface transfer rate in MB/sec	Optional	
Link to vendor/model information	Optional but recommended	

4. Enter that information into the `new_disk_models.txt` file. See [Content of the new_disk_models.txt file](#) for format, order, and examples.

Content of the new_disk_models.txt file

The `new_disk_models.txt` file has required and optional fields. The fields are comma separated, so do not use commas *within* the fields.

All fields are required except for seek time, transfer rates and `additional_info`. If available, include the vendor/model website link in the `additional_info` field.

Using a text editor, enter the following information in this order, separated by commas, for each new disk model you wish to add:

1. **key**: use the model number (required)
2. **vendor**: name (required)
3. **model number**: full number (usually the same value as in "key") (required)
4. **rpm of the disk**: for example 10000 or 15000 (required)
5. **size**: capacity in GB (required)
6. **interface type**: ATA, SATA, FC, SAS, FATA, SSD, OTHER (required)
7. **seek time**: in ms (optional)
8. **potential transfer rate**: the potential transfer rate in MB/sec. Maximum transfer rate of the disk itself. (optional)
9. **interface transfer rate**: the rate to and from the host in MB/sec (optional).
10. **Additional Info**: Any additional information you want to capture. Best practice is to enter the link to the vendor page where the specs are found, for reference (optional)

For any optional fields left blank, be sure to include the comma.

Examples (each on one line with no spaces):

```
ST373405,Seagate,ST373405,10000,73,FC,5.3,64,160,http://www.seagate.com/staticfiles/support/disc/manuals/enterprise/cheetah/73(LP)/100109943e.pdf
```

```
SLR5B-M400SS,HITACHI,SLR5B-M400SS,1000000,400,SSD,,,,,
```

```
X477_THARX04TA07,TOSHIBA,X477_THARX04TA07,7200,4000,SATA,9.5,,,https://storage.toshiba.eu/export/sites/toshiba-sdd/media/products/datasheets/MG03ACAxxy.pdf
```

Monitoring your environment

Insight helps you to prevent problems in your environment and troubleshoot potential problems quickly.

Asset page data

Asset pages provide performance troubleshooting data and present summary information about a base asset (such as a virtual machine or a volume) and the related assets it uses (such as storage pools, storage nodes, and connected switch ports), with links to additional information.

Beginning with OnCommand Insight 7.3.1, all asset pages have a **Main** page and an **Additional data** page. On the Main page are a summary of the asset and different sections for charts, topology and other information. The **Additional data** page allows you to configure a customizable dashboard page for the current asset type.

A solid red circle next to a line or message on the asset page main tab indicates potential issues with the monitored environment.

Types of asset pages

Asset pages summarize the current status of an asset and contain links to additional information about the asset and its related assets.

OnCommand Insight provides asset pages for the following assets:

- Virtual machine
- Volume
- Internal volume
- Physical host
- Storage pool
- Storage
- Datastore
- Hypervisor
- Application
- Storage node
- Qtree
- Disk
- VMDK
- Port
- Switch
- Fabric
- Object storage (for example, Atmos, Centera, Amazon S3)
- Zone

Mapping and Masking information can be viewed in tables on Zone, Volume, VM, and Host/Hypervisor asset pages.



Summary information is available for object storage assets; however, you can only access this information from the Data sources detail page.

Searching your environment for specific assets

You can locate information about specific assets by using the search facility. For example, if a system user contacts the storage administrator with a complaint about a particular server, the administrator can search the server name and display an asset page summarizing the status and supplying additional linked information.

Steps

1. Open the OnCommand Insightweb UI.
2. On the toolbar, click .

The **Search assets** box is displayed.

3. Enter the name of an asset or a portion of the name.
4. Select the resource you want from the search results.

The asset page for that resource is displayed.

Advanced search techniques

Multiple search techniques can be used to search for data or objects in your monitored environment.

Wildcard search

You can perform multiple character wildcard search using the * character. For example, *applic*n* would return application.

Phrases used in search

A phrase is a group of words surrounded by double quotation marks; for example, "PAW VNX LUN 5". You can use double quotes to search for documents that contain spaces in their names or attributes.

Boolean Operators

Using Boolean operators, you can combine multiple terms to form a more complex query.

- **OR**

- The OR operator is the default conjunction operator.

If there is no Boolean operator between two terms, the OR operator is used.

- The OR operator links two terms and finds a matching document if either of the terms exists in a document.

For example, "storage OR netapp" searches for documents that contain either "storage" or "netapp".

- High scores are given to documents that match most of the terms.

- **AND**

You can use the AND operator to find documents in which both the search terms exist in a single document. For example, "aurora AND netapp" searches for documents that contain both "storage" and "netapp".

You can use the symbol && instead of the word AND.

- **NOT**

When you use the NOT operator, all the documents that contain the term after NOT are excluded from the search results. For example, "storage NOT netapp" searches for documents that contains only "storage" and not "netapp".

You can use the symbol ! instead of the word NOT.

Prefix and suffix search

- As soon as you start typing a search string, the search engine does a prefix and suffix search to find the best match.
- Exact matches are given a higher score than a prefix or suffix match. The score is calculated based on the distance of the search term from the actual search result. For example, we have three storages: “aurora”, “aurora1”, and “aurora11”. Searching for “aur” will return all three storages. However, the search result for “aurora” will have the highest score because it has the closest distance to the prefix search string.
- The search engine also searches for terms in reverse order, which allows you to perform a suffix search. For example, when you type “345” in the search box, the search engine searches for “345”.
- Search is case-insensitive.

Search using indexed terms

Searches that match more of the indexed terms result in higher scores.

The search string is split into separate search terms by space. For example, the search string “storage aurora netapp” is split into three keywords: “storage”, “aurora”, and “netapp”. The search is performed using all three terms. The documents that match most of these terms will have the highest score. The more information you provide, the better are the search results. For example, you can search for a storage by its name and mode.

The UI displays the search results across categories, with the three top results per category. If you did not find a document that you were expecting, you can include more terms in the search string to improve the search results.

The following table provides a list of indexed terms that can be added to the search string.

Category	Indexed terms
Storage	<ul style="list-style-type: none">• “storage”• name• vendor• model
StoragePool	<ul style="list-style-type: none">• “storagepool”• name• name of the storage• IP addresses of the storage• serial number of the storage• storage vendor• storage model• names for all associated internal volumes• names for all associated disks

Internal Volume	<ul style="list-style-type: none"> • “internalvolume” • name • name of the storage • IP addresses of the storage • serial number of the storage • storage vendor • storage model • name of the storage pool • names of all associated shares • names of all associated applications and business entities
Volume	<ul style="list-style-type: none"> • “volume” • name • label • names of all internal volumes • name of the storage pool • name of the storage • IP addresses of the storage • serial number of the storage • storage vendor • storage model
Storage Node	<ul style="list-style-type: none"> • “storagenode” • name • name of the storage • IP addresses of the storage • serialnumber of the storage • storage vendor • storage model
Host	<ul style="list-style-type: none"> • “host” • name • IP addresses • names of all associated applications and business entities

<p>Datastore</p>	<ul style="list-style-type: none"> • “datastore” • name • virtual center IP • names of all volumes • names of all internal volumes
<p>Virtual Machines</p>	<ul style="list-style-type: none"> • “virtualmachine” • name • DNS name • IP addresses • name of the host • IP addresses of the host • names of all datastores • names of all associated applications and business entities
<p>Switches (regular and NPV)</p>	<ul style="list-style-type: none"> • “switch” • IP address • wwn • name • serial number • model • domain ID • name of the fabric • wwn of the fabric
<p>Application</p>	<ul style="list-style-type: none"> • “application” • name • tenant • line of business • business unit • project
<p>Tape</p>	<ul style="list-style-type: none"> • “tape” • IP address • name • serial number • vendor

Port	<ul style="list-style-type: none"> • “port” • wwn • name
Fabric	<ul style="list-style-type: none"> • “fabric” • wwn • name

Changing the time range of displayed data

By default, an asset page displays the last 24 hours of data; however, you can change the segment of data displayed by selecting another fixed time or a custom range of time to view less or more data.

About this task

You can change the time segment of displayed data by using an option that is located on every asset page, regardless of asset type.

Steps

1. Log in to the OnCommand Insightweb UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it. The asset page displays.
3. In the upper left corner of the page, click any of the following time icons to change the segment of data displayed:
 - **3h**
Displays the last three hours of data.
 - **24h**
Displays the last 24 hours of data.
 - **3d**
Displays the last three days of data.
 - **7d**
Displays the last seven days of data.
 - **30d**
Displays the last thirty days of data.

- **Custom**

Displays a dialog box that enables you to choose a custom range of time. You may display up to 31 days of data at a time.

4. If you chose **Custom**, do the following:
 - a. Click the date field, and select a month, day, and year for the begin date.
 - b. Click the time list, and select a begin time.
 - c. Repeat both steps a and b for the end data and time.
 - d. Click .

Determining data source acquisition status

Because data sources are the primary source of information for Insight, it is imperative that you ensure that they remain in a running state.

The ability to see the data source acquisition status is available on every asset page for all assets that are directly acquired. Either of the following acquisition scenarios can occur, in which the status is displayed in the upper right corner of the asset page:

- Acquired successfully from data source

Displays the status “Acquired `xxxxx`”, where `xxxxx` indicates the most recent acquisition time of the asset’s data sources.

- There is an acquisition error.

Displays the status “Acquired `xxxxx`”, where `xxxxx` indicates the most recent acquisition time of the asset’s one or more data sources with . When you click , a window displays each data source for the asset, the data source’s status, and the last time data was acquired. Clicking a data source displays the data source’s detail page.

If an asset is not directly acquired, no status is displayed.

Asset page sections

An asset page displays several sections containing information relevant to the asset. The sections that you see depend on the type of asset.

Summary

The Summary section on an asset page displays a summary of information about the particular asset and shows issues related to the asset, indicated by a red circle, with hyperlinks to additional information about related assets and to any performance policies assigned to the asset.

The following example shows some of the types of information available in the Summary section of an asset page for a virtual machine. Any item with a solid red circle next to it indicates potential issues with the monitored environment.

Summary

Power state:	On
Guest state:	Running
Datstore:	DS_SP1_1
CPU:	41.05%
Memory:	● 51% (1,047 / 2,048 MB)
Capacity:	10% (19.5 / 195.3 GB)
Latency:	1.93 ms (6.00 ms max)
IOPS:	1,317.33 IO/s (4,964.00 IO/s max)
Throughput:	38.79 MB/s (142.00 MB/s max)
DNS name:	VM_Cs_travBookcomp.com
IP:	10.97.133.23
OS:	Microsoft Windows Server 2008 R2(64-bit)
Processors:	4
FC Fabrics Connected:	1
Performance Policies:	VM Latency-Critical VM Latency-Warning Comp Corp.Customer Support SLA latency ● Exchange SLO

Using the Summary section

You can view the Summary section to see general information about an asset. Specifically, it is helpful to see if any metrics (for example, memory, capacity, and latency) or any performance policies are cause for concern, which OnCommand Insight indicates by displaying a red circle next to the metric or performance policy.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it. The asset page displays.



The information that displays in the Summary section depends on the type of asset page you are viewing.

3. You can click any of the asset links to view their asset pages.

For example, if you are viewing a storage node, you can click a link to view the asset page of the storage it is associated with or click to view the asset page of the HA partner.

4. You can view the metrics associated with the asset.

A red circle next to a metric indicates that you might need to diagnose and resolve potential problems.



You may notice that volume capacity might show greater than 100% on some storage assets. This is due to metadata related to the capacity of the volume being part of the consumed capacity data reported by the asset.

5. If applicable, you can click a performance policy link to view the performance policy or policies associated with the asset.

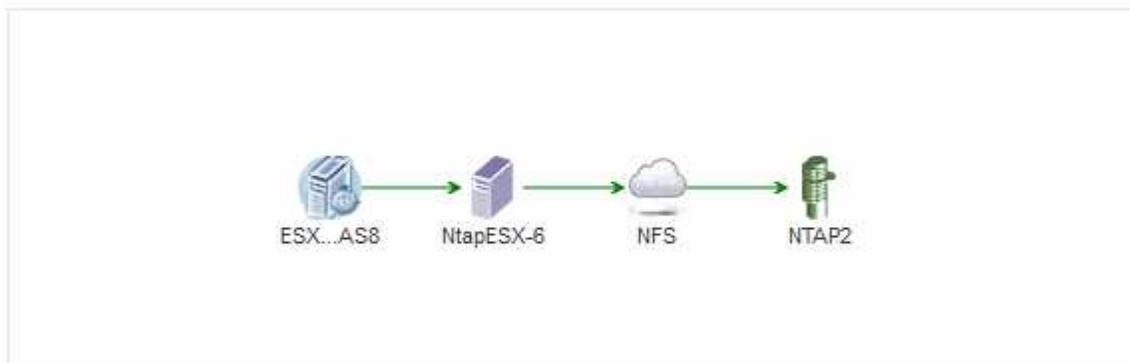
If a red circle appears next to a performance policy, this indicates an asset has crossed the performance policy's defined threshold. You should examine the performance policy to further diagnose the issue.

Topology

The Topology section, if applicable to an asset, enables you to see how a base asset is connected to its related assets.

The following shows an example of what might display in the Topology section of a virtual machine asset page.

Topology



If the topology for the asset is larger than will fit in the section, the **Click link to see the topology** hyperlink is displayed instead.

Using the Topology section

The Topology section enables you to view how the assets in your network are connected to each other and display information about related assets.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it. The asset page displays. You can find the Topology section in the upper right-hand corner of the asset page.

If the topology for the asset is larger than will fit in the section, click the **Click link to see the topology** hyperlink.

3. To view more information about the base asset's related assets, position your cursor over a related asset in the topology and click its name, which displays its asset page.

User Data

The User Data section of an asset page displays and enables you to change any user-defined data such as applications, business entities, and annotations.

The following shows an example of what might display in the User Data section of a virtual machine asset page when an application, business entity, and annotation are assigned to the asset:

User Data

Application(s):	Concur
Business Entities:	Hybridsoft Corporation.Sales.Wes...
Birthday:	01/30/2016  
+ Add	

Using the User Data section to assign or modify applications

You can assign applications running in your environment to certain assets (host, virtual machines, volumes, internal volumes, and hypervisors). The User Data section enables you to change the application assigned to an asset or assign an application or additional applications to an asset.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it. The asset page displays.
3. You can do the following:
 - To view the asset page for the application, click the application's name.
 - To change the application assigned or to assign an application or additional applications, position your cursor over the application name, if an application is assigned, or over **None**, if no application is assigned, click , type to search for an application or select one from the list, and then click .

If you choose an application that is associated with a business entity, the business entity is automatically assigned to the asset. In this case, when you place your cursor over the business entity name, the word *derived* displays. If you want to maintain the entity for only the asset and not the associated application, you can manually override the assignment of the application.

- To remove an application, click .

Using the User Data section to assign or modify business entities

You can define business entities to track and report on your environment data at a more granular level. The User Data section in an asset page enables you to change the business entity assigned to an asset or remove a business entity from an asset.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it. The asset page displays.
3. You can do the following:
 - To change the entity assigned or to assign an entity, click  and select an entity from the list.
 - To remove a business entity, click .



You cannot remove an entity that is derived from an application that is assigned to the asset.

Using the User Data section to assign or modify annotations

When customizing OnCommand Insight to track data for your corporate requirements, you can define specialized notes, called *annotations*, and assign them to your assets. The User Data section of an asset page displays annotations assigned to an asset and also enables you to change the annotations assigned to that asset.

Steps

1. Log in to the OnCommand Insight web UI.
 2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it. The asset page displays.
 3. In the **User Data** section of the asset page, click .
- The Add Annotation dialog box displays.
4. Click **Annotation** and select an annotation from the list.
 5. Click **Value** and do either of the following, depending on type of annotation you selected:
 - If the annotation type is list, date, or Boolean, select a value from the list.
 - If the annotation type is text, type a value.
 6. Click **Save**.

The annotation is assigned to the asset. You can later filter assets by annotation using a query.

7. If you want to change the value of the annotation after you assign it, click  and select a different value.

If the annotation is of list type for which the **Add values dynamically upon annotation assignment** option is selected, you can type to add a new value in addition to selecting an existing value.

Expert view

The Expert View section of an asset page enables you to view a performance sample for the base asset based on any number of applicable metrics in context with a chosen time period (3 hours, 24 hours, 3 days, 7 days, or a custom time period) in the performance chart and any assets related to it.

The following is an example of the Expert View section in a volume asset page:



You can select the metrics you want to view in the performance chart for the time period selected.

The Resources section shows the name of the base asset and the color representing the base asset in the performance chart. If the Top Correlated section does not contain an asset you want to view in the performance chart, you can use the **Search assets** box in the Additional resources section to locate the asset and add it to the performance chart. As you add resources, they appear in the Additional resources section.

Also shown in the Resources section, when applicable, are any assets related to the base asset in the following categories:

- Top correlated

Shows the assets that have a high correlation (percentage) with one or more performance metrics to the base asset.

- Top contributors

Shows the assets that contribute (percentage) to the base asset.

- Greedy

Shows the assets that take away system resources from the asset through sharing the same resources,

such as hosts, networks, and storage.

- Degraded

Shows the assets that are depleted of system resources due to this asset.

Expert View metric definitions

The Expert View section of an asset page displays several metrics based on the time period selected for the asset. Each metric is displayed in its own performance chart. You can add or remove metrics and related assets from the charts depending on what data you want to see.

Metric	Description
BB credit zero Rx, Tx	Number of times the receive/transmit buffer-to-buffer credit count transitioned to zero during the sampling period. This metric represents the number of times the attached port had to stop transmitting because this port was out of credits to provide.
BB credit zero duration Tx	Time in milliseconds during which the transmit BB credit was zero during the sampling interval.
Cache hit ratio (Total, Read, Write) %	Percentage of requests that result in cache hits. The higher the number of hits versus accesses to the volume, the better is the performance. This column is empty for storage arrays that do not collect cache hit information.
Cache utilization (Total) %	Total percentage of cache requests that result in cache hits
Class 3 discards	Count of Fibre Channel Class 3 data transport discards.
CPU utilization (Total) %	Amount of actively used CPU resources, as a percentage of total available (over all virtual CPUs).
CRC error	Number of frames with invalid cyclic redundancy checks (CRCs) detected by the port during the sampling period
Frame rate	Transmit frame rate in frames per second (FPS)
Frame size average (Rx, Tx)	Ratio of traffic to frame size. This metric enables you to identify whether there are any overhead frames in the fabric.

Frame size too long	Count of Fibre Channel data transmission frames that are too long.
Frame size too short	Count of Fibre Channel data transmission frames that are too short.
I/O density (Total, Read, Write)	Number of IOPS divided by used capacity (as acquired from the most recent inventory poll of the data source) for the Volume, Internal Volume or Storage element. Measured in number of I/O operations per second per TB.
IOPS (Total, Read, Write)	Number of read/write I/O service requests passing through the I/O channel or a portion of that channel per unit of time (measured in I/O per sec)
IP throughput (Total, Read, Write)	Total: Aggregated rate at which IP data was transmitted and received in megabytes per second. Read: IP Throughput (Receive): Average rate at which IP data was received in megabytes per second. Write: IP Throughput (Transmit): Average rate at which IP data was transmitted in megabytes per second.
Latency (Total, Read, Write)	Latency (R&W): Rate at which data is read or written to the virtual machines in a fixed amount of time. The value is measured in megabytes per second. Latency: Average response time from the virtual machines in a data store. Top Latency: The highest response time from the virtual machines in a data store.
Link failure	Number of link failures detected by the port during the sampling period.
Link reset Rx, Tx	Number of receive or transmit link resets during the sampling period. This metric represents the number of link resets that were issued by the attached port to this port.
Memory utilization (Total) %	Threshold for the memory used by the host.

Partial R/W (Total) %	<p>Total number of times that a read/write operation crosses a stripe boundary on any disk module in a RAID 5, RAID 1/0, or RAID 0 LUN. Generally, stripe crossings are not beneficial, because each one requires an additional I/O. A low percentage indicates an efficient stripe element size and is an indication of improper alignment of a volume (or a NetApp LUN).</p> <p>For CLARiiON, this value is the number of stripe crossings divided by the total number of IOPS.</p>
Port errors	Report of port errors over the sampling period/given time span.
Signal loss count	Number of signal loss errors. If a signal loss error occurs, there is no electrical connection, and a physical problem exists.
Swap rate (Total Rate, In rate, Out rate)	Rate at which memory is swapped in, out, or both from disk to active memory during the sampling period. This counter applies to virtual machines.
Sync loss count	Number of synchronization loss errors. If a synchronization loss error occurs, the hardware cannot make sense of the traffic or lock onto it. All the equipment might not be using the same data rate, or the optics or physical connections might be of poor quality. The port must resynchronize after each such error, which impacts system performance. Measured in KB/sec.
Throughput (Total, Read, Write)	Rate at which data is being transmitted, received, or both in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Timeout discard frames - Tx	Count of discarded transmit frames caused by timeout.
Traffic rate (Total, Read, Write)	Traffic transmitted, received, or both received during the sampling period, in mebibytes per second.
Traffic utilization (Total, Read, Write)	Ratio of traffic received/transmitted/total to receive/transmit/total capacity, during the sampling period.
Utilization (Total, Read, Write) %	Percentage of available bandwidth used for transmission (Tx) and reception (Rx).

Write pending (Total)	Number of write I/O service requests that are pending.
-----------------------	--

Using the Expert View section

The Expert view section enables you to view performance charts for an asset based on any number of applicable metrics during a chosen time period, and to add related assets to compare and contrast asset and related asset performance over different time periods.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it. The asset page displays. By default, the performance chart shows two metrics for time period selected for the asset page. For example, for a storage, the performance chart shows latency and total IOPS by default. The Resources section displays the resource name and an Additional resources section, which enables you to search for assets. Depending on the asset, you might also see assets in the Top correlated, Top contributor, Greedy, and Degraded sections.
3. You can click **Select metrics to show**, and select a metric to add a performance chart for a metric.

A performance chart is added for the selected metric. The chart displays the data for the selected time period. You can change the time period by clicking on another time period in the top left-hand corner of the asset page.

You can perform the step again, and click to clear a metric. The performance chart for the metric is removed.

4. You can position your cursor over the chart and change the metric data that displays by clicking either of the following, depending on the asset:
 - **Read or Write**
 - **Txor Rx**
Total is the default.
5. You can drag your cursor over the data points in the chart to see how the value of the metric changes over the time period selected.
6. In the **Resources** section, you can do any of the following, if applicable, to add any related assets to the performance charts:
 - You can select a related asset in the Top correlated, Top contributors, Greedy, or Degraded sections to add data from that asset to the performance chart for each selected metric. Assets must have a minimum 15% correlation or contribution to be shown.

After you select the asset, a color block appears next to the asset to denote the color of its data points in the chart.

- For any asset shown, you can click the asset name to display its asset page, or you can click the percentage that the asset correlates or contributes to the base asset to view more information about the assets relation to the base asset.

For example, clicking the linked percentage next to a top correlated asset displays an informational message comparing the type of correlation that asset has with the base asset.

- If the Top correlated section does not contain an asset you want to display in a performance chart for comparison purposes, you can use the **Search assets** box in the Additional resources section to locate other assets.

After you select an asset, it displays in the Additional resources section. When you no longer want to view information about the asset, click  .

Related Assets

If applicable, an asset page displays a Related Assets section. For example, a volume asset page might show information about assets like Storage Pools, Connected switch ports, and Compute Resources. Each section comprises a table that lists any of the related assets in that category, with links to their respective asset pages, and several performance statistics related to the asset.

Using the Related Assets section

The Related Assets section enables you to view any of the assets that are related to the base asset. Each related asset is displayed in a table along with pertinent statistics for the asset. You can export the asset information, view the asset statistics in the Expert View performance charts, or show a chart that displays statistics for only related assets.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate an asset page by doing either of the following:
 - On the Insight toolbar, click  , type in the name of the asset, and then select the asset from the list.
 - Click **Dashboards**, select **Assets Dashboard**, locate an asset name, and click it. The asset page displays.
3. To control how assets display in the table:
 - Click the name of any asset to display its asset page.
 - Use the **filter** box to show only specific assets.
 - Click a page number to browse through the assets by page if there are more than five assets in the table.
 - Change the sort order of the columns in a table to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header.
 - Add a related asset to any performance chart in the Expert View section by placing your cursor over the related asset and clicking .
4. To export the information displayed in the table to a .CSV file:
 - a. Click  .
 - b. Click **Open with** and then **OK** to open the file with Microsoft Excel and save the file to a specific location, or click **Save file** and then **OK** to save the file to your Downloads folder.

All of the object attributes for the columns currently selected for display are exported to the file. Only the attributes for the displayed columns will be exported. Note that only the first 10,000 rows of the

table are exported.

- To display the related asset information in a chart below the table, click  and do any of the following:
 - Click **Read**, **Write**, or **Total** to change the metric data that displays. **Total** is the default.
 - Click  to select a different metric.
 - Click  to change the chart type. **Line chart** is the default.
 - Move your cursor over the data points in the chart to see how the value of the metric changes over the time period selected for each related asset.
 - Click a related asset in the chart legend to add it to or remove it from the chart.
 - Click a page number in the related asset table to view other related assets in the chart.
 - Click  to close the chart.

Violations

You can use the Violations section of an asset page to see the violations, if any, that occur in your environment as a result of a performance policy assigned to an asset. Performance policies monitor your network thresholds and enable you to detect a violation of a threshold immediately, identify the implication, and analyze the impact and root cause of the problem in a manner that enables rapid and effective correction.

The following example shows a Violations section that displays on an asset page for a hypervisor:

Violations filter...

Time	Description
06/05/2015 5:00:00 pm	Port balance index of 74 on esx1 exceeds the threshold of 50
06/12/2015 8:59:54 am	2 violations for esx2 with 'Swap out rate' > 3
06/12/2015 12:04:54 pm	esx1 violation with 'Swap out rate' > 3.00 KB/s (value of 86.85 KB/s)
06/12/2015 12:29:54 pm	esx1 violation with 'Swap in rate' > 3.00 KB/s (value of 59.90 KB/s)
06/12/2015 1:04:54 pm	7 violations for ds-30 with 'Latency - Total' > 50

Showing 1 to 5 of 32 entries < 1 2 3 4 5 >

Using the Violations section

The Violations section enables you to view and manage any of the violations that occur in your network as the result of a performance policy assigned to an asset.

Steps

- Log in to the OnCommand Insight web UI.
- 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").
- You can perform any of the following optional tasks:
 - Use the **filter** box to show only specific violations.

- Click a page number to browse through the violations by page if there are more than five violations in the table.
- Change the sort order of the columns in a table to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header.
- Click the asset name in any description to display its asset page; a red circle indicates issues that need further investigation.

You can click the performance policy, which displays the Edit Policy dialog box, to review the performance policy and make changes to the policy if necessary.

- Click  to remove a violation from the list if you determine the issue is no longer a cause for concern.

Customizable asset page

Additional data can be displayed in customizable widgets on each asset page. Customizing the page for an asset applies the customization to the pages for all assets of that type.

You customize asset page widgets by performing the following actions:

1. Add a widget to the page
2. Create a query or expression for the widget to showcase desired data
3. Choose a filter if desired
4. Choose a rollup or grouping method
5. Save the widget
6. Repeat for all desired widgets
7. Save the asset page

You can also add variables to the custom asset page that can be used to further refine your showcased data in widgets. In addition to regular variables, each asset type can use a set of "\$this" variables to quickly identify resources directly related to the current asset, for example, all virtual machines hosted by the same hypervisor that hosts the current virtual machine.

This custom asset page is unique for each user as well as for each asset type. For example, if User A creates a custom asset page for a virtual machine, that custom page will display for any virtual machine asset page, for that user.

Users can only view, edit, or delete custom asset pages that they create.

Custom asset pages are not included in Insight's export/import functionality.

Understanding "\$this" variables

Special variables on an asset's "Additional data" customizable page allow you to easily showcase additional information that is directly related to the current asset.

About this task

To use the "\$this" variables in widgets on your asset's customizable landing page, follow the steps below. For this example, we will add a table widget.



“\$this” variables are only valid for an asset’s customizable landing page. They are not available for other Insight dashboards. The available “\$this” variables varies according to asset type.

Steps

1. Navigate to an asset page for an asset of your choosing. For this example, let’s choose a Virtual Machine (VM) asset page. Query or search for a VM and click on the link to go to that VM’s asset page.

The asset page for the VM opens.

2. Click on the **Change view:** > **Additional Virtual Machine data** drop-down to go to that asset’s customizable landing page.
3. Click on the **Widget** button and choose **Table widget**.

The Table widget opens for editing. By default, all storages are shown in the table.

4. We want to show all virtual machines. Click on the asset selector and change **Storage** to **Virtual Machine**.

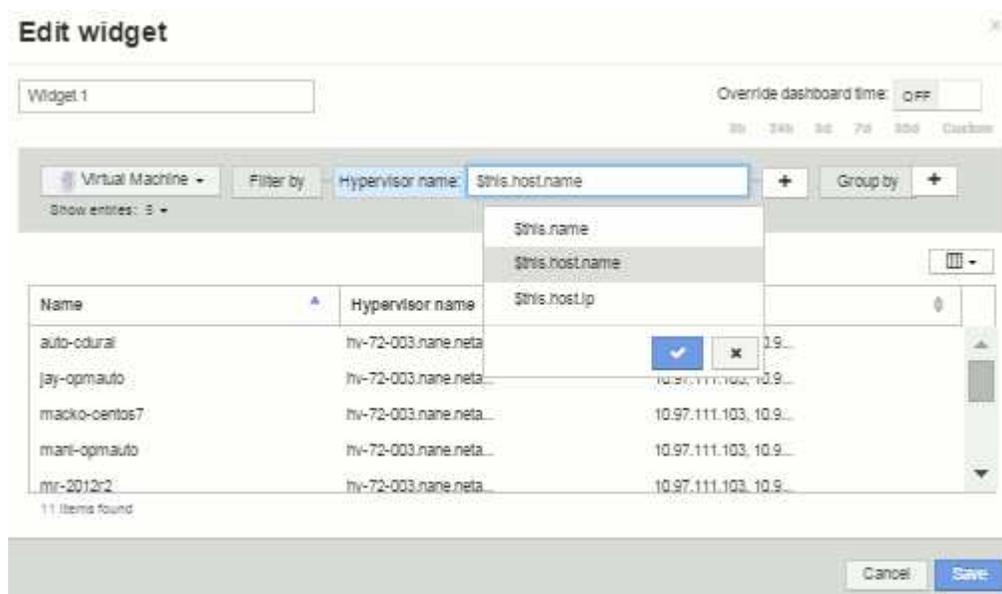
All virtual machines are now shown in the table.

5. Click on the **Column selector** button  and add the **hypervisor name** field to the table.

The hypervisor name is shown for each VM in the table.

6. We only care about the hypervisor that hosts the current VM. Click on the **Filter by** field’s+button and select **hypervisor name**.

7. Click on **Any** and select the **\$this.host.name** variable. Click the check button to save the filter.



The screenshot shows the 'Edit widget' interface. At the top, there's a text input 'Widget 1' and an 'Override dashboard time' dropdown set to 'OFF'. Below that, a 'Virtual Machine' selector is active. The 'Filter by' dropdown is set to 'Hypervisor name' with a search input containing '\$this.host.name'. A dropdown menu is open showing options: '\$this.name', '\$this.host.name' (highlighted), and '\$this.host.ip'. The table below has columns for Name, Hypervisor name, and IP. The table contains 11 items, with the first few rows visible: 'auto-cdural', 'jay-opmauto', 'macko-centos7', 'mani-opmauto', and 'mr-2012r2'. All have the same hypervisor name: 'hvy-72-003.nane.neta...'. The 'IP' column shows values like '10.97.111.103, 10.9...'. At the bottom, there are 'Cancel' and 'Save' buttons.

8. The table now shows all the VM’s hosted by the current VM’s hypervisor. Click **Save**.

Results

The table that you created for this virtual machine asset page will be displayed for any VM asset page you display. The use of the **\$this.host.name** variable in the widget means that only the VM’s owned by the current assets’s hypervisor will be displayed in the table.

Balancing network resources

To resolve balancing issues, use the asset pages to find the problems and identify high capacity volumes that are underused.

Steps

1. Open the Assets Dashboard in your browser.
2. In the Virtual Machines IOPS heat map, you notice the name of a VM in very large print that often reports problems.
3. Click the VM name to display the asset page.
4. Check for error messages in the summary.
5. Check the performance charts and particularly the top correlated resources to locate any volumes that might be in contention.
6. Add volumes to the performance chart to compare the patterns of activity and display more asset pages for other resources involved in the problem.
7. Scroll to the bottom of the asset page to see lists of all of the resources associated with the VM. Note any VMDKs running at high capacity. This is likely causing the contention.
8. To resolve the balancing problem, identify a resource that is under-utilized to receive the load from an over-utilized resource or remove a less demanding application from the heavily used resource.

Examining network performance

You can examine your storage environment performance and identify under-utilized and over-utilized resources and identify risks before they turn into problems.

Insight helps you to resolve or prevent performance and availability problems that are revealed through the collected storage data.

You can use Insight to perform these performance management tasks:

- Monitor performance across your environment
- Identify resources influencing the performance of other devices

The Importance of Ports

The Insight Server and Data Warehouse (DWH) server may require a number of TCP ports to be free in order to operate reliably. Some of these ports are only utilized for processes bound to the localhost adapter (127.0.0.1), but are still required for core services to operate reliably. The number of ports required is a superset of what ports are used across the network.

Insight Server Ports

Insight Servers can have software firewalls installed. The "holes" that would need to be opened would be as described below.

Inbound HTTPS 443 - assuming you have the Insight WebUI running on TCP 443, you must expose that as to allow any and all of the following consumers:

- Insight users of the WebUI

- Remote Acquisition Units seeking to connect to the Insight server
- OCI DWH servers with connectors to this Insight server.
- Any programmatic interactions with the Insight REST API

Our general recommendation for anyone looking to implement Insight server host-level firewalling is to allow HTTPS access to all corporate network IP blocks.

Inbound MySQL (TCP 3306). This port only needs to be exposed to any Insight DWH server with a connector

While Insight has dozens of data collectors, they are all poll-based - Insight will cause its Acquisition Units (AUs) to initiate outbound communication to various devices. As long as your host based firewall is "stateful" such that it allows return traffic to be allowed through the firewall, host based firewalls on the Insight Server should not impact data acquisition.

Data Warehouse Ports

For Insight DWH servers:

Inbound HTTPS 443 - assuming you have the Insight WebUI running on TCP 443, you must expose that as to allow the following consumers:

- Insight administrative users of the DWH admin portal

Inbound HTTPS (TCP 9300) - this is the Cognos reporting interface. If you will have users interacting with the Cognos reporting interface, this must be exposed remotely.

We can imagine environments where the DWH may not need to be exposed - perhaps the report authors just make RDP connections to the DWH server, and craft and schedule reports there, while having all reports scheduled to be delivered via SMTP, or written to a remote file system.

Inbound MySQL (TCP 3306). This port only needs to be exposed if your organization has any MySQL-based integrations with DWH data - are you extracting data out of the various DWH data marts for ingestion into other applications like CMDBs, chargeback systems, etc.

Analyzing slow PC performance

If you receive calls from network users complaining that their computers are running slowly, you need to analyze host performance and identify the affected resources.

Before you begin

In this example, the caller gave the host name.

Steps

1. Open Insight in your browser.
2. Enter the host name in the **Search assets** box and click the host name in the search results.

The *asset page* for the resource opens.

3. On the asset page for the host, examine the performance charts in the center of the page. You might want to show different types of data in addition to the Latency and IOPS that are usually pre-selected. Click the check boxes for other types of data, such as Throughput, Memory, CPU, or IP Throughput depending on

the device type.

4. To display a description of a point on a chart, position the mouse pointer over the point.
5. You might also want to change the time range with the selection at the top of the page to be 3 hours up to 7 days or All of the available data.
6. Examine the list of **Top correlated resources** to see if there are other resources with the same pattern of activity as the base resource.

The first resource in the list is always the base resource.

- a. Click a linked percentage beside a correlated resource to see if the correlated activity pattern is for IOPS or CPU for the base resource and another resource.
 - b. Click the check box for a correlated resource to add its data to the performance charts.
 - c. Click the linked name of the correlated resource to display its asset page.
7. For a VM, as in this example, locate the storage pool in the **Top correlated resources** and click the storage pool name.

Analyzing correlated resources

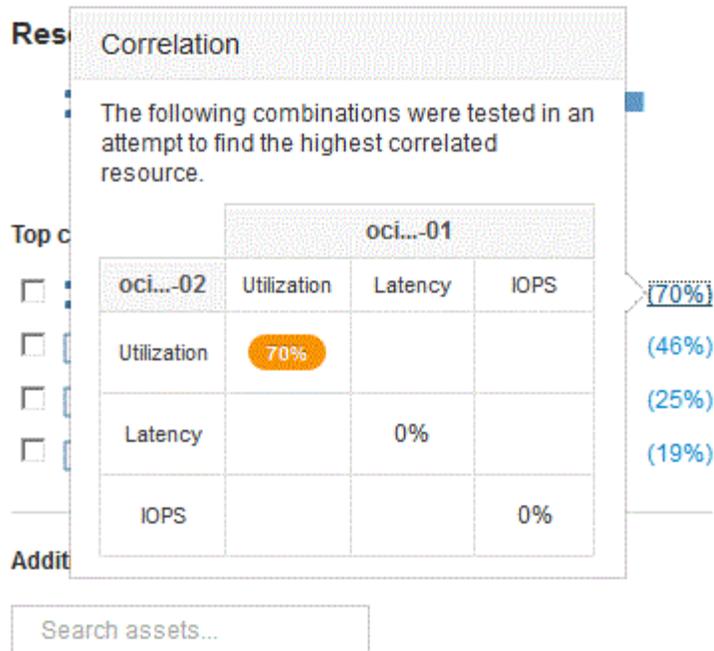
When you are researching performance problems and you open the *asset page* for a device, you should use the Top correlated resources list to refine data displayed in the performance charts. A resource with a high percentage indicates that resource has similar activity to the base resource.

About this task

You are investigating a performance problem and opened the asset page for a device.

Steps

1. In the **Top correlated resources** list, the first resource is the base resource. The correlated resources in the list are ranked by percentage of correlated activity to the first device. Click the linked percentage of correlation to see the details. In this example, the 70% correlation is in Utilization, so both the base resource and this correlated resource have equally high utilization.



- To add a correlated resource to the performance charts, select the check box in the **Top correlated resources** list for the resource you want to add. By default each resource provides the Total data available, but you can select only Read or only Write data from the menu on the check box.

Each resource in the charts has a different color so that you can compare the performance measurements for each resource. Only the appropriate type of data is plotted for the selected measurement metrics. For example, CPU data does not include Read or Write metrics, so only Total data is available.

- Click the linked name of the correlated resource to display its asset page.
- If you do not see a resource listed in the Top correlated resources that you believe should be considered in the analysis, you can use the **Search assets** box to find that resource.

Fibre Channel environment monitoring

Using OnCommand Insight's Fibre Channel asset pages, you can monitor the performance and inventory of the fabrics in your environment and be aware of any changes that might cause issues.

Fibre Channel asset pages

Insight's asset pages present summary information about the resource, its topology (the device and its connections), performance charts, and tables of associated resources. You can use the fabric, switch, and port asset pages to monitor your Fibre Channel environment. Particularly helpful when troubleshooting a Fibre Channel issue is the performance chart for each port asset, which shows the traffic for the selected top contributor port. Additionally, you can also show buffer-to-buffer credit metrics and port errors in this chart, with Insight displaying a separate performance chart for each metric.

Performance policies for port metrics

Insight enables you to create performance policies to monitor your network for various thresholds and to raise alerts when those thresholds are crossed. You can create performance policies for ports based on available port metrics. When a violation of a threshold occurs, Insight detects and reports it in the associated asset page

by displaying a red solid circle; by email alert, if configured; and in the Violations Dashboard or any custom dashboard that reports violations.

Time-to-live (TTL) and downsampled data

Starting with OnCommand Insight 7.3, data retention or time-to-live (TTL) has been increased to from 7 to 90 days. Because that means much more data is processed for charts and tables and the potential for tens of thousands of datapoints, data is downsampled before being displayed.

Downsampling provides a statistical approximation of your data in charts, giving you an efficient overview of data without having to display every data point, while maintaining an accurate view of your collected data.

Why is downsampling needed?

Insight 7.3 increases the time-to-live (TTL) for data to 90 days. This means an increase in the amount of processing needed to prepare data for display in charts and graphs. In order to allow charts to display quickly and efficiently, data is downsampled in a manner that keeps the overall shape of a chart without needing to process every single data point for that chart.



No actual data is lost during downsampling. You can choose to view actual data for your chart instead of downsampled data by following the steps illustrated below.

How downsampling works

Data is downsampled under the following conditions:

- When your selected time range includes 7 days of data or less, no downsampling occurs. Charts display actual data.
- When your selected time range includes more than 7 days of data but less than 1,000 data points, no downsampling occurs. Charts display actual data.
- When your selected time range includes more than 7 days of data and more than 1,000 data points, data is downsampled. Charts display approximated data.

The following examples show downsampling in action. The first illustration shows latency and IOPS charts on a Datastore asset page for a 24-hour period, as shown by selecting **24h** on the asset page's time selector. You can also see the same data by selecting **Custom** and setting the time range to the same 24-hour period.

Since we are choosing a time range of less than 7 days and we have less than 1,000 data points to chart, the data displayed is actual data. No downsampling occurs.



However, if you are viewing data by choosing either **30d** on the asset page time selector, or by setting a custom time range of more than 7 days (or in the event that Insight has collected more than 1,000 data samples for the time period chosen), the data is downsampled before being displayed. When you zoom in on a downsampled chart, the display continues to show the approximated data.



When you zoom in on a downsampled chart, the zoom is a digital zoom. The display continues to show the approximated data.

You can see this in the following illustration, where the time range is first set to 30d, and the chart is then zoomed in to show the same 24-hour period as above.



The downsampled charts are showing the same 24-hour period as the "actual" charts above, so the lines follow the same general shape, allowing you to quickly spot interesting peaks or valleys in your performance data.



Due to the way data is approximated for downsampling, chart lines may be off slightly when comparing downsampled vs. actual data, to allow for better alignment in the graphs. However, the difference is minimal and does not affect the overall accuracy of the data displayed.

Violations on downsampled charts

When viewing downsampled charts, be aware that violations are not shown. To see violations, you can do one of two things:

- View the actual data for that time range by selecting Custom in the asset page time selector, and entering a range of time less than 7 days. Hover over each red dot. The tooltip will show the violation that occurred.
- Note the time range and find the violation(s) in the Violation Dashboard.

Pruning of inventory history

Starting with version 7.3.2, Insight keeps inventory (foundation) change history for 90 days. Previous versions of Insight kept all inventory change history from the time of installation. Following an upgrade from an older version of Insight, older inventory history is pruned down to and then kept at 90 days.

After upgrading to the current version of OnCommand Insight, history is pruned to the most recent 90 days. Insight prunes the history in 30-day chunks occurring once a day, starting with the oldest, until 90 days' worth of history remains. Then, history is pruned daily, to keep only 90 days' worth of inventory change history.

NAS path for VMs

OnCommand Insight 7.3 supports NAS paths for Virtual Machines to storage shares. These paths are similar to NAS paths for hosts to storage shares. When a VM's IP address is allowed to access a share, a NAS path is created.

NAS paths for virtual machines are displayed on the Internal Volumes landing page. This page contains a Guest Mounted Storage Resources widget which identifies the Internal Volumes that VMs have access to.

- NAS paths are created when virtual machines have access to the backend shares. There is no acknowledgment of whether the virtual machines access the shares or not.
- Correlation calculation is based on latencies and IOPs, and do not include cases where VMs have NAS paths to the backend storage.
- User can query the share by initiator IP address, but querying by path is not supported.

The Compute Resources table of the Internal Volume now also displays VM's with NAS paths. For each VM, CPU and memory, utilization and performance data is provided.

Data warehouse impact

Changes to the Data Warehouse that are present after upgrading to OnCommand Insight 7.3 include the following:

- The `dwh_inventory.nas_logical` table is removed from the Inventory data mart and replaced with a view.

Any Insight 7.2.x reports containing the NFS path table are preserved.

- The `dwh_inventory.nas_cr_logical` table is added to the Inventory data mart and includes the following:
 - Compute resource
 - Internal volume
 - Storage
 - NAS share

Capacity as Time Series

With OnCommand Insight 7.3.1, capacity information is reported and charted as time series data.

Previously, capacity information acquired from data sources has been exclusively "point-in-time" (PIT) data, meaning it could not be used in charts as time series data. Now, capacity values for assets can be used as time series data in the following ways:

- Graphed in tables, widgets, expert views, and any place where time series data is displayed
- Applied to performance thresholds with violations using existing semantics
- Used in expressions with other performance counters where appropriate

Note that if you upgrade from a previous version of Insight, previous PIT capacity values used in queries or in filters for custom dashboards will be replaced with time series capacity data. This may result in small changes in the way that capacity data is reported or filtered when compared to the equivalent data in previous Insight versions.

Data Warehouse administration

Welcome to OnCommand Insight Data Warehouse

The OnCommand Insight Data Warehouse is a centralized repository that stores data from multiple OnCommand Insight servers and transforms data into a common, multidimensional data model for querying and analysis.

The OnCommand Insight Data Warehouse enables access to an open database consisting of several data marts that let you generate custom capacity and performance reports such as chargeback reports, trending reports with historical data, consumption analyses, and forecasting reports.

Data Warehouse features

The OnCommand Insight Data Warehouse is an independent database made up of several data marts.

Data Warehouse includes the following features:

- Current and historical configuration and inventory data that enables you to create trending reports useful for forecasting and planning
- Several multidimensional historical data marts and an additional current-only inventory data mart
- An optimized database for predefined queries or user-defined queries
- A platform for integration with third-party reporting and business intelligence engines, including:
 - Configuration management databases
 - Financial accounting systems
 - Asset management systems

Data Warehouse components

Data Warehouse contains several components.

- Data Warehouse Portal
- OnCommand Insight Reporting Portal
- Report authoring tools

What you can do using the Data Warehouse Portal

The Data Warehouse Portal is a web-based user interface that you use to configure options and set up fixed schedules to retrieve data. From the Data Warehouse Portal, you can also access the OnCommand Insight reporting portal.

Using Data Warehouse portal, you can do the following:

- Access the OnCommand Insight reporting portal to view predesigned reports or to create custom reports using report authoring tools.

- Consolidate multiple OnCommand Insight databases.
- Manage connections to OnCommand Insight servers.
- Check the status of current jobs or queries that are running.
- Schedule Data Warehouse builds.
- Edit the site name.
- View Data Warehouse version and upgrade history, including specific information such as module versions, sites, and licenses.
- Import annotations.
- Configure a build from history.
- View Data Warehouse documentation and the database schema.
- Reset the Data Warehouse database.
- Back up and restore the Data Warehouse database.
- Troubleshoot Data Warehouse issues.
- Manage user accounts.

Data Warehouse software components

OnCommand Insight Data Warehouse includes several software components.

- MySQL database

The back-end repository for data mart tables

- IBM Cognos

The reporting engine for OnCommand Insight

- Apache Derby Database

Used for storing Cognos configuration and content

- WildFly

The Java Enterprise application server that hosts OnCommand Insight components

Data Warehouse processes

Data Warehouse performs many types of processes.

- **ETL process**

The Extract Transform and Load (ETL) process retrieves data from multiple OnCommand Insight databases, transforms the data, and saves it into the data mart. The Data Warehouse build process is an ETL process.

- **Jobs**

Data Warehouse performs and reports on jobs such as these: inventory, dimensions, capacity, port

capacity, VM capacity, file system utilization, performance, capacity efficiency, licenses, history build, dynamic annotations, connector removal, skipped build, ASUP option, and maintenance jobs.

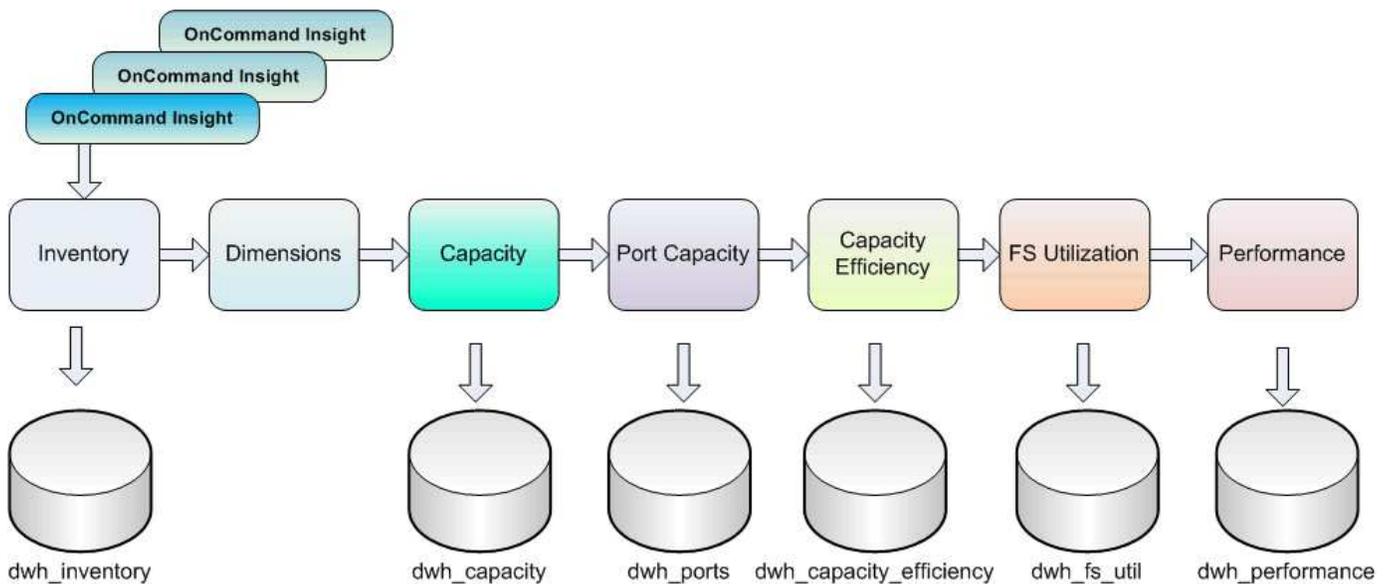
- **Consolidation process**

Data Warehouse supports the consolidation of multiple OnCommand Insight servers into the same Data Warehouse database. In many configurations it might happen that the same object is reported from multiple connectors (that is, the same switch exists in two OnCommand Insight instances). In that case, Data Warehouse consolidates the multiple objects into one (a primary connector is chosen and the object's data is taken from that connector only).

How Data Warehouse extracts data

The Extract, Transform, and Load (ETL) process retrieves data from multiple OnCommand Insight databases, transforms the data, and saves it into the data marts.

OnCommand Insight connectors invoke a series of batch jobs to extract data from multiple OnCommand Insight mySQL databases and publish the data in various data marts, as shown in the following diagram.



The ETL process includes these individual processes:

- **Extract**

This process takes data from multiple OnCommand Insight databases, transforms the data, and saves it into the data mart. The process is performed against each OnCommand Insight instance at the same time. To ensure that data cleansing and deduplication is performed, it is not possible to split the ETL process into multiple scheduled ETL operations.

- **Transform**

This process applies business logic rules or functions to extract the data from the OnCommand Insight database.

- **Load**

This process loads the transformed data into public data marts.

ETL frequency and date data

You should run the Extract, Transform, and Load (ETL) process at least once per day; however, you choose to run ETL numerous times if needed.

By default, the Cognos reporting engine treats all capacity and performance facts as additive. As a result, there is a risk of double counting capacity data if the ETL process is run multiple times per day without the proper time filters.

Two date data elements in the Date dimension are related to the daily ETL process. The Date dimension, which is used in several data models, includes the following data elements that are affected by the ETL:

- **Is Day Representative**

The "Is Day Representative" data element is set to a value of 1 (true) during the first ETL process run during any given day. If the first ETL process is run at 1:00 a.m., Is Day Representative is set to 1 for all of the data loaded during the 1:00 a.m. ETL process. If a second ETL is scheduled later (for example, 1:00 p.m.), Is Day Representative is set to 0 (false) for the data loaded during that ETL process.

- **Is Latest**

The "Is Latest" member is set to a value of 1 (true) after each ETL process completes. If the first ETL process is run at 1:00 a.m., Is Latest is set to 1 for all of the data loaded during the 1:00 a.m. ETL process. If another ETL process is scheduled later (for example, 1:00 p.m.), Is Latest is set to 1 for data loaded during the 1 p.m. ETL process. The ETL process also sets the 1:00 a.m. ETL load's Is Latest entry to 0 (false).

How historical data is retained in Data Warehouse

Data is maintained in Data Warehouse according to a schedule. As data gets older, the data record retention is reduced.

Data Warehouse retains historical data based on the data marts and granularity of the data, as shown in the following table.

Data mart	Measured object	Granularity	Retention period
Performance marts	Volumes and internal volumes	Hourly	14 days
Performance marts	Volumes and internal volumes	Daily	13 months
Performance marts	Application	Hourly	13 months
Performance marts	Host	Hourly	13 months
Performance marts	Switch performance for port	Hourly	5 weeks

Performance marts	Switch performance for host, storage, and tape	Hourly	13 months
Performance marts	Storage node	Hourly	14 days
Performance marts	Storage node	Daily	13 months
Performance marts	VM performance	Hourly	14 days
Performance marts	VM performance	Daily	13 months
Performance marts	Hypervisor performance	Hourly	14 days
Performance marts	Hypervisor performance	Daily	13 months
Performance marts	VMDK performance	Hourly	14 days
Performance marts	VMDK performance	Daily	13 months
Performance marts	Disk performance	Hourly	14 days
Performance marts	Disk performance	Daily	13 months
Capacity marts	All (except individual volumes)	Daily	13 months
Capacity marts	All (except individual volumes)	Monthly representative	14 months and beyond
Inventory marts	Individual volumes	Current state	1 day (or until next ETL)

After 13 months (which is configurable), Data Warehouse retains only one record per month instead of one record per day for capacity, performance, and resource data in the following fact tables:

- Chargeback fact table (dwh_capacity.chargeback_fact)
- File System Utilization fact table (dwh_fs_util.fs_util_fact)
- Host fact table (dwh_sa.sa_host_fact)
- Internal Volume Capacity fact table (dwh_capacity.internal_volume_capacity_fact)
- Ports fact table (dwh_ports.ports_fact)
- Qtree Capacity fact table (dwh_capacity.qtree_capacity_fact)
- Storage and Storage Pool Capacity fact table (dwh_capacity.storage_and_storage_pool_capacity_fact)
- Volume Capacity fact table (dwh_capacity.vm_capacity_fact)
- Storage Node Hourly Performance (storage_node_hourly_performance_fact) and Storage Node Daily Performance (storage_node_daily_performance_fact) fact tables

Data retention, ETL, and time periods

OnCommand Insight Data Warehouse retains data obtained from the Extract, Transform, and Load (ETL) process for different time periods based on the different data marts and time granularity of the data.

Performance Marts and hourly granularity for volumes and internal volumes

The OnCommand Insight Data Warehouse records the hourly averages, hourly maximums, and access bit for each hour of the day (24 data points) for 14 days. The access bit is a Boolean value that is true if the volume is accessed or false if the volume is not accessed during the hourly interval. All 24 data points for the preceding day are obtained during the first ETL process of the day.

You do not need to run one ETL process per hour to gather the hourly data points. Running additional ETL processes during the day does not obtain any performance information from the OnCommand Insight Servers.

Performance Marts and daily granularity for volumes and internal volumes

Each day when the ETL is processed, the daily averages for the preceding day are calculated and populated within Data Warehouse. The daily average is a summary of the 24 data points for the previous day. The performance data marts retain daily summaries for volumes and internal volumes for 13 months.

Capacity marts and daily granularity

The Capacity marts provide daily measurements for various capacity facts on a daily basis for a period of 13 months. The capacity facts in Data Warehouse are current as of the last data source acquisition for the device prior to the ETL.

Capacity marts and monthly granularity

Data Warehouse retains daily capacity data for 13 months. After the 13-month threshold is reached, the capacity data is summarized on a monthly basis. The monthly data is based on the values reflected by the date that is the month representative date.

The following table shows which monthly data is included in the monthly summary:

Date	Is Month Representative value	Allocated capacity
Jan 1	1 (True)	50 TB
Jan 2	0 (False)	52 TB
...
Jan 31	0 (False)	65 TB
Feb 1	1 (True)	65 TB

Based on the table, a monthly report would show 50 TB allocated for January and 65 TB allocated for February. All of the other capacity values for January would not be included in the monthly summary.

Inventory mart

The Inventory data mart is not historical. Each time an ETL process is run, the Inventory mart is erased and rebuilt. Therefore, any reports generated out of the Inventory mart do not reflect historical inventory configuration.

Getting started with Data Warehouse

OnCommand Insight Data Warehouse enables you to configure options needed before generating reports that include your data. Data Warehouse contains many features; however, you need to use only a few of them to get started. To set up Data Warehouse, you use options in the Data Warehouse Portal.

About this task

To set up OnCommand Insight Data Warehouse, a storage administrator should complete the following procedures:

- Accessing the Data Warehouse portal
- Connecting Data Warehouse to OnCommand Insight servers
- Building the database from history
- Setting up backup and restore processes

Additionally, a storage administrator might want to complete the following procedures.

- Accessing MySQL using the command line interface
- Scheduling daily builds
- Setting up multiple tenancy in reporting
- Troubleshooting setup issues
 - Why can't I see my annotations?
 - What should I do with failing historical build points?

If this is the first time you are using the Data Warehouse Portal, you must set up Data Warehouse before any information can appear on the Jobs page. You also need to repeat this setup process after resetting the Data Warehouse database.

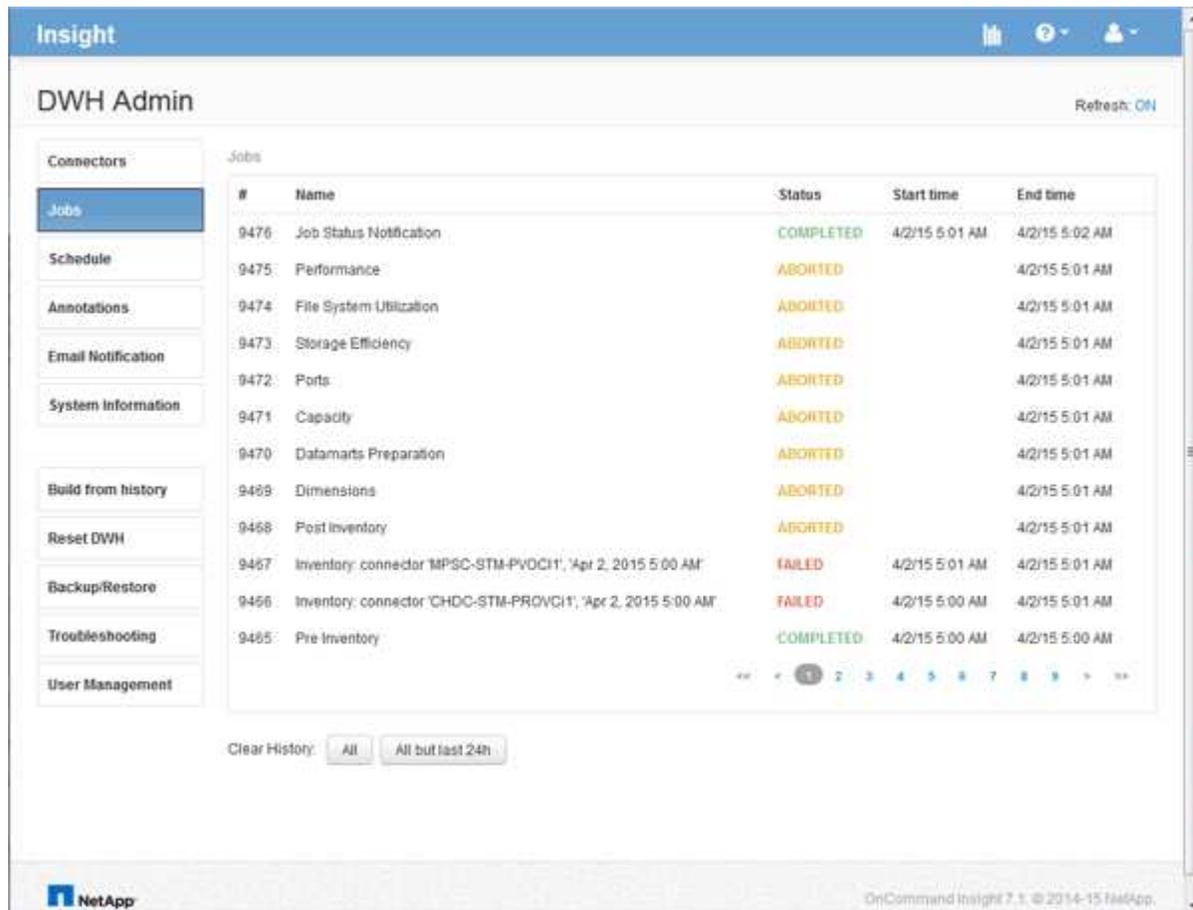
Accessing the Data Warehouse portal

The OnCommand Insight Data Warehouse portal is a web-based user interface that you can use to update connector information, view job queues, schedule daily builds, select annotations, set up email notifications, view system information, build the database, reset Data Warehouse, back up and restore the database, troubleshoot issues, manage Data Warehouse and Reporting portal user accounts, and access documentation and schema diagrams.

Steps

1. Log in to the Data Warehouse portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. Enter your user name and password.
3. Click **Login**.

The Data Warehouse portal opens:



Managing Data Warehouse and Reporting user accounts

User accounts, user authentication, and user authorization for the OnCommand Insight reporting tools are defined and managed from the Data Warehouse (DWH). Based on these configurations, users and administrators gain access to some or all of the available OnCommand Insight reports.

Access to the User Management in the Data Warehouse requires an account with System Administrator privileges. This includes:

- Full administrative capabilities for the Data Warehouse
- Configuration and maintenance of all user accounts
- Read access to the database
- Capability to set up connectors in the ETL, schedule Data Warehouse jobs, reset the database, assign or

change roles, and add and remove user accounts

Accessing the Data Warehouse and Reporting portal

The Data Warehouse portal provides access to administration options. From the Data Warehouse portal, you can also access the Reporting portal.

Steps

1. Log in as an administrator to the Data Warehouse portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. On the Data Warehouse toolbar, click  to open the Reporting portal.

Reporting user roles

Each user account is assigned a role with a set of permissions. The number of users is limited by the number of Reporting licenses attached to each role.

Each role can perform the following actions:

- **Recipient**

Views OnCommand Insight Reporting portal reports and sets personal preferences such as those for languages and time zones.



Recipients cannot create reports, run reports, schedule reports, export reports, nor perform administrative tasks.

- **Business Consumer**

Runs reports and performs all Recipient options.

- **Business Author**

Views scheduled reports, runs reports interactively, creates stories, in addition to performing all Business Consumer options.

- **Pro Author**

Creates reports, creates packages and data modules, in addition to performing all Business Author options.

- **Administrator**

Performs reporting administrative tasks such as the import and export of report definitions, configuration of reports, configuration of data sources, and the shutdown and restart of reporting tasks.

The following table shows the privileges and the maximum number of users allowed for each role:

Feature	Recipient	Business Consumer	Business Author	Pro Author	Admin

View reports in the Team Content tab	Yes	Yes	Yes	Yes	Yes
Run reports	No	Yes	Yes	Yes	Yes
Schedule reports	No	Yes	Yes	Yes	Yes
Upload external files	No	No	Yes	Yes	No
Create stories	No	No	Yes	Yes	No
Create reports	No	No	Yes	Yes	No
Create Packages and Data Modules	No	No	No	Yes	No
Perform administrative tasks	No	No	No	No	Yes
Number of users	Number of OnCommand Insight users	20	2	1	1

When you add a new Data Warehouse and Reporting user, if you exceed the limit in a role, the user is added as “deactivated,” and you need to deactivate or remove another user with that role to give a new user membership.



Report authoring capabilities require Insight Plan license. You can add additional Business Author and Pro Author users by purchasing the ARAP (Additional Report Authoring Package). Contact your OnCommand Insight representative for assistance.

These reporting user roles do not affect direct database access. These reporting user roles do not impact your ability to create SQL queries using the data marts.

Adding a Reporting user

You must add a new user account for each person who requires access to the Reporting portal. Having a different user account for each person provides a way of controlling access rights, individual preferences, and accountability.

Before you begin

Before adding a Reporting user, you must have allocated a unique user name, determined what password to use, and verified the correct user role or roles. These roles are specialized in the Reporting portal.

Steps

1. Log in as an administrator to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **User Management**.
3. In the **User Management** window, click **Add New User**.
4. Enter the following information for the new Reporting user:

- **User name**

User name (alphanumeric, including a-z, A-Z, and 0-9) for the account

- **E-mail Address**

Email address associated with the user account and required if the user subscribes to any reports

- **Password**

Password to log in to OnCommand Insight with this user account, which is typically selected by the user and confirmed in the interface

- **Insight role**

Roles available to the user with appropriate permissions



The options for the OnCommand Insight role are shown only if OnCommand Insight is installed on the same machine as the reporting facilities, which is not typical.

- **Reporting roles**

Reporting role for this user account (for example, Pro Author)



The Administrator role is unique. You can add this role to any user.

5. Click **Add**.

Managing user accounts

You can configure user accounts, user authentication, and user authorization from the Data Warehouse portal. Each user account is assigned a role with one of the following permission levels. The number of users is limited by the number of Reporting licenses attached to each role.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **User Management**.

User Management

Name	OnCommand Insight roles			Reporting roles					E-mail				
	Guest	User	Administrator	Recipient	Business Consumer	Business Author	Pro Author	Administrator					
guest	X									Edit	Delete	Change password	Deactivate
user	X	X								Edit	Delete	Change password	Deactivate
admin	X	X	X				X	X		Edit		Change password	
oadmin	X	X	X							Edit		Change password	Deactivate

LDAP Configuration Add New User Change DWH User password

The following table shows the privileges for each reporting role:

Feature	Recipient	Business Consumer	Business Author	Pro Author	Administrator
View reports (in Public Folder tab, My Folders)	Yes	Yes	Yes	Yes	Yes
Run reports	No	Yes	Yes	Yes	Yes
Schedule Reports	No	Yes	Yes	Yes	Yes
Create reports in Query Studio	No	No	Yes	Yes	No
Create reports in Workspace (Standard)	No	Yes	Yes	Yes	No
Create reports in Workspace (Advanced)	No	No	Yes	Yes	No
Create reports in Report Studio	No	No	No	Yes	No
Perform administrative tasks	No	No	No	No	Yes

3. Do one of the following:
 - To edit an existing user, select the row for the user and click **Edit**.
 - To change a user’s password, select the row for the user and click **Change password**.
 - To delete a user, select the row for the user and click **Delete**
4. To activate or deactivate a user, select the row for the user and click **Activate** or **Deactivate**.

Configuring LDAP for Reporting

From the Data Warehouse portal, the Administrator can configure LDAP usage for Data Warehouse and Reporting.

Before you begin

You must log in to Insight as an Administrator to perform this task.

For all Secure Active Directory (i.e. LDAPS) users, you must use the AD server name exactly as it is defined in the certificate. You can not use IP address for secure AD login.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system on which OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **User Management**.
3. Click **LDAP Configuration**.
4. Select **Enable LDAP** to start the LDAP user authentication and authorization process.
5. Make whatever changes are necessary to configure LDAP.

The majority of the fields contain default values. The default settings are valid for the Active Directory.

- **User principal name attribute**

Attribute that identifies each user in the LDAP server. Default is `userPrincipalName`, which is globally unique. OnCommand Insight attempts to match the contents of this attribute with the username

that has been supplied above.

- **Role attribute**

LDAP attribute that identifies the user's fit within the specified group. Default is `memberOf`.

- **Mail attribute**

LDAP attribute that identifies the user's email address. Default is `mail`. This is useful if you want to subscribe to reports available from OnCommand Insight. Insight picks up the user's email address the first time each user logs in and does not look for it after that.



If the user's email address changes on the LDAP server, be sure to update it in Insight.

- **Distinguished name attribute**

LDAP attribute that identifies the user's distinguished name. default is `distinguishedName`.

- **Referral**

Indicates whether to follow the path to other domains if there are multiple domains in the enterprise. You must always use the default `follow` setting.

- **Timeout**

Length of time to wait for a response from the LDAP server before timing out, in milliseconds. default is 2,000, which is adequate in all cases and should not be modified.

- **LDAP servers**

This is the IP address or DNS name to identify the LDAP server. To identify a specific port, where `ldap-server-address` is the name of the LDAP server, you can use the following format:

```
ldap://ldap-server-address:port
```

To use the default port, you can use the following format:

```
ldap://ldap-server-address
```



When entering multiple LDAP servers in this field, separate entries with a comma, and ensure that the correct port number is used in each entry.

+

To import the LDAP certificates, click **Import Certificates** and automatically import or manually locate the certificate files.

- **Domain**

LDAP node where OnCommand Insight should start looking for the LDAP user. Typically this is the top-level domain for the organization. For example:

```
DC=<enterprise>,DC=com
```

- **Insight server admins group**

LDAP group for users with Insight Server Administrator privileges. Default is `insight.server.admins`.

- **Insight administrators group**

LDAP group for users with Insight Administrator privileges. Default is `insight.admins`.

- **Insight users group**

LDAP group for users with Insight User privileges. Default is `insight.users`.

- **Insight guests group**

LDAP group for users with Insight Guest privileges. Default is `insight.guests`.

- **Reporting administrators group**

LDAP group for users with Insight Reporting administrator privileges. Default is `insight.report.admins`.

- **Reporting pro authors group**

LDAP group for users with Insight Reporting pro authors privileges. Default is `insight.report.proauthors`.

- **Reporting business authors group**

LDAP group for users with Insight Reporting business authors privileges. Default is `insight.report.business.authors`.

- **Reporting business consumers group**

LDAP group for users with Insight Reporting business consumers privileges. Default is `insight.report.business.consumers`.

- **Reporting recipients group**

LDAP group for users with Insight Reporting recipient privileges. Default is `insight.report.recipients`.

6. Enter values in the **Directory lookup user** and **Directory lookup user password** fields if you made any changes.

If you do not enter the revised values in these fields, your changes are not saved.

7. Retype the directory lookup user password in the **Confirm directory lookup user password** field, and

click **Validate Password** to validate the password on the server.

8. Click **Update** to save the changes. Click **Cancel** to remove changes.

Connecting Data Warehouse to OnCommand Insight servers

Connectors establish connections from the OnCommand Insight Data Warehouse to the OnCommand Insight servers. You can connect Data Warehouse with one or more OnCommand Insight servers. You can add or remove connections to or from OnCommand Insight databases.

About this task

Data Warehouse assigns a global unique ID to the connector that is used in conjunction with the connector name. After adding a connector, Data Warehouse queries the OnCommand Insight database for the OnCommand Insight site name and version.

You can choose to connect to a data source with or without SSL. Choosing the secure data source forces the connection to use SSL when communicating with the OnCommand Insight remote database.

Data Warehouse can provide a consolidated view of data from multiple OnCommand Insight installations. This consolidated database provides the following information:

- Globally Unique Identifiers

Each object is assigned a globally unique ID that is independent of the IDs used by individual sites, to avoid conflicting IDs and enable duplicate detection. These IDs are shared between all the data marts. This ID is the Globally Unique ID (GUID) in the Comment column of the Inventory data mart tables.

- No duplication

Entities that exist in multiple OnCommand Insight databases are registered only once in the consolidated database.

- Current record

The data in the consolidated database (Inventory data mart) is always the most up-to-date possible.

When you add or edit a connection, you can also test the connection. The test does the following:

- Verifies the host IP address, user name, and password and ensures that a connection can be established.

Invalid connections appear in red.

- Compares the OnCommand Insight version to the Data Warehouse version.

If the versions are not compatible, an error message appears.

- Verifies that the OnCommand Insight database has not been changed or restored to a different database as seen by the last Data Warehouse processing. If there has been a change, an error message appears.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the

system where OnCommand Insight Data Warehouse is installed.

2. From the navigation pane on the left, click **Connectors**.

The Connectors table appears blank at first and shows connector information after you add a connector.

3. Click **New** to add a new connector.

4. Enter the following:

- **Encryption**

To enable Data Warehouse requests to be made using SSL encryption, select `Enabled`.

- **Name**

A connector name that will identify the connector on the Connectors view.

- **Host**

Host IP address

- **User name**

"inventory"



Using this user name and password, you can log in to the remote OnCommand Insight database and perform queries on the data.

- **Password**

"sanscreen"

5. To specify the port to use for TCP connections to the host, click **Advanced** and enter the TCP port number.
6. To specify the port (other than the default port) to use for HTTPS connections to the host, click **Advanced** and enter the port number.
7. Click **Test**.

Data Warehouse tests the connection.

8. Click **Save**.

If you enter multiple connections for multiple installations, Data Warehouse invokes independent build processes, one for each database from which data should be extracted. Each such build process extracts data from an OnCommand Insight database and loads it into the consolidated database.

Data Warehouse database build from history overview

You can build the Data Warehouse database using historical data in your OnCommand Insight server. Data Warehouse extracts data from the OnCommand Insight servers and builds the Data Warehouse data marts according to the build from history schedule.

This option does not require a special license and inventory data is included in the build. However, to build capacity information, the OnCommand Insight Plan and OnCommand Insight Perform licenses are required.

If any build (from history or current) has already been performed, the build cannot be done on dates before the last job. This means if you perform a current build, you cannot build from history. More specifically, if you performed builds from history that ended on Jan 1, 2012, you cannot perform any build on the year 2011.

If the history build does not include a day or two of any unsuccessful ETL processes, do not try building history for just these few days. Historical data is for longer periods and a day or two is not going to change trending significantly. If you do want to rebuild from history, rebuild the entire history.

The Build from History view displays all build jobs from all connectors. For example, the view might display an inventory job for every connector, a port capacity job for every build run, and an annotations job.

Before you configure the Build from History, the following must occur:

- Connectors must be configured.
- Annotations should be entered in OnCommand Insight and can be manually updated using the **Force Update of Annotations for DWH option** in the old OnCommand Insight Portal or will be automatically updated 15 minutes after they are set.

Adding a job that builds a Data Warehouse database from history

You can build the Data Warehouse database using historical data that is kept in your OnCommand Insight server, which enables you to run projection reports.

Before you begin

You must have updated annotations in the OnCommand Insight server and forced an update of annotation information for Data Warehouse.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Build from History**.

Build From History

Target time	Start running	Status
3/13/15 12:00 AM	3/25/15 9:28 AM	COMPLETED
3/14/15 12:00 AM	3/25/15 9:34 AM	COMPLETED
3/15/15 12:00 AM	3/25/15 9:39 AM	COMPLETED
3/16/15 12:00 AM	3/25/15 9:45 AM	COMPLETED
3/17/15 12:00 AM	3/25/15 9:51 AM	COMPLETED
3/18/15 12:00 AM	3/25/15 9:57 AM	COMPLETED
3/19/15 12:00 AM	3/25/15 10:03 AM	COMPLETED
3/20/15 12:00 AM	3/25/15 10:09 AM	COMPLETED
3/21/15 12:00 AM	3/25/15 10:16 AM	COMPLETED
3/22/15 12:00 AM	3/25/15 10:23 AM	COMPLETED
3/23/15 12:00 AM	3/25/15 10:30 AM	COMPLETED
3/24/15 12:00 AM	3/25/15 10:38 AM	COMPLETED
3/25/15 12:00 AM	3/25/15 10:44 AM	COMPLETED

« < 1 2 3 > »

Cancel Pending Jobs

Configure

Run

Skip history build failures:

3. Click **Configure**.

Configure Build From History

Start time:	11	February		
	2015	...		
End time:	02	April		
	2015	...		
Interval:	<input checked="" type="radio"/> Daily	<input type="radio"/> Weekly	<input type="radio"/> Monthly	<input type="radio"/> Quarterly
Hour:	12:00 AM	▼		
<input type="button" value="Save"/> <input type="button" value="Reset"/> <input type="button" value="Cancel"/>				

4. Enter the start and end times.

To display a calendar from which you can select these dates, click the down arrow near the month name.

The time format depends upon the locale of the Data Warehouse server.

The start and end times must be within the range of history contained in all the OnCommand Insight servers to which Data Warehouse is connected, as set in the Data Warehouse portal Connectors option. The default start and end times reflect the maximum valid period. The Data Warehouse build job runs automatically at the time you specify.



Configuring a non-realistic schedule such as “Daily for 4 years” results in 1460 build cycles, which could take 10 days to complete.

5. Choose the interval.

If you select a monthly or weekly interval, the Day field appears. If you selected monthly, then Day is a date. If you selected weekly, Day is Sunday through Saturday.

6. Choose the hour when the build will take place.

7. Optionally, to return the options to default settings, click **Reset**.

8. Click **Save**.

9. From the **Build from History** page, to perform a build outside of the automatic schedule build, click **Run**.

The Target Time column displays the time that this entry was built. The Status column displays whether the build was completed or failed.

Canceling a build from history job

You can cancel all planned jobs. The job status becomes “Aborted”.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Build from History**.
3. Click **Cancel**.

Backing up the Data Warehouse database

You can back up the Data Warehouse database, which also includes a Cognos backup, to a file and later restore it using the Data Warehouse portal. Such a backup enables you to migrate to a different Data Warehouse server or upgrade to a new Data Warehouse version.

Steps

1. Log in to the Data Warehouse Portal at `https://fqdn/dwh`.
2. From the navigation pane on the left, select **Backup/Restore**.
3. Click **Backup** and select your backup configuration:
 - a. All Datamarts except Performance Datamart
 - b. All Datamarts

This operation can take 30 minutes or more.

+

Data Warehouse creates a backup file and displays its name.

4. Right-click the backup file and save it to a location you want.

You might not want to change the file name; however, you should store the file outside the Data Warehouse installation path.

The Data Warehouse backup file includes the DWH instance's MySQL; custom schemas (MySQL DBs) and tables; LDAP configuration; the data sources that connect Cognos to the MySQL database (not the data sources that connect the Insight server to devices to acquire data); import and export tasks that imported or exported reports; reporting security roles, groups, and namespaces; user accounts; any modified Reporting Portal reports; and any custom reports, regardless of where they are stored, even in the My Folders directory. Cognos system configuration parameters, such as SMTP server setting, and Cognos custom memory settings are not backed up.

The default schemas where custom tables are backed up include the following:

dwh_capacity
dwh_capacity_staging
dwh_dimensions
dwh_fs_util
dwh_inventory
dwh_inventory_staging
dwh_inventory_transient
dwh_management
dwh_performance
dwh_performance_staging
dwh_ports
dwh_reports
dwh_sa_staging

Schemas where custom tables are excluded from backup include the following:

information_schema
acquisition

cloud_model
host_data
innodb
inventory
inventory_private
inventory_time
logs
management
mysql
nas
performance
performance_schema
performance_views
sansscreen
scrub
serviceassurance
test
tmp
workbench

In any backup initiated manually, a `.zip` file is created that contains these files:

- A daily backup `.zip` file, which contains Cognos report definitions
- A reports backup `.zip` file, which contains all the reports in Cognos, including those in the My Folders directory

- A Data Warehouse database backup file

In addition to manual backups, which you can perform at any time, Cognos creates a daily backup (automatically generated each day to a file called `DailyBackup.zip`) that includes the report definitions. The daily backup includes the top folders and packages shipped with the product. The My Folders directory and any directories that you create outside the product's top folders are not included in the Cognos backup.



Due to the way Insight names the files in the `.zip` file, some unzip programs show that the file is empty when opened. As long as the `.zip` file has a size greater than 0 and does not end with a `.bad` extension, the `.zip` file is valid. You can open the file with another unzip program like 7-Zip or WinZip®.

Backing up custom reports and report artifacts

If you created custom reports in a version of Insight earlier than 7.0, and you want to upgrade to the newest Insight version, you should back up your reports and report artifacts before the upgrade installation and restore them after the upgrade installation. You should also pay attention to the folders that you are using to store report artifacts.

About this task

If you made changes to the predesigned reports, create your own copies of those reports in a separate folder. That way, when you update the predesigned artifacts, you do not overwrite your changes.

If you have reports in the My Folders area, you should copy them to the Custom Reports folders so that they are not lost.

Restoring the Data Warehouse database

You can restore a Data Warehouse database by using the `.zip` file that was created when you backed up that Data Warehouse database.

About this task

When you restore a Data Warehouse database, you have the option to restore user account information from the backup as well. User management tables are used by the Data Warehouse report engine in a Data Warehouse only installation.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system on which OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Backup/Restore**.
3. In the **Restore Database and Reports** section, click **Browse**, and locate the `.zip` file that contains the Data Warehouse backup.
4. If you want to restore reports or user account data, select one or both of the following check boxes:

- **Restore database**

Includes Data Warehouse settings, data marts, connections, and user account information.

- **Restore reports**

Includes custom reports, predesigned reports, changes that you made to predesigned reports, and reporting settings that you created in the Reporting Portal.



If your database backup contains a custom report that has a slash (/) or an open bracket ([]) in its name (for example, US IT Center Switch Port Boston/July), the restore operation renames the report, replacing the slash or open bracket with an underscore (for example, US IT Center Switch Port Boston_July).

5. Click **Restore**.

After the restore process is completed, a message is displayed below the Restore button. If the restore process is successful, the message indicates success. If the restore process fails, the message reports the specific exception that caused the failure. If an exception occurs and the restore process fails, the original database is automatically reset.

Setting up multiple tenancy in reporting

OnCommand Insight Data Warehouse accommodates multiple tenancy (often shortened to “multi-tenancy” or “multitenancy”) in reporting by allowing you to associate users with one or more business entities. With this feature, administrators can separate data or reports according to user attributes or user affiliation.

Business entities use a hierarchy for the purposes of capacity chargeback using the following values:

- Tenant: Primarily used by service providers to associate resources with a customer, for example, NetApp.
- Line of Business (LOB): A line of business within a company, for example "Hardware" or "Software."
- Business Unit: A traditional business unit such as "Sales" or "Marketing."
- Project: A project to which you might want to assign capacity chargeback.

The process of configuring multiple tenancy involves the following major steps:

- Configure a Data Warehouse user account.
- Create a group in Reporting Portal.
- Assign users to one or more groups, which represent business entities.
- Assign users to one or more business entities. For example, users associated with "NetApp" obtain access to all business entities that have “NetApp” as a tenant.
- Test that users can see only those reports that they should see.

The following points summarize how users access reporting data:

- A user, not assigned to any group, gets access to all the data.
- A user, assigned to any group, will not be able to get access to records without business entity.

For example, you might have the following departments and need to separate reports for users within these departments.

User	Engineering	Support	Finance	Legal
------	-------------	---------	---------	-------

User1	X	X		
User2			X	X
User3		X		

Configuring user accounts

You must complete several steps to configure user accounts.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **User Management**.
3. Configure each user account.

Assigning users to business entities

You must complete a series of steps to assign users to business entities. Data Warehouse accommodates multiple tenancy (often shortened to “multi-tenancy” or “multitenancy”) in reporting by allowing you to associate users with one or more business entities. This enables administrators to separate data or reports according to user attributes or user affiliation.

Steps

1. Log in to the Data Warehouse Portal as administrator at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. On the Data Warehouse toolbar, click  to open the Reporting Portal.
3. Enter your user name and password and click **Login**.
4. From the Launch menu, select **IBM Cognos Administration**.
5. Click the **Security** tab:
6. In the Directory, select **Cognos**.
7. Create a new subfolder in the Cognos folder called “BEs”, for business entities.
8. Open the BEs folder.
9. Click the **New Group** icon to add groups that correspond to different permission levels.

These permission levels can be either the full name of the business entity (for example, NetApp.N/A) or a prefix (for example, NetApp.N/A.Finance). Either of these formats enables access to all projects within the business entity (NetApp.N/A.Finance).

The New Group wizard displays.

10. Complete the pages of the wizard.
11. Select a business entity and click **More**.

12. Click **Set members**.
13. Click **Add**.
14. Select the SANscreen directory.
15. From the list of users, select each user that you want to include in the Business Entity and add the user to the Selected Entries box.
16. Click **OK**.
17. Repeat the process of adding members to each of the business entity groups.

Troubleshooting setup issues

There are several common issues with annotations, builds, and reports that you may face during setup. You can troubleshoot these issues by following the steps outlined.

Why I cannot see my annotations

If you cannot see annotations in Data Warehouse, you might need to force an update of annotations and then initiate a Data Warehouse build.

Missing annotations affect the way data is imported into Data Warehouse and is displayed in the reports. For example, if the annotation “Tier” is not available, you will not be able to group storage systems by tier in Data Warehouse reports.

Forcing an update of annotations for Data Warehouse

You can initiate an update of annotations from OnCommand Insight to Data Warehouse.

About this task

You can update annotations using one of two options:

- Including deleted objects: This includes data about devices that no longer exist such as hosts, storage arrays, or switches that were removed. This is needed if you want to build Data Warehouse data with historical data points.
- Not including deleted objects: Choose this option if you want to exclude deleted objects.

Steps

1. Log in to the OnCommand Insight Portal as administrator `https://hostname`, where `hostname` is the name of the system where OnCommand Insight is installed.
2. Click on **Admin > Troubleshooting**. At the bottom of the page, click on **Advanced Troubleshooting**.
3. In the **Actions** tab, click **Update DWH Annotations (include deleted)**.

Generating a manual Data Warehouse build

After forcing an annotations update (running transient data) in OnCommand Insight, you need to initiate a Data Warehouse build. You can wait until the next scheduled build or initiate a build now.

Steps

1. Log in as an administrator to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Schedule**.
3. Click **Build now**.

Importing user-defined annotations into Data Warehouse

After forcing an annotation update in OnCommand Insight, you need to select the annotations you want in Data Warehouse and initiate a Data Warehouse build. You can wait until the next scheduled build or initiate a build now.

Steps

1. Log in as an administrator to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Annotations**.

Annotations

Annotation	Column Name	Target Object	Published
Compute_Resource_Group	Compute_Resource_Group	Virtual Machine	
Data_Center	dataCenter	Host	✓
Data_Center	dataCenter	Storage	✓
Data_Center	dataCenter	Switch	✓
Note	Note	Switch	
Switch_Level	switchLevel	Switch	✓
Tier	Tier	Internal Volume	
Tier	Tier	Qtree	
Tier	Tier	Storage	
Tier	Tier	Storage Pool	
Tier	Tier	Volume	

Edit

The list displays a row for every annotation type and a target object to which the annotation can be assigned. A check mark in the Published column indicates that the annotation was already selected for the particular target object and is already available through the Data Warehouse data marts.

3. Click **Edit** to edit how annotations will be imported from OnCommand Insight.

Edit Annotations

Annotation	Column Name	Target Object	Published All / None	Init With Current All / None
Compute_Resource_Group	Compute_Resource_Group	Virtual Machine	<input type="checkbox"/>	<input type="checkbox"/>
Data_Center	dataCenter	Host	<input type="checkbox"/>	<input type="checkbox"/>
Data_Center	dataCenter	Storage	<input type="checkbox"/>	<input type="checkbox"/>
Data_Center	dataCenter	Switch	<input type="checkbox"/>	<input type="checkbox"/>
Note	Note	Switch	<input type="checkbox"/>	<input type="checkbox"/>
Switch_Level	switchLevel	Switch	<input type="checkbox"/>	<input type="checkbox"/>
Tier	Tier	Internal Volume	<input type="checkbox"/>	<input type="checkbox"/>
Tier	Tier	Qtree	<input type="checkbox"/>	<input type="checkbox"/>
Tier	Tier	Storage	<input type="checkbox"/>	<input type="checkbox"/>
Tier	Tier	Storage Pool	<input type="checkbox"/>	<input type="checkbox"/>
Tier	Tier	Volume	<input type="checkbox"/>	<input type="checkbox"/>

4. To edit the annotation process, do the following:

- Select **Published** to add annotations retrieved from OnCommand Insight into the Data Warehouse database. Click **All** to select all annotations on all objects. Click **None** to ensure that all options are not selected.



Uncheck this option to remove the annotation column from the specific object's inventory table and associated data marts. If any custom-designed reports use annotation data, the reports do not run successfully.

- Check **Init with Current** to initialize historical data in Data Warehouse dimension tables with the current annotation value. Click **All** to select all annotations on all objects. Click **None** to ensure that all options are not selected. This check box is disabled after an annotation is published; the check box is enabled for annotations that are not published.
For example, if a host is annotated with annotation type "floor" and gets the value "1", and there are 3 rows for that host in the host_dimension table, then selecting **Init with Current** associates the value "1" in the "floor" column for all 3 rows in the host_dimension table. If **Init with Current** is not selected, then only the latest row for that host will have the value "1" in the floor column.

5. Click **Save**.

A warning message appears indicating that this will cause changes to the structure of the data or data loss, if you are removing annotations.

6. To continue, click **Yes**.

Data Warehouse initiates an asynchronous annotations job that applies the requested changes. You can see the job in the Jobs page. You can also see the changes in the Data Warehouse database schema.

What to do with failing historical build points

You can build from history, omitting any failed builds by enabling the **Skip history build failures** option.

If you do this, the build from history continues. If a build fails and this option is enabled, Data Warehouse continues building and ignores any failed builds. In such cases, there is no data point in the historical data for any skipped builds. If you do not enable this option and the build fails, all subsequent jobs are aborted.

Administrative tasks you can perform using Data Warehouse

OnCommand Insight Data Warehouse is a web-based user interface that enables users to configure and troubleshoot data in OnCommand Insight Data Warehouse and to set up schedules to retrieve data from OnCommand Insight.

Using the Data Warehouse portal, you can perform the following administrative tasks:

- Check the status of current jobs or queries that are running
- Manage annotations
- Configure email notifications
- Access and create custom reports
- Review Data Warehouse documentation and database schema
- Edit the site name
- Identify the Data Warehouse version and upgrade history
- Build the Data Warehouse data from history
- Reset the Data Warehouse database
- Back up and restore the Data Warehouse database
- Troubleshoot Data Warehouse issues and look at OnCommand Insight logs
- Manage user accounts

Managing jobs

You can see a list of current jobs and their status. The first job in a build cycle is in bold type. The build that Data Warehouse performs for each connector and for each data mart is considered a job.

About this task

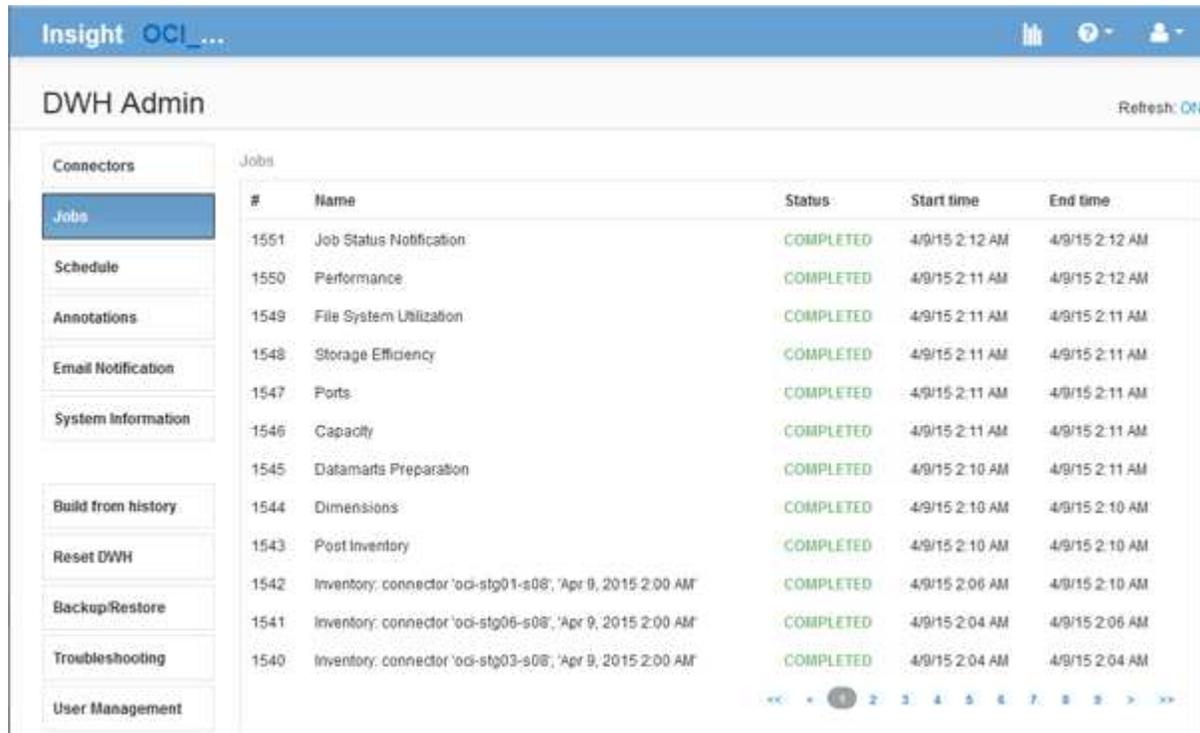
You can cancel any pending job you have scheduled or begun. You can also clear the history of previously executed jobs. You can clear history of jobs that are not pending, running, or in the process of aborting. You can clear all history or all history except the previous 24 hours to remove all but the last day's entries.

You can see information about the following types of jobs: License, Pre Inventory, Inventory, Post Inventory, Dimensions, Datamarts Preparation, Capacity, Ports, Storage Efficiency, File System Utilization, Performance, Job Status Notification, History build, Dynamic annotations, Connector removal, Skipped build, Phone Home, and Maintenance.

A maintenance job runs weekly and uses MySQL tools to optimize the database.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Jobs**.



The screenshot shows the 'DWH Admin' interface with the 'Jobs' tab selected in the left navigation pane. The main area displays a table of jobs with columns for ID, Name, Status, Start time, and End time. All jobs shown are in a 'COMPLETED' status.

#	Name	Status	Start time	End time
1551	Job Status Notification	COMPLETED	4/9/15 2:12 AM	4/9/15 2:12 AM
1550	Performance	COMPLETED	4/9/15 2:11 AM	4/9/15 2:12 AM
1549	File System Utilization	COMPLETED	4/9/15 2:11 AM	4/9/15 2:11 AM
1548	Storage Efficiency	COMPLETED	4/9/15 2:11 AM	4/9/15 2:11 AM
1547	Ports	COMPLETED	4/9/15 2:11 AM	4/9/15 2:11 AM
1546	Capacity	COMPLETED	4/9/15 2:11 AM	4/9/15 2:11 AM
1545	Datamarts Preparation	COMPLETED	4/9/15 2:10 AM	4/9/15 2:11 AM
1544	Dimensions	COMPLETED	4/9/15 2:10 AM	4/9/15 2:10 AM
1543	Post Inventory	COMPLETED	4/9/15 2:10 AM	4/9/15 2:10 AM
1542	Inventory: connector 'oci-stg01-s08', 'Apr 9, 2015 2:00 AM'	COMPLETED	4/9/15 2:06 AM	4/9/15 2:10 AM
1541	Inventory: connector 'oci-stg06-s08', 'Apr 9, 2015 2:00 AM'	COMPLETED	4/9/15 2:04 AM	4/9/15 2:06 AM
1540	Inventory: connector 'oci-stg03-s08', 'Apr 9, 2015 2:00 AM'	COMPLETED	4/9/15 2:04 AM	4/9/15 2:04 AM

If a Pending status appears, a cancel link appears.

3. To cancel a pending job, click **cancel**.
4. To remove the job history, click **All** or **All but last 24h**.

Monitoring Data Warehouse health

The Data Warehouse (DWH) includes a health monitor that displays information about the state of DWH. Alarm messages are displayed on the **Connectors** and **Jobs** pages of the DWH, as well as sent to the connected Insight server, where they are displayed on the **Admin > Health** page.

DWH collects metrics every ten minutes, and displays an alarm under the following conditions:

- Connection to the Insight server is down
- Disk utilization is greater than 90%
- Reporting (Cognos) service is down
- A query holds a lock on any table for a prolonged time
- A maintenance job is disabled
- Automatic backup is disabled
- Security risk: default encryption keys detected

Health monitor warnings in the Data Warehouse can be suppressed for up to 30 days.

When email notification is enabled, these events are also reported by email. Note that the email does not contain any attachments.

These events are logged in the `dwh_troubleshoot.log` file in the following locations:

- Windows: `<install dir>\SANSscreen\Wildfly\Standalone\Logs`
- Linux: `/var/log/netapp/oci/wildfly/`

Scheduling daily builds

Although you can manually build Data Warehouse by using the Build now control at any time, it is best practice to schedule automatic builds, defining when and how often to build the Data Warehouse database. Data Warehouse performs a build job for each connector and for each data mart. Data Warehouse performs a build job for each connector for licenses and inventory and all other build jobs (for example, capacity) are performed on the consolidated database.

About this task

Whenever the Data Warehouse is built, it performs an inventory job for every connector. After the inventory jobs are complete, Data Warehouse performs jobs for dimensions, capacity, and the remaining data marts.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Edit Schedule**.

Automatic Schedule

Enabled:	<input type="text" value="yes"/>	<input type="button" value="Edit schedule"/>
Schedule:	<input type="text" value="Daily at 2:00 AM,7:00 PM"/>	
Next run:	<input type="text" value="4/2/15 7:00 PM"/>	

3. In the **Build Schedule** dialog box, click **Edit** to add a new schedule.

Type:

Enabled:

Run at:

<input type="checkbox"/> 12:00 AM	<input type="checkbox"/> 1:00 AM	<input checked="" type="checkbox"/> 2:00 AM	<input type="checkbox"/> 3:00 AM	<input type="checkbox"/> 4:00 AM	<input type="checkbox"/> 5:00 AM	<input type="checkbox"/> 6:00 AM	<input type="checkbox"/> 7:00 AM	<input type="checkbox"/> 8:00 AM	<input type="checkbox"/> 9:00 AM	<input type="checkbox"/> 10:00 AM	<input type="checkbox"/> 11:00 AM
<input type="checkbox"/> 12:00 PM	<input type="checkbox"/> 1:00 PM	<input type="checkbox"/> 2:00 PM	<input type="checkbox"/> 3:00 PM	<input type="checkbox"/> 4:00 PM	<input type="checkbox"/> 5:00 PM	<input type="checkbox"/> 6:00 PM	<input checked="" type="checkbox"/> 7:00 PM	<input type="checkbox"/> 8:00 PM	<input type="checkbox"/> 9:00 PM	<input type="checkbox"/> 10:00 PM	<input type="checkbox"/> 11:00 PM

- Choose the frequency - weekly.
- Choose the time of day for each day you want the job to run.
- Chose N/A for days you do not want to run the build.
- To enable the schedule, select **Enabled**.



If you do not check this, the schedule build does not occur.

- Click **Save**.
- To build Data Warehouse outside of the automatic scheduled build, click **Build now**.

Configuring a weekly schedule

Although you can manually build Data Warehouse by using the Build now control at any time, it is best practice to schedule automatic builds, defining when and how often to build the Data Warehouse database. Data Warehouse performs a build job for each connector and for each data mart. Data Warehouse performs a build job for each connector for licenses and inventory and all other build jobs (for example, capacity) are performed on the consolidated database. With a weekly schedule, you can specify the time you want the build to run for each day of the week.

Steps

- Log in to the Data Warehouse Portal at <https://hostname/dwh>, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
- From the navigation pane on the left, click **Edit Schedule**.
- Choose the frequency - weekly.
- Choose the time of day for each day you want the job to run.
- Chose N/A for days you do not want to run the build.
- To enable the schedule, select **Enabled**.



If you do not check this, the schedule build does not occur.

- Click **Save**.
- To build Data Warehouse outside of the automatic scheduled build, click **Build now**.

Scheduling daily backups

Although you can manually back up Data Warehouse by using the Backup/Restore control at any time, it is best practice to schedule automatic backups, defining when and how often to back up the Data Warehouse database and Cognos content store. Backups offer protection from data loss, allowing you to restore the Data Warehouse database if needed. You also use a backup when migrating to a new Data Warehouse server or when upgrading to a new Data Warehouse version.

About this task

Scheduling backups during times when the Data Warehouse server is not busy improves backup performance and reduces the impact on users.

Steps

- Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
- From the navigation pane on the left, click **Schedule**.
- In the **Backup Schedule** dialog box, click **Edit** to add a new schedule.

Backup Enabled:

Backup Location:

Select Backup Configuration:

Run every:

<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>				
Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday

Run at hour:

Cleanup:

- To enable the scheduled backups, select **Backup Enabled**.
- Specify the location where you want to store the backup files.
- Specify the data you want backed up.
- Specify the day or days that you want the backup to be performed.
- Specify what time of day you want the backup started.
- Specify how many past backup copies you want to retain.
- Click **Save**.

Running custom scripts in Data Warehouse

Data Warehouse allows customers to create jobs that can run custom scripts that prepare customized data in Data Warehouse.

Before you begin

To prevent the custom script from being deleted during a Data Warehouse upgrade, you must not store the script in the SANscreen directory.

About this task

The job can specify only one script. You can run multiple scripts and commands from one script.

Steps

1. In Data Warehouse, select **DWH Admin > Schedule**.
2. Select the **Script enabled** check box.
3. Enter the absolute path of the script name in the **Script Location** text box.
4. Click **Save**.

Results

The Data Warehouse job engine schedules the task to run a “Custom scripting” job. The job is scheduled to run after an ETL and avoiding other conflicting background processes. The job is not run by a “Build from history” operation.

What you can do using annotations

Annotations provide a method for defining information that relates to objects in your environment and then allows you to track objects based on the annotation. For example, you could add building or floor number annotations to devices in your environment and then create a query that returns all of the devices on the first floor of a data center.

Additionally, you might want to look at all devices in a specific data center or business entity and determine which business entity is using the most tier 1 storage. To do this, you assign a data center, business entity, or tier annotation to the device using the OnCommand Insight web UI. Then, you can bring selected user-defined annotations from OnCommand Insight into Data Warehouse. You want to do this to see the annotation values assigned to objects appear in your custom reports.

You can specify which user-defined annotations propagate to Data Warehouse. Annotations are added as additional columns to the object table in the inventory, and to the relevant dimension table in the data marts. When you update the annotations on resources using the OnCommand Insight user interface and initiate or wait for the next Data Warehouse build, you see the results in the following tables:

- `dwh_inventory.annotation_value`
- `dwh_inventory.object_to_annotation`

Ensuring annotations entered in OnCommand Insight are included in Data Warehouse requires the following major processes:

- Before you import annotations into Data Warehouse, you must ensure that they are prepared in OnCommand Insight.

To do this, you can manually run the **Troubleshooting > Force Update of Annotations for Data Warehouse** option or wait until the next scheduled transient data run process. When you force the update of annotations, you force the OnCommand Insight server to calculate and place the transient data (such as annotation values) into database tables so that the Data Warehouse ETL process can read the data. The update of annotations data occurs automatically every fifteen minutes; however, you can force it to happen more frequently.

- You then import annotations into Data Warehouse by using the Data Warehouse **Annotations** option.
- If you want to include annotations in reports that you create by using the OnCommand Insight Reporting Portal report authoring tools, you must update the OnCommand Insight reporting metadata model.

When you upgrade Data Warehouse, the annotations job runs automatically during the database restore process. The annotations job runs automatically also when WildFly starts up.



WildFly is an application server where the OnCommand Insight Java code runs and is needed for both for the OnCommand Insight server and for Data Warehouse.

Preparing annotations in OnCommand Insight

Annotations must be prepared in OnCommand Insight before they can be imported into Data Warehouse.

Steps

1. Log in to the OnCommand Insight Portal as administrator `https://hostname`, where `hostname` is the name of the system where OnCommand Insight is installed.
2. Click on **Admin > Troubleshooting**. At the bottom of the page, click on **Advanced Troubleshooting**.
3. In the **Actions** tab, click **Update DWH Annotations (include deleted)**.

Importing user-defined annotations into Data Warehouse

After forcing an annotation update in OnCommand Insight, you need to select the annotations you want in Data Warehouse and initiate a Data Warehouse build. You can wait until the next scheduled build or initiate a build now.

Steps

1. Log in as an administrator to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Annotations**.

Annotations

Annotation	Column Name	Target Object	Published
Compute_Resource_Group	Compute_Resource_Group	Virtual Machine	
Data_Center	dataCenter	Host	✓
Data_Center	dataCenter	Storage	✓
Data_Center	dataCenter	Switch	✓
Note	Note	Switch	
Switch_Level	switchLevel	Switch	✓
Tier	Tier	Internal Volume	
Tier	Tier	Qtree	
Tier	Tier	Storage	
Tier	Tier	Storage Pool	
Tier	Tier	Volume	

The list displays a row for every annotation type and a target object to which the annotation can be assigned. A check mark in the Published column indicates that the annotation was already selected for the particular target object and is already available through the Data Warehouse data marts.

3. Click **Edit** to edit how annotations will be imported from OnCommand Insight.

Edit Annotations

Annotation	Column Name	Target Object	Published All / None	Init With Current All / None
Compute_Resource_Group	Compute_Resource_Group	Virtual Machine	<input type="checkbox"/>	<input type="checkbox"/>
Data_Center	dataCenter	Host	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Data_Center	dataCenter	Storage	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Data_Center	dataCenter	Switch	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Note	Note	Switch	<input type="checkbox"/>	<input type="checkbox"/>
Switch_Level	switchLevel	Switch	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Tier	Tier	Internal Volume	<input type="checkbox"/>	<input type="checkbox"/>
Tier	Tier	Qtree	<input type="checkbox"/>	<input type="checkbox"/>
Tier	Tier	Storage	<input type="checkbox"/>	<input type="checkbox"/>
Tier	Tier	Storage Pool	<input type="checkbox"/>	<input type="checkbox"/>
Tier	Tier	Volume	<input type="checkbox"/>	<input type="checkbox"/>

4. To edit the annotation process, do the following:

- Select **Published** to add annotations retrieved from OnCommand Insight into the Data Warehouse database. Click **All** to select all annotations on all objects. Click **None** to ensure that all options are not selected.



Uncheck this option to remove the annotation column from the specific object's inventory table and associated data marts. If any custom-designed reports use annotation data, the reports do not run successfully.

- Check **Init with Current** to initialize historical data in Data Warehouse dimension tables with the current annotation value. Click **All** to select all annotations on all objects. Click **None** to ensure that all

options are not selected. This check box is disabled after an annotation is published; the check box is enabled for annotations that are not published.

For example, if a host is annotated with annotation type “floor” and gets the value “1”, and there are 3 rows for that host in the host_dimension table, then selecting **Init with Current** associates the value “1” in the “floor” column for all 3 rows in the host_dimension table. If **Init with Current** is not selected, then only the latest row for that host will have the value “1” in the floor column.

5. Click **Save**.

A warning message appears indicating that this will cause changes to the structure of the data or data loss, if you are removing annotations.

6. To continue, click **Yes**.

Data Warehouse initiates an asynchronous annotations job that applies the requested changes. You can see the job in the Jobs page. You can also see the changes in the Data Warehouse database schema.

Viewing the Annotations job in the Jobs list

You can view the Annotations job in the Jobs list and apply the annotation changes to Data Warehouse data marts.

Steps

1. Log in as an administrator to the Data Warehouse Portal at <https://hostname/dwh>, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Jobs**.

Displaying annotation changes in the database schema

The database schema reflects the changes in the specific table.

About this task

For example, if you add annotations to a storage array, they appear in the storage or switch table in the inventory or other data marts.

If you update the annotations on resources using the OnCommand Insight user interface and initiate or wait for the next Data Warehouse build, you see a new column added or removed in the corresponding object in inventory (`dwh_inventory`) and in the corresponding dimension table as well (in the appropriate data mart). You see the results in the following tables:

- `dwh_inventory.annotation_value`
- `dwh_inventory.object_to_annotation`

Steps

1. Click  on the Data Warehouse tool bar and select **Documentation**.
2. Select **Database Schema**.
3. In the **Database Schema** pane on the left, scroll to the **DWH_INVENTORY** section and click **switch**.

Column	Type	Nullable	Description
id	int(11)	false	GUID of the switch.
fabricId	int(11)	true	GUID of the fabric on which this switch is configured to operate. References: <ul style="list-style-type: none"> id in dwh_inventory.fabric
identifier	varchar (255)	false	Identifier of the device.
wwn	varchar (255)	false	WWN of the switch.
ip	varchar (255)	false	IP address of the switch.
Name	varchar (255)	false	Name of the switch.
Manufacturer	varchar (255)	true	Manufacturer of the switch
Model	varchar (255)	true	Manufacturer's model of the switch.
Firmware	varchar (255)	true	Firmware version running on the switch.

4. The `dwh_inventory.switch` table reflects the changes:

Column	Type	Nullable	Description
tk	int(11)	false	TK of this storage array row.
name	varchar (255)	false	Name of the storage array.
identifier	varchar (255)	false	Identifier of the device.
ip	varchar (255)	false	IP address of the storage array.
model	varchar (255)	true	Manufacturer's model of the storage array.
manufacturer	varchar (255)	true	Manufacturer of the storage array.
serialNumber	varchar (255)	true	Serial number for the storage array.
microcodeVersion	varchar (255)	true	Version of the firmware running on the storage array.
family	varchar (255)	true	Family name of the storage array (e.g. Clariion, Symmetrix etc).
id	int(11)	true	GUID of the storage array in <code>dwh_inventory.storage</code> .

The dataCenter annotation column appears in the storage_dimensions table.

Setting email notifications

You can have Data Warehouse send email to a specific email address when Data Warehouse jobs do not complete successfully.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.

2. From the navigation pane on the left, click **Email Notification**.

3. Enter the following:

- SMTP server address

Specifies the server that is acting as the SMTP server in your organization, identified using either a hostname or an IP address using the `nnn.nnn.nnn.nnn` format. If you specify a host name, ensure that DNS can resolve it.

- SMTP server username and password

Specifies the user name to access the email server and is required only if your SMTP server requires a user to log into the server. This is the same user name you use to log in to the application and access your email.

- Notifications enabled

Yes enables the notifications; **No** disables the notifications.

- Sender's Email

Specifies the email address that is used to send the notifications. This must be a valid email address in your organization.

- Recipient's Email

Specifies the email address or addresses of the person or people who will always receive the email. Separate multiple addresses with commas.

- Email subject

Specifies the subject for the notification.

- Email signature

Specifies the information that displays at the bottom of the email, for example, the department name.

Accessing the Reporting Portal

From the Data Warehouse Portal, you can access the Reporting Portal, where you can create custom reports using report authoring tools such as Workspace Advanced and

Report Studio.

Steps

1. On the Data Warehouse toolbar, click  to open the Insight Reporting Portal.
2. Enter your user name and password and click **Login**.

Viewing the Data Warehouse database schema documentation

You can review Data Warehouse database schema information.

Steps

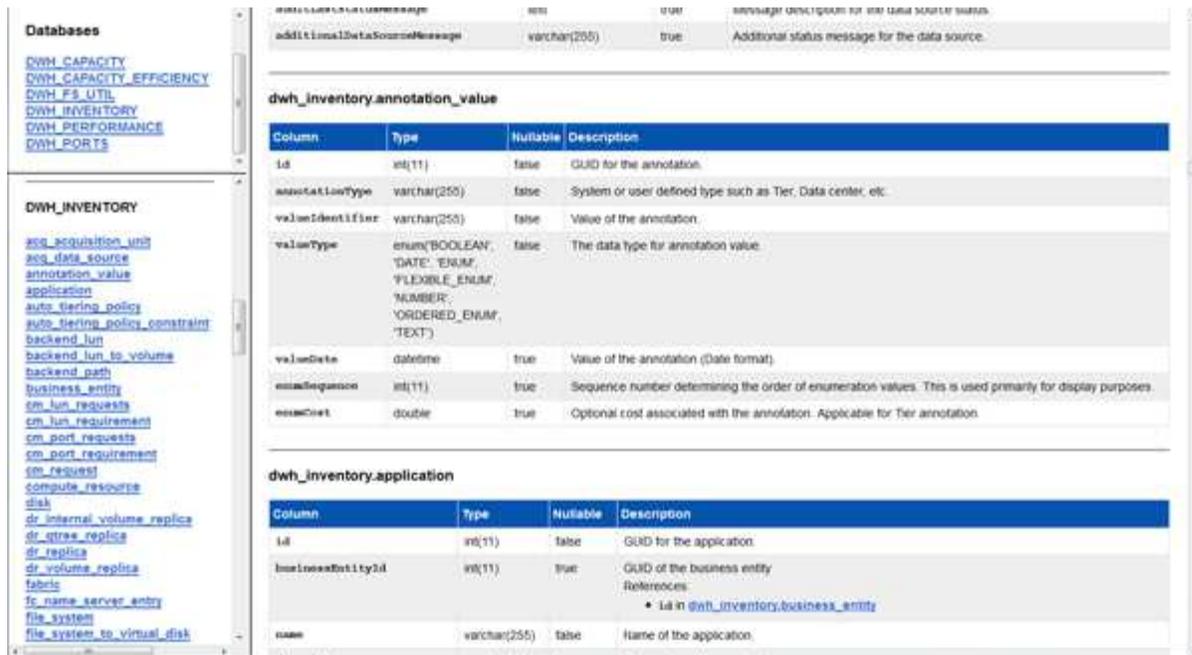
1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. On the Data Warehouse toolbar, click  and select **Schema**.

Viewing the Data Warehouse database schema

You might want to view the database schema to understand how to use the data in another API or to develop SQL queries. The schema option lists all databases, tables, and columns in the schema. You can also review the database schema diagrams showing the table relationships.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. Click  on the Data Warehouse toolbar, and select **Documentation**.
3. Select **Database Schema**.
4. For example, in the **Databases** pane, click **DWH_INVENTORY**.
5. In the **All tables** pane, scroll down to **DWH_INVENTORY** section, and click the **annotation_value** table.



The dwh_inventory.annotation table appears.

Viewing system information

You can view system, module, license, and Data Warehouse upgrade information.

Steps

1. Log in to the Data Warehouse Portal at <https://hostname/dwh>, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **System Information**.
3. On the **System** tab, review the system information and edit it the site name, if needed, by doing the following:
 - a. Click **Edit Site Name**
 - b. Enter the new site name and click **Save**.
4. To see application information (application name, module, version, and install date), click the **Application Info** tab.
5. To see license information (protocol, code, expiration date, and quantity), click the **Licenses** tab.
6. To see application upgrade information (application name, from date, to date, time, user, and file size), click **Upgrade History**.

Advanced options

Data Warehouse includes various advanced options.

Skipping failed builds

After your first build, sometimes you might encounter an unsuccessful build. To ensure that all the jobs after an unsuccessful build complete successfully, you can enable the

Skip history build failures option.

About this task

If a build fails and the **Skip history build failures** option is enabled, Data Warehouse continues building and ignores any failed builds. If this occurs, there will not be a data point in the historical data for any skipped builds.

Use this option only if the build is not successful.

If a build fails in Build from History and the **Skip history build failures** check box is not selected, all subsequent jobs are aborted.

Steps

1. Log in to the Data Warehouse Portal at <https://hostname/dwh>, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Build from history**.

Build From History

Target time	Start running	Status
3/13/15 12:00 AM	3/25/15 9:28 AM	COMPLETED
3/14/15 12:00 AM	3/25/15 9:34 AM	COMPLETED
3/15/15 12:00 AM	3/25/15 9:39 AM	COMPLETED
3/16/15 12:00 AM	3/25/15 9:45 AM	COMPLETED
3/17/15 12:00 AM	3/25/15 9:51 AM	COMPLETED
3/18/15 12:00 AM	3/25/15 9:57 AM	COMPLETED
3/19/15 12:00 AM	3/25/15 10:03 AM	COMPLETED
3/20/15 12:00 AM	3/25/15 10:09 AM	COMPLETED
3/21/15 12:00 AM	3/25/15 10:16 AM	COMPLETED
3/22/15 12:00 AM	3/25/15 10:23 AM	COMPLETED
3/23/15 12:00 AM	3/25/15 10:30 AM	COMPLETED
3/24/15 12:00 AM	3/25/15 10:38 AM	COMPLETED
3/25/15 12:00 AM	3/25/15 10:44 AM	COMPLETED

« < 1 2 3 > »

Cancel Pending Jobs Configure Run

Skip history build failures:

3. Click **Configure**.
4. Configure the build.
5. Click **Save**.
6. To skip failed builds, check **Skip history build failures**.

You can see this check box only if the **Run** button is enabled.

7. To perform a build outside of the automatic scheduled build, click **Run**.

Resetting the Data Warehouse database or Reporting server

You can delete the contents of the Data Warehouse data marts and delete all configured connectors. You might want to do this if an installation or upgrade did not complete successfully and it left the Data Warehouse database in an intermediate state. You can also delete only the Inventory data model or the Cognos Reporting data model.

Steps

1. Log in to the Data Warehouse Portal at <https://hostname/dwh>, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Reset DWH database**.
3. Click one of the following options:

- **Reset DWH Database**

This deletes the contents of all Data Warehouse data marts and all configured connectors and places the Data Warehouse to the default installed state without any custom configurations. You might choose this option, for example, if you changed your connected servers, but restored a different Data Warehouse database accidentally on your server and need to return to a default installed state. This does not delete any reports. (Reports are saved in the Cognos Content Store.)

- **Reset Inventory Only**

This deletes the contents of the Inventory data model only. This does not delete any historical data.

- **Reset Reporting Content**

This resets the content of the reporting server. This deletes any custom reports you may have. Backup your reports before you choose this option.

A warning message displays.

4. To continue, click **Yes**.

Restoring and upgrading reports for versions prior to 6.3

If you are upgrading an Insight version prior to 6.3, you must manually restore your reporting artifacts.

Before you begin

Follow the instructions in the "Upgrading the Data Warehouse (DWH)" and "Backing up custom reports and reporting artifacts" topics.

Steps

1. To restore Reporting artifacts from releases prior to version 6.3, copy the Export Backup.zip file you created and stored in your `<install>\cognos\c10_64\deployment` directory.
2. Open a browser and go to <http://<server>:<port>/reporting> for the server and port you used

during installation.

3. Enter your user name and password and click **Login**.
4. From the **Launch** menu, select **Insight Reporting Administration**.
5. Click the **Configuration** tab.

Due to changes in the data model, the reports in the old packages may not run and need to be upgraded.

6. Click **Content Administration**.
7. Click the **New Import** button.
8. Make sure that archive you copied to the deployment directory (for example, `backup6.0.zip`) is selected, and click **Next**.
9. If you entered a password to protect the archive, enter the password and click **OK**.
10. Change the name `Export...` to `Import Backup` and click **Next**.
11. Click on the pencil icon next to each package name and enter a new target name if necessary. For example, add a `_original` suffix to the existing name. Then click **OK**.
12. After you rename the target package names for all packages, select all blue folders and click **Next** to continue.
13. Accept all default values.
14. Click **Finish** and then select **Run**.
15. Check for the details of this import and click **OK**.
16. Click **Refresh** to view the status of the import.
17. Click **Close** after the import is complete.

Results

Two sets of packages appear in the Public Folders tab. For example, one with a `7.0` suffix (for the newer version) and one with a `_original` (or whatever you entered during the backup/restore procedure) suffix which contains your old reports. Due to changes in the data model, the reports in the old packages may not run and need to be upgraded. Your portal tabs now point to the current version of the portal pages.

Accessing MySQL using the command-line interface

In addition to accessing Data Warehouse data elements through the report authoring tools, you can obtain access to Data Warehouse data elements directly by connecting as a MySQL user. You might want to connect as a MySQL user to use the data elements in your own applications.

About this task

There are many ways to connect. The following steps show one way.

When accessing MySQL, connect to the MySQL database on the machine where Data Warehouse is installed. The MySQL port is 3306 by default; however, you can change it during installation. The user name and password is `dwhuser/netapp123`.

Steps

1. On the machine where Data Warehouse is installed, open a command-line window.
2. Access the MySQL directory in the OnCommand Insight directory.
3. Type the following user name and password: `mysql -udwhuser -pnetapp123`

The following is displayed, depending on where Data Warehouse is installed:

```
c:\Program Files\SANscreen\mysql\bin> mysql -udwhuser -pnetapp123
```

```
Welcome to the MySQL monitor. Commands end with ; or \g.
Your MySQL connection id is 882
Server version: 5.1.28-rc-community MySQL Community Server (GPL)
```

```
Type 'help;' or '\h' for help. Type '\c' to clear the buffer.
```

4. Show the Data Warehouse databases: `show databases;`

The following is displayed:

```
mysql> show databases;
+-----+
| Database                |
+-----+
| information_schema      |
| dwh_capacity            |
| dwh_capacity_efficiency |
| dwh_fs_util             |
| dwh_inventory           |
| dwh_performance        |
| dwh_ports               |
+-----+
```

Troubleshooting Data Warehouse

You can do various tasks related to troubleshooting Data Warehouse.

- Use OnCommand Insight ASUP.
- View OnCommand Insight logs.
- Resolve issues related to upgrades and business entities.
- Resolve issues related to the consolidation of multiple OnCommand Insight servers.

You can consolidate multiple OnCommand Insight servers into the same Data Warehouse database. Many configurations may report the same object from multiple connectors (that is, the same switch exists in two OnCommand Insight instances). In such cases, Data Warehouse consolidates the multiple objects into one (a primary connector is chosen and the object's data will be taken from that connector only).

The storage administrator can use the Troubleshooting page to solve problems related to consolidation issues.

Issue resolution with ASUP

You can send ASUP logs to technical support for assistance in troubleshooting. ASUP for Data Warehouse is configured to run automatically. In the Data Warehouse Portal, you can disable the automatic send process, choose to include a backup of the Data Warehouse database, or initiate a transmittal to ASUP.

The information in the logs is forwarded to technical support by using HTTPS protocol. No data is forwarded using ASUP unless you first configure it on the Insight Server.

Data Warehouse sends the logs to the OnCommand Insight Server that is the first connector listed in the Data Warehouse Portal Connectors page. The automatic process sends the following files:

- Data Warehouse logs, which includes the following:
 - boot.log (including backups)
 - dwh.log (including backups such as dwh.log.1)
 - dhw_troubleshoot.log
 - dwh_upgrade.log (including backups)
 - WildFly.log (including backups)
 - ldap.log (including backups)
 - SQL dump of the Data Warehouse management database
 - mysql: my.cnf, .err and slow query logs
 - full innodb status

- Cognos logs, which include the following:

- cognos-logs.zip

Contains the Cognos log files from the <install>\cognos\c10_64\logs directory. It also contains the logs generated by Cognos as well as the OnCommand InsightAP.log file that contains all logging from users logging in to and out of OnCommand Insight reporting.

- DailyBackup.zip

Contains the backup of the reporting artifacts in the Public Folders. The contents of My Folders is not included in this.

- cognos_version_site_name_content_store.zip

Contains a full backup of the Cognos Content Store.

You can generate a troubleshooting report manually. The Troubleshooting Report .zip file contains the following Data Warehouse information:

- boot.log (including backups)
- dwh.log (including backups such as dwh.log.1)
- dwh_upgrade.log (including backups)
- wildfly.log (including backups)
- ldap.log (including backups)
- dump files in c:\Program Files\SANscreen\wildfly\standalone\log\dwh\
- SQL dump of the Data Warehouse management database
- mysql: my.cnf, .err and slow query logs
- full innodb status



ASUP does not automatically send a backup of the OnCommand Insight database to technical support.

Disabling automatic ASUP transmissions

All NetApp products are equipped with automated capabilities to provide the best possible support to troubleshoot issues that occur in your environment. ASUP periodically sends predefined, specific, information to Customer Support. By default, ASUP is enabled for Data Warehouse; however, you can disable it if you no longer want the information sent.

Steps

1. From the navigation pane on the left, click **Troubleshooting**.
2. Click **Disable** to prevent ASUP from sending a daily report.

A message displays saying ASUP is disabled.

Including a backup of the Data Warehouse database

By default, ASUP sends only the Data Warehouse log files to technical support for assistance in troubleshooting; however, you can also choose to include a backup of the Data Warehouse database and select the type of data that is sent.

Steps

1. Log in to the Data Warehouse portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Troubleshooting**.
3. To specify that ASUP should include a backup of the Data Warehouse database, click the **Include DWH Database Backup** list and select one of the following options for the type of data that the backup should include:
 - All (including Performance)
 - All Except Performance
 - Inventory Only

4. Click **Update**.

Sending Insight logs to ASUP

You can send ASUP logs to technical support for assistance in troubleshooting. ASUP for Data Warehouse is configured to run automatically. In the Data Warehouse portal, you can disable the automatic send process, choose to include a backup of the Data Warehouse database, or initiate a transmittal to ASUP. When you request an ASUP report, the report request appears as a job in the Data Warehouse portal Jobs page.

About this task

The job is managed by the job queue similar to the processing of other jobs. If an ASUP job is in a Pending or Running state already, an error message appears indicating that the ASUP report request cannot be added to the job request, because the job queue contains pending or running requests.

Steps

1. Log in to the Data Warehouse portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Troubleshooting**.
3. In the **OnCommand Insight ASUP** section of the **Troubleshooting** page, click **Download DWH Troubleshooting Report** to retrieve the troubleshooting report.
4. To send the report to the OnCommand Insight Server listed as the first connector in the Data Warehouse Portal **Connectors** page, click **Send Now**.

Viewing OnCommand Insight logs

You can view various Data Warehouse and Cognos logs in OnCommand Insight.

About this task

You can examine troubleshooting and status information in Cognos and Data Warehouse log files.

Steps

1. Log in to the Data Warehouse portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. In the navigation pane on the left, click **Troubleshooting**.
3. In the **Logs** section, click **Log Files**.

The following log files are displayed:

dwh.log
Lists the status of Data Warehouse jobs
wildfly.log

Provides information about the WildFly application server
dwh_upgrade log
Provides information about the upgrade on Data Warehouse
ldap.log
Logs messages related to LDAP authentication
dwh_troubleshoot.log
Logs messages that can help troubleshooting DWH problems
sansscreenap.log
Provides information about connection to the server, authentication and access to the Cognos repository, and information about other processes
cognosserver.log
Cognos log

4. Click on the name of the log file you want to view.

Multiple server chassis consolidation issues

You can view the connectors that report on hosts and adapters and SAN switches and storage arrays. You can also see the various connectors that report on an object and identify the primary connector, which is the connector that was chosen for the object.

Viewing hosts and adapters consolidation issues

The reported data for hosts and their associated adapters is derived from the Inventory data mart.

Steps

1. Log in to the Data Warehouse Portal at <https://hostname/dwh>, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. In the navigation pane to the left, click **Troubleshooting**.
3. In the **Chassis Consolidation** section, click **Hosts and Adapters**.



The configuration shown in this example is not a valid configuration. Principal and Available connectors on localhost suggest that the Insight server and DWH are both installed on the same server. The intent of this example is to familiarize you with the consolidation table.

Hosts and Adapters Consolidation

Host GUID	Host Name	Host IP	Adapter GUID	Adapter WWN	Principal Connector	Available Connectors	Insight ID	Insight Change Time
288	Agassi	192.1.168.71			localhost (1)	localhost (1)	9927	11/18/10 1:36 PM
			576	40:A0:00:00:00:00:84	localhost (1)	localhost (1)	9928	11/18/10 1:36 PM
			577	40:A0:00:00:00:00:85	localhost (1)	localhost (1)	9930	11/18/10 1:36 PM
305	AI_Host1	192.1.168.88			localhost (1)	localhost (1)	12254	11/18/10 1:38 PM
			597	40:A0:00:00:00:00:01:05	localhost (1)	localhost (1)	12255	11/18/10 1:38 PM
306	AI_Host2	192.1.168.89			localhost (1)	localhost (1)	12257	11/18/10 1:38 PM
			598	40:A0:00:00:00:00:01:06	localhost (1)	localhost (1)	12258	11/18/10 1:38 PM
307	AI_Host3	192.1.168.90			localhost (1)	localhost (1)	12260	11/18/10 1:38 PM

For all hosts and adapters there is a row for each connector that reports on them, as well as the Primary Connector from which the host and adapter are taken. For hosts and adapters only, a host that is reported by one connector may have its adapters reported by a different connector.

You can also see the OnCommand Insight change time of a host/adapter for each connector. Using this parameter, you can discover when an update has occurred in OnCommand Insight for the host/adapter and when the same host/adapter has been updated in other OnCommand Insight servers.

- Optionally, filter data in this view by typing a portion of the text and clicking **Filter**. To clear the filter, delete the text in the **Filter** box and click **Filter**. You can filter by host name, host IP, adapter WWN, or OnCommand Insight object ID.

The filter is case sensitive.

- Review the following data:

- **Host GUID**

Global Unique Identifier for this type of consolidated device (hosts)

- **Host Name**

Name of the consolidated host as it appears in the data warehouse

- **Host IP**
IP address of the consolidated host
- **Adapter GUID**
Global Unique identifier for the host adapter
- **Adapter WWN**
WWN of the host adapter
- **Principal Connector**
Name of the OnCommand Insight connector that was the actual source of the data
- **Available Connectors**
All OnCommand Insight connectors where the consolidated host / adapter reside
- **Insight ID**
OnCommand Insight ID of the consolidated host/adapter for the relevant reporting connector
- **Insight Change Time**
When an update has occurred in OnCommand Insight for the host/adapter and when the same host/adapter has been updated in other OnCommand Insight servers

6. To obtain detail about the connector, click on the connector.

You can see the following information for the connector:

- Host name
- The last time a Data Warehouse job was run on that connector
- The last time a change was received from that connector
- The version of the OnCommand Insight server pointed to by that connector

Viewing storage arrays consolidation issues

The reported data for storage arrays is derived from the Inventory data mart. For all storage arrays, there is a row for each connector that reports on them, as well as the Primary Connector from which each array is taken.

Steps

1. Log in to the Data Warehouse Portal at <https://hostname/dwh>, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Troubleshooting**.
3. In the **Chassis Consolidation** section, click **SAN Storage Arrays**.
4. Optionally, to filter data in this view, type a portion of the text in the Filter box and click **Filter**. To clear the filter, delete the text in the Filter box and click **Filter**. You can filter by storage name, storage IP, vendor

model, or OnCommand Insight object ID.

The filter is case sensitive.

5. Review the following data:

- **GUID**

Global Unique Identifier for this type of consolidated device (storage array)

- **Name**

Name of the consolidated storage array as it appears in the Data Warehouse

- **IP**

IP address of the consolidated storage array

- **Vendor and Model**

Name of the vendor who sells the consolidated storage array and the manufacturer's model number

- **Principal Connector**

Name of the OnCommand Insight connector that was the actual source of the data

- **Available Connectors**

All OnCommand Insight connectors where the consolidated storage array resides

- **Insight ID**

ID of the consolidated storage array on the OnCommand Insight chassis where the Principal Connector resides

- **Insight Change Time**

When an update has occurred in OnCommand Insight for the storage array and when the same storage array has been updated in other OnCommand Insight servers

Viewing switches consolidation issues

The reported data for switches is derived from the Inventory data mart. For all switches, there is a row for each connector that reports on them, as well as the Primary Connector from which each switch is taken.

Steps

1. Log in to the Data Warehouse Portal at <https://hostname/dwh>, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Troubleshooting**.
3. In the **Chassis Consolidation** section, click **SAN Switches**.
4. Optionally, filter data in this view by typing a portion of the text and clicking **Filter**. To clear the filter, clear the Filter box and click **Filter**. You can filter by switch name, switch IP, vendor model, or OnCommand

Insight object ID.

The filter is case sensitive.

5. Review the following data:

- **GUID**

Global Unique Identifier for this type of consolidated device (storage array)

- **Name**

Name of the consolidated storage array as it appears in the data warehouse

- **IP**

IP address of the consolidated storage array

- **Vendor and Model**

Name of the vendor who sells the consolidated storage array and the manufacturer's model number

- **WWN**

WWN for the consolidation switch

- **Principal Connector**

Name of the OnCommand Insight connector that was the actual source of the data

- **Available Connectors**

All OnCommand Insight connectors where the consolidated storage array resides

- **Insight ID**

ID of the consolidated storage array on the OnCommand Insight chassis where the Principal Connector resides

- **Insight Change Time**

When an update has occurred in OnCommand Insight for the storage array and when the same storage array has been updated in other OnCommand Insight servers

Resolving multiple server annotation consolidation issues

The Annotation Consolidation view in the Data Warehouse Troubleshooting view displays a table that contains all the available Annotation Types and the Object Types to which they can be applied.

About this task

The consolidation of annotation values is based on the value of the Annotation Type. A storage array could have two different tier values, each coming from a different connector. Thus, if in one connector there is a tier defined by the name gold and in a second connector a tier is defined with the name goldy, this information

appears in Data Warehouse as two separate tiers.

Because some Annotation Types allow assignment of multiple annotation values to the same object, Data Warehouse allows objects (for example, "host") to have multiple annotation values assigned to them (for example, "data center 1" and "data center 2" could be assigned to the same host).

Tier annotation on volumes functions somewhat differently from the general annotation tables. Potentially, there could be a very large number of volumes in the environment and displaying all of them in the Data Warehouse would affect the usability of the information. Therefore, the Annotations Consolidation view displays only the volumes that have multiple tier values assigned to them, and the storage containing each such volume.

Steps

1. Log in to the Data Warehouse Portal at `https://hostname/dwh`, where `hostname` is the name of the system where OnCommand Insight Data Warehouse is installed.
2. From the navigation pane on the left, click **Troubleshooting**.
3. In the **Annotation Consolidation** section, click **Show** in the row for the object.

The following shows an example of the annotations for `Data_Center`:

Troubleshooting Annotations Consolidation

Annotation Type: Data_Center

Object Type: Host

Host GUID	Host Name	Host Natural Key	Data_Center Value	Connector
305	AI_Host1	192.1.168.88	New York	localhost (1)
306	AI_Host2	192.1.168.89	New York	localhost (1)
307	AI_Host3	192.1.168.90	New York	localhost (1)

Reporting

Welcome to OnCommand Insight reporting

OnCommand Insight reporting is a business intelligence tool that enables you to view pre-defined reports or create custom reports. OnCommand Insight reporting generates reports from the Data Warehouse (DWH) data.

With OnCommand Insight reporting you can perform the following tasks:

- Run a pre-defined report
- Create a custom report
- Customize the report format and delivery method
- Schedule reports to run automatically
- Email reports
- Use colors to represent thresholds on data

Pre-defined reports are the standard OnCommand Insight reports. This guide describes the pre-defined reports that are available with all of the product licenses.

Accessing the OnCommand Insight Reporting Portal

You can access the OnCommand Insight Reporting Portal directly from a web browser, from the Data Warehouse, or from the Insight server . You use the Reporting Portal to access predefined reports or to create your own reports using Data Warehouse data.

Access the reporting portal from a web browser

Steps

1. Open a web browser.
2. Enter the following URL: `https://server-name:9300/bi`

9300 represents the default port that was specified during installation. If another port was specified, you must change the port.

3. Enter your user name and password, and then click **OK**.

Accessing the reporting portal from the Insight server

Steps

1. Open a web browser.
2. Enter the following URL to access the the Insight server: `https://server-name`
3. Enter your user name and password, and then click **OK**.
4. In the Insight toolbar, click .

5. In the login page that appears, enter your user name and password, and then click **OK**.

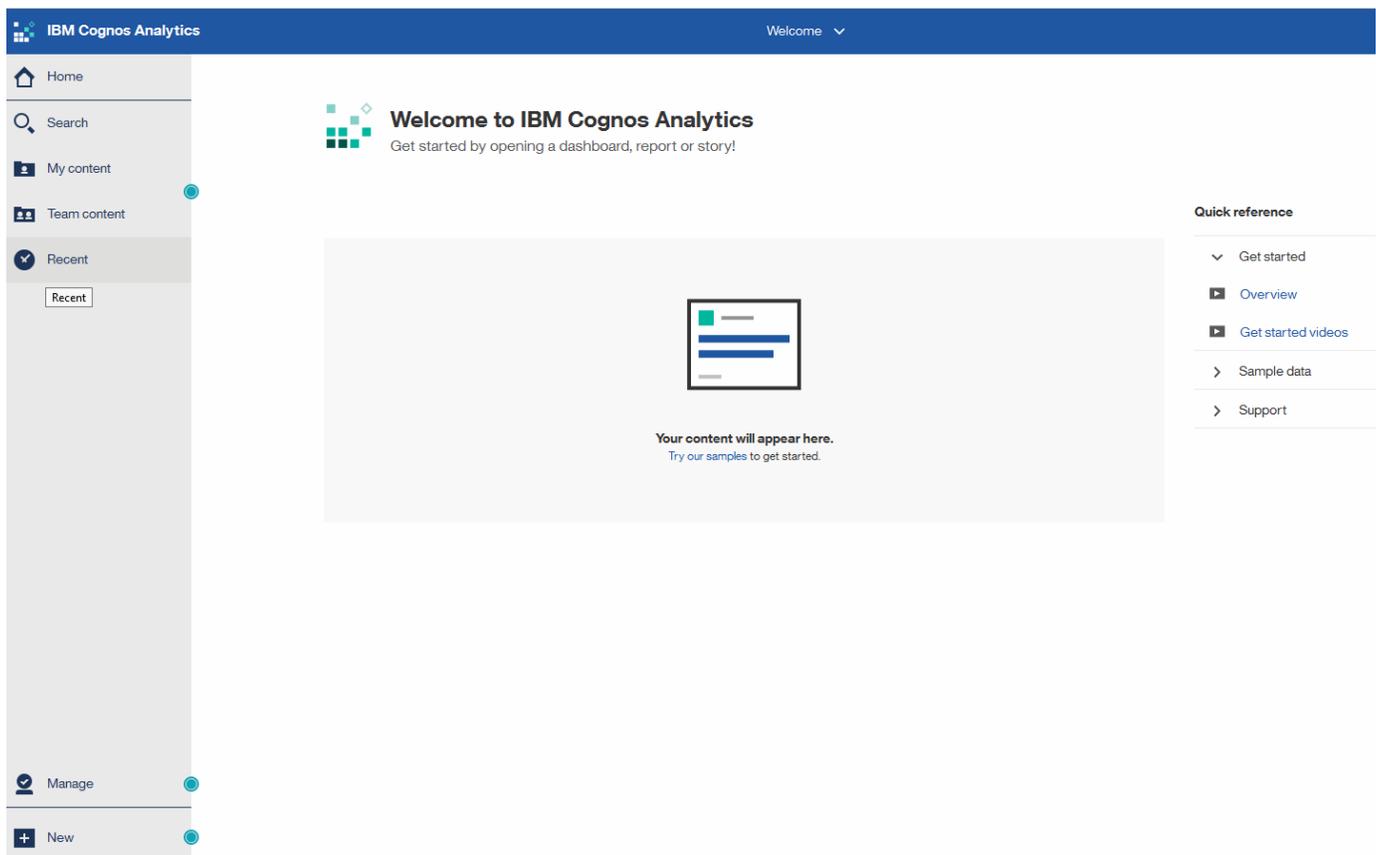
Accessing the reporting portal from the Data Warehouse

Steps

1. Open a web browser.
2. Enter the following URL to access the Data Warehouse: `https://server-name/dwh`
3. Enter your user name and password, and then click **OK**.
4. In the Data Warehouse toolbar, click .
5. In the login page that appears, enter your user name and password, and then click **OK**.

Results

The IBM Cognos Analytics welcome page is displayed. This is the default landing page of the OnCommand Insight Reporting Portal.



The screenshot shows the IBM Cognos Analytics interface. The top navigation bar is blue with the IBM Cognos Analytics logo on the left and a 'Welcome' dropdown on the right. Below the navigation bar is a sidebar with the following items: Home (house icon), Search (magnifying glass icon), My content (person icon), Team content (two people icon), Recent (checkmark icon), and a 'Recent' list. At the bottom of the sidebar are 'Manage' (gear icon) and 'New' (plus icon). The main content area has a white background with the IBM Cognos Analytics logo and the text 'Welcome to IBM Cognos Analytics' and 'Get started by opening a dashboard, report or story!'. Below this is a large light gray box with a bar chart icon and the text 'Your content will appear here. Try our samples to get started.' On the right side, there is a 'Quick reference' sidebar with a list of links: Get started, Overview, Get started videos, Sample data, and Support.

Variations due to installed licenses

Data in the OnCommand Insight reports is based upon the OnCommand Insight licenses that you have purchased. For example without the Plan license, you get point in time data (today) in the Inventory datamart for capacity and performance, but you do not have the ability to trend (report over a time period) the capacity or performance data for any device.

The absence of a Plan license removes the ability to create new reports or edit existing reports. You might see differences between the available reports in your OnCommand Insight system compared to the illustrations in the documentation. These variations are due to differences between the installed licenses on your system and the licenses on the system used to create the illustrations.

For more information about Licenses, see the OnCommand Insight Installation guide.

Reporting user roles

Each user account is assigned a role with a set of permissions. The number of users is limited by the number of Reporting licenses attached to each role.

Each role can perform the following actions:

- **Recipient**

Views OnCommand Insight Reporting portal reports and sets personal preferences such as those for languages and time zones.



Recipients cannot create reports, run reports, schedule reports, export reports, nor perform administrative tasks.

- **Business Consumer**

Runs reports and performs all Recipient options.

- **Business Author**

Views scheduled reports, runs reports interactively, creates stories, in addition to performing all Business Consumer options.

- **Pro Author**

Creates reports, creates packages and data modules, in addition to performing all Business Author options.

- **Administrator**

Performs reporting administrative tasks such as the import and export of report definitions, configuration of reports, configuration of data sources, and the shutdown and restart of reporting tasks.

The following table shows the privileges and the maximum number of users allowed for each role:

Feature	Recipient	Business Consumer	Business Author	Pro Author	Admin
View reports in the Team Content tab	Yes	Yes	Yes	Yes	Yes
Run reports	No	Yes	Yes	Yes	Yes

Schedule reports	No	Yes	Yes	Yes	Yes
Upload external files	No	No	Yes	Yes	No
Create stories	No	No	Yes	Yes	No
Create reports	No	No	No	Yes	No
Create Packages and Data Modules	No	No	No	Yes	No
Perform administrative tasks	No	No	No	No	Yes
Number of users	Number of OnCommand Insight users	20	2	1	1

When you add a new Data Warehouse and Reporting user, if you exceed the limit in a role, the user is added as “deactivated,” and you need to deactivate or remove another user with that role to give a new user membership.



Report authoring capabilities require Insight Plan license. You can add additional Business Author and Pro Author users by purchasing the ARAP (Additional Report Authoring Package). Contact your OnCommand Insight representative for assistance.

These reporting user roles do not affect direct database access. These reporting user roles do not impact your ability to create SQL queries using the data marts.

Enabling Security Headers

HTTP headers can be configured to enhance the overall security of the Cognos Analytics web application.

To add the response headers:

- Log in to the Cognos Analytics UI and navigate to **Manage -> Configuration -> System -> Advanced Settings**
- Add the following Key/Value and apply:
 - Key: BIHeaderFilter.responseHeaders
 - Value: [{"name": "X-FRAME-OPTIONS", "value": "SAMEORIGIN"}, {"name": "X-XSS-Protection", "value": "1"}, {"name": "X-Content-Type-Options", "value": "nosniff"}]
- Refresh your browser to enable the headers.

Reporting made easy

You can generate pre-defined reports from the OnCommand Insight Reporting Portal, email them to other users, and even modify them. Several reports enable you to filter by device, business entity, or tier. The reporting tools use IBM Cognos as a foundation and give you many data presentation options.

- The OnCommand Insight pre-defined reports show your inventory, storage capacity, chargeback, performance, storage efficiency, and cloud cost data. You can modify these pre-defined reports and save your modifications.

The report data available to you is controlled by several things, including the following:

- Login access to the OnCommand Insight Reporting Portal, which is defined by roles.
- The setup of the OnCommand InsightData Warehouse, which stores the data for the reports.

You can generate reports in various formats, including HTML, PDF, CSV, XML, and Excel.

OnCommand Insight accommodates multiple tenancy in reporting by enabling you to associate users with business units. With this feature, administrators can separate data or reports according to the attributes of a user or his affiliation.



With Cognos version 11.1.2 onward, reporting URLs are not considered "stable" and are subject to change. If you have bookmarked reporting URLs, these bookmarks may subsequently fail. More information can be found here: <http://queryvision.com/ibm-analytics-11-x-urls-they-are-a-changing/>



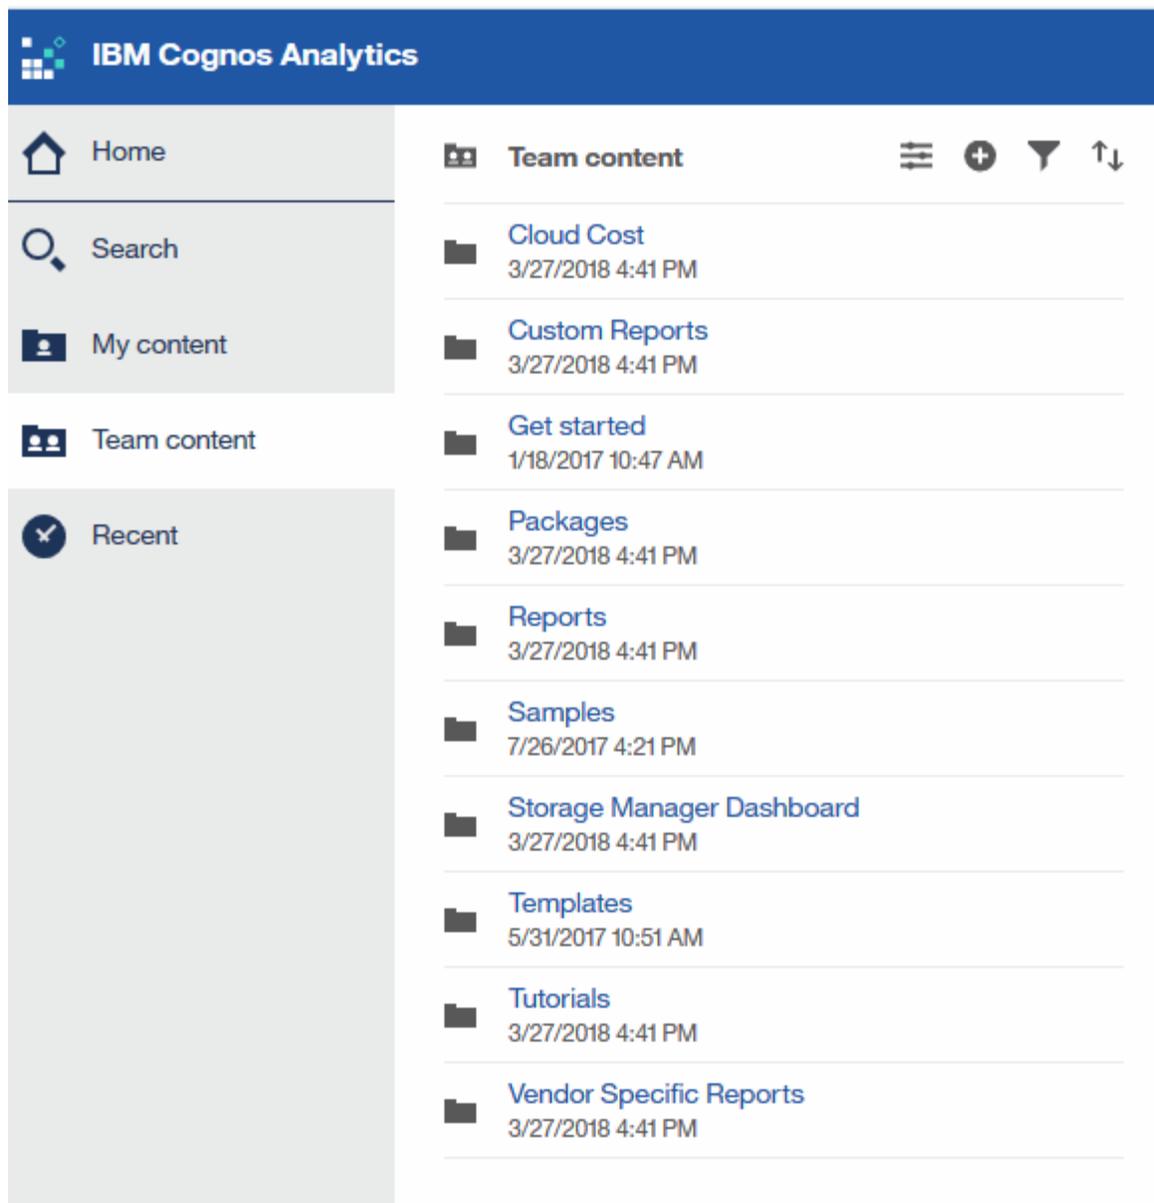
OnCommand Insight does not support any Dashboards created using Packages in IBM Cognos, unless using the new Data Module feature.

Navigating to pre-defined OnCommand Insight reports

When you open the Reporting Portal, the Team content folder is the starting point for you to select the type of information that you require in the OnCommand Insight reports.

Steps

1. In the left navigation pane, click **Team content** and select the information category that you want to use.



2. Click **Reports** to access the pre-defined reports.
3. Click **Get Started**, **Samples** or **Tutorials** to learn how to create reports.

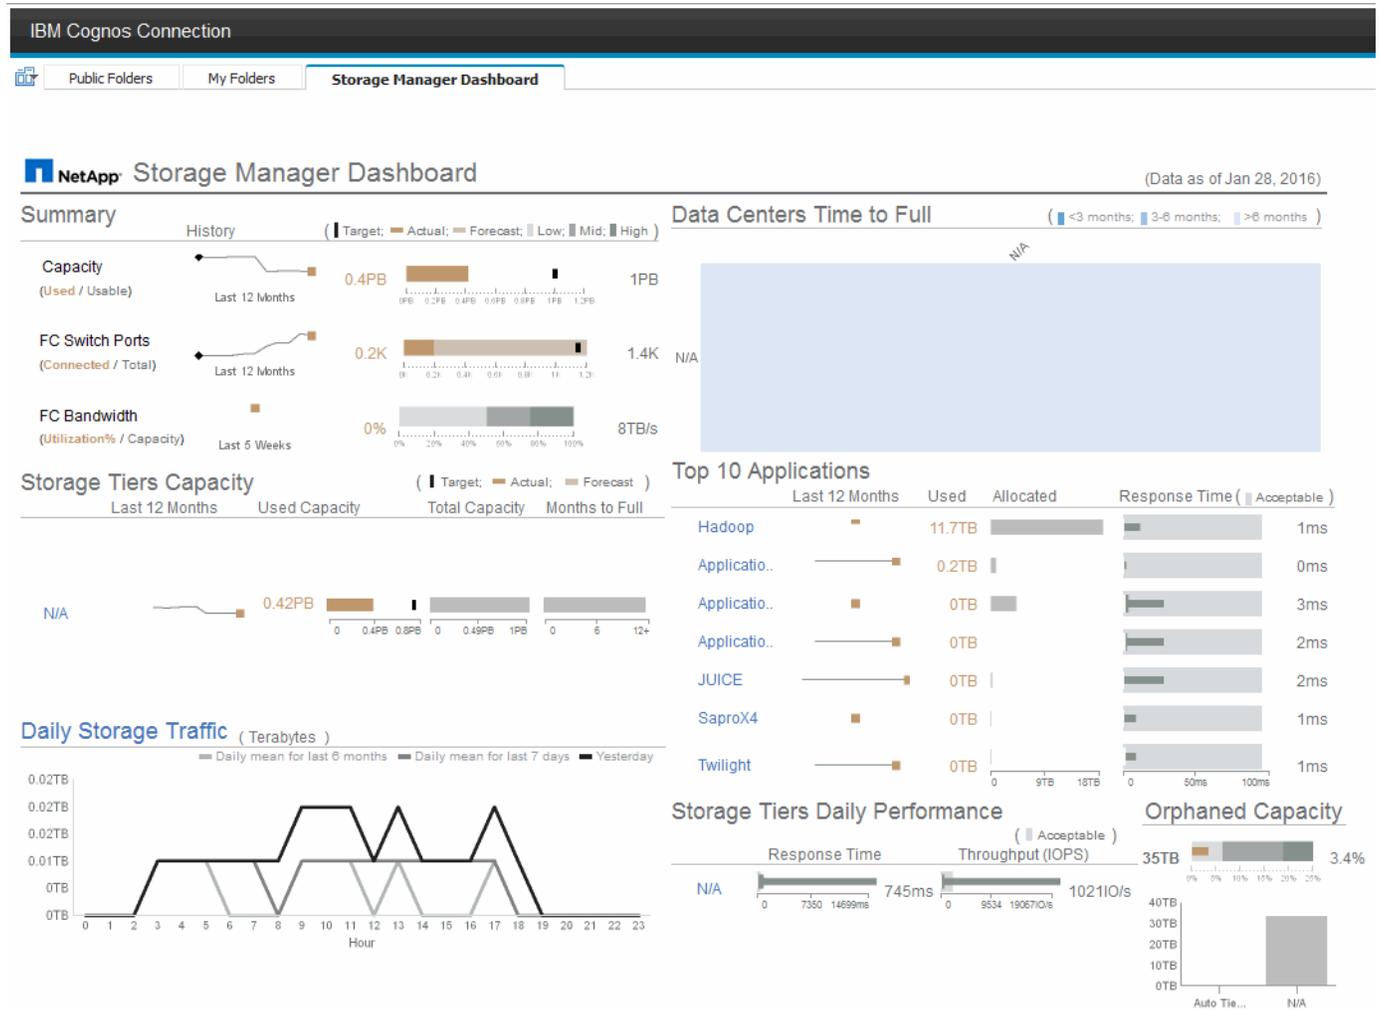
What the Storage Manager Dashboard enables you to do

You can use the Storage Manager Dashboard for the daily management of your storage services.

The Storage Manager Dashboard provides you with a centralized visualization that enables you to compare and contrast resource usage over time against the acceptable ranges and previous days of activity. Showing only the key performance metrics for your storage services, you can make decisions about how to maintain your data centers.

The dashboard comprises seven components that contain contextual information on certain aspects of your storage environment. You can drill down on the aspects of your storage services to perform an in-depth of analysis of a section that interests you most.

Summary



This component shows the used versus usable storage capacity, total switch ports versus the number of switch ports connected, and total connected switch port utilization versus the total bandwidth, and how each of these trend over time. You can view the actual utilization compared against the low, mid, and high ranges, which enables you to compare and contrast usage between Insight projections and your desired actuals, based on a target. For capacity and switch ports, you can configure this target. The forecast is based on an extrapolation of the current growth rate and the date you set. When the forecasted used capacity, which is based on future usage projection date, exceeds the target, an alert (solid red circle) appears next to Capacity.

Storage Tiers Capacity

This component shows the tier capacity used versus the capacity allocated to the tier, which indicates how the used capacity increases or decreases over a 12-month period and how many months are remaining to full capacity. Capacity usage is shown with values provided for actual usage, the usage forecast by Insight, and a target for capacity, which you can configure. When the forecasted used capacity, which is based on future usage projection date, exceeds the target capacity, an alert (solid red circle) appears next to a tier.

You can click any tier to display the Storage Pools Capacity and Performance Details report, which shows free versus used capacities, number of days to full, and performance (IOPS and Response Time) details for all the pools in the selected tier. You can also click any storage or storage pool name in this report to display the asset page summarizing the current state of that resource.

Daily Storage Traffic

This component shows how the environment is performing, if there is any large growth, changes, or potential issues compared to the previous six months. It also shows the average traffic versus the traffic for the previous seven days, and for the previous day. You can visualize any abnormalities in the way the infrastructure is performing because it provides information that highlights both cyclical (previous seven days) and seasonal variations (previous six months).

You can click the title (**Daily Storage Traffic**) to display the Storage Traffic Details report, which shows the heat map of the hourly storage traffic for the previous day for each storage system. Click any storage name in this report to display the asset page summarizing the current state of that resource.

Data Centers Time to Full

This component shows all the data centers versus all of the tiers and how much capacity remains in each data center for each tier of storage based on Insight forecasted growth rates. Tier capacity level is shown in blue; the darker the color, the lesser time the tier at the location has left before it is full.

You can click a section of a tier to display the Storage Pools Days to Full Details report, which shows total capacity, free capacity, and number of days to full for all the pools in the selected tier and the data center. Click any storage or storage pool name in this report to display the asset page summarizing the current state of that resource.

Top 10 Applications

This component shows the top 10 applications based on the used capacity. Regardless of how the tier organizes the data, this area displays the current used capacity and share of the infrastructure. You can visualize the range of user experience for the previous seven days to see if consumers experience acceptable (or, more importantly, unacceptable) response times.

This area also shows trending, which indicates if the applications meet their performance service level objectives (SLO). You can view the previous week's minimum response time, the first quartile, the third quartile, and the maximum response time, with a median shown against an acceptable SLO, which you can configure. When the median response time for any application is out of the acceptable SLO range, an alert (solid red circle) appears next to the application. You can click an application to display the asset page summarizing the current state of that resource.

Storage Tiers Daily Performance

This component shows a summary of the tier's performance for response time and IOPS for the previous seven days. This performance is compared against a SLO, which you can configure, enabling you to see if there is opportunity to consolidate tiers, realign workloads delivered from those tiers, or identify issues with particular tiers. When median response time or median IOPS is out of the acceptable SLO range, an alert (solid red circle) appears next to a tier.

You can click a tier name to display the Storage Pools Capacity and Performance Details report, which shows free versus used capacities, number of days to full, and performance (IOPS and response time) details for all the pools in the selected tier. Click any storage or storage pool in this report to display the asset page summarizing the current state of that resource.

Orphaned Capacity

This component shows the total orphaned capacity and orphaned capacity by tier, comparing it against acceptable ranges for total usable capacity and showing the actual capacity that is orphaned. Orphaned capacity is defined by configuration and by performance. *Storage orphaned by configuration* describes a

situation in which there is storage allocated to a host. However, the configuration has not been performed properly and the host cannot access the storage. *Orphaned by performance* is when the storage is correctly configured to be accessed by a host. However, there has been no storage traffic.

The horizontal stacked bar shows the acceptable ranges. The darker the gray, the more unacceptable the situation is. The actual situation is shown with the narrow bronze bar that shows the actual capacity that is orphaned.

You can click a tier to display the Orphaned Storage Details report, which shows all the volumes identified as orphaned by configuration and performance for the selected tier. Click any storage, storage pool, or volume in this report to display the asset page summarizing the current state of that resource.

Using predefined reports to answer common questions

OnCommand Insight includes predefined reports that address a number of common reporting requirements, providing critical insight that stakeholders need to make informed decisions about their storage infrastructure.

The following predefined reports are available in **Team content > Reports** or **Team content > Vendor Specific Reports**.

Newer versions of reports might be available at the NetApp Storage Automation Store. You should check the Automation Store regularly for reports.

- **AWS Cloud Cost Data**

The Cloud cost report provides a consolidated view of all assets so you can track, analyze and optimize usage and cost of cloud-based as well as on-prem services as they dynamically scale in your environment.

The report provides infrastructure-to-cost correlation, giving clear and actionable reporting to ensure right-sizing through focused capacity planning and waste detection.

- **Application Service Level Capacity and Performance**

The Application Service Level Capacity and Performance report provides a high level overview of your applications. You can use this information for capacity planning or for a migration plan.

- **Chargeback**

The Chargeback report provides storage capacity chargeback and accountability information by hosts, application, and business entities, and includes both current and historical data.

To prevent double counting do not include ESX servers, only monitor the VMs.

An updated version of this report is available at the NetApp Storage Automation Store.

- **Data Sources**

The Data Sources report shows all the data sources that are installed on your site, the status of the data source (success/failure), and status messages. The report provides information about where to start troubleshooting data sources. Failed data sources impact the accuracy of Insight reporting and the general usability of the product.

- **ESX vs VM Performance**

The ESX vs VM Performance report provides a comparison of ESX servers and VMs, showing average and peak IOPs, throughput, and latency and utilizations for ESX servers and VMs. To prevent double counting, exclude the ESX servers; only include the VMs.

An updated version of this report is available at the NetApp Storage Automation Store.

- **Fabric Summary**

The Fabric Summary report identifies switches and switch information, including port counts, firmware versions, and license status. The report does not include NPV switch ports.

- **Host HBAs**

The Host HBAs report provides an overview of the hosts in the environment and provides the vendor, model, and firmware version of HBAs, and the firmware level of the switches to which they are connected. This report can be used to analyze firmware compatibility when planning a firmware upgrade for a switch or an HBA.

- **Host Service Level Capacity and Performance**

The Host Service Level Capacity and Performance report provides an overview of storage utilization by host for block only applications.

- **Host Summary**

The Host Summary report provides an overview of storage utilization by each selected host with information for Fibre Channel and iSCSI hosts. The report enables you to compare ports and paths, the Fibre Channel and iSCSI capacity, and violation counts.

- **License Details**

The License Details report shows the entitled quantity of resources you are licensed for across all sites with active licenses. The report also shows a summation of actual quantity across all the sites with active licenses. The summation may include overlaps of storage arrays managed by multiple servers.

- **Mapped but not Masked Volumes**

The Mapped but not Masked Volumes report lists the volumes whose logical unit number (LUN) has been mapped for use by a particular host, but is not masked to that host. In some cases these could be decommissioned LUNs that have been unmasked. Unmasked volumes can be accessed by any host, making them vulnerable to data corruption.

- **NetApp Capacity and Performance**

The NetApp Capacity and Performance report provides global data for allocated, utilized, and committed capacity with trending and performance data for NetApp capacity.

- **OCI Scorecard**

The OCI Scorecard report provides a summary and general status of all assets discovered by OnCommand Insight. Status is indicated with green, yellow, and red flags:

- Green indicates normal condition
- Yellow indicates a potential issue in the environment
- Red indicates an issue that requires attention

All of the fields in the report are described in the Data Dictionary provided with the report.

- **Storage Summary**

The Storage Summary report provides a global summary of used and unused capacity data for raw, allocated, storage pools, and volumes. This report provides an overview of all of the storage discovered.

A newer version of this report is available at the NetApp Storage Automation Store.

- **VM Capacity and Performance**

Describes the virtual machine (VM) environment and its capacity usage. VM tools must be enabled to view some data, such as when VMs were powered down.

- **VM Paths**

The VM Paths report provides data store capacity data and performance metrics for which virtual machine is running on which host, which hosts are accessing which shared volumes, what the active access path is, and what comprises capacity allocation and usage.

- **HDS Capacity by Thin Pool**

The HDS Capacity by Thin Pool report shows the amount of usable capacity on a storage pool that is thin provisioned.

- **NetApp Capacity by Aggregate**

The NetApp Capacity by Aggregate report shows raw total, total, used, available, and committed space of aggregates.

- **Symmetrix Capacity by Thick Array**

The Symmetrix Capacity by Thick Array report shows raw capacity, useable capacity, free capacity, mapped, masked, and total free capacity.

- **Symmetrix Capacity by Thin Pool**

The Symmetrix Capacity by Thin Pool report shows raw capacity, useable capacity, used capacity, free capacity, used percentage, subscribed capacity, and subscription rate.

- **XIV Capacity by Array**

The XIV Capacity by Array report shows used and unused capacity for the array.

- **XIV Capacity by Pool**

The XIV Capacity by Pool report shows used and unused capacity for storage pools.

Creating a report using Cognos 11

Creating reports with Cognos 11 differs from previous versions of Cognos. Use this procedure to create a report using the pre-defined OnCommand Insight reports.

About this task

Use the following steps to generate a simple report on physical capacity of storage and storage pools in a number of data centers.

Steps

1. In the toolbar, click 

2. Click **Report**

3. Click **Templates > Blank**

4. Click **Themes > Cool Blue > OK**

The Source and Data tabs is displayed

5. Click **Source >** 

6. In the Open file dialog, click **Team content > Packages**

A list of available packages is displayed.

7. Click **Storage and Storage Pool Capacity > Open**

8. Click 

The available styles for your report are displayed.

9. Click **List**

Add appropriate names for List and Query

10. Click **OK**

11. Expand **Physical Capacity**

12. Expand to the lowest level of **Data Center**

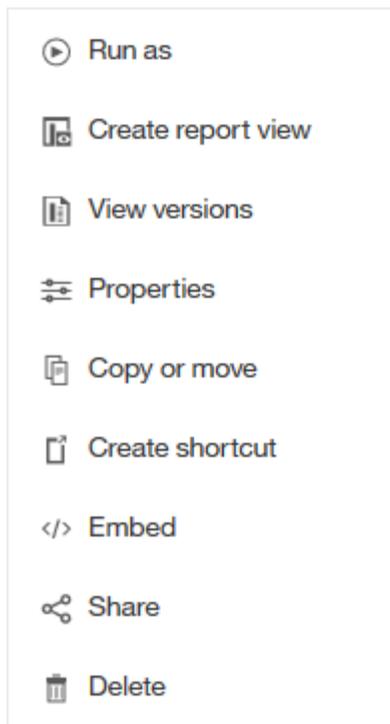
13. Drag  **Data Center** to the Reporting palate.

14. Expand **Capacity (MB)**

15. Drag **Capacity (MB)** to the Reporting palate.

16. Drag **Used Capacity (MB)** to the Reporting palate.

17.



Run the report, by clicking  and selecting an output type.

Results

A report similar to the following is created:

	Data Center	Capacity (MB)	Used Capacity (MB)
	Asia	122,070,096.00	45,708,105.00
	BLR	100,709,506.00	54,982,204.00
	Boulder	22,883,450.00	12,011,075.00
	DC01	1,707,024,715.00	1,407,609,686.00
	DC02	732,370,688.00	732,370,688.00
	DC03	314,598,162.00	65,448,975.00
	DC04	573,573,884.00	282,645,615.00
	DC05	89,245,458.00	62,145,011.00
	DC06	19,455,433,799.00	11,283,487,744.00
	DC08	100,709,506.00	44,950,171.00
	DC10	112,916,718.00	43,346,818.00
	DC14	23,565,735,054.00	17,357,431,924.00
	DC56	137,549,084.00	10,657,793.00
	Europe	743,942,208.00	240,369,325.00
	HIO	9,823,036,853.00	4,216,750,338.00
	London	0.00	0.00
	N/A	9,049,939,023.00	5,887,911,992.00
	RTP	12,386,326,262.00	5,638,948,477.00
	SAC	9,269,642,330.00	6,197,549,437.00
	↑ Top ↑ Page up ↓ Page down ↓ Bottom		

Managing reports

For each report, you can select the **More** link in the Actions column, and access all of the report operations, such as setting report properties, scheduling reports, or emailing reports. Administrators have more management options available than other users.

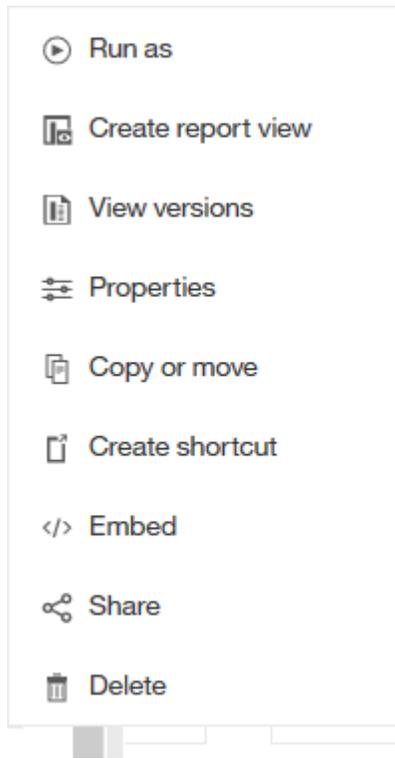
Administrators can set permissions for other report users according to their OnCommand Insight roles.

Customizing a report's output format and delivery

You can customize the format and delivery method of reports.

Steps

1. Open the OnCommand Insight Reporting Portal and select the report you want to customize and click [...].



2. Click **Properties** > **Schedule**

< Back

Create schedule

Period

Start

 2018-04-06

 1:49 PM

End

 2018-07-06

 1:49 PM

No end date

Run every

1 week(s)

On day(s)

M

T

W

T

F

S

S

Daily time interval

Options

Format

 HTML

>

Delivery

 Save

>

Prompts

Set values

>

Languages

English (United States)

>

3. You can set the following options:

- **Schedule** when you want reports to run.
- **Format** the report output.
- **Delivery** print, save, or email the report.
- **Languages** define the language the report is delivered in.

4. Click **Create** to produce the report using the selections you made.

Copying a report to the clipboard

Use this process to copy a report to the clipboard.

Steps

1. Open the Cognos 11 Reporting Portal: <https://server-name:9300/bi/>
2. In the toolbar, click 

3. Click **Report**

4. Click the **Pages** icon 

The **Report** icon  **Report** is displayed

5. Left click the **Report** icon

Report options are displayed.

6. Click **Copy Report to Clipboard**.

Opening reports (xml) from the clipboard

You can open a report specification that was previously copied to the clipboard.

About this task

You enter the Reporting user interface by creating a new report or opening an existing report

Steps

1. Open the Cognos 11 Reporting Portal: <https://server-name:9300/bi/>

2. In the toolbar, click 

3. Click **Report**

4. Click the **Pages** icon 

The **Report** icon  **Report** is displayed

5. Left click the **Report** icon

Report options are displayed.

6. Click **Open report from clipboard**.

Creating custom ad hoc reports

You can use the report authoring tools to create custom reports. After creating reports, you can save them and run them on a regular schedule. The results of reports can be automatically sent by email to yourself and others.

The examples in this section show the following process, which can be used for any of the OnCommand Insight data models:

- Identifying a question to be answered with a report
- Determining the data needed to support the results
- Selecting data elements for the report

What you need to do before you design your report

Before you design your custom report, you need to complete some prerequisite tasks. If you do not complete these, reports could be inaccurate or incomplete.

For example, if you do not finish the device identification process, your capacity reports will not be accurate. Or, if you do not finish setting annotations (such as tiers, business units, and data centers), your custom reports might not accurately report data across your domain and might show "N/A" for some data points.

Before you design your reports, complete the following tasks:

- Configure all data sources. For details, see the *OnCommand Insight Configuration and Administration Guide*.
- Enter annotations (such as tiers, data centers, and business units) on devices and resources in your environment. It is beneficial to have annotations stable before generating reports, because OnCommand Insight Data Warehouse collects historical information.
- Configure OnCommand Insight Data Warehouse to accept the data from the OnCommand Insight server in the Extract, Transform, and Load (ETL) process.

Process of creating reports

The process of creating ad hoc reports involves several tasks.

- Plan the results of your report.
- Identify data to support your results.
- Select the data model (for example, Chargeback data model, Inventory data model, and so on) that contains the data.
- Select data elements for the report.
- Optionally format, sort, and filter report results.

How to plan the results of your custom report

Before you open the report design tools, you might want to plan the results you want from the report. With report authoring tools, you can create reports easily and might not need a great deal of planning; however, it is a good idea to get a sense from the report requestor about the report requirements.

- Identify the exact question you want to answer. For example:
 - How much capacity do I have left?
 - What are the chargeback costs per business unit?
 - What is the capacity by tier to ensure that business units are aligned at the proper tier of storage?
 - How can I forecast power and cooling requirements? (Add customized metadata by adding annotations to resources.)
- Identify the data elements that you need to support the answer.
- Identify the relationships between data that you want to see in the answer. Do not include illogical relationships in your question, for example, "I want to see the ports that relate to capacity."
- Identify any calculations needed on data.

- Determine what types of filtering are needed to limit the results.
- Determine if you need to use current or historical data.
- Determine if you need to set access privileges on reports to limit the data to specific audiences.
- Identify how the report will be distributed. For example, should it be emailed on a set schedule or included in the Team content folder area?
- Determine who will maintain the report. This might affect the complexity of the design.
- Create a mockup of the report.

Tips for designing reports

Several tips might be helpful when you are designing reports.

- Determine whether you need to use current or historical data.

Most reports only need to report on the latest data available in the Data Warehouse.

- Data Warehouse provides historical information on capacity and performance, but not on inventory.
- Everybody sees all data; however, you might need to limit data to specific audiences.

To segment the information for different users, you can create reports and set access permissions on them.

Reporting data model

Your enterprise can benefit from the data that is discovered and stored in OnCommand Insight Data Warehouse. The OnCommand Insight Data Warehouse is a centralized repository that stores data from multiple information sources and transforms them into a common, multidimensional data model for efficient querying and analysis.

From this repository, you can generate custom reports such as chargeback, consumption analysis, and forecasting reports that answer questions such as the following:

- What inventory do I have?
- Where is my inventory?
- Who is using our assets?
- What is the chargeback for allocated storage for a business unit?
- How much headroom do I have on switch ports?
- How long until I need to acquire additional storage capacity?
- Are business units aligned along the proper storage tiers?
- How is storage allocation changing over a month, quarter, or year?

Using the data model provided with OnCommand Insight Reporting, you can use report authoring tools to design and schedule reports.

Data model overview

OnCommand Insight provides several data models for use in report development. Each

data model is an aggregation that summarizes data so that it can be queried and searched. For example, reports about capacity planning use the Capacity data model.

The OnCommand Insight enterprise reporting data models provide data elements and interactive relationships among data elements that yield business views of the data. Using the data elements and relationships, you can create reports using the IBM Cognos Analytics report generation tools recommended by NetApp.

OnCommand Insight also provides data marts that can be used to develop your own SQL queries. There is a distinction between these SQL query data marts and the data models used in reporting. The individual OnCommand Insight reporting data models use the underlying OnCommand Insight database schema provided in the data marts; however, the data models use additional tables and sometimes new elements in the tables. For instance, the data model includes a Monthly Capacity Fact table in the Storage Capacity data model that is based on the Capacity Fact table from the database schema and data mart. The data model filters out the values from the database schema table to show only month information.

Another example of a difference between the database schema used in data marts and the data model is in the Violation table and the Violation Type column. The data model translates programmatic-named values in the database to match the text displayed in the OnCommand Insight web UI.

OnCommand Insight data models

OnCommand Insight includes several data models from which you can either select predefined reports or create your own custom report.

Each data model contains a simple data mart and an advanced data mart:

- Simple data mart provides quick access to the most commonly used data elements and includes only the last snapshot of Data Warehouse data; it does not include historical data.
- Advanced data mart provides all values and details available from the simple data mart and includes access to historical data values.
- **Capacity data model**

Enables you to answer questions about storage capacity, file system utilization, internal volume capacity, port capacity, qtree capacity, and virtual machine (VM) capacity. The Capacity data model is a container for several capacity data models. You can create reports answering various types of questions using this data model:

- **Storage and Storage Pool Capacity data model**

Enables you to answer questions about storage capacity resource planning, including storage and storage pools, and includes both physical and virtual storage pool data. This simple data model can help you answer questions related to capacity on the floor and the capacity usage of storage pools by tier and data center over time.

If you are new to capacity reporting, you should start with this data model because it is a simpler, targeted data model. You can answer questions similar to the following using this data model:

- What is the projected date for reaching the capacity threshold of 80% of my physical storage?
- What is the physical storage capacity on an array for a given tier?
- What is my storage capacity by manufacturer and family as well as by data center?
- What is the storage utilization trend on an array for all of the tiers?

- What are my top 10 storage systems with the highest utilization?
- What is the storage utilization trend of the storage pools?
- How much capacity is already allocated?
- What capacity is available for allocation?

- **File System Utilization data model**

Enables you to answer questions about file system utilization. This data model provides visibility about capacity utilization by hosts at the file system level. Administrators can determine allocated and used capacity per file system, determine the type of file system, and identify trending statistics by file system type. You can answer the following questions using this data model:

- What is the size of the file system?
- Where is the data kept and how is it accessed, for example, local or SAN?
- What are the historical trends for the file system capacity? Then, based on this, what can we anticipate for future needs?

- **Internal Volume Capacity data model**

Enables you to answer questions about internal volume used capacity, allocated capacity, and capacity usage over time:

- Which internal volumes have a utilization higher than a predefined threshold?
- Which internal volumes are in danger of running out of capacity based on a trend?
- What is the used capacity versus the allocated capacity on our internal volumes?

- **Port Capacity data model**

Enables you to answer questions about switch port connectivity, port status, and port speed over time. You can answer questions similar the following to help you plan for purchases of new switches:

- How can I create a port consumption forecast that predicts resource (port) availability (according to data center, switch vendor and port speed)?
- Which ports are likely to run out of capacity, providing data speed, data center, vendor and number of Host and storage ports?
- What are the switch port capacity trends over time?
- What are the port speeds?
- What type of port capacity is needed and which organization is about to run out of a certain port type or vendor?
- What is the optimal time to purchase that capacity and make it available?

- **Qtree Capacity data model**

Enables you to trend qtree utilization (with data such as used versus allocated capacity) over time. You can view the information by different dimensions—for example, by business entity, application, tier, and service level. You can answer the following questions using this data model:

- What is the used capacity for qtrees versus the limits set per application or business entity?
- What are the trends of our used and free capacity so that we can do capacity planning?
- Which business entities are using the most capacity?

- Which applications consume the most capacity?

- **VM Capacity data model**

Enables you to report your virtual environment and its capacity usage. This data model lets you report on changes in capacity usage over time for VMs and data stores. The data model also provides thin provisioning and virtual machine chargeback data.

- How can I determine capacity chargeback based on capacity provisioned to VMs and data stores?
- What capacity is not used by VMs and which portion of unused is free, orphaned, or other?
- What do we need to purchase based on consumption trends?
- What are my storage efficiency savings achieved by using storage thin provisioning and deduplication technologies?
Capacities in the VM Capacity data model are taken from virtual disks (VMDKs). This means that the provisioned size of a VM using the VM Capacity data model is the size of its virtual disks. This is different from the provisioned capacity in the Virtual Machines view in OnCommand Insight, which shows the provisioned size for the VM itself.

- **Volume Capacity data model**

Enables you to analyze all aspects of the volumes in your environment and organize data by vendor, model, tier, service level, and data center. You can view the capacity related to orphaned volumes, unused volumes, and protection volumes (used for replication). You can also see different volume technologies (iSCSI or FC), and compare virtual volumes to non-virtual volumes for array virtualization issues. You can answer questions similar to the following with this data model:

- Which volumes have a utilization higher than a predefined threshold?
- What is the trend in my data center for orphan volume capacity?
- How much of my data center capacity is virtualized or thin provisioned?
- How much of my data center capacity must be reserved for replication?

- **Chargeback data model**

Enables you to answer questions about used capacity and allocated capacity on storage resources (volumes, internal volumes, and qtrees). This data model provides storage capacity chargeback and accountability information by hosts, application, and business entities, and includes both current and historical data. Report data can be categorized by service level and storage tier.

You can use this data model to generate chargeback reports by finding the amount of capacity that is used by a business entity. This data model enables you to create unified reporting of multiple protocols (including NAS, SAN, FC, and iSCSI).

- For storage without internal volumes, chargeback reports show chargeback by volumes.
- For storage with internal volumes:
 - If business entities are assigned to volumes, chargeback reports show chargeback by volumes.
 - If business entities are not assigned to volumes but assigned to qtrees, chargeback reports show chargeback by qtrees.
 - If business entities are not assigned to volumes and not assigned to qtrees, chargeback reports show the internal volume.
 - The decision whether to show chargeback by volume, qtree or internal volume is made per each internal volume, so it is possible for different internal volumes in the same storage pool to show

chargeback at different levels.

Capacity facts are purged after a default time interval. For details, see Data Warehouse processes.

Reports using the Chargeback data model might display different values than those reports using the Storage Capacity data model.

- For storage arrays that are not NetApp storage systems, the data from both data models is the same.
- For NetApp and Celerra storage systems, the Chargeback data model uses a single layer (of volumes, internal volumes, or qtrees) to base its charges, while the Storage Capacity data model uses multiple layers (of volumes and internal volumes) to base its charges.

• **Inventory data model**

Enables you to answer questions about inventory resources including hosts, storage systems, switches, disks, tapes, qtrees, quotas, virtual machines and servers, and generic devices. The Inventory data model includes several submarts that enable you to view information about replications, FC paths, iSCSI paths, NFS paths, and violations. The Inventory data model does not include historical data. Questions you can answer with this data mart could include the following:

- What assets do I have and where are they?
- Who is using the assets?
- What types of devices do I have and what are components of those devices?
- How many hosts per OS do I have and how many ports exist on those hosts?
- What storage arrays per vendor exist in each data center?
- How many switches per vendor do I have in each data center?
- How many ports are not licensed?
- What vendor tapes are we using and how many ports exist on each tape?
- Are all the generic devices identified before we begin working on reports?
- What are the paths between hosts and storage volumes or tapes?
- What are the paths between generic devices and storage volumes or tapes?
- How many violations of each type do I have per data center?
- For each replicated volume, what are the source and target volumes?
- Do I have any firmware incompatibilities or port speed mismatches between Fibre Channel host HBAs and switches?

• **Performance data model**

Enables you to answer questions about performance for volumes, application volumes, internal volumes, switches, applications, VMs, VMDKs, ESX versus VM, hosts, and application nodes. Using this data model, you can create reports that answer several types of performance management questions:

- What volumes or internal volumes have not been used or accessed during a specific period?
- Can we pinpoint any potential misconfiguration for storage for an application (unused)?
- What was the overall access behavior pattern for an application?
- Are tiered volumes assigned appropriately for a given application?
- Could we use cheaper storage for an application currently running without impact to application performance?

- What are the applications that are producing more accesses to currently configured storage?
When you use the switch performance tables, you can obtain the following information:
- Is my host traffic through connected ports balanced?
- Which switches or ports are exhibiting a high number of errors?
- What are the most used switches based on port performance?
- What are the underutilized switches based on port performance?
- What is the host trending throughput based on port performance?
- What is the performance utilization for last X days for one specified host, storage system, tape, or switch?
- Which devices are producing traffic on a specific switch (for example, which devices are responsible for use of a highly utilized switch)?
- What is the throughput for a specific business unit in our environment?
When you use the disk performance tables, you can obtain the following information:
- What is the throughput for a specified storage pool based on disk performance data?
- What is the highest used storage pool?
- What is the average disk utilization for a specific storage?
- What is the trend of usage for a storage system or storage pool based on disk performance data?
- What is the disk usage trending for a specific storage pool?

When you use VM and VMDK performance tables, you can obtain the following information:

- Is my virtual environment performing optimally?
- Which VMDKs are reporting the highest workloads?
- How can I use the performance reported from VMDs mapped to different datastores to make decisions about re-tiering.

The Performance data model includes information that helps you determine the appropriateness of tiers, storage misconfigurations for applications, and last access times of volumes and internal volumes. This data model provides data such as response times, IOPs, throughput, number of writes pending, and accessed status.

• **Storage Efficiency data model**

Enables you to track the storage efficiency score and potential over time. This data model stores measurements of not only the provisioned capacity, but also the amount that is used or consumed (the physical measurement). For example, when thin provisioning is enabled, OnCommand Insight indicates how much capacity is taken from the device. You can also use this model to determine efficiency when deduplication is enabled. You can answer various questions using the Storage Efficiency data mart:

- What is our storage efficiency savings as a result of implementing thin provisioning and deduplication technologies?
- What are the storage savings across data centers?
- Based on historical capacity trends, when do we need to purchase additional storage?
- What would be the capacity gain if we enabled technologies such as thin provisioning and deduplication?
- Regarding storage capacity, am I at risk now?

Data model fact and dimension tables

Each data model includes both fact and dimension tables.

- Fact tables: Contain data that is measured, for example, quantity, raw and usable capacity. Contain foreign keys to dimension tables.
- Dimension tables: Contain descriptive information about facts, for example, data center and business units. A dimension is a structure, often composed of hierarchies, that categorizes data. Dimensional attributes help describe the dimensional values.

Using different or multiple dimension attributes (seen as columns in the reports), you construct reports that access data for each dimension described in the data model.

For explanations of all data elements used in creating reports, see the Data Glossary.

Colors used in data model elements

Colors on data model elements have different indications.

- Yellow assets: Represent measurements.
- Non-yellow assets: Represent attributes. These values do not aggregate.

Using multiple data models in one report

Typically, you use one data model per report. However, you can write a report that combines data from multiple data models.

To write a report that combines data from multiple data models, choose one of the data models to use as the base, then write SQL queries to access the data from the additional data marts. You can use the SQL Join feature to combine the data from the different queries into a single query that you can use to write the report.

For example, say you want the current capacity for each storage array and you want to capture custom annotations on the arrays. You could create the report using the Storage Capacity data model. You could use the elements from the Current Capacity and dimension tables and add a separate SQL query to access the annotations information in the Inventory data model. Finally, you could combine the data by linking the Inventory storage data to the Storage Dimension table using the storage name and the join criteria.

FAQ

General Questions

This FAQ answers common general questions about OnCommand Insight.

When was OnCommand Insight (OCI) introduced?

OCI is one of the most mature infrastructure monitoring products in the industry today with over a decade in active development. Formerly known as Onaro or SANscreen, the SANscreen name was changed when joining the OnCommand Portfolio suite of products and is now referred to as OnCommand Insight, or more commonly Insight or OCI.

How long will OCI take to deploy in my environment?

OCI is simply a software download. Software is installed on two dedicated virtual or physical servers. Typical installations can be performed in as little as 2 hours and inventory, capacity and performance data will begin to be provided almost immediately. Any additional performance and best practice policies, user annotation, and cost awareness setup will require additional planning discussions.

Does OCI require agents, collectors, or probes?

OCI is 100% agentless and does not require the use of agents, taps or probes. All device discovery is read only, performed out of band, and over IP.

How does OCI discover and connect to devices?

OCI setup leverages the native APIs and protocols often already present in the data center environment, with no need of agents or probes. SSH, HTTP, SMIS and CLI are just a few examples. Where device element managers already exist (such as EMC's Unisphere, for example), OCI will communicate to the element manager(s) to capture the existing environmental data. Most device discoveries require only an IP address and read-only username and password. These device discoveries can be "one-to-many", such as with OCI's VMware data source. By discovering the VMware vCenter, OCI in turn discovers all of its ESXi hosts and their associated VM's, all with a single IP address and credential.

Does OCI require Professional Services? Is that available, and what do they offer?

For moderately-sized environments we recommend Professional Services for deployment, configuration, and integrations, as well as a wide variety of custom reporting and data validation possibilities. A short discussion with the OCI team and account engagement manager can help determine what services will benefit you the most.

How often does OCI release updates for new features and improvements?

Product updates and Service Packs are available for multiple versions of OCI. Major or minor releases are typically provided every few months, with service packs including new device support and firmware released more frequently. Both are available on the support.netapp.com download site. Certain updates such as new disk models that come out more frequently from manufacturers are pushed out automatically to the OCI software. Additionally, OCI data source device collection can be patched on site immediately after a development fix or update.

How does the OCI management team prioritize requests for new data sources?

OCI's Product Management team actively tracks all customer enhancement and interoperability feature requests (IFR's). Each request is detailed, evaluated for feasibility and prioritized based on customer demand and overall strategic business impact. Once accepted, requests are sized based on level of effort and scheduled for future development. The agile nature of the OCIs development process routinely allows for new data sources to be made available outside regular scheduled release cycles. NetApp account representatives can assist in customer inquiries and in submitting new requests on your behalf. Data sources can be patched on site, without the need to upgrade OCI.

My company runs completely on Linux. Will OCI work on Linux?

Yes, OCI supports several flavors of Linux as well as Windows. Be aware that Cognos (IBM's reporting tool used by OCI in conjunction with the Data Warehouse) is only supported on Windows, so if you are using OCI for reporting, you will need to run the reporting tool on a Windows server. The OCI Installation Guide lists the server requirements and supported operating systems for each OCI component.

Is OCI suitable for secure environments without internet access?

Yes, OCI is used by the top 10 Fortune 500 companies and by leading banking, healthcare, research and government agencies around the world today. OCI provides support for US military common access cards (CAC) and offers solutions for geographically-dispersed or heavily-firewalled environments.

I keep hearing that OnCommand Unified Manager (OCUM) is the management solution for cDOT. Can you help me understand why I would also use OCI?

OnCommand Unified Manager operates at the storage array "device management" layer, providing in-depth incident and event-based analysis of Clustered Data ONTAP (cDOT) arrays and their cluster interconnects. OCI provides a holistic view of on-premise and globally-dispersed environments consisting of 7-mode, Clustered Data ONTAP and other 3rd party arrays. Its end-to-end visibility, from VM to spindle, allows for historical trending and forecasting of capacity, performance and cost modeling that promotes a proactive service quality approach to data center management.

What is the OnCommand Insight Secondary ETL mentioned on the Automation Storefront?

The "Secondary ETL" requirement referenced in some OnCommand Insight Automation Storefront report downloads refers to a developed professional services implementation used for invoking additional Extract, Transform and Load (ETL) of captured data, for population into the OnCommand Insight data warehouse.

The Secondary ETL process primary purpose is to prefetch "batch" data allowing for more complex reports to generate faster, or to be scheduled to run on a daily basis.

This Secondary ETL is in addition to the recommended "once per day" ETL detailed in the OnCommand Insight data warehouse administration guide.

NetApp Professional Services is qualified to configure Secondary ETL scripting to avoid impact to existing OnCommand Insight report schedules, automated backups, scalability, or other system performance activities. For additional information regarding ETL scripting or data validation needs, please contact your NetApp Sales Representative and discuss how NetApp's Professional Services can assist you.

OnCommand Insight Licensing

Answers to common questions about OnCommand Insight licensing.

OCI Licensing Overview

OCI is licensed by capacity. Customers must purchase a license for each module they want to enable:

Discover is a prerequisite for Assure, Perform, and Plan and is not offered on its own. Discover is licensed by TB of managed capacity.

Assure is licensed by TB of managed capacity (as a single unit of charge for all storage infrastructure: FC, NAS, iSCSI, FCoE).

Perform is licensed by TB of managed capacity.

Plan is licensed by TB of managed capacity.

“Managed capacity” is defined as the raw capacity of the physical disks, virtual disks, and tapes prior to formatting. This is applicable to all storage discovered by Insight, both on-premises and in the cloud.

Most data sources are looking at disk raw base 2 capacity. There is no consideration for the disk role, such as a spare disk, unassigned disk, or RAID disk.

There are two types of Insight licenses available: **Perpetual** and **Subscription**.

Perpetual licenses allow you to indefinitely use the specific version/release of the software obtained subject to applicable license terms. If you have purchased a Software Support Plan (SSP), NetApp provides access to commercially available software updates through the NetApp Support site when and if updates are available in accordance with its Support Services terms. NetApp also provides access to special patches as determined by the NetApp Technical Support Center.

Subscription is a fixed term license of software which grants the right to:

- Use the software on-premises for a limited period only (most commonly 12 months) subject to applicable license terms
- Receive Software Support (previously referred to as SSP) for the period of the term
- While in effect the Licensee may use the most current commercially available version, release, or update, should any be made available as well as receive support for the software

At the end of each fixed term (most commonly 12 months), the license may be renewed for an additional fixed term (most commonly 12 months). If the license is not renewed, Licensee will no longer have the rights to use the software, will no longer be entitled to the benefits of SSP, and Licensee must destroy all copies of the software.

More about OCI License Modules

OCI has 4 core license modules to meet today's datacenter environment needs. These modules are **Discover, Perform, Assure and Plan**. Discover is the base module and is required for all other module purchases.

The **Discover** module enables OCI to locate the assets in the datacenter and dynamically map the device service paths. Information such as capacity, vendor information, model, firmware and serial numbers are provided.

Perform is OCI's performance collection module. Perform captures IOPS, throughput, latency, and CPU and memory information as well as provides other analytics.

Assure is positioned toward Fibre Channel environments and efficiency technologies. It helps identify and manage risks in fibre channel and iSCSI environments. Assure also helps with information on identification, mapping and alerting of masking, mapping and zoning service path entries and efficiency best practice policies such as fabric redundancy, switch hops, fan-out ratios and thin provisioning.

Plan provides the ability to identify and forecast trends across compute, fabric and various types of storage (cDOT, 7-mode, 3rd party) in hybrid on premise and globally-dispersed Data Center environments. It allows for longer retention times. The Data Warehouse consists of a built-in Intelligence to allow report authoring and avoids double counting of metrics in enterprise shared storage environments. It has the ability to generate and schedule a compliment of "out of the box" productized reports, or create your own reports using its "drag and drop" integrated report authoring tools.

Configuration and Supported devices

This FAQ answers common questions about OnCommand Insight configuration and supported devices.

Does OCI make changes to my environment?

No. OCI is a read-only tool that gathers information about your environment. OCI never makes any changes to your assets or configurations.

What permission-level access does OCI need to my devices?

In most cases where the device supports it, a read-only access is all that is required. There are some solutions that do not permit read-only access and thus would require the appropriate elevated permissions.

How often does OCI collect information?

OCI typically collects performance data every 5 minutes and discovery of logical and physical constructs every ½ hr. OCI sets the default polling intervals according to suggested best practices and scalability but does permit the user complete control over these intervals.

What is OCI's impact to my Environment?

OCI's agentless, out-of-band and passive IP communications help minimize setup, maintenance and impact to the data center ecosystem. OCI's performance development team takes great measures to minimize any impact to the Data center's performance in activities of monitoring performance itself. Impact is considered negligible in normal operating environments and can be relaxed or tightened in highly utilized or underperforming technology platforms. See the OnCommand Insight Installation Guide for more information.

How can I list all the hosts/VM's in OCI?

OCI's compliment of widgets and query-listing possibilities can be used to provide inventory style listings for Data Center assets. Listings of Virtual Machines down to the spindles and numerous constructs in between can all be made available to queries, widgets, dashboards, and data warehouse reporting, and are accessible through the RESTful API.

Does OCI provide the same type of support for related non-hypervisor hosts (i.e. physical servers)?

Hypervisors such as VMware provide detailed information on the ESXI hosts and their associated virtual machines (VMs). For physical servers, OCI collects metrics up to the host HBA. OCI employs a unique method in which it discovers physical servers using a patent-pending technology. Once storage and/or switches are discovered, host names for physical servers are contained within the fabric alias information. OCI selects these host names, matches them in DNS, and automatically brings the hosts into OCI. This technique greatly minimizes the need for manual entry updates and tool inventory maintenance.

Does OCI provide the same device metric depth (parity) across heterogenous environments?

There are varying degrees of standardizing, commonality and nomenclature across 3rd party platforms and vendor technologies. OCI attempts to normalize capacity and performance information into a consistent framework. Some capacity and performance metrics are provided natively from the device's counters, such as IOPs, latency and raw capacity. When counters are not provided, OCI can attempt to summarize the values (for example, by totaling the IOPs or capacities of underlying volumes), and in cases where neither is available, OCI will attempt to derive the metric values through various computational algorithms. OCI provides a generic SNMP integration capability to incorporate additional metrics not currently collected by OCI today.

Does OCI support Fibre Channel switches?

Yes, In addition to gathering data from your storage assets, OCI also acquires Inventory and Performance data from Cisco, Brocade and QLogic switches in your environment.

Are topology views of the whole infrastructure available? Does OCI show “end-to-end visibility”?

Yes, OCI dynamically discovers and maps the logical and physical constructs, providing an interactive end-to-end topology view of Compute, Fabrics, Virtualizers and back-end Storage. Topology icons allow quick launch navigation to impacted resources and aid in identification of workloads & violations in shared storage environments.

Scale and Ease of Use

This FAQ answers common questions about OnCommand Insight scaling and ease of use.

How does OCI scale?

OCI is a leader with respect to interoperability and the number of assets it can acquire with a minimal footprint. At its core, OCI requires 2 virtual or physical servers: one for the Operational Server which discovers all the data center assets, and one for its consolidated Data Warehouse for long term historical reporting. OCI's enterprise coverage supports hundreds of arrays, tens of thousands of Virtual Machines, 100,000 Fibre Channel paths and 10,000+ fibre channel ports, all in a single server instance.

How many people are needed to manage the OCI application?

OCI can be managed by a single person. But OCI has capabilities that can be used by multiple personas within the business environment, each with different roles, each with different reporting, troubleshooting or analytics needs. All efforts are made to minimize tool maintenance—from health and notification menus displaying

configuration problems, to automatic discovery of Physical Hosts attached to a fabric. Flexible annotations bring business context to the ecosystem data for all types of users. From Storage, Fabric, and Virtualization Administrators to Capacity Planners, Business Analysts and Executives, OCI brings the sharing of information across business silos and technologies together in a single pane of glass.

Does OCI support custom reporting?

Yes. OCI provides reporting via the IBM Cognos business intelligence tool, which allows you to create your own fully-customized reports from data collected in OCI's Data Warehouse.

How easy is it to create custom reports?

OCI reporting offers features for both novice and advanced users. OCI provides a number of report authoring capabilities including "drag and drop" report authoring and SQL query-based reporting for the more advanced user or professional service engagement. OCI's built-in business Intelligence solution (IBM Cognos) avoids common mistakes such as double counting capacity. With a complement of out of the box reports, widgets, queries and dashboards there are offerings to fit anyone's reporting requirements.

Customers can also find downloadable reporting templates available from the OCI community store.

Can OCI show performance and availability with "traffic light" simplicity?

Yes. OCI Data Warehouse and Reporting allow for reports with color enhancements—e.g. red/yellow/green "conditional styling" of values. Generating a colored font or background in a report can be implemented both by end users and Professional Services. OCIs widget libraries allow business specific performance metrics to be displayed in dashboards.

Performance troubleshooting

This FAQ answers common questions about OnCommand Insight performance troubleshooting.

How can I create a list of all the greedy resources in my environment?

OCI's correlation analytics help with identification of greedy and degraded resources for a specified service path. The correlation feature's generated analysis is performed in real time while viewing each object. The analysis provided greatly reduces the time necessary for troubleshooting performance issues and identifying root cause. Exploring generated violations of defined performance policies are one point of entry to discovering greedy or degraded resources. Both widgets and dashboards using the latest query capability help to filter, sort and visualize resources with higher than expected IOPS (greedy), Utilization or Latency.

Can OCI give one place to diagnose performance problems?

Yes. Performance Troubleshooting in OCI can be approached in multiple ways. OCI has a number of alerting methods possible. SNMP, Syslog and emailed Alerts are used commonly. Emailed Alerts allow users to quickly click and launch to the impacted resources within OCI. A global search window allows administrators to simply type in a resource name to begin analyzing the situation.

OCI's Violation Dashboard allows users to prioritize efforts based on the number of events, the duration and the time of day. An example of various alerting types would be Latency, IOPS, Utilization, Severity, business unit or even associated application.

OCI's correlation analytics helps administrators compare objects associated with the impacted resource and determines their impact to IOPS, Latency, Utilization, CPU and BB credits.

OCI's Query technology and Widget dashboards allows for pinpoint specifics in organized views that targets problem areas within the Datacenter.

Can OCI help with my 7-mode to cDOT migrations?

Yes, OCI provides an invaluable understanding for existing workload demand and post migration validations. OCI's role in modernizing today's datacenter allows for change management simulations, pre-migration optimization planning and defining the right tier of service. OCI effortlessly collects and correlates the business impact across thousands of NFS shares and Fibre channel paths in multi-vendor environments with just a few clicks. From migration to tech refreshes, OCI is providing a pathway to reliable, right-sized migrations and mitigating unplanned service disruptions.

How "real time" is OCI performance monitoring?

OCI is considered **near-real-time** for both on-premises and hybrid cloud data center management. While polling of data sources can be configured to occur more often, most users don't get significant analytical benefit from having a performance collection interval for most devices of less than 5 minutes. More frequent collection can put unnecessary burden on the objects under management and the analyses performed. Of course, there may be circumstances where a more granular collection is required, and fortunately OCI allows complete flexibility including configurable device inventory and performance polling intervals to suit your specific data center environment needs.

Why is my "Total" different from my "Read" plus "Write"?

In some instances, you may notice that the *Total* for a counter is not equal to the sum of *Reads* plus *Writes* for that counter. There are a few instances where this could happen.

IOPS: In addition to reads and writes, a storage array or other asset will process internal operations unrelated to the workload data flow. These are sometimes referred to as "system", "metadata", or simply "other" operations, and can be attributed to internal processes such as snapshots, deduplication, or space reallocation. In these cases, to find the amount of system operations for a given asset, subtract the sum of *Read* and *Write* IOPS from the *Total* IOPS. The sum of Read plus Write IOPS is the total IOPS directly related to your data flow.

Latency: The total response time (latency) for an operation can sometimes be reported as *less than* the write response time, because the total response time is a time-weighted average. I/O workloads will often consist of more read than write operations, with the writes typically observing larger latencies. For example, if a workload performed 10 read operations with an average latency of 5ms, and 5 write operations with an average latency of 10ms, the total weighted average latency will be calculated as the number of reads times the average read latency, plus the number of writes times the average write latency, divided by the total number of I/O operations, e.g. $(10 * 5 + 5 * 10) / (10 + 5) = 6.33\text{ms}$.

Why do OCI and OCUM show different values for overcommitted space?

The OnCommand Unified Manager (OCUM) notion of "Provisioned" space may include autogrow limits to which FlexVols (OnCommand Insight internal volumes) may grow. The OCI "Capacity" does not reflect those autogrow limits. As such, in an environment where autogrow Flexvols exist, the OCUM Provisioned capacity total will exceed the OCI storage level "Over-committed Capacity" total - the delta will be the difference between the Flexvols capacity and their autogrow capacity.

Managing your environment

This FAQ answers common questions about managing OnCommand Insight environments.

Can I give access to OCI to a specific user, while restricting the view to only certain resources (ie. SVM and related volumes, VMs, servers)?

OCI provides role-based access. For example, access to Reporting is controlled through OCI's Data Warehouse reporting. Reports can be scheduled, emailed as PDF, HTML or CSV, or to a file share or even a URL requiring the user to authenticate before viewing. User-based access is granted in the form of Admins, Users and Guests. Active directory/ldap support is also available.

Integrating Insight with other tools

This FAQ answers common questions about integration of OnCommand Insight with other tools.

Can OCI integrate with other tools and what integration points are available?

Yes, OCI is an extensible (wide open) solution allowing for Integrations with 3rd party orchestration, business management, change control and ticketing systems as well as custom CMDB integrations. OCI's fully published RESTful API and open MySQL database primary points of integration allow simple and effective movement of data and allow users to gain seamless access to their data.

Insight's Swagger-based API documentation is found in the product under **(?) Help > REST API Documentation**.

What is the Insight BMC Connector?

The OnCommand Insight Connector for BMC integrates the OnCommand Insight Data Warehouse (DWH) and the BMC Atrium Configuration Management Database (CMDB). The Insight Connector for BMC maps physical and logical stored data about network storage systems (for example, storage units, host storage services, VS Storage Service, and VM Storage Service) and their relationships with devices (hosts, storage switches, and tapes) and imports them into the BMC CMDB as configuration items and relationships. You can find more information about the OnCommand Insight Connector for BMC on the NetApp Support Site.

Does OCI work with SCOM or VROPs?

Yes, OCI complements a number of business management solutions and is considered an authoritative source for Storage, Compute, Hypervisor and Fabric information for the data center. OCI customers leverage OCI's RESTful API and Extensible MySQL database to enhance numerous applications like BMC's Remedy, ServiceNow, SCOM, Vrops, and Splunk, to name a few. OCI extends integrations by importing information from almost any source of record and/or by sending the captured environmental metrics to popular 3rd party Monitoring, Ticketing, CMDB billing & orchestration systems.

Can OCI work with cloud services I already use or am considering using?

Yes, OCI's management of both traditional on-premise and agile hybrid cloud environments provides visibility when determining the best, most cost-effective platform for your business service's needs. OCI can be leveraged for pre- and post-migration analysis, helping identify workloads that are suitable for the cloud.

Historical capacity trending, performance and cost are all necessary in order to select the appropriate cloud service. Service Design Workshops leveraging OCI's I/O density and other metrics can also help you answer questions like whether you optimized your environment and if the cloud makes sense. OCI continues to expand its coverage with support for NetApp Private Storage, Cloud ONTAP, Amazon S3 and Openstack KVM. OCI continues to play a vital role in NetApp's Cloud management campaign, especially in areas where visibility into Capacity Planning, Performance, Service Quality and Chargeback are important.

Can OCI open incidents in our incident management solution?

Yes, OCI violation events can be triggered and sent via SNMP as a trap or via Syslog to a server, and some by RESTful API. Details contained within the provided events can be interpreted by many 3rd party incident management and ticketing solutions.

Can you allocate resources to a business unit or departments?

Yes, OCI incorporates a method of metadata tagging called Annotations. Business Units, Lines of Business, Tenants, and Projects can be assigned to data center resources for richer business context around assets, capacity planning, troubleshooting and reporting.

Does OCI work with Work Flow Automator (WFA)?

OCI's integration capabilities with 3rd party CMDB, Billing and Orchestration technologies are a key value to its success, and WFA is no exception. NetApp's Professional Services have performed a number of successful Integrations that exist today with WFA workflows and OCI . There is a WFA connector available for download for OCI on the NetApp Automation Storefront.

How long are the OCI retention times for performance data?

The OCI server holds 90 days of near-real-time performance as well as the current (point in time) inventory (Logical and Physical constructs).

OCI performance polling intervals are user-configurable. Storage performance is typically configured for every 5 minutes for most vendors. Each day, performance/inventory data is sent to the OCI data warehouse (DWH) for long term historical and forecast reporting. DWH transforms this data into summarized data (Hourly, Daily, Monthly rollup data). Our ability to track "changes" e.g. monitored environmental history for Storage/Compute /fabric configuration/mappings, has no defined limit today.

Data Warehouse retains historical data based on the data marts and granularity of the data.

Are there any performance planning reports?

Yes, There are several reports provided with OCI and there are many others available in our Professional Services catalog, based on use case. The Data Warehouse module also comes with a suite of Cognos report authoring tools that allow users to create their own reports. There is also a complement of community-generated reporting templates and other downloads available from the NetApp Automation Storefront.

Data ONTAP Storage IOPS

This FAQ answers common questions about how IOPS numbers are derived from Data ONTAP storage systems.

How storage IOPS are derived from Data ONTAP storage systems

- Storage Array level IOPS are aggregated from the Internal Volumes IOPS
- Storage Node level IOPS include meta-data OPS
- Storage Pool level IOPS excludes meta-data OPS; only measures the disks
- Internal Volume level IOPS include Read + Write OPS (operations) + Other OPS

Question - How can the Aggregate IOPS be sometimes higher than Node IOPS?

Before CDOT 8.3.1 Node IOPS are made up of protocol IOPS. In CDOT 8.3.1. and later, they are made up of system constituents metrics. They include 'only' requests for data, request that come through the front door, but do not include backend tasks like snapmirrors, dedupe, and so on. On the other hand these tasks do produce disk IOPS, therefore Aggregate IOPS. Hence you might see Aggregate IOPS higher than the Node IOPS.

Question - How is Meta data or Other OPS calculated

Other OPS = Total - (Read + Write)

How-To guides

Getting Started with Insight

After OnCommand Insight has been installed and properly licensed, there are a number of tasks you should do to begin preparing your environment to showcase the data that is important to you.

Some of the tasks performed in a typical environment include the following:

1. **Annotating your assets** to prepare them for querying and reporting. Useful initial annotations typically include data Center, Tier, and Service Level.
2. **Creating queries** to show important data and help troubleshooting
3. **Assigning applications and business entities** to assets
4. **Creating performance policies and alerts** for violations against those policies
5. **Creating custom dashboards** to highlight data according to need or user role

Setting up notifications

You can configure Insight to send notifications on trigger events such as performance policy, global path, or capacity violations via email, SNMP, or Syslog. You can also configure Insight to send email notifications on system-level events such as data source errors or acquisition unit failures.

These are basic instructions. For more detail on notifications, see Configuration and administration > Insight configuration and administration > Setting up Insight.

Setting up email for notifications

Insight can send email notifications on trigger events, such as performance policy violations.

About this task

Follow these basic steps to configure email notifications:

Steps

1. Click **Admin > Notifications** and go to the **Email** section.
2. In the **Server** box, enter the name of your SMTP server. You can enter either a fully-qualified domain name or an IP address.
3. Enter your SMTP user name and (if required by your SMTP server) password.
4. In the **Sender email** box, enter the sender email account that will be identified as the sender on the notifications.

This account must be a valid email account within your organization.

5. In the **Email signature** box, enter any text that you want to be inserted in every email that is sent.

6. In the **Recipients** box, click **+** to enter an email address, and click **OK**.
7. Click **Save**.

To edit or remove an email address, or to send a test email, select the address and click the appropriate button that appears.

Note that you can configure Insight to send email notifications for specified performance policy violations to specific individuals or groups. For example, you might send cloud asset violations to one group, and physical host events to another. Go to **Manage > Performance policies** to configure individual policy notifications.

Setting up Syslog for logging

Insight can send syslog events for capacity or path violations and performance alerts.

About this task

Follow these basic steps to configure syslog notification in Insight:

Steps

1. Click **Admin > Notifications** and go to the **Syslog** section.
2. Place a check in the **Syslog enabled** checkbox..
3. In the **Server** field, enter the IP address of the log server.
4. In the **Facility** field, select the facility level that corresponds to the type of program that is logging the message.
5. Click **Save**.

Setting up SNMP for notifications

Insight can send SNMP notifications on trigger events, such as violations or when data source thresholds are exceeded.

About this task

Follow these basic steps to configure SNMP in Insight:

Steps

1. Click **Admin > Notifications** and go to the **SNMP** section.
2. Click **Actions** and select **Add trap source**.
3. In the **Add SNMP trap recipients** dialog box, enter the **IP** address and **Port** to which you want SNMP trap messages sent. For **Community String**, use "public" for SNMP trap messages.
4. Click **Save**.

Preparing assets: Annotating

Annotating allows you to associate specific tags or labels to the assets you choose, which aids in managing and reporting on those assets.

Creating annotations for your enterprise

This guide describes how to create and customize annotations for your environment that can be used for querying, filtering, alert notifications and reporting.

An annotation is a note or tag that you associate with specific assets in your environment. OnCommand Insight provides several annotations that you can configure for your assets as needed, or you can create your own custom annotations based on your business needs.

The examples that follow are those that are typically configured first in new customer environments, to serve as a baseline for additional actions. Your own annotation needs may vary, but the steps described herein can be used as a guide to configuring any annotations you may need on the assets you desire.

This guide is based on the following assumptions:

- You have OnCommand Insight Server installed and properly licensed.
- You want to explore best practices, not every available option.
- You understand that these are examples only and that your specific needs may vary.

This guide walks you through modifying existing annotations as well as creating custom ones

In our example environment, we wish to be able to list assets according to Data Center, Tier, Service Level and Environment.

Configuring Data Center annotations

The Data Center annotation is typically used to associate a storage array, switch, or physical host asset with a data center location. You may choose to associate the Data Center annotation with other assets in your environment as well.

Steps:

- Log in to Insight as a user with administrative permissions.
- Select **Manage > Annotations**.
- Choose the **Data Center** annotation and click on the **Edit** icon.
- Click **+Add** and add the Name and Description of your first data center to the annotation list.
- Do the same for your other data centers.
- When finished, click **Save**.

Example Data Center annotations:

Name	Description
DC1_SVL	Sunnyvale Bldg 1
DC2_SVLb3	SVL Bldg3 ENG
DC3_NY	New York

DC4_London	London
...	

Insight comes with several out-of-the-box annotation types that allow users to define or modify values to fit their needs. These default annotation types will always be available to the Insight web UI as well as reporting. Newly-created custom annotations are visible in the Insight web UI but require additional measures to make them available to reporting. For information on including custom annotations in reports, contact your NetApp Customer Support representative.



Some users may be inclined to use the Country annotation to set asset locations, as opposed to, or in conjunction with, the Data Center annotation. However, you should bear in mind that the Country annotation is treated as a custom annotation type in the Insight data warehouse and therefore may not show up in reporting at the same granularity as Data Center.

Configure Tier annotations

The Tier annotation is used to associate assets with their respective Tiers, for use in cost accounting, for example. Insight comes with a number of default Tier annotations; you may modify these according to your tiering naming conventions or create your own Tiers as needed.

When setting Tier annotations, keep the following in mind:

- Cost is per gigabyte.
- Tier 1, 2, 3 are default tiers configured at a storage array level, by disk type. However, many customers will have multiple disk types within an array, or across arrays of the same type.
- Best practice is to create Tier annotations based on disk type and/or disk speed. This is a typical Tier methodology; your own needs may vary.

Steps:

- Choose the **Tier** annotation and click on the **Edit** icon.
- If desired, click **+Add** and add the Name and Description of your first tier to the annotation list.
- Do the same for your other tiers.
- When finished, click **Save**.

Example Tier annotations:

Name	Description	Cost per Gb
Auto Tier	Automatic Storage Tiering Tier	0.5
Tier 1 SSD	All Flash Array	0.5
Tier 2 SAS	SAS	0.25

Tier 3 SATA	SATA	0.1
...		

Configure Service Level annotations

The Service Level annotation is used to associate assets with their respective service levels.

Service Level annotations are typically only set in customer environments that use auto-tiering. In the Insight data warehouse, Tier is preferred. However, best practice is to use Service Level when you want to detail Provisioned Cost vs. Customer Cost. When both are present in Data Warehouse, Service Level will supersede Tier.

Steps:

- Choose the **Service Level** annotation and click on the **Edit** icon.
- Click **+Add** and add the Name and Description of your first service level to the annotation list.
- Do the same for your other service levels.
- When finished, click **Save**.

Example Service Level annotations:

Name	Description	Cost per Gb
Service Level 1	FAS controllers with FC or SAS, local & remote mirror and tape	0.93
Service Level 2	FAS controllers with FC or SAS, local & remote mirrors	0.85
Service Level 3	FAS controllers with SATA and local mirror	0.48
...		

Configure custom Environment annotations

The Environment annotation is a custom annotation for associating assets with their respective environmental location or use, for example, Lab, R&D, Production, etc. By creating the Environment annotation and setting it onto these assets, you can easily find, filter and report on your Lab assets separately from your Production assets, for example.

Steps:

- Select **Manage > Annotations**.
- Click the **+Add** button at the top of the page.

- For **Name**, enter ‘**Environment**’.
- For **Description**, enter ‘**Asset environment type**’.
- For **Type**, select **List**. New fields are displayed for you to create your list.
- For now, leave **Add new assets on the fly** unchecked. You will check this on if you want to be able to add new Environments to the list of choices at the same time you are associating them with assets.
- Enter the Name and Description of your first environment.
- Click **+Add** and do the same for your other environments.
- When finished, click **Save**.

Example Environment annotations:

Name	Description
Lab	Lab
Dev	Development
Prd	Production
...	

Finding assets: Querying

You can easily find and display assets in your environment using powerful queries.

Using queries to annotate your assets

Now that you have created your initial annotations, let’s take a look at how to associate those annotations with specific assets.

In the examples that follow, we will apply these annotations to specific assets. For example, we will create a query to list all the storage arrays that reside in a specific data center, and mark those with the appropriate annotation. Then we’ll do the same for assets belonging to a specific tier and service level.

Querying and annotating data centers

You use queries to associate your annotations with the appropriate assets in your environment. In this example, we will associate Data Center annotations with selected assets.

During data source acquisition, Insight gathers (among other information) the names of each asset it discovers. For this example, we will assume that all of your storage arrays have been named according to the data center in which they reside, such as “SVL_NN_<label>” for arrays residing in Sunnyvale. Insight queries make annotating these assets simple.

- Log in to Insight as a user with administrative permissions
- Select **Queries** > **+New Query**

- Drop down the **Search for...** field and select **Storage**. A list of all of your storage arrays is displayed.
- In the **Name** filter field, type “SVL” and click the  button (or press Enter). The query results list is now updated to show only those arrays that contain the string “SVL”.
- When filtering, you can use any of the following characters alone or combined to refine your search in any text box on the Query page:
 - An asterisk enables you to search for everything. For example, “vol*rhel” displays assets that start with “vol” and end with “rhel”.
 - The question mark enables you to search for a specific number of characters. For example, filtering for “SVL-PRD??-S12” displays SVL-PRD12-S12, SVL-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.
- Select the storage arrays you wish to associate with this data center. When all of your desired arrays are selected, click the **Actions** button and select **Edit annotation**.
- In the **Add Annotation** dialog, select the **Data Center** annotation.
- Choose the desired **Value**, for example, “DC1_SVL”.
- Click **Save**.
- If the Data Center column is not visible on the Query results page, select it by dropping down the **Columns** button and choosing **Data Center**.
- If desired, you can save the query for future use by clicking the **Save** button in the upper right corner of the Query page and giving it a unique and explicit name. For example, “Storage Arrays - SVL data center”.

If you wish to associate the “SVL” annotation with other assets, create a new query and follow these steps for each asset type you desire.

Repeat these steps for assets in each of your data centers.

Querying and annotating tiers

You use queries to associate your annotations with the appropriate assets in your environment. Here, we will associate those tiers with the appropriate assets.

Earlier, we set up annotations for your Tiers. For this example, we will associate tiers with storage pools, and will assume your tier annotations are configured as follows:

Value	Description	Cost per Gb
Tier 1 SSD	All Flash Array	0.5
Tier 2 SAS	SAS	0.25
Tier 3 SATA	SATA	0.1

Let’s search for all SSD disks in your environment, and associate the “Tier 1 SSD” annotation with them.

- Log in to Insight as a user with administrative permissions
- Select **Queries > +New Query**

- Drop down the **Search for...** field and select **Storage Pool**. A list of all of your storage pools is displayed.
- The **Name** field may not be helpful this time, so let's use another field. Click the **More** drop-down and select "Least performing disk type". This field lists the disk types that we are interested in. Enter "SSD" in the field and click the  button. The query results list shows only your SSD storage pools.
- You may filter further by clicking the **More** drop-down and selecting additional fields.
- Select the storage pools you wish to associate with this tier. When all of your desired storage pools are selected, click the **Actions** button and select **Edit annotation**.
- In the **Add Annotation** dialog, select the **Tier** annotation.
- Choose the desired **Value** from the list. For this example, choose "Tier 1 SSD".
- Click **Save**.
- If the Tier column is not visible on the Query results page, select it by dropping down the **Columns** button and choosing **Tier**. You should see the appropriate annotation now associated with your assets.
- Save the query by clicking the **Save** button in the upper right corner of the Query page and giving it a unique and explicit name. For example, "Storage Pools - Tier 1 SSD".

If you wish to associate the "Tier 1 SSD" annotation with other assets, create a new query and follow these steps for each asset type you desire.

Repeat these steps for the assets in each of your remaining tiers.

Service Level and Environment annotations

Add Service Level and Environment annotations to the appropriate assets using the steps and concepts you've learned.

To add Service Level and Environment annotations to the appropriate assets in your environment, follow the steps noted above, choosing the desired assets and appropriate Service Level or Environment annotations. You can have multiple annotations associated with the same assets, and in fact this practice will allow you greater flexibility in managing your environment through Insight.

Now that you have created queries to annotate your assets, you can use these annotations in many different ways, such as:

- Performance policies to alert you when events occur on desired assets
- Custom dashboards and widgets to monitor activity
- Reporting

Your corporate structure: setting up business entities and applications

Understanding the elements of your corporate structure helps you to keep track of asset usage and report on costs.

Configuring business entities for your company

Understanding the business elements of your corporate structure helps you to keep track of asset usage and report on costs. Here we will configure your company business entities.

About this task

OnCommand Insight allows you to define business entities in a hierarchy that includes up to four levels of granularity.

- **Tenant**

Primarily used by service providers to associate resources with a customer. Tenant level is needed if your company is an ISP and you want to track customer usage of resources.

- **Line of Business (LOB)**

A line of business or product line within a company, for example, Data Storage. Line of Business is needed in the hierarchy if the data for different product lines needs to be tracked.

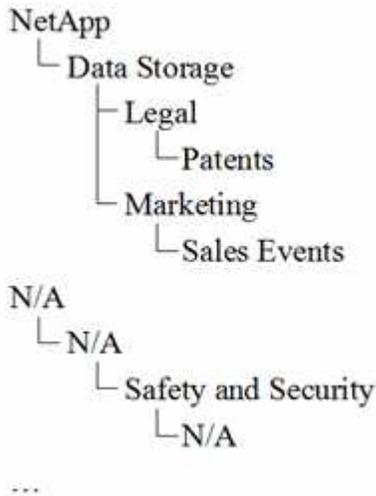
- **Business Unit**

Represents a traditional business unit such as Legal or Marketing. 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**

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.

An example of a business entity hierarchy would be:



Best practice: Create a table with each row showing one full business entity in your hierarchy:

Tenant	Line of Business	Business Unit	Project
NetApp Inc.	Data Storage	Legal	Patents
NetApp Inc.	Data storage	Marketing	Sales Events

N/A	N/A	Safety and Security	N/A
...			



You are not required to use all of the levels in the design of your corporate hierarchy. You can choose “N/A” for levels that you do not use.

To create a business entity hierarchy in Insight:

Steps

1. Log in to Insight as a user with administrative permissions.
2. Select **Manage > Business entities**.
3. Click the **+Add** button
4. Click in the **Tenant** box and type your tenant name.

If you have already entered tenants for your environment, a list of existing tenants will appear from which you may choose. You may also choose N/A if tenant does not apply for this business entity.

5. Repeat for **Line of Business, Business Unit, and Project**.
6. Click **Save**

After you finish

Best Practices:

- Map out your business hierarchy in a table, and check that the names in the hierarchy will be self-explanatory in Insight views and reports.
- Create your business entities in Insight before creating applications.
- Identify and list all applications that will be associated with each business entity.

Configuring applications for your company

Understanding the applications used in your company’s environment helps you to keep track of asset usage and report on costs. Here we will configure your company’s applications and associate them with the appropriate assets.

About this task

In the *Configuring business entities for your company* section, we created some business entities, and recommended that you list out all the applications you associate with each business entity. OnCommand Insight allows us to then track data associated with those applications for things like usage or cost reporting.

Before you can track data associated with the applications running in your environment, you must first define those applications and associate them with the appropriate assets. You can associate applications with the following assets: hosts, virtual machines, volumes, internal volumes, qtrees, shares, and hypervisors.

In this walkthrough, we want to track the usage of virtual machines that the Marketing Team uses for its Exchange email. You will remember the following table we created while defining our business entities. Let’s

add a column to this worksheet listing the applications used by each business entity. (This table is a worksheet example only. You will not see an “Applications” column in the business entities table in Insight.)

Tenant	Line of Business	Business Unit	Project	Applications
NetApp	Data Storage	Legal	Patents	Oracle Identity Manager, Oracle On Demand, PatentWiz
NetApp	Data storage	Marketing	Sales Events	Exchange, Oracle Shared DataBase, BlastOff Event Planner
N/A	N/A	Safety and Security	N/A	N/A
...				

Creating applications in Insight:

Steps

1. Log in to Insight as a user with administrative permissions.
2. Select **Manage > Applications**
3. Click the **+Add** button
4. Enter the name of the application (for our example, enter “Exchange”)
5. Select a priority for the application
6. If you wish to associate the application to a business entity, select one from the **Business Entity** drop-down. Otherwise, you can leave this as “None”.
7. If you want to ensure each host has access to the same volumes in a cluster, ensure that the **Validate volume sharing** box is checked. 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. If you do not use volume sharing, clear the **Validate volume sharing** box. This requires the Assure license.
8. Click Save.
9. Repeat for all other Applications in your environment.

After you finish

We see that the Marketing Team uses the Exchange application. We want to see their virtual machine utilization for Exchange, in order to predict when we will need to add more storage. Let’s associate the Exchange application with all of Marketing’s VM’s. The easiest way to accomplish this is through a query.

By following these steps, you can associate each of your applications with their appropriate assets.

Associating Applications to assets:

Now that you have created your applications (and tied them to business entities, as desired), we can now

associate those applications with assets in your environment. In this example, we will associate the Exchange application with a number of virtual machines in your company. The easiest way to do this is what a query.

1. Select **Queries** > **+New query**.
2. In the **Select Resource Type** drop-down, choose *Virtual Machine*
3. We will assume the Marketing team names their assets with the string “_mktg_”. In the Name filter box, enter “_mktg_”(without quotes) and click the apply (checkmark) button.
4. The list of all VM’s with the “_mktg_”string is shown.
5. If desired, click the **More** drop-down and add additional filters.
6. Select the VM’s used for Exchange by clicking in the checkbox beside each desired VM Name, or select all VMs by clicking the checkbox at the top of the column.
7. When the desired VM’s are selected, click the **Actions** button and choose **Add Application**.
8. In the Assign Application dialog, click the **Application** drop-down and select “Exchange”.
9. Click **Save**.
10. Repeat as necessary to associate the Exchange application with other assets (hosts, volumes, etc.)

Creating performance policies for alerting

Performance Policies allow you to monitor and send alerts when specific conditions are met.

About this task

Now that we have annotated our assets, let’s create a Performance Policy that we can use to alert us when latency is higher than 2ms in any storage array at our Sunnyvale (DC1_SVL) data center. When these conditions occur, we will send an email to selected recipients.

Steps

1. Select **Manage** > **Performance policies**.

The Performance policies page opens. There are several default policies already set, which you can modify to suit your needs. We will create a new policy, however.

2. Click the **+Add** button.

The **Add Policy** dialog opens.

3. In the **Policy name** field, enter “SVL Data Center Latency 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.

4. For **Apply to objects of type**, select **Storage**.
5. In the **With annotation** field, select **Data Center** is “DC1_SVL” (or choose the name of your desired data center here).
6. Apply after a window of **First occurrence** .

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.

7. From the **With severity** list, select **Warning**.
8. Under **Email recipients**, click to override the global recipient list. Click **+** to add the email address of your first desired alert recipient, and click **OK**. Repeat for any additional desired email recipients.
9. Leave the default choice to Create alert if **any** of the following are true. This will send an alert if any one of the set thresholds is met. You can also choose to alert only if **all** of the set thresholds are met.
10. To set your first threshold, select **Latency - Total** in the drop-down and set it to greater than 2 ms.
11. If desired, add additional thresholds to alert on by clicking the **Add threshold** button. When the policy is customized the way you want it, click **Save**.
12. You can also choose to **Stop processing further policies if alert is generated**. This will halt additional policy alerting if this policy's conditions are met.
13. You may add as many new policies as you like, setting alerts for other recipients based on different conditions, according to business need. Any policies configured without specific recipients will send alerts to the global recipient list set in the **Admin > Notifications** page

After you finish

Each new policy is automatically activated when it is saved, and recipients will start receiving alerts when the policy's conditions are met (known as a *violation*). You can also monitor these violations in the **Dashboards > Violations Dashboard**.

Highlighting data using dashboards

Now that you have your assets annotated and have configured performance policies to alert for violations, you can create dashboards to highlight specific data that you want to target.

About this task

In this example we will provide a high-level view of dashboard creation by creating a dashboard with a single widget highlighting VM Performance data. You can add as many widgets as you need on a single dashboard, and you can create as many dashboards as you need. Widgets can be resized and moved as desired.

More information on Dashboards and Widgets can be found throughout the OnCommand Insight documentation.

Steps

1. Log in to Insight as a user with administrative permissions.
2. From the **Dashboards** menu, select **+New dashboard**.

The New dashboard page opens.

3. Best practice: Name and save your dashboard as soon as you create it. Click the **Save** button and enter a unique name for the dashboard in the **Name** field. For example "VM Performance Dashboard". Click **Save**.
4. If necessary, slide the **Edit** switch to "On" to enable Edit mode. This allows you to begin adding widgets to your dashboard.

5. Click the **+Widgets** button and select **Table** to add a new table widget to the dashboard.

The Edit Widget dialog opens.

6. In the Name field, delete “Widget 1” and enter “Virtual Machine Performance table”.
7. Click the asset type drop-down and change **Storage** to **Virtual Machine**.

The table data changes to show all virtual machines in your environment.

8. To add additional columns to the table, click the *Columns*  button and select the desired columns, for example *Data Center*, *Storage name*, and *Tier*. You can sort the table by any of these columns.
9. You can set filters as needed to highlight the data that is important to you for this dashboard, for example, you might choose to only show only virtual machines with “Tier 1 - SSD” annotation. Click the “+” button next to **Filter by** and select *Tier*. Click on **Any** and enter “Tier 1 - SSD”. Click the  button to save the filter.

The table now shows only virtual machines in the “SSD” tier.

10. You can group results by clicking the “+” button next to **Group by** and selecting a field to group by, such as *Data Center*. Grouping is automatically applied to the table.
11. When you have customized the widget to your satisfaction, click the **Save** button.

The table widget is saved to the dashboard.

12. You can resize the widget on the dashboard by dragging the lower-right corner.
13. To add more widgets, click the **+Widget** button. Each widget is added to the dashboard when it is saved.
14. When you have made all the changes desired, be sure to click **Save** to save the dashboard.
15. You can create additional dashboards to highlight different data.

Creating custom dashboards

OnCommand Insight 7.3 includes enhanced custom dashboard capabilities to give users an operational view of the data that is important to them and provide a single-stop view of that data.

OnCommand Insight provides users the flexibility to create operational views infrastructure data across IT platforms, by allowing you to create custom dashboards with a variety of widgets, each of which provides extensive flexibility in displaying and charting your data. In this How-To, we will create an example dashboard to highlight VM Performance.

This How-To should serve as an example only and does not cover every possible scenario. The concepts and steps herein can be used to create your own custom dashboards to highlight the data specific to your particular needs.

Overview

You create a custom dashboard by either of the following methods:

- **Dashboards > +New dashboard**

- **Dashboards > Show all dashboards** and click **+Add**

The New Dashboard screen has several controls:

- **Time selector:** allows you to view dashboard data for a range of time from 3 hours up to 90 days by use of the custom date range selector. You can choose to override this global time range in individual widgets.
- **Edit** button: Selecting “On” will enable Edit mode, which allows you to make changes to the dashboard. New dashboards open in Edit mode by default.
- **Save** button: Allows you to save, rename or delete the dashboard.
- **Variable** button: Variables can be added to dashboards. Changing the variable updates all of your widgets at once. For more information on variables, see [Custom Dashboard concepts](#)
- **Widget** button, which allows you to add any number of tables, charts, or other widgets to the dashboard.

Widgets can be resized and relocated to different positions within the dashboard, to give you the best view of your data according to your current needs.

Widget types

You can choose from the following widget types:

Table widget: A table displaying data according to filters and columns you choose. Table data can be combined in groups that can be collapsed and expanded.

Line, Spline, Area, Stacked Area charts: These are time-series chart widgets on which you can display performance and other data over time.

Single value widget: A widget allowing you to display a single value that can be derived either directly from a counter or calculated using a query or expression. For example, you can display the sum of total IOPS for all storage in your environment as a single value at the top of your dashboard.

Bar chart: A chart to display top or bottom 5, 10, 20, or 50 values.

Box Plot chart: A plot of the min, max, median, and the range between lower and upper quartile of data in a single chart.

Scatter Plot chart: Plots related data as points, for example, IOPS and latency. In this example, you would quickly see assets with high latency and low IOPS.

Additionally, there are a number of legacy widgets that you can choose. In the **Widgets** drop-down, select **Show more...** to see these widgets.

Custom Dashboard concepts

Custom dashboards and widgets allow great flexibility in how data is displayed. Here are some concepts to help you get the most from your custom dashboards. Each concept is explained in greater detail in the following sections.

Variables

Variables allow you to change the data displayed in some or all widgets on a dashboard at once. By setting each widget to use a common variable, changes made in one place cause the data displayed in each widget to update automatically.

Multiple queries and/or expressions

Each time series widget (line, spline, area, or stacked area charts) can have up to five queries and/or expressions to determine what data to display, allowing you to compare different sets of data on a single chart. For example, you can have a line chart showing IOPS for both Storage and VM's, or a single chart comparing Throughput and Latency for all Storage Pools.

Rollup and Grouping

Data displayed in each widget is rolled up from the data points collected. You can choose to roll up this data in one of several ways:

- Avg: rolls up data as the average of the underlying data
- Max: rolls up data to the maximum of the underlying data
- Min: rolls up data to the minimum of the underlying data
- Sum: rolls up data as the sum of the underlying data

By default, all of the underlying data is rolled up into a single entry (All) on the chart or table. You can choose to roll up data for a specific attribute instead, such as Data Center or Tier, to distribute the underlying data into desired groups. Your widget will display data for only the attributes you select.

You can group data in a table widget according to the attribute you choose. For example, you might choose to group your table by Data Center. Groups can be expanded or collapsed at will. Performance data in a table is rolled up in the group header according to the rollup method you set in the widget (average, max, min, or sum).

Table widgets can be sorted by any column, and columns can be moved or resized as needed.

Top / Bottom

Use this to limit the result set in chart widgets, to select whether to display the top N results in your widget, or the bottom N results. You can choose this option when data is either not rolled up or is rolled up by a specific attribute.

Override dashboard time

By default, most widgets you add to a dashboard show data according to the dashboard's time range setting (3h, 24h, 3d, 7d, 30d or custom range). However, you can override this time setting in individual widgets to force them to show data in a specific time context, regardless of the dashboard's time setting.

These concepts are explained in greater detail in the following section.

Dashboard variables

Dashboard variables allow you to filter data across multiple widgets on a dashboard quickly and easily.

Before you begin

This example requires the **City** annotation (also called City attribute) to be set on multiple storage assets.

For best results, set different cities on different storages.

About this task

Variables provide a quick and simple way of filtering the data shown in some or all of the widgets on a custom dashboard. The following steps will guide you to creating widgets that use variables, and show you how to use them on your dashboard.

Steps

1. Log in to Insight as a user with administrative permissions
2. Click on **Dashboards > +New Dashboard**.
3. Before adding widgets, we want to define the variables we will use to filter the dashboard data. Click on the **Variable** button.

The list of attributes is displayed.

4. Let's say we want to set the dashboard to filter based on City. Select the **City** attribute from the list.

The \$city variable field is created and added to the dashboard.

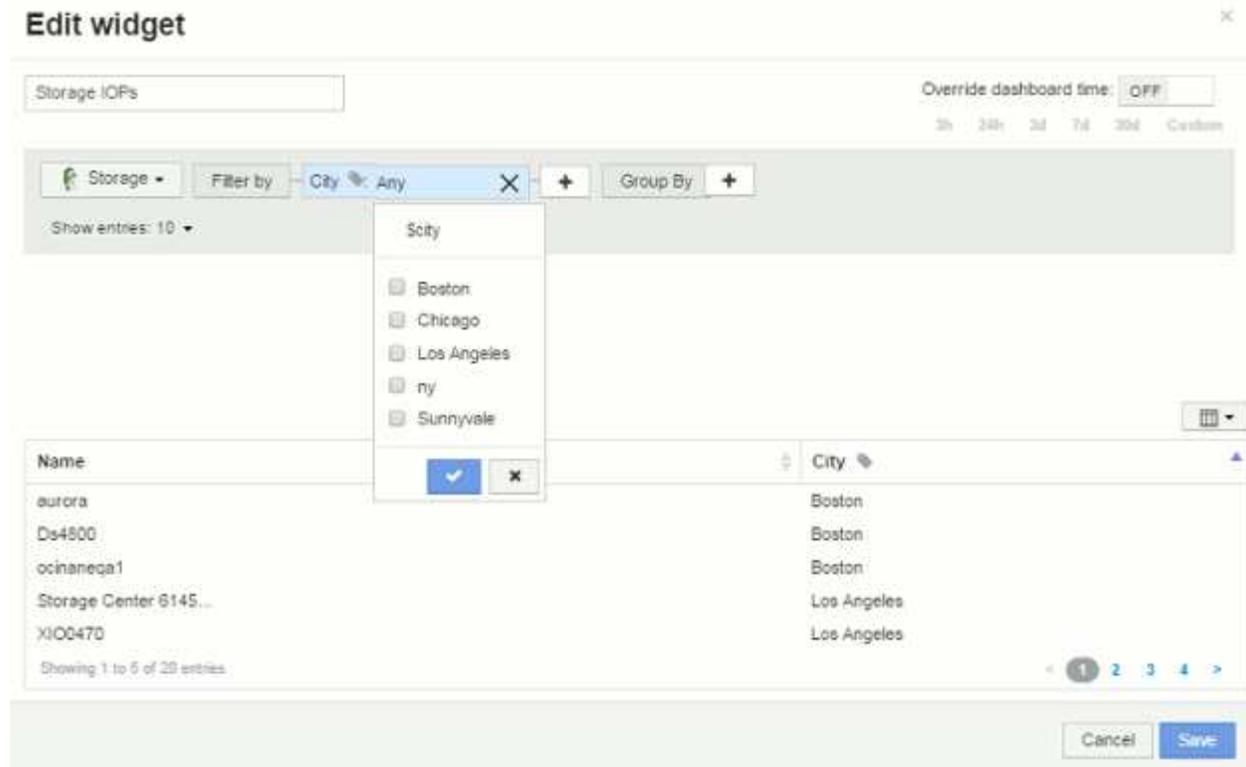
5. Next, we must tell our widgets to use this variable. The simplest way to illustrate this is to add a table widget showing the City column. Click on the **Widget** button and select the **Table** widget.

6. First, add the City field to the table by selecting it from the column picker  button.

City is a list-type attribute, so it contains a list of previously-defined choices. You may also choose text, boolean, or date-type attributes.

7. Next, click the **Filter by +** button and choose **City**.
8. Click **Any** to view the possible filter choices for City. Notice that the list now includes "\$city" at the top, in addition to any previously-available choices. Select "\$city" to use this dashboard variable.

The "\$city" choice only appears here if it was defined previously on the main dashboard page. If the variable was not previously defined, only the existing choices for the filter will be shown. Only variables that are applicable to the selected attribute type will be displayed in the drop-down for that filter.



9. **Save** the widget.

10. On the dashboard page, click on **Any** next to the \$city variable, and select the city or cities you want to see.

Your table widget updates to show only the cities you selected. You can change the values in the \$city variable at will, and all widgets on your dashboard that are set to use the \$city variable will refresh automatically to show only data for the values you selected.

11. Be sure to **Save** your dashboard when you have it configured as you want it.

More on dashboard variables

Dashboard variables come in several types, can be used across different fields, and must follow rules for naming. These concepts are explained here.

Variable types

A variable can be one the following types:

Text: Alphanumeric string. This is the default variable type.

Numerical: a number or range of numbers.

Boolean: Use for fields with values of True/False, Yes/No, 0/1, etc. For the boolean variable, the choices are *Yes*, *No*, *None*, *Any*.

Date: A date or range of dates.

“Generic” variables

You can set a generic or universal variable by clicking the **Variable** button and selecting one of the types listed above. These types are always shown at the top of the drop-down list. The variable is given a default name, for example “\$var1”, and is not tied to a specific annotation or attribute.

Configuring a generic variable allows you to use that variable in widgets to filter for *any* field of that type. For example, if you have a table widget showing *Name*, *Alias*, and *Vendor* (which are all text-type attributes), and “\$var1” is a text-type variable, you can set filters for each of those fields in the widget to use the \$var1 variable. You can set other widgets to use \$var1 for those or any text fields.

On your dashboard page, setting \$var1 to a value (for example “NetApp”) will filter *all* of those fields in *all* widgets that are set to use that variable. In this way, you can update multiple widgets at once to highlight dashboard data you choose at will.

Because generic variables can be used for any field of that type, you can change the name of a generic variable without changing its functionality.



All variables are treated as "generic" variables, even those you create for a specific attribute, because all configured variables of a type are shown when you set a filter for any attributes or annotations of that type. However, best practice is to create a generic variable when you will use it to filter for a value across multiple fields, as in the *Name/Alias/Vendor* example above.

Variable naming

Variables names:

- Must always be prefixed with a “\$”. This is added automatically when you configure a variable.
- Cannot contain any special characters; only the letters a-z and the digits 0-9 are allowed.
- Cannot be longer than 20 characters, including the “\$” symbol.
- Are not case-sensitive: \$CityName and \$cityname are the same variable.
- Cannot be the same as an existing variable name.
- Cannot be only the “\$” symbol.

Widgets that use variables

Variables can be used with the following widgets:

- Area Chart
- Bar Chart
- Box Plot Chart
- Line Chart
- Scatter Plot Chart
- Single Value Widget
- Spline Chart
- Stacked Area Chart
- Table Widget

Displaying widget legends

Widgets in dashboards can be displayed with or without legends.

Legends in widgets can be turned on or off on a dashboard by either of two methods:

1. When creating or editing the widget itself, check the Legends checkbox and save the widget.
2. With the dashboard in Edit mode, click the Options button on the widget and check the Legends checkbox in the menu.

As you edit and change the data displayed in the widget, the legend for that widget is updated dynamically.

When legends are displayed, if the landing page of the asset indicated by the legend can be navigated to, the legend will display as a link to that asset page.

Dashboard widget queries and filters

The Query in a dashboard widget is a powerful tool for managing the display of your data. Here are some things to note about widget queries.

Some widgets can have up to five queries. Each query will plot its own set of lines or graphs in the widget. Setting rollup, grouping, top/bottom results, etc. on one query does not affect any other queries for the widget.

You can click on the eye icon to temporarily hide a query. The widget display updates automatically when you hide or show a query. This allows you to check your displayed data for individual queries as you build your widget.

The following widget types can have multiple queries:

- Area chart
- Stacked area chart
- Line chart
- Spline chart
- Single value widget

The remaining widget types can have only a single query:

- Table
- Bar chart
- Box plot
- Scatter plot

Filtering in dashboard queries

You can filter using any of the following to refine your search in any **text field** in the query:

- An asterisk enables you to search for everything. For example, `vol*rhel` displays all resources that start with “vol” and end with “rhel”.
- The question mark enables you to search for a specific number of characters. For example, `BOS-PRD??-S12` displays BOS-PRD12-S12, BOS-PRD13-S12, and so on.

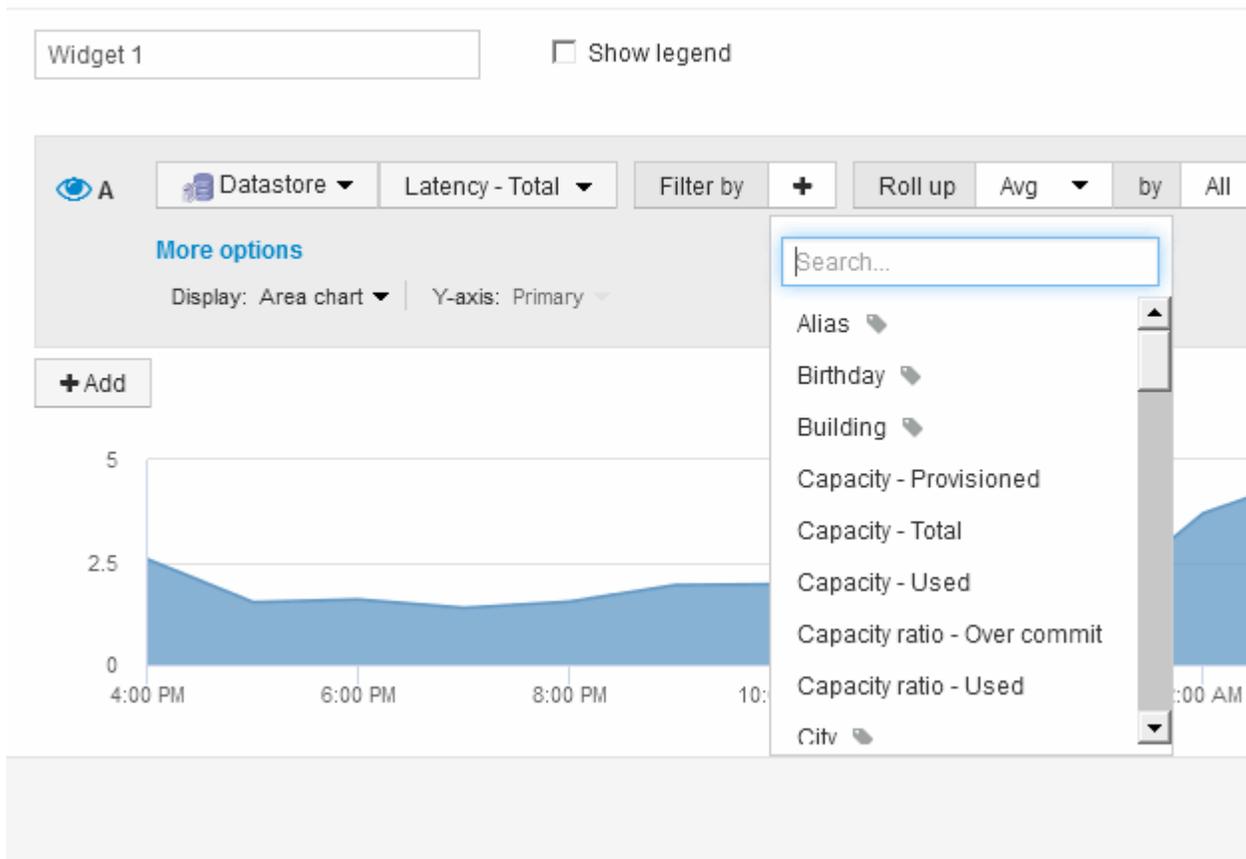
- The OR operator enables you to specify multiple entities. For example, `FAS2240 OR CX600 OR FAS3270` finds multiple storage models.
- The NOT operator allows you to exclude text from the search results. For example, `NOT EMC*` finds everything that does not start with “EMC”. You can use `NOT *` to display fields that contain null values.

If you enclose a filter string in double quotes, Insight treats everything between the first and last quote as an exact match. Any special characters or operators inside the quotes will be treated as literals. For example, filtering for “*” will return results that are a literal asterisk; the asterisk will not be treated as a wildcard in this case. The operators AND, OR, and NOT will also be treated as literal strings when enclosed in double quotes.

Identifying objects returned by queries and filters

The objects returned by queries and filters look similar to those shown in the following illustration. Objects with 'tags' assigned to them are annotations while the objects without tags are performance counters or object attributes.

Edit widget



Roll up and Aggregation

Data displayed in dashboard widgets is rolled up from acquired data points, allowing flexibility and conciseness in your dashboards.

Data displayed in each widget is rolled up from the underlying data points collected during acquisition. For example, if you have a line chart widget showing Storage IOPS over time, you might want to see a separate line for each of your data centers, for a quick comparison. You can choose to roll up this data in one of several ways:

- **Avg**: displays each line as the *average* of the underlying data.
- **Max**: displays each line as the *maximum* of the underlying data.
- **Min**: displays each line as the *minimum* of the underlying data.
- **Sum**: displays each line as the *sum* of the underlying data.

To do this, in your widget's query, first choose an asset type (for example, *Storage*) and metric (such as *IOPS - Total*). For **Roll up**, choose a rollup method (such as *Avg*) and select an attribute or annotation by which to roll up the data (for example, *Data Center*). The widget updates automatically and shows a line for each of your data centers.

You can also choose to roll up *all* of the underlying data into the chart or table. In this case, you will get a single line for each query in the widget, which will show the average, min, max or sum of the chosen metric for all of the underlying assets.

If you have set a filter for the query, the data that is rolled up is based on the filtered data.

Note that when you choose to roll up a widget by any field (for example, *Model*), you will still need to **Filter by** that field in order to properly display the data for that field properly on the chart or table.

Aggregating data: You can further align your time-series charts (line, area, etc.) by aggregating data points into minute, hour, or day buckets before that data is subsequently rolled up by attribute (if chosen). You can choose to aggregate data points according to their Avg, Max, Min, or Sum, or by the Last data point collected during the chosen interval. To choose an aggregation method, click on **More options** in the widget's query section.

The minimum allowed interval is ten minutes. A small interval combined with a long time range may result in a "Aggregation interval resulted in too many data points." warning. You might see this if you have a small interval and increase the dashboard time frame to 7 days. In this case, Insight will temporarily increase the aggregation interval to 1 hour until you select a smaller time frame.

You can also aggregate data in bar chart widget and single-value widget.

Most asset counters aggregate to *Avg* by default. Some counters aggregate to *Max*, *Min*, or *Sum* by default. For example, port errors aggregate to *Sum* by default, where storage IOPS aggregate to *Avg*.

Showing top/bottom results in dashboard widgets

In a chart widget on a custom dashboard, you can show either the Top or Bottom results for rolled up data, and choose the number of results shown. In a table widget, you can select the number of rows displayed and sort by any column.

Chart widget top/bottom

In a chart widget, when you choose to rollup data by a specific attribute, you have the option of viewing either the top N or bottom N results. Note that you cannot choose the top or bottom results when you choose to rollup by *all* attributes.

You can choose which results to display by choosing either **Top** or **Bottom** in the query's **Show** field, and selecting a value from the list provided.

Table widget show entries

In a table widget, you can select the number of results shown in the table results. You can choose from 5, 10,

20, or 50 results. You are not given the option to choose top or bottom results because the table allows you to sort ascending or descending by any column on demand.

You can choose the number of results to show in the table on the dashboard by selecting a value from the query's **Show entries** field.

Note that the more results you choose to display, the taller your widget will be when you save it to the dashboard. You will not be able to resize the widget smaller than the number of rows displayed.

Grouping in table widgets

Data in a table widget can be grouped by any available attribute, allowing you to see an overview of your data, and to drill-down into it for more detail. Metrics in the table are rolled up for easy viewing in each collapsed row.

Table widgets allow you to group your data based on the attributes you set. For example, you might want your table to show Total Storage IOPS grouped by the data centers in which those storages live. Or you might want to display a table of Virtual machines grouped according to the hypervisor that hosts them. From the list, you can expand each group to view the assets in that group.

Grouping is only available in the **Table** widget type.

Performance data roll up

If you include a column for performance data (for example, *IOPS - Total*) in a table widget, when you choose to group the data you can then choose a roll up method for that column. The default roll up method is to display the *average* of the underlying data in the group row. You can also choose to display the *sum*, *minimum*, or *maximum* of the data..

Grouping example (with rollup explained)

Table widgets allow you to group data for easier display.

About this task

In this example, we will create a table widget showing all VMs grouped by Data Center.

Steps

1. Create or open a dashboard, and add a **Table** widget.
2. Select **Virtual Machine** as the asset type for this widget.
3. Click on the Column Selector  and choose *Hypervisor name* and *IOPS - Total*.

Those columns are now displayed in the table.

4. Let's disregard any VM's with no IOPS, and include only VMs that have total IOPS greater than 1. Click the **Filter by +** button and select **IOPS - Total**. Click on **Any**, and in the **from** field, type 1. Leave the **to** field empty. Click the check button to apply the filter.

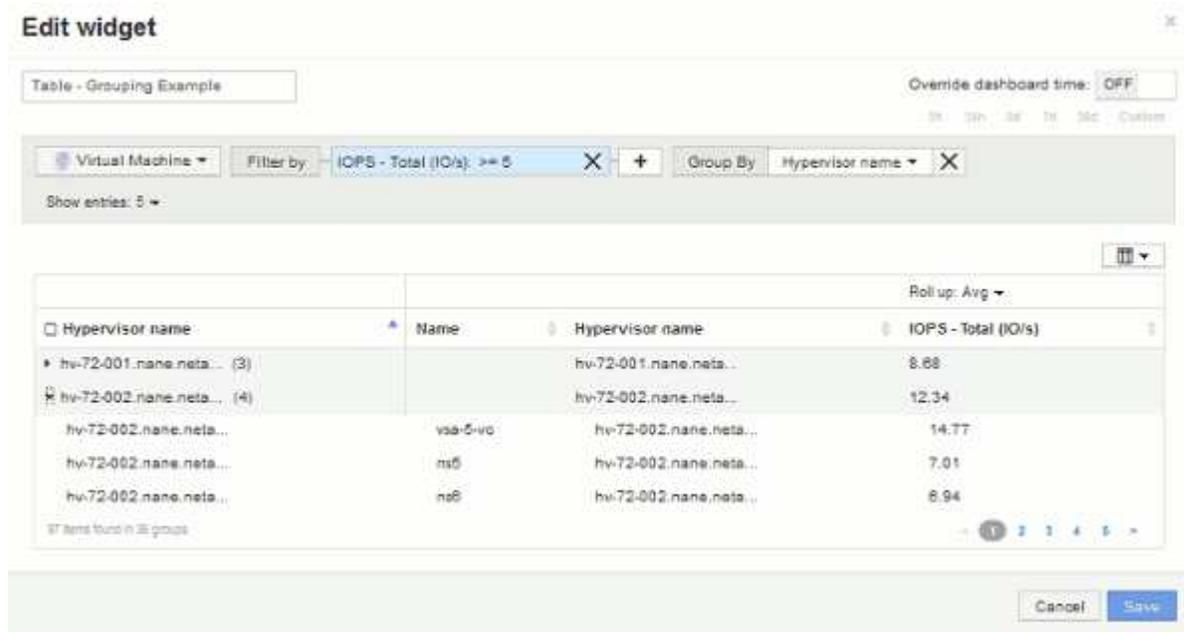
The table now shows all VMs with Total IOPS greater than or equal to 1. Notice that there is no grouping in the table. All VMs are shown.

5. Click the **Group by +** button.

Because **All** is selected as the grouping method by default, all VMs are moved into a single group named “All”.

- Above the *IOPS - Total* column is now a **Roll up** option. The default roll up method is *Avg*. This means that the number shown for the group is the average of all the Total IOPS reported for each VM inside the group. You can choose to roll this column up by *Avg*, *Sum*, *Min* or *Max*. Each column that you display that contains performance metrics can be rolled up individually.
- Click **All** and select **Hypervisor name**.

The VM list is now grouped by Hypervisor. You can expand each hypervisor to view the VMs hosted by it.



- Click **Save** to save the table to the dashboard. You can resize the widget.
- Click **Save** to save the dashboard.

Overriding dashboard time for individual widgets

You can override the main dashboard time frame setting in individual widgets. These widgets will display data based on their set time frame, not the dashboard time frame.

To override the dashboard time and force a widget to use a its own time frame, in the widget’s edit mode set the **Override dashboard time** to **On**, and select a time frame for the widget. **Save** the widget to the dashboard.

The widget will display its data according to the time frame set, regardless of the time frame you select on the dashboard itself.

The time frame you set for one widget will not affect any other widgets on the dashboard.

Primary and Secondary axis explained

The secondary axis makes it easier to view data from two different sets of values that use different units of measurement.

About this task

Different metrics use different units of measurements for the data they report in a chart. For example, when looking at IOPS, the unit of measurement is the number of I/O operations per second of time (IO/s), while Latency is purely a measure of time (milliseconds, microseconds, seconds, etc.). When charting both metrics on a single line chart using a single set of values for the Y-Axis, the latency numbers (typically a handful of milliseconds) are charted on the same scale with the IOPS (typically numbering in the thousands), and the latency line gets lost at that scale.

But it is possible to chart both sets of data on a single meaningful graph, by setting one unit of measurement on the primary (left-side) Y-axis, and the other unit of measurement on the secondary (right-side) Y-axis. Each metric is charted at its own scale.

Steps

1. Create or open a dashboard. Add a **line chart**, **spline chart**, **area chart** or **stacked area chart** widget to the dashboard.
2. Select an asset type (for example **Storage**) and choose **IOPS - Total** for your first metric. Set any filters you like, and choose a roll-up method if desired.

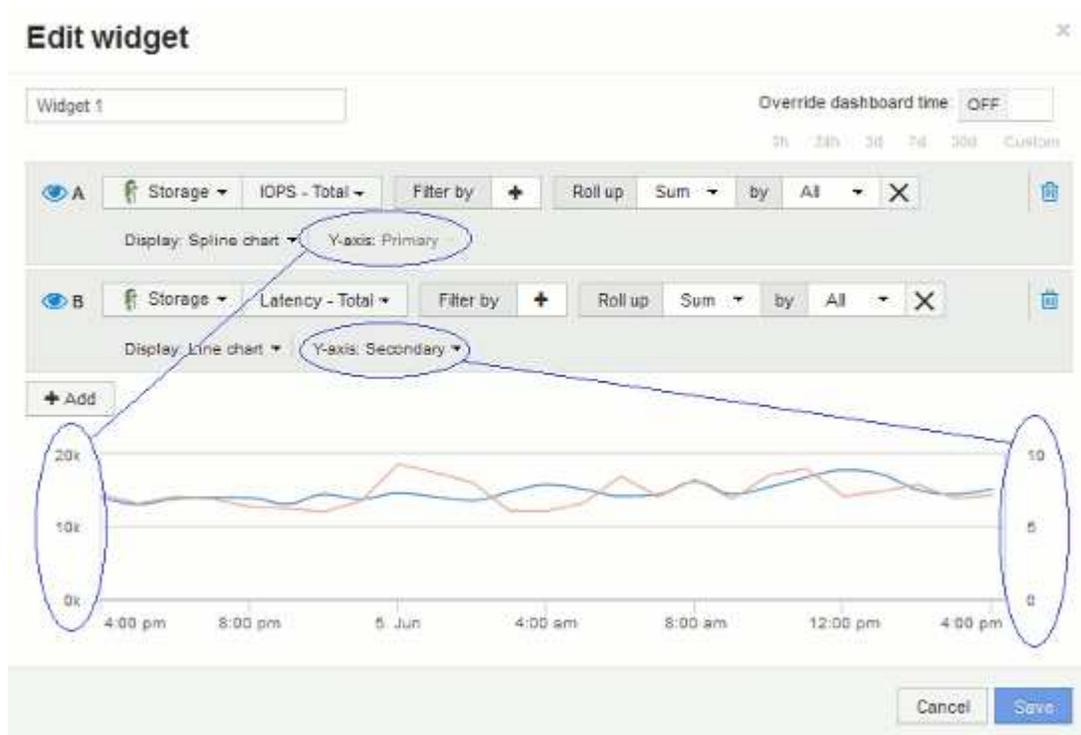
The IOPS line is displayed on the chart, with its scale shown on the left.

3. Click **+Add** to add a second line to the chart. For this line, choose **Latency - Total** for the metric.

Notice that the line is displayed flat at the bottom of the chart. This is because it is being drawn at the same scale as the IOPS line.

4. In the Latency query, select **Y-Axis: Secondary**.

The Latency line is now drawn at its own scale, which is displayed on the right side of the chart.



Expressions in dashboard widgets

Expressions in time series widgets allow you to display data based on calculations with metrics of your choosing.

In a dashboard, any time series widget (line, spline, area, stacked area) allows you to build expressions from metrics you choose, and show the result of those expressions in a single graph. The following examples use expressions to solve specific problems. In the first example, we want to show Read IOPS as a percentage of Total IOPS for all storage assets in our environment. The second example gives us visibility into the "system" or "overhead" IOPS that occur in our environment—those IOPS that are not from reading or writing data.

Expressions example: Read IOPS percentage

Using expressions, you can view metrics by alternate means, such as percentage of total.

About this task

In this example we want to show Read IOPS as a percentage of Total IOPS. You can think of this as the following formula:

- Read Percentage = (Read IOPS / Total IOPS) x 100

This data can be shown in a line graph on your dashboard. To do this, follow these steps:

Steps

1. Create a new dashboard, or open an existing dashboard in **edit mode**.
2. Add a widget to the dashboard. Choose **Area chart**.

The widget opens in edit mode. By default, a query is displayed showing **IOPS - Total** for **Storage** assets. If desired, select a different asset type.

3. Click the **Convert to Expression** button.

The current query is converted to Expression mode. Notice that you cannot change the asset type while in Expression mode. Notice that while you are in Expression mode, the button changes to **Revert to Query**. Click this if you wish to switch back to Query mode at any time. Be aware that switching between modes will reset fields to their defaults.

For now, stay in **Expression** mode.

4. The **IOPS - Total** metric is now in the alphabetic variable field "a". In the "b" variable field, click **Select** and choose **IOPS - Read**.

You can add up to a total of five alphabetic variables for your expression by clicking the **+** button following the variable fields. For our Read Percentage example, we only need Total IOPS ("a") and Read IOPS ("b").

5. In the **Expression** field, you use the letters corresponding to each variable to build your expression. We know that *Read Percentage = (Read IOPS / Total IOPS) x 100*, so we would write this expression as: (b / a) * 100
6. The **Label** field identifies the expression. Change the label to "Read Percentage", or something equally meaningful for you.
7. Change the **Units** field to "%" or "Percent".

The chart displays the IOPS Read percentage over time for the chosen storage devices. If desired, you can set a filter, or choose a different rollup method. Be aware that if you select **Sum** as the rollup method, all percentage values are added together, which potentially may go higher than 100%.

8. Click **Save** to save the chart to your dashboard.

You can also use expressions in **Line chart**, **Spline chart**, or **Stacked Area chart** widgets.

Expressions example: "System" I/O

Expressions give you the freedom to chart data that can be calculated from other metrics.

About this task

Example 2: OnCommand Insight acquires many metrics from data sources. Among them are read, write, and total IOPS. However, the total number of IOPS reported by acquisition sometimes includes "system" IOPS, which are those IO operations that are not a direct part of data reading or writing. This system I/O can also be thought of as "overhead" I/O, necessary for proper system operation but not directly related to data operations.

To show these system I/Os, you can subtract read and write IOPS from the total IOPS reported from acquisition. The formula might look like this:

- System IOPS = Total IOPS - (Read IOPS + Write IOPS)

This data can then be shown in a line graph on your dashboard. To do this, follow these steps:

Steps

1. Create a new dashboard, or open an existing dashboard in **edit mode**.
2. Add a widget to the dashboard. Choose **Line chart**.

The widget opens in edit mode. By default, a query is displayed showing **IOPS - Total** for **Storage** assets. If desired, select a different asset type.

3. Click the button to create a copy of the query.

A duplicate of the query is added below the original.

4. In the second query, click the **Convert to Expression** button.

The current query is converted to Expression mode. Click **Revert to Query** if you wish to switch back to Query mode at any time. Be aware that switching between modes will reset fields to their defaults.

For now, stay in **Expression** mode.

5. The **IOPS - Total** metric is now in the alphabetic variable field "a". Click on **IOPS - Total** and change it to **IOPS - Read**.
6. In the "b" variable field, click **Select** and choose **IOPS - Write**.
7. In the **Expression** field, you use the letters corresponding to each variable to build your expression. We would write our expression simply as: a + b. In the **Display** section, choose **Area chart** for this expression.
8. The **Label** field identifies the expression. Change the label to "System IOPS", or something equally meaningful for you.

The chart displays the total IOPS as a line chart, with an area chart showing the combination of read and write IOPS below that. The gap between the two shows the IOPS that are not directly related to data read or write operations.

9. Click **Save** to save the chart to your dashboard.

Custom Dashboard: Virtual Machine Performance

OnCommand Insight's custom dashboards and widgets help provide operational views into inventory and performance trends.

About this task

There are many challenges facing IT operations today. Administrators are being asked to do more with less, and having full visibility into your dynamic data centers is a must. In this example, we will show you how to create a custom dashboard with widgets that give you operational insights into the virtual machine performance in your environment. By following this example, and creating widgets to target your own specific needs, you will be able to visualize backend storage performance compared to frontend virtual machine (VM) performance, or view VM latency versus I/O demand.

Custom dashboards allow you to prioritize efforts and identify resource availability. You can respond to the ebb and flow of workloads and minimize the time to detect and remediate emerging issues. Custom dashboards allow you the flexibility to create prioritized views into business-critical infrastructure, and are useful for identifying performance availability across multi-vendor technologies.

Here we will create a Virtual Machine Performance dashboard containing the following:

- a table listing VM names and performance data
- a chart comparing VM Latency to Storage Latency
- a chart showing Read, Write and Total IOPS for VMs
- a chart showing Max Throughput for your VMs

This is just a basic example. You can customize your dashboard to highlight and compare any performance data you choose to target for your own operational best practices.

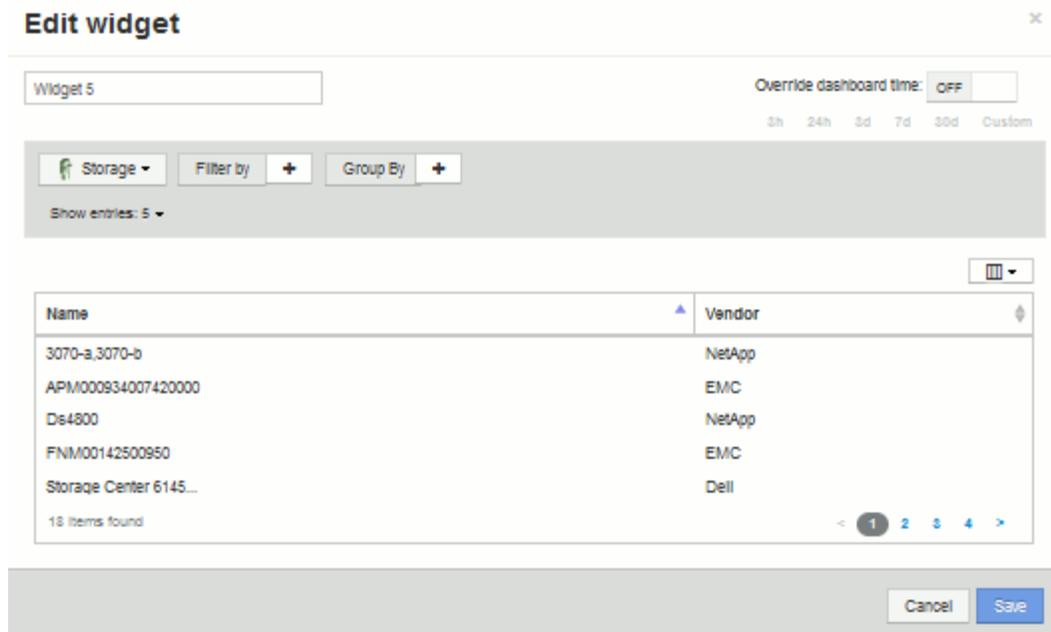
Steps

1. Log in to Insight as a user with administrative permissions.
2. From the **Dashboards** menu, select **+New dashboard**.

The New dashboard page opens.

3. Let's give our dashboard a meaningful name. Click **Save**. In the **Name** field, enter a unique name for the dashboard, for example "VM Performance by Application".
4. Click **Save** to save the dashboard with the new name.
5. Let's start adding our widgets. If necessary, slide the **Edit** switch to "On" to enable Edit mode.
6. Click the **Widget** button and select **Table widget** to add a new table widget to the dashboard.

The Edit Widget dialog opens. The default name is "Widget 1" and the default data displayed is for all storages in your environment.



7. We can customize this widget. In the Name field, delete “Widget 1” and enter “Virtual Machine Performance table”.
8. Click the asset type drop-down and change **Storage** to **Virtual Machine**.

The table data changes to show all virtual machines in your environment. For now, the table only shows the VM names. Let’s add a few columns to the table.

9. Click the *Columns*  button and select *Data Center*, *Storage name*, and *IOPS - Total*. You can also try typing the name into the search to quickly display the desired field(s).

These columns are now displayed in the table. You can sort the table by any of these columns. Note that the columns are displayed in the order in which they were added to the widget.

10. For this exercise we will exclude VMs that are not actively in use, so let’s filter out anything with less than 10 total IOPS. Click the “+” button next to **Filter by** and select *IOPS - Total (IO/s)*. Click on **Any** and enter “10” in the **from** field. Leave the **to** field empty. Click the  button to save the filter.

The table now shows only VMs with 10 or more total IOPS.

11. We can further collapse the table by grouping results. Click the “+” button next to **Group by** and select a field to group by, such as Application or Cluster. Grouping is automatically applied.

The table rows are now grouped according to your setting. You can expand and collapse the groups as needed. Grouped rows show rolled up data for each of the columns. Some columns allow you to choose the roll up method for that column.

Edit widget

Virtual Machine Performance Table

Override dashboard time: OFF

Virtual Machine Filter by: IOPS - Total (IO/s) >= 10 Group By: Application

Show entries: 5

Application	Name	Data Center	Storage name	IOPS - Total (IO/s)
N/A (1)	MAP aomshen Dev Rb...	NANE	vifsane05,vifsana...	15.75
vm2 app (57)		NANE		85.29
applicationT7 (53)		NANE		84.22
Application T7 (lar... (52)		NANE		85.60
application T11 (45)		NANE		80.10

218 items found in 2 groups

Cancel Save

12. When you have customized the table widget to your satisfaction, click the **Save** button.

The table widget is saved to the dashboard.

13. You can resize the widget on the dashboard by dragging the lower-right corner. Make the widget wider to show all the columns clearly. Click **Save** to save the current dashboard.

14. Next we will add some charts to show our VM Performance. Let's create a line chart comparing VM latency with Storage latency.

15. If necessary, slide the **Edit** switch to "On" to enable Edit mode.

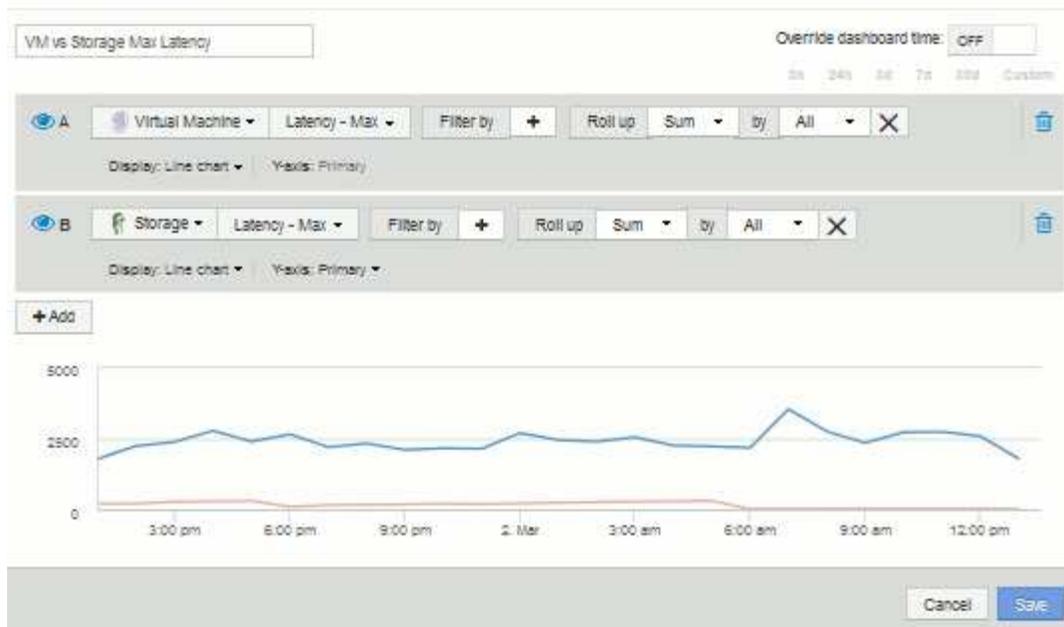
16. Click the **Widget** button and select **Line Chart** to add a new line chart widget to the dashboard.

The Edit Widget dialog opens. Click the **Name** field and name this widget "VM vs Storage Max Latency"

17. Select **Virtual Machine** and choose **Latency - Max**. Set any filters you wish, or leave **Filter by** empty. For **Roll up**, choose "Sum" by "All". Display this data as a **Line Chart**, and leave Y-Axis as **Primary**.

18. Click the **+Add** button to add a second data line. For this line, select **Storage** and **Latency - Max**. Set any filters you wish, or leave **Filter by** empty. For **Roll up**, choose "Sum" by "All". Display this data as a **Line Chart**, and leave Y-Axis as **Primary**.

Edit widget



19. Click **Save** to add this widget to the dashboard.
20. Next we will add a chart showing VM Read, Write and Total IOPS in a single chart.
21. Click the **Widget** button and select **Area Chart** to add a new area chart widget to the dashboard.

The Edit Widget dialog opens. Click the **Name** field and name this widget "VM IOPS"

22. Select **Virtual Machine** and choose **IOPS - Total**. Set any filters you wish, or leave **Filter by** empty. for **Roll up**, choose "Sum" by "All". Display this data as a**Area Chart**, and leave Y-Axis as **Primary**.
23. Click the +Add button to add a second data line. For this line, select **Virtual Machine** and choose **IOPS - Read**. Leave Y-Axis as **Primary**.
24. Click the +Add button to add a third data line. For this line, select **Virtual Machine** and choose **IOPS - Write**. Leave Y-Axis as **Primary**.

Edit widget

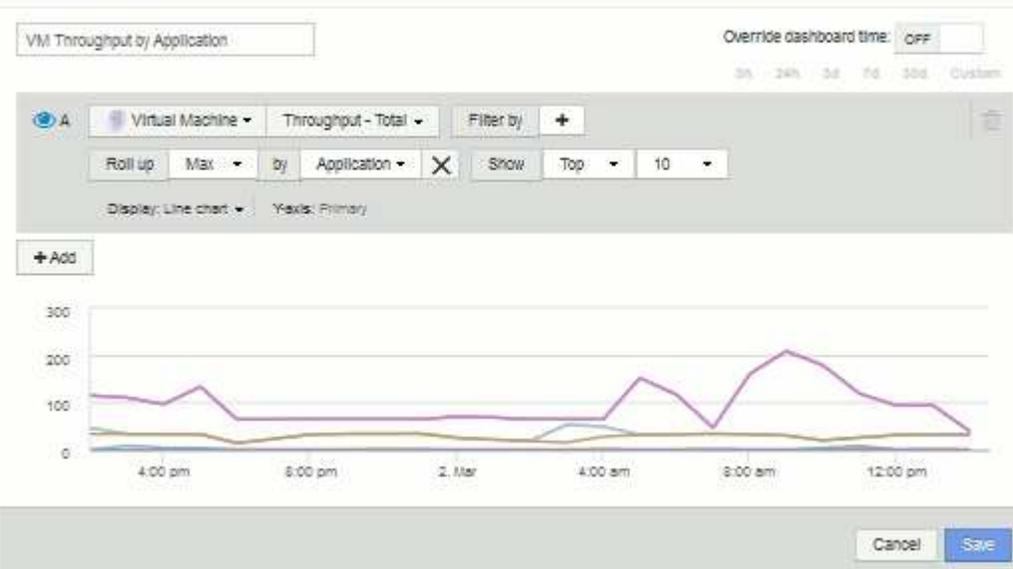


25. Click **Save** to add this widget to the dashboard.
26. Next we will add a chart showing VM Throughput for each Application associated with the VM. We will use the Roll Up feature for this.
27. Click the **Widget** button and select **Line Chart** to add a new line chart widget to the dashboard.

The Edit Widget dialog opens. Click the **Name** field and name this widget “VM Throughput by Application”

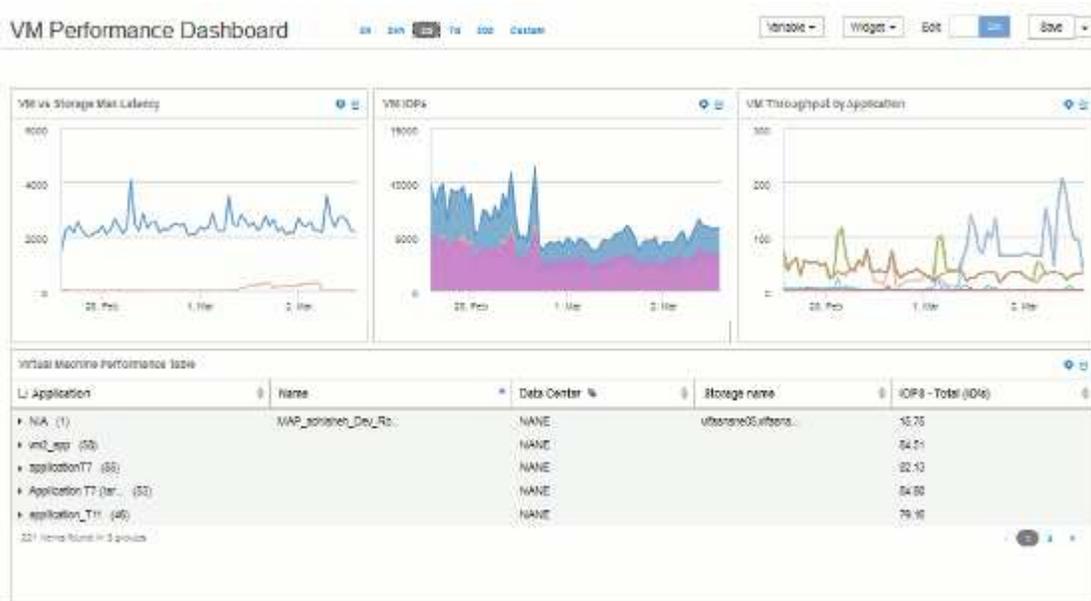
28. Select **Virtual Machine** and choose **Throughput - Total**. Set any filters you wish, or leave **Filter by** empty. For **Roll up**, choose “Max” and select by “Application” or “Name”. Show the **Top 10** applications. Display this data as a **Line Chart**, and leave Y-Axis as **Primary**.

Edit widget



29. Click **Save** to add this widget to the dashboard.
30. You can move widgets by holding down the mouse button anywhere in the top of the widget and dragging to a new location. You can resize widgets by dragging the lower-right corner. Be sure to **Save** the dashboard after you make your changes.

Your final VM Performance Dashboard will look like this:



Example storage node utilization dashboard with variables

Create a custom dashboard for Storage Analysis which has variables for storage, storage pool, node, tier, utilization and latency.

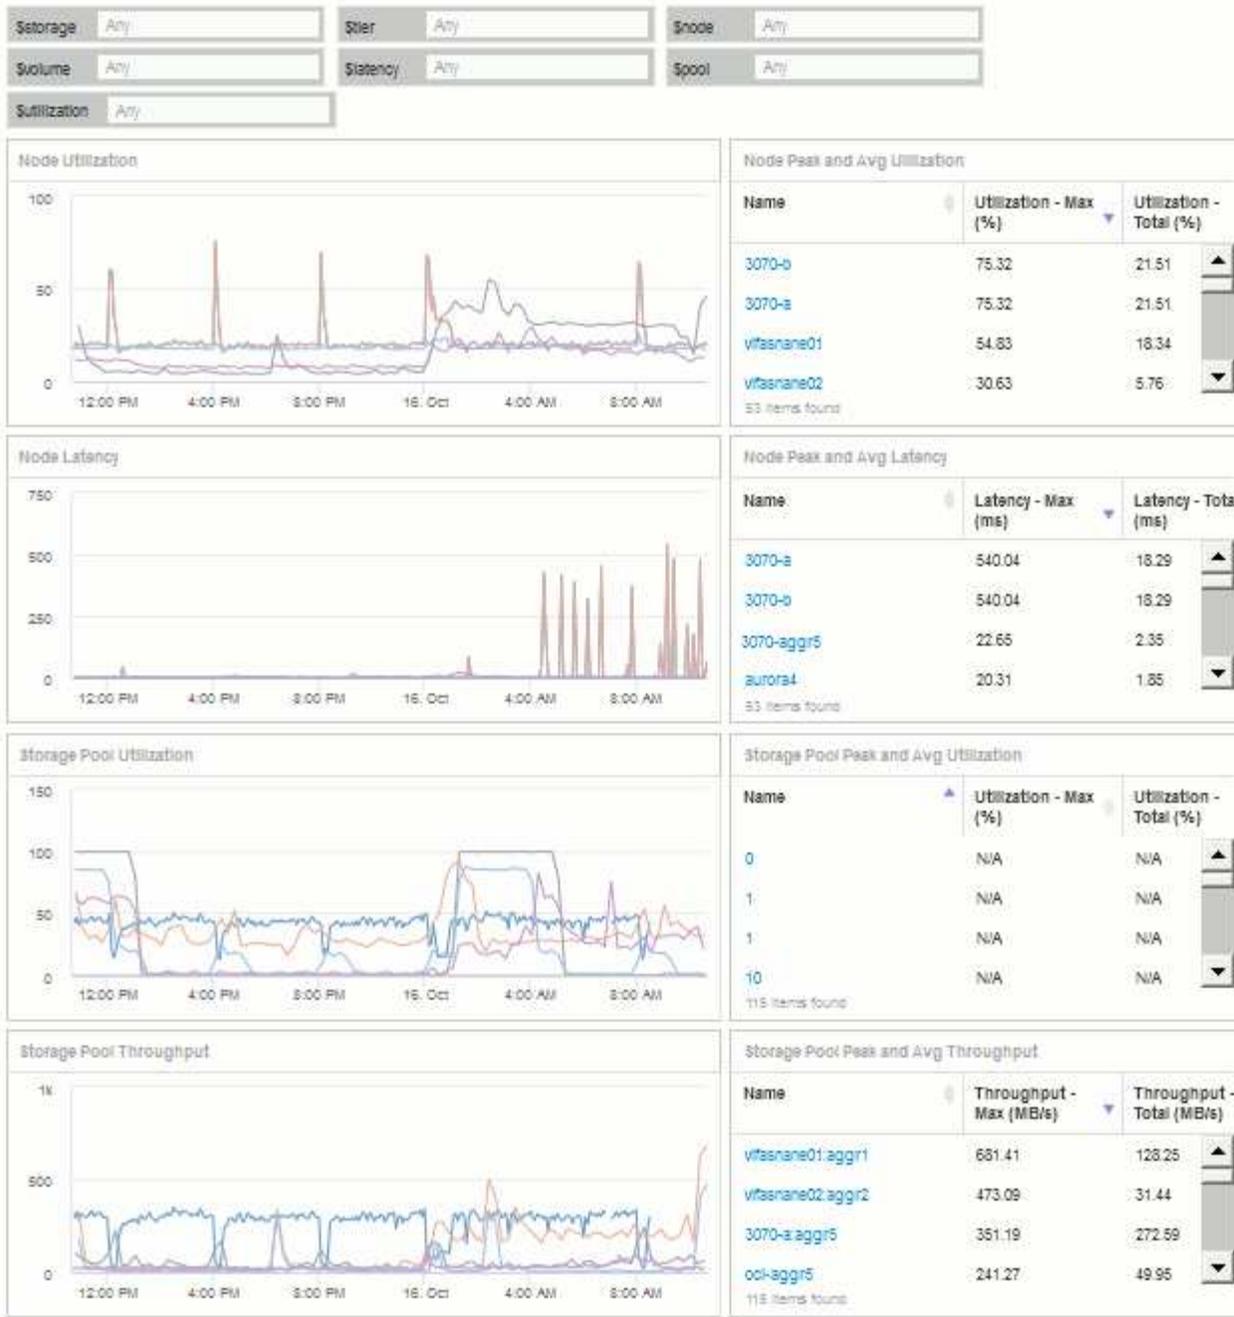
Before you begin

Familiarity with dashboards in Insight is recommended but not required.

About this task

The following procedure will create a custom Storage Analysis Overview dashboard which uses variables for storage, storage pool, node, tier, utilization and latency. Variables in the example below will be used to filter the displayed assets or metrics across one or many widgets available on the dashboard. The widgets that use these variables as filters will be updated with filtered content on-demand according to the values entered in the variable fields on the dashboard, allowing you to quickly filter multiple charts and graphs to drill down to a specific area of interest.

By following the steps in this example, you will create a dashboard like the following. You can change these widgets, or add any number of additional widgets, to highlight any data you choose.



Steps

1. Create a new dashboard, and name it "Analysis: Storage Overview", or something equally descriptive.
2. Click on the **Variable** drop-down and select **Text** variable type. By default, the variable is named `$var1`. Click on `$var1` to edit the name, and change it to `$storage`, then click the check mark to save the variable. Repeat to create text variables for `$node`, `$pool`, and `$volume`.
3. Repeat the above process to create **Number**-type variables named `$utilization` and `$latency`.
4. Click on the **Variable** drop-down and search for the `Tier` annotation. Select it to create a variable named `$tier`.

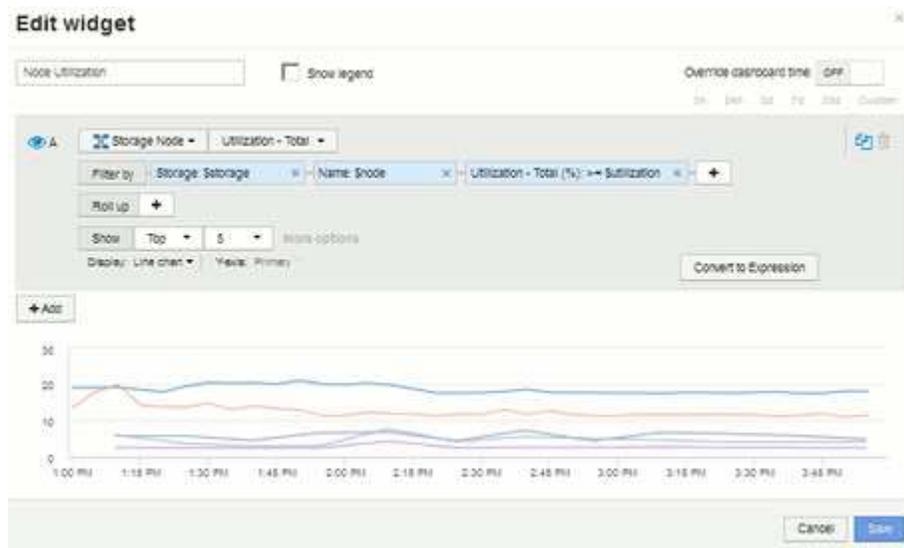
You can add variables at any time, however it is easier to create them up front and therefore make them available to all widgets as you create them.

5. Add a widget by clicking on the **Widget** drop-down and selecting either a **line chart** or **area chart** widget. Name the widget “Node Utilization”. Click on the **Storage** asset type and change it to **Storage Node**. Select **Utilization - Total** for the chart data.
6. Click on the **Filter by +** button to add a filter. Search for and select **Storage**, then click on **Any** and select the $\$storage$ variable.
7. Click the **+button** to add another filter for **Name**. Set the variable to $\$node$.

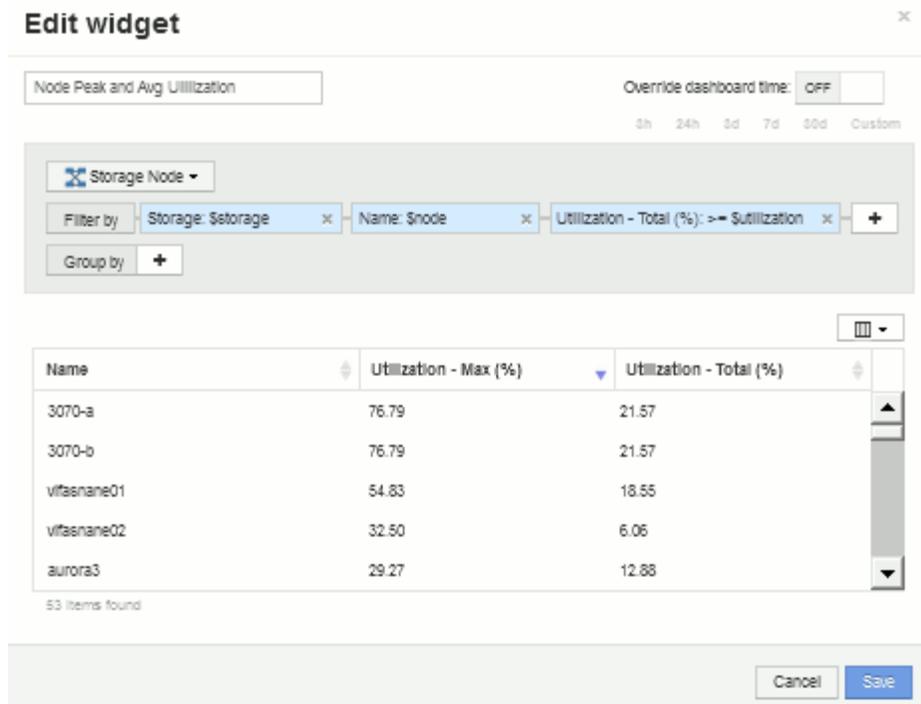
Different variables can be assigned to the annotation name filter. Use the name/variable pair at the lowest level depending on the object in the widget. For example:

- You can assign the $\$node$ variable to the **Name** filter for a Node-focused widget.
 - You can assign the $\$pool$ variable to the **Name** filter for a Pool-focused widget.
8. Click the **+button** to add another filter for **Utilization - Total (%)**. Set the variable to $\geq \$utilization$.
 9. Click the **X** after the **Roll up** field to collapse the field.
 10. Select **Show Top 5** and click **Save** to save the widget and return to your Dashboard.

Your widget should look something like this:



11. Add another line or area chart widget to your dashboard. Select **Storage Node** as the asset type and **Latency - Total** as the metric to chart.
12. Click on the **Filter by +** button to add filters for **Storage: \$storage** and **Name: \$node**.
13. Add a filter for **Latency - Total** and select the $\$latency$ variable.
14. Name the widget “Node Latency” and save it.
15. You can add supporting tables to show more details for the charts you created, for example, Max or Avg node utilization. Add a **Table widget** to the dashboard and select **Storage Node** as the asset type, and create filters for **Storage: \$storage**, **Name: \$node**, and **Utilization - Total: \$utilization**.
16. Add columns to the table for **Utilization - Max**, **Utilization - Total**, or any other desired columns.
17. Name the widget “Node Peak and Avg Utilization” and save it.



- Repeat the steps to create a table for Node Latency, showing **Latency - Max**, **Latency - Total**, or other columns as desired.
- To round out your dashboard, you might add additional chart and table widgets for some or all of the following:

Chart	Table
Storage Pool Utilization	Storage Pool Peak and Avg Utilization
Storage Pool Throughput	Storage Pool Peak and Avg Throughput
Volume Latency	Volume Peak and Avg Latency
Volume IOPS	Volume Peak and Avg IOPS

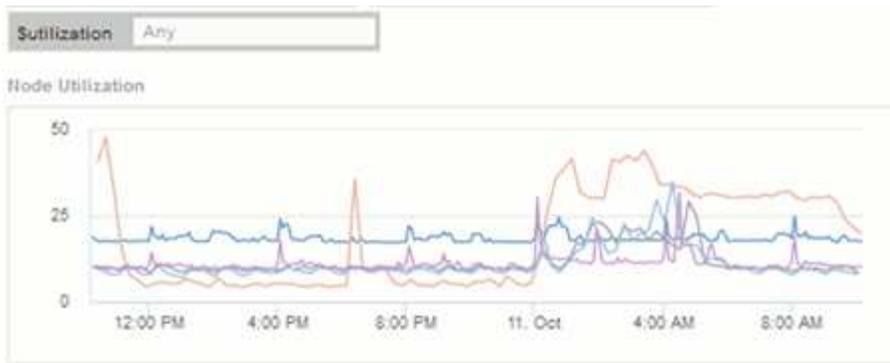
- You can move and resize the widgets in whatever placement you want on your dashboard. When finished, be sure to **Save** the dashboard.

Your final dashboard will look something like this:



- You can use the variables to focus on specific assets in your dashboard. As you enter values into the variable fields, your widgets update automatically to reflect those variables. For example, by entering "15" in the \$utilization variable field, the widgets using that variable update to show only assets with total utilization >= 15%.

Node utilization widget showing top 5 of all nodes:



Node utilization widget showing nodes with 15% or greater utilization:



3. Keep in mind the following when creating your widgets:

- The \$tier variable will only impact resources that are annotated with the **Tier** annotation.
- Not all filters will impact all widgets, depending on whether the widget is designed to accept the variable(s) specified.
- Number variables are applied as “greater than or equal to” the value specified.
 Note that any variable can be used as a filter on any widget at any level in a storage hierarchy, as long as the variable is valid for the asset against which the widget is running. As you move down from a Node level to Storage Pool to a Volume widget, more variables are present for assignment as filters. For example, at a Storage Node level widget, the *Storage* and *Name* variables can be assigned as filters. At a Storage Pool level, *Storage*, *Nodes*, *Storage Pools* and *Name* are all available. Assign your variables as appropriate and use the \$name variable at the lowest level in the stack. Doing this will allow your \$name variable to filter on the actual name of the asset against which the widget is running.

Node dashboard example widget settings

Widget settings for the node dashboard with variables example.

Following are the settings for each of the widgets in the storage node dashboard example.

Node utilization:

Edit widget

Node Utilization Show legend Override dashboard time: OFF

3h 24h 3d 7d 30d Custom

Storage Node Utilization - Total

Filter by Storage: \$storage Name: \$node

Utilization - Total (%) >= \$utilization +

Roll up +

Show Top 5 More options

Display: Line chart Y-axis: Primary

+ Add

Edit widget

Node Peak and Avg Utilization Override dashboard time: OFF

3h 24h 3d 7d 30d Custom

Storage Node

Filter by Storage: \$storage Name: \$node Utilization - Total (%) >= \$utilization +

Group by +

Name	Utilization - Max (%)	Utilization - Total (%)
3070-a	76.79	21.57
3070-b	76.79	21.57
vifasname01	54.83	18.55
vifasname02	32.50	6.06
aurora3	29.27	12.88

53 items found

Node latency:

Edit widget

Node Latency Show legend Override dashboard time: OFF

3h 24h 3d 7d 30d Custom

A Storage Node Latency - Total

Filter by: Storage: \$storage Name: \$node

Latency - Total (ms) => \$latency

Roll up: +

Show: Top 5 More options

Display: Line chart Y-axis: Primary

Edit widget

Node Peak and Avg Latency Override dashboard time: OFF

3h 30h 3d 7d 30d Custom

Storage Node

Filter by: Storage: \$storage Name: \$node Latency - Total (ms) => \$latency

Group by: +

Name	Latency - Max (ms)	Latency - Total (ms)
vfasname04	9.05	7.70
vfasname05	2.25	0.41
vfasname02	1.62	0.90
vfasname01	1.42	1.03
vfasname06	0.97	0.64

0 items found

Storage pool utilization:

Edit widget

✕

Storage Pool Utilization Show legend Override dashboard time:

3h 24h 3d 7d 30d Custom

A 🔗 📄

Storage Pool Utilization - Total

Filter by: Storage: \$storage x Nodes: \$node x Name: \$pool x

Utilization - Total (%) >= \$utilization x Tier: \$tier x +

Roll up: +

Show: Top 5 More options

Display: Line chart Y-axis: Primary Convert to Expression

+ Add

Edit widget

✕

Storage Pool Peak and Avg Utilization Override dashboard time:

3h 24h 3d 7d 30d Custom

Storage Pool

Filter by: Storage: \$storage x Nodes: \$node x Name: \$pool x

Utilization - Total (%) >= \$utilization x Tier: \$tier x +

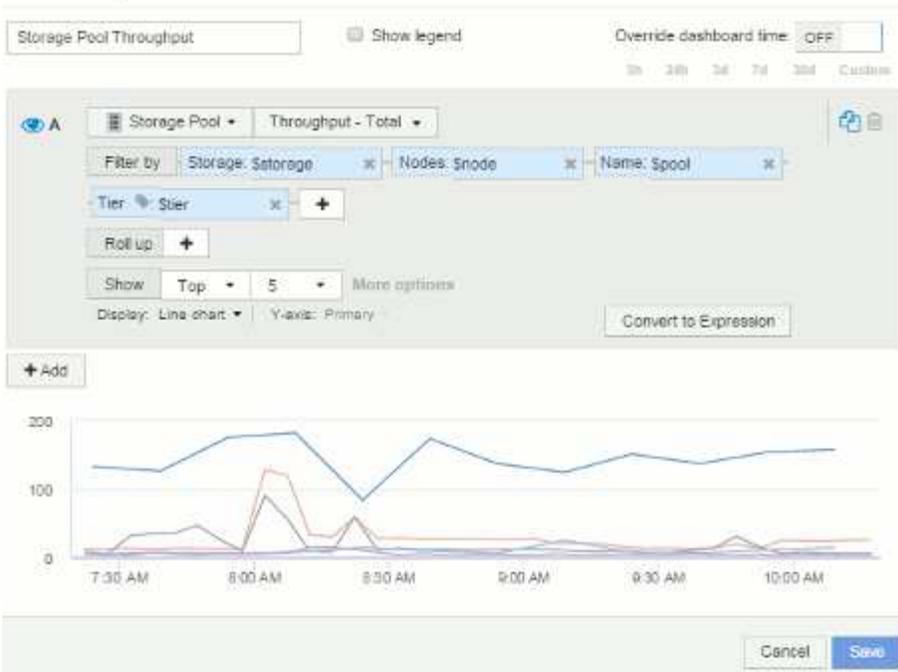
Group by: +

Name	Utilization - Max (%)	Utilization - Total (%)
vfasname01:aggr1	15.85	8.52
vfasname01:vfassn...	16.19	4.71
vfasname02:aggr2	9.28	3.65
vfasname02:vfassn...	4.96	1.63
vfasname03:aggr3	1.04	0.68

14 items found

Storage pool throughput:

Edit widget



Edit widget

Storage Pool Peak and Avg Throughput Override dashboard time: OFF

3h 24h 3d 7d 30d Custom

Storage Pool

Filter by Storage: \$storage Nodes: \$node Name: \$pool

Tier: \$tier

Group by: +

Name	Throughput - Max (MB/s)	Throughput - Total (MB/s)
vfasname01:aggr1	181.17	143.62
vfasname06:aggr1	127.19	26.75
vfasname05:aggr1	89.83	18.20
vfasname02:aggr2	24.57	9.70
vfasname05:aggr_opm1	14.61	4.75

14 items found

Cancel Save

Volume latency:

Edit widget

Volume Latency Show legend Override dashboard time: OFF

3h 24h 3d 7d 30d Custom

Volume Latency - Total

Filter by: Storage: \$storage Nodes: \$node

Storage pools: \$pool Name: \$volume Tier: \$tier

Roll up: +

Show: Top 5 More options

Display: Line chart Y-axis: Primary

Convert to Expression

+ Add

Cancel Save

Edit widget

Volume Peak and Avg Latency Show legend Override dashboard time: OFF

3h 24h 3d 7d 30d Custom

Volume

Filter by: Storage: \$storage Nodes: \$node Storage pools: \$pool

Name: \$volume Latency - Total (ms) => \$latency Tier: \$tier

Group by: +

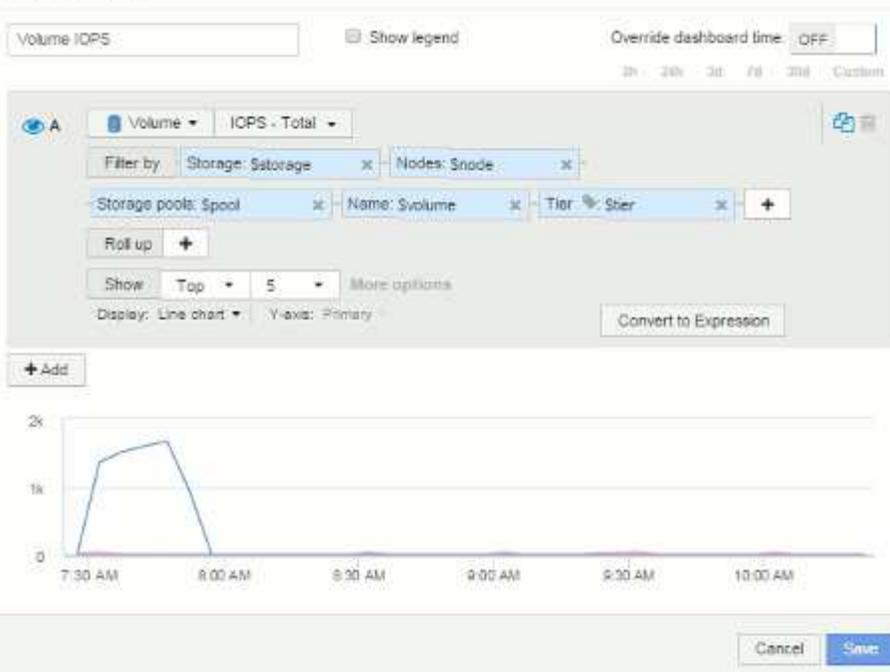
Name	Latency - Max (ms)	Latency - Total (ms)
vfasname05/vol/bo...	0.00	0.00
vfasname05/vol/bo...	0.19	0.06
vfasname05/vol/bo...	0.00	0.00
vfasname05/vol/bo...	0.00	0.00
vfasname05/vol/bo...	0.00	0.00

51 items found

Cancel Save

Volume IOPS:

Edit widget



Edit widget

Volume Peak and Avg IOPS

Override dashboard time: OFF

2h 24h 3d 7d 30d Custom

Volume

Filter by: Storage: Sstorage x Nodes: Snode x Storage pools: Spool x

Name: Svolume x Tier: Stier x +

Group by +

Name	IOPS - Max (IO/s)	IOPS - Total (IO/s)
vfasname05/vol/vi...	1,889.31	198.97
vfasname05/vol/vi...	50.03	19.18
vfasname05/vol/bo...	1.51	1.11
vfasname05/vol/bo...	0.00	0.00
vfasname06/vol/bo...	0.00	0.00

51 items found

Cancel Save

Best Practices for Dashboards and Widgets

Tips and tricks to help you get the most out of the powerful features of dashboards and widgets.

Best Practice: finding the right metric

OnCommand Insight acquires counters and metrics using names that sometimes differ from data source to data source.

When searching for the right metric or counter for your dashboard widget, keep in mind that the metric you want could be under a different name from the one you are thinking of. While drop-down lists in OnCommand Insight are usually alphabetical, sometimes a term may not show up in the list where you think it should. For example, terms like "raw capacity" and "used capacity" do not appear together in most lists.

Best practice: Use the search feature in fields such as **Filter by** or places like the column selector  to find what you are looking for. For example, searching for "cap" will show all metrics with "capacity" in their names, no matter where it occurs. You can then easily select the metrics you want from that short list.

Here are a few alternative phrases you can try when searching for metrics:

When you want to find:	Try also searching for:
CPU	Processor
Capacity	Used capacity Raw capacity Provisioned capacity Storage pools capacity <other asset type> capacity Written capacity
Disk Speed	Lowest disk speed Least performing disk type
Host	Hypervisor Hosts
Hypervisor	Host's hypervisor
Microcode	Firmware
Name	Alias Hypervisor name Storage name <other asset type> name Simple name Resource name Fabric Alias

Read / Write	Partial R/W Pending writes IOPS - Write Written capacity Latency - Read Cache utilization - read
Virtual Machine	VMIs virtual

This is not a comprehensive list. These are examples of possible search terms only.

Best Practice: finding the right assets

The Insight assets you can reference in widget filters and searches vary from asset type to asset type.

In dashboards, the asset type around which you are building your widget determines the other asset type counters for which you can filter or add a column. Keep the following in mind when building your widget:

This asset type / counter:	Can be filtered for under these assets:
Virtual Machine	VMDK
Datastore(s)	Internal Volume VMDK Virtual Machine Volume
Hypervisor	Virtual Machine
Is hypervisor	Host
Host(s)	Internal Volume Volume
Cluster	Host Virtual Machine
Fabric	Port

This is not a comprehensive list.

Best practice: If you are filtering for a particular asset type that does not appear in the list, try building your query around an alternate asset type.

Scatterplot Example: knowing your axis

Changing the order of counters in a scatterplot widget changes the axes on which the data is displayed.

About this task

This example will create a scatter plot that will allow you to see under-performing VMs that have high latency compared to low IOPS.

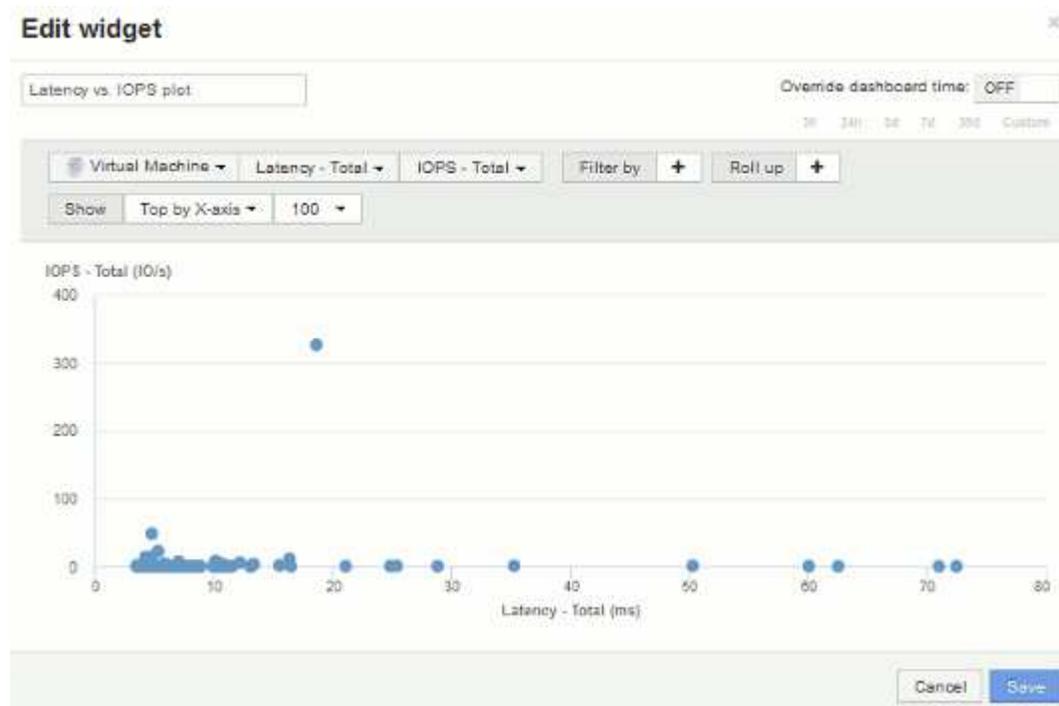
Steps

1. Create or open a dashboard in edit mode and add a **Scatter Plot Chart** widget.
2. Select an asset type, for example, **Virtual Machine**.
3. Select the first counter you wish to plot. For this example, select **Latency - Total**.

Latency - Total is charted along the X-axis of the chart.

4. Select the second counter you wish to plot. For this example, select **IOPS - Total**.

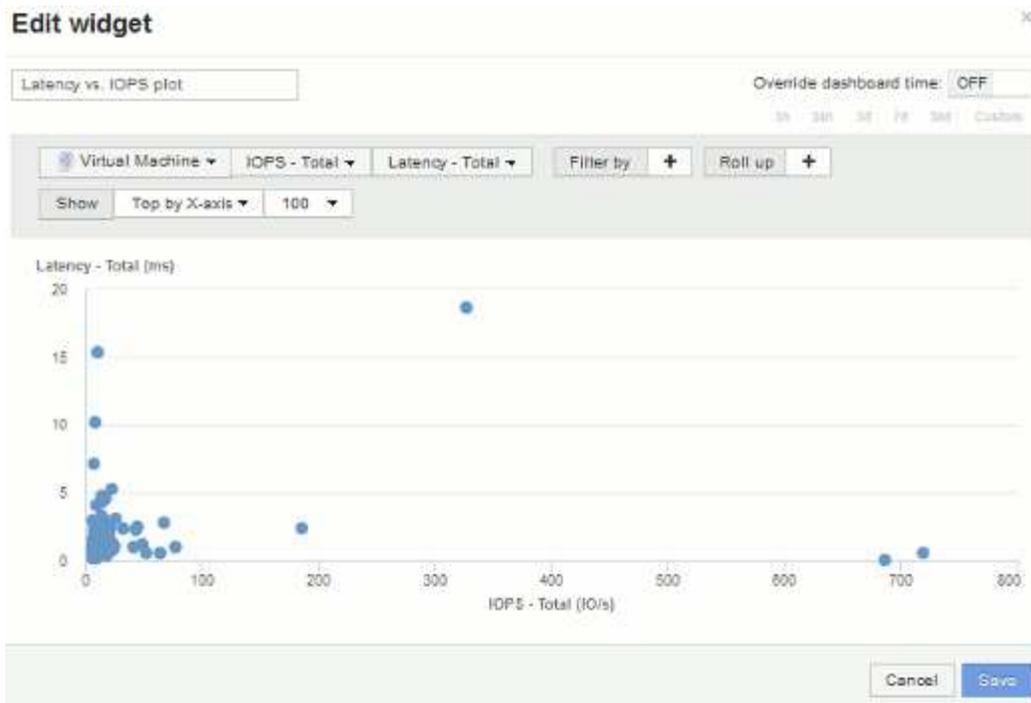
IOPS - Total is charted along the Y-axis in the chart. VMs with higher latency display on the right side of the chart. Only the top 100 highest-latency VMs are displayed, because the **Top by X-axis** setting is current.



5. Now reverse the order of the counters by setting the first counter to **IOPS - Total** and the second to **Latency - Total**.

latency- Total is now charted along the Y-axis in the chart, and *IOPS - Total* along the X-axis. VMs with higher IOPS now display on the right side of the chart.

Note that because we haven't changed the **Top by X-Axis** setting, the widget now displays the top 100 highest-IOPS VMs, since this is what is currently plotted along the X-axis.



6. You can choose for the chart to display the Top N by X-axis, Top N by Y-axis, Bottom N by X-axis, or Bottom N by Y-axis. In our final example, the chart is displaying the Top 100 VMs that have the highest *total IOPS*. If we change it to Top by Y-axis, the chart will once again display the top 100 VMs that have the highest *total latency*.

Note that in a scatterplot chart, you can click on a point to open the asset page for that resource.

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

Performance Policies

[Add new policy](#)

Database policies

Policy Name	Severity	Annotations	Time Window	Thresholds
Latency	Warning		First occurrence	Latency - Total = 200 ms
Database_0	Warning		First occurrence	IOPS - Total > 0 I/Os or Latency - Total = 0 ms

Showing 1 of 2 entries

Internal volume policies

Policy Name	Severity	Annotations	Time Window	Thresholds
Atmos Service Level	Critical	Service_Level = Atmos	First occurrence	Latency - Total > 100 ms or IOPS - Total > 100 I/Os or Throughput - Total = 200 MB/s
Global	Critical		First occurrence	Latency - Total > 200 ms or IOPS - Total > 1 I/Os or Throughput - Total = 300 MB/s

Showing 1 of 2 entries

Storage policies

Policy Name	Severity	Annotations	Time Window	Thresholds
Storage_Storage	Warning		First occurrence	IOPS - Read > 10 I/Os
Storage_0	Warning		First occurrence	Throughput - Total = 0 MB/s or IOPS - Total = 0 I/Os

Showing 1 of 2 entries

displayed.

Policies are organized by object, and are evaluated in the order in which they appear in the list for that object.

3. Click **Add new policy**.

The Add Policy dialog box is displayed.

4. In the **Policy name** field, enter a name for the policy.

You must use a name that is different from all the other policy names for the object. For example, you cannot have two policies named “Latency” for an internal volume; however, you can have a “Latency” policy for an internal volume and another “Latency” policy for a different volume. The best practice is to always use a unique name for any policy, regardless of the object type.

5. From the **Apply to objects of type** list, select the type of object to which the policy applies.
6. From the **With annotation** list, select an annotation type, if applicable, and enter a value for the annotation in the **Value** box to apply the policy only to objects that have this particular annotation set.
7. If you selected **Port** as the object type, from the **Connected to** list, select what the port is connected to.
8. From the **Apply after a window of** list, select when an alert is raised to indicate a threshold violation.

The First occurrence option triggers an alert when a threshold is exceeded on the first sample of data. All other options trigger an alert when the threshold is crossed once and is continuously crossed for at least the specified amount of time.

9. From the **With severity** list, select the severity for the violation.
10. By default, email alerts on policy violations will be sent to the recipients in the global email list. You can override these settings so that alerts for a particular policy are sent to specific recipients.
 - Click the link to open the recipients list, then click the **+** button to add recipients. Violation alerts for that policy will be sent to all recipients in the list.
11. Click the **any** link in the **Create alert if any of the following are true** section to control how alerts are triggered:

- **any**

This is the default setting, which creates alerts when any of the thresholds related to a policy are crossed.

- **all**

This setting creates an alert when all of the thresholds for a policy are crossed. When you select **all**, the first threshold that you create for a performance policy is referred to as the primary rule. You must ensure that the primary rule threshold is the violation that you are most concerned about for the performance policy.

12. In the **Create alert if** section, select a performance counter and an operator, and then enter a value to create a threshold.
13. Click **Add threshold** to add more thresholds.
14. To remove a threshold, click the trash can icon.
15. Select the **Stop processing further policies if alert is generated** check box if you want the policy to stop processing when an alert occurs.

For example, if you have four policies for datastores, and the second policy is configured to stop processing when an alert occurs, the third and fourth policies are not processed while a violation of the second policy is active.

16. Click **Save**.

The Performance Policies page displays, and the performance policy appears in the list of policies for the object type.

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**.

Monitoring the violations in your network

When Insight generates violations due to the thresholds set in performance policies, you can view them using the Violations Dashboard. The dashboard lists all the violations that occur in your network and enables you to locate and address issues.

Steps

1. Open OnCommand Insight in your browser.
2. On the Insight toolbar, click **Dashboards** and select **Violations Dashboard**.

The Violations Dashboard displays.

3. You can use the **Violations By Policies** pie chart in the following ways:

- You can position your cursor over any slice of a chart to display the percentage of the total violations that occurred for a particular policy or metric.
- You can click a slice of a chart to “enlarge” it, which enables you to emphasize and study more carefully that slice by moving it away from the rest of the chart.
- You can click the  icon in the upper-right corner to display the pie chart in full screen mode, and click  again to minimize the pie chart.

A pie chart can contain a maximum of five slices; thus, if you have six policies that generate violations, Insight combines the fifth and sixth slices into an “Others” slice. Insight assigns the most violations to the first slice, the second most violations to the second slice, and so on.

4. You can use the **Violations History** chart in the following ways:

- You can position your cursor over the chart to display the total number of violations that occurred at a particular time and the number that occurred out of the total for each specified metric.
- You can click a legend label to remove the data associated with the legend from the chart.

Click on the legend to display the data again.

- You can click the  icon in the upper-right corner to display the chart in full screen mode, and click  again to minimize the pie chart.

5. You can use the **Violations Table** in the following ways:

- You can click the  icon in the upper-right corner to display the table in full screen mode, and click  again to minimize the pie chart.

If your window size is too small, then the Violations Table displays only three columns; however, when you click , additional columns (up to seven) display.

- You can display violations for a particular time period (**1h**, **3h**, **24h**, **3d**, **7d**, and **30d**), with Insight showing a maximum number of 1000 violations for the selected time period.
- You can use the **filter** box to show only the violations you want.
- You can change the sort order of the columns in a table to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header; to return to the default sort order, click any other column header.

By default, the table displays the violations in descending order.

- You can click a violation in the ID column to display the asset page for the duration of the violation.

- You can click the resource links (for example, storage pool and storage volume) in the Description column to display the asset pages associated with those resources.
- You can click the performance policy link in the Policy column to display the Edit Policy dialog box.

You might want to adjust the thresholds for a policy if you feel it generates too few or too many violations.

- You can click a page number to browse through data by page if there is more data than fits on a single page.
- You can click  to dismiss the violation.

Troubleshooting Fibre Channel BB credit 0 errors

Fibre Channel uses buffer-to-buffer credits (BB credits) to control transmission flow. The credit value is decremented when a frame is sent from a port and the credit value is replenished when the port receives a response. If the BB credits in the port are not replenished, the transmission flow can be impacted. Ports need memory, or buffers, to temporarily store frames until they are assembled in sequence, and delivered. The number of buffers is the number of frames a port can store and is called a Buffer Credit.

As the available credits for a given port approach zero, an error warns that the port will stop receiving transmissions when zero is reached and will not resume until the BB credits are replenished.

Insight performance policies allow you to set thresholds on the following port metrics.

BB credit zero - Rx
Number of times the receive buffer-to-buffer credit count transitioned to zero during the sampling period
BB credit zero - Tx
Number of times the transmit buffer-to-buffer credit count transitioned to zero during the sampling period
BB credit zero - Total
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

BB Credit errors might be caused by some of the following scenarios:

- If a given implementation has a high percentage of FC frames of sizes significantly less than the maximum size, then more BB_Credits might be required.
- Workload changes to your environment that could be impacting ports or devices that are connected to them, such as storage nodes.

You can use the fabric, switch, and port asset pages to monitor your Fibre Channel environment. Port asset pages present summary information about the resource, its topology (the device and its connections), performance charts, and tables of associated resources. When troubleshooting Fibre Channel issues the performance chart for each port asset is helpful because it shows traffic for the selected top contributor port. Port asset pages also show buffer-to-buffer credit metrics and port errors in this chart, with Insight displaying a separate performance chart for each metric.

Creating performance policies and thresholds for ports

You can create performance policies with thresholds for metrics associated with a port. By default, performance policies apply to all devices of the specified type when you create them. You can create an annotation to include only a specific device or a set of devices in the performance policy. For simplicity, an annotation is not used in this procedure.

Before you begin

If you want to use an annotation with this performance policy, you must create the annotation before you create the performance policy.

Steps

1. From the Insight toolbar, click **Manage > Performance policies**

The existing policies are displayed. If a policy exists for switch ports, you can edit the existing policy, adding the new policies and thresholds.

2. Edit an existing port policy or create a new port policy

- Click the pencil icon at the far right of the existing policy. Add the thresholds described in steps “d” and “e”.
- Click **+Add** to add a new policy
 - a. Add a “Policy Name”: Slow Drain Device
 - b. Select port as object type
 - c. Enter First occurrence for “Apply after window” of
 - d. Enter threshold: BB credit zero - Rx > 1,000,000
 - e. Enter threshold: BB credit zero - Tx > 1,000,000
 - f. Click “Stop processing further policies if alert is generated”
 - g. Click “Save”

The policy you create monitors the thresholds you set over a period of 24 hours. If the threshold is exceeded a violation is reported.

3. Click **Dashboards > Violations Dashboard**

The system displays all of the violations that have occurred on the system. Search or sort the violations to view the “Slow Drain Device” violations. The Violations Dashboard shows all of the ports that experienced BB Credit 0 errors exceeding the thresholds that were set in the performance policy. Each switch port identified in the violations dashboard is a highlighted link to the Port Landing Page.

4. Click on a highlighted port link to display the Port Landing Page.

The port landing page is displayed and includes information useful for BB Credit 0 troubleshooting:

- Devices the port is connected to
- Identification of the port reporting the violation, which is a Fiber Channel Switch port.
- The port's speed
- The associated node and port name

The screenshot shows the 'SWITCH Port' landing page for 'fc4'. At the top, there are time range filters: 3h, 24h, 3d, 7d, and a 'Custom' button. Below the filters is a 'Summary' section containing the following details:

WWN:	10:B0:00:00:00:00:38:97
Node WWN:	20:00:44:7F:EA:40:AC:32
Device:	Cisco-MDS9148A.prod.srvc.com
Fabric:	20:D0:00:00:00:00:01:24
Speed:	4 Gb/s
State:	Online
Port Index:	4
Role:	N/A
Type:	F port
Connected To:	BSM-prod-A, BSM-prod-B port BSM-prod-A.01
Blade:	N/A
Class of service:	0
GBIC:	N/A
Performance Policies:	Slow Drain Device
Traffic - Tx (MB/s):	0.05 (0.10 max)
Traffic - Rx (MB/s):	0.05 (0.10 max)
Port Errors:	No errors
BB Credit Zero:	20,984,868

5. Scroll down to view the port metrics. Click **Select metrics to show > BB credit zero** to display the BB credit zero graph.

The screenshot shows the 'SWITCH Port' metrics page for 'fc4'. The top navigation includes time range filters (3h, 24h, 3d, 7d, Custom) and the text 'Acquired Never'. On the left, a dropdown menu 'Select metrics to show' is open, with 'BB credit zero' selected. The main area features a line graph showing BB credit zero over time from 6:00 am to 9:00 pm. The graph shows several sharp peaks, indicating periodic violations. On the right, the 'Resources' section shows 'fc4' with a blue square icon. Below it, 'Top correlated' resources are listed, with 'BSM-prod-A' checked and showing a correlation of 97%. An 'Additional resources' section with a search box is also visible.

6. Click **Top correlated**

The top correlated resource analysis shows the connected controller node the port is servicing as the resource most correlated with performance. This step compares the IOPS metrics of the port activity with the overall node activity. The display shows the Tx and Rx BB Credit Zero metrics and the controller node's IOPS. The display shows the following:

- The controller IOs are highly correlated with the port traffic
- The performance policy is violated when the port is transmitting IO to the server.
- Given that our port performance violation is occurring in conjunction with a high IOPS load on the storage controller it is likely that the violation is due to the workload on the storage node.



7. Return to the Port Landing Page and access the storage controller node's landing page to analyze the workload metrics.

The Node shows a utilization violation and metrics show high "cache reads replaced" that correlates to buffer-to-buffer zero credit states.

Storage:	BSM-prod-A, BSM-prod-B
HA partner:	BSM-prod-B
State:	N/A
Model:	FAS6070
Version:	8.0.5 7-Mode
Serial number:	700001181351
Memory:	98,304 MB
Utilization:	● 21.26% (94.56% max)
IOPS:	232.73 IO/s (1,153.00 IO/s max)
Latency:	7.07 ms (15.00 ms max)
Throughput:	22.44 MB/s (106.00 MB/s max)
Processors:	12
Performance Policies:	● Node Utilization Node Read Latency

8. From the Node landing page, you can compare the BB credit zeroes by selecting the port from the correlated resources list, and select utilization data, including Cache utilization data, for our node from the metrics menu.



This data makes it clear that the Cache hit Ratio is inversely correlated to our other metrics. Instead of being able to respond to the server load from cache, the storage node is experiencing high cache reads replaced. It is likely that having to retrieve most of the data from disk rather than cache is causing the delay in the port's transmission of data to the server. The cause of the performance problem appears likely to be a workload generated change in IO behavior, and that the node cache, and its configuration, are the cause. The problem might be solved by increasing the node's cache size or change the caching algorithm's behavior.

Analyzing your infrastructure

The procedures in this topic are ones that you might use to perform an analysis of parts of the infrastructure in your environment. The steps, views, and data you gather in this exercise use virtual computing objects as an example. Analysis of other assets in your environment will follow similar steps using relevant counters for each specific asset. The intent of this exercise is to familiarize you with the variety of options Insight offers to monitor and understand the characteristics of the assets in your data center.

About this task

Some of the actions you can take to analyze the state of your infrastructure might include the following:

- Observe an object's behavior over time
- Compare an object's metrics against the metrics of the top 10 like objects
- Compare numbers for objects
- Compare the top 10 objects against the average
- Compare metrics A vs. B for many objects to show categories and anomalies
- Compare a range of objects against other objects
- Use an expression to display metrics not available in the web UI

You can create all of these views of objects in your infrastructure in a dashboard using widgets for each analysis you perform. The dashboards can be saved to provide quick access to current data on your infrastructure.

Observe an object's behavior over time

You can observe the behavior of a single object to determine if the object is operating within expected operational levels.

Steps

1. Use a query to identify the VM that will be the subject of analysis: **Query > + New query > Virtual machine > "name"**

Leaving the name field blank returns all VMs. Select the VM that you want to use in this exercise. You can select it by scrolling through the list of VMs.

2. Create a new dashboard for the information you want to collect. From the toolbar, click **Dashboards > +New Dashboard**.

3. In the new Dashboard, select **Variable > Text**.
 - a. Add the VM name from your query as the `$var1` value.
 - b. Click the check box.

The variable is used to easily swap between different sets of objects you want to analyze. In other steps of your analysis, you may reuse this variable for additional analysis against the single VM initially chosen. Variables become more useful when identifying multiple objects.

4. Add a line chart widget to the new dashboard: **Widget > Line chart**.
 - a. Change the default asset type to virtual machine: Click **Virtual machine > Latency-Total**.
 - b. Click **Filter by > Name > \$var1**.
 - c. Change the time period on the dashboard: **Override dashboard time > On > 7 days**.

You can change the duration of the display using any of the preset selections or by specifying a custom time range.

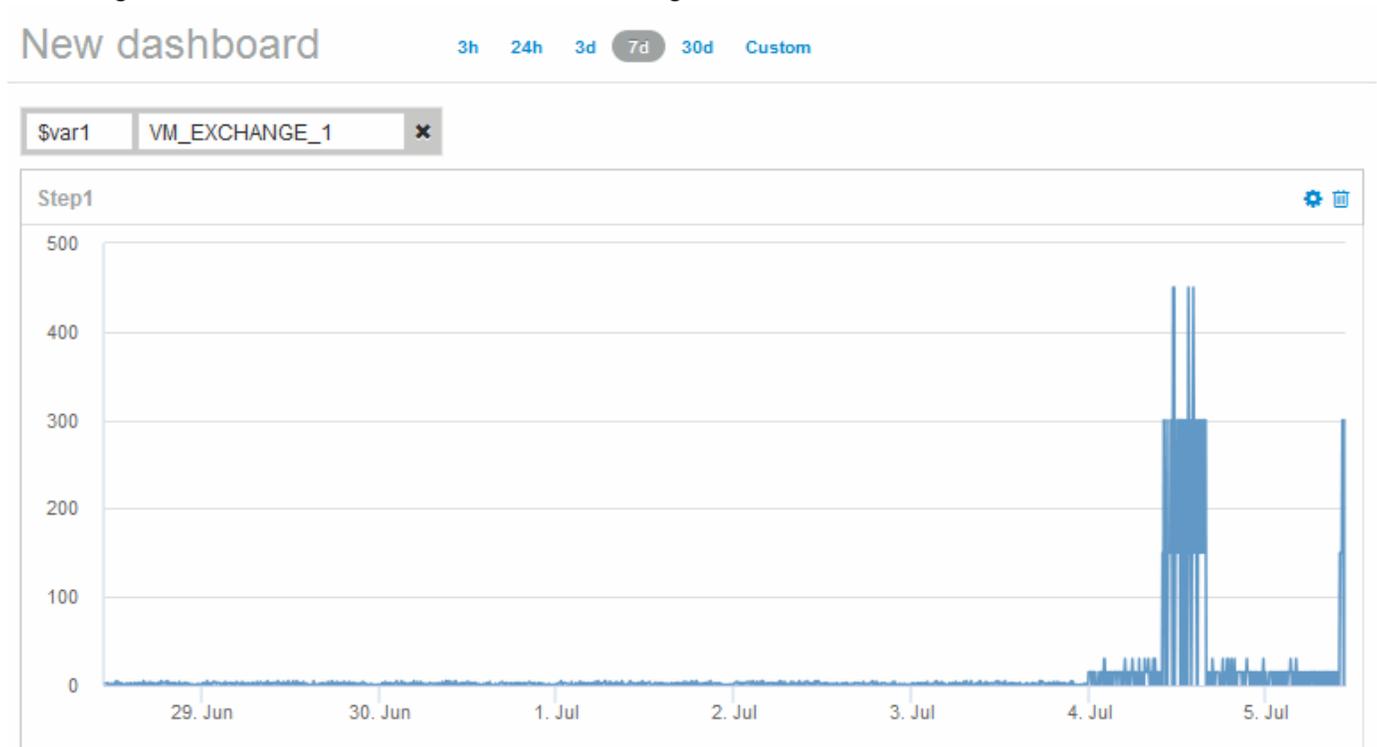
+

The dashboard displays the VM's **IOPS-Total** for the period of time you specify.

5. Assign a name to the widget and save the widget.

Results

Your widget should contain data similar to the following:



The VM shows a period of abnormally high latency for a short period of time in the 7 days that are displayed.

Compare objects with the top 10 latency total to the average latency for all like objects

You might want to compare the VMs with the top 10 latency total to the average latency total to identify any that are extremely out of the average range. This information could help in decisions to balance workloads on VMs.

Steps

1. Add a widget with a stacked area chart to the to the new dashboard: **Widget > Stacked Area Chart**

a. Change the default device to Virtual machine: Click **Storage > Virtual machine > Latency total**

The widget displays the Latency Total, for all VMs, for 24 hours in a stacked area chart.

b. Create a second display in this widget that shows Latency Total averaged for all VMs: **Widget > Line chart**

c. Change the default device to Virtual machine: Click **Virtual machine > Latency-total**

The widget displays the Latency Total for the default 24 hour period of time using a line chart.

d. Click **X** on the **Roll up** bar and select **Show > Top > 10**

The system displays the Top 10 VMs based on Latency Total.

2. To compare the average Latency Total for all VMs to the Top 10 IOPS total use the following steps:

a. Click **+Add**

b. Change the default device to Virtual machine: Click **Storage > Virtual machine > IOPS total**

c. Click **X** on the **Roll up** bar and select **Show > Top > 10**

The system displays the 10 objects with high latency and shows the average latency in a line chart.

+

image::../media/analytics-top10-avg.gif[]

+

The average latency is 1.6 ms, while in the top ten, the are VMs experiencing latency of over 200 ms.

Compare one object's latency total to the latency total of the top 10 objects

The following steps compare a single VM's Latency Total to the VMs reporting the Top 10 Latency Total in the entire virtual infrastructure.

Steps

1. Add a widget with a line chart to the to the new dashboard: **Widget > Line Chart**

a. Change the default device to Virtual machine: Click **Storage > Virtual machine > Latency-total**

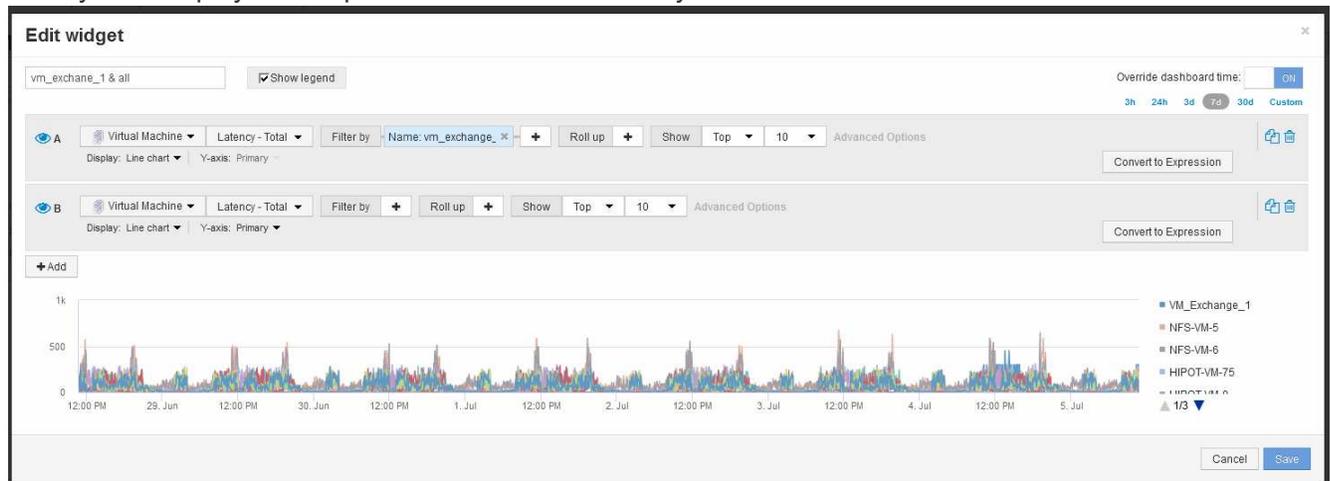
The widget displays the total Latency, for all VMs, for the default 24 hours in an area chart.

- b. Create a second display in this widget that shows Latency Total averaged for all VMs: **Widget > Line chart**
- c. Change the default device to Virtual machine: Click **Storage > Virtual machine > Latency-Total**

The widget displays the Latency total for the default 24 hour period of time using a line chart.

- d. Click **X** on the **Roll up** bar and select **Show > Top > 10**

The system displays the Top 10 VMs based on Latency - Total.



2. Add the VM that you want to compare to the Top 10:
 - a. Click **+Add**
 - b. Change the default device to Virtual machine: Click **Storage > Virtual machine > Latency total**
 - c. Click **Filter by > Name > \$var1**
3. Click **Show legend**

Results

A legend identifies each of the VMs under analysis. You can easily identify VM_Exchange_1 and determine if it is experiencing latency similar to the top ten VMs in the environment.

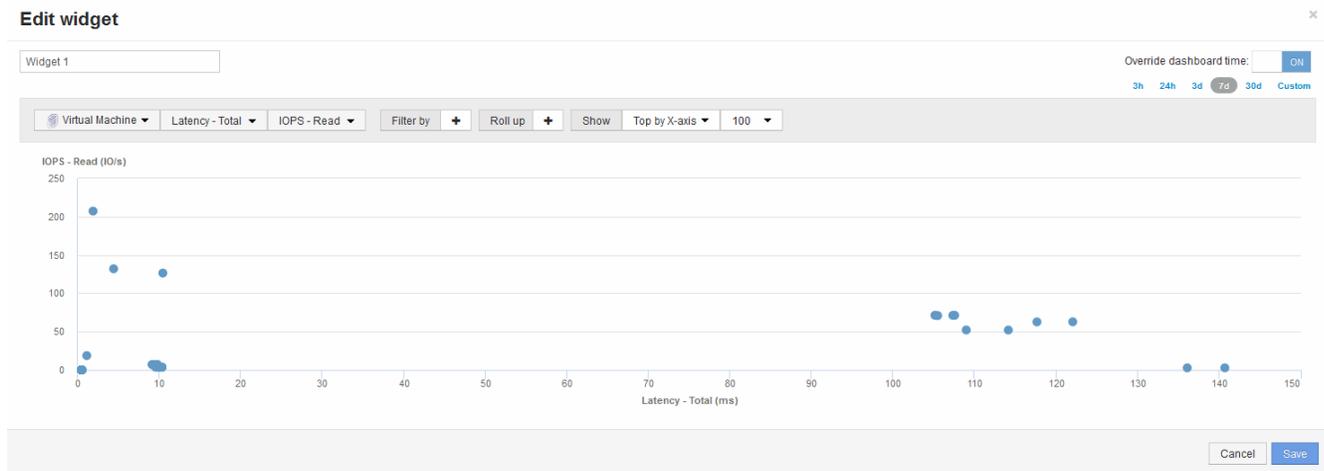
Compare metrics-A against metrics-B to show categories and anomalies

You can use a scatter plot to show two sets of data for each object. For example, you can specify IOPS Read and Latency Total to be displayed for each object. Using this chart you can identify the object you consider troublesome based on both the IOPS and the Latency combined.

Steps

1. Add a widget with a scatter plot chart to the to the new dashboard: **Widget > Scatter Plot Chart**
2. Change the default device to Virtual machine: Click **Storage > Virtual machine > Latency total > IOPS Read**

The system displays a scatter plot similar to the following:



Use an expression to identify alternate metrics

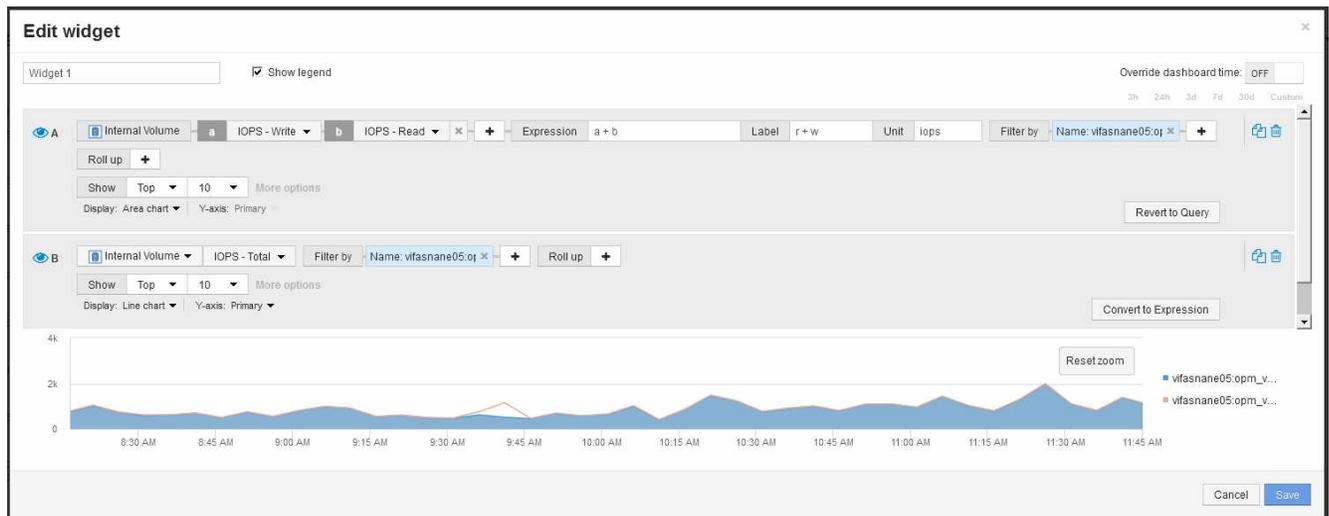
You can use expressions to view metrics not provided by the web UI, such as the IOPS that are system overhead generated.

About this task

You might want use an expression to show total IOPS generated by non-read or non-write operations, such as overhead operations for an internal volume.

Steps

1. Add a widget to the dashboard. Choose **Area chart**.
2. Change the default device to Internal volume: Click **Storage > Internal volume > IOPS Write**
3. Click the **Convert to Expression** button.
4. The **IOPS - Write** metric is now in the alphabetic variable field “a”.
5. In the “b” variable field, click **Select** and choose **IOPS - Read**.
6. In the **Expression** field, type **a + b**. In the **Display** section, choose **Area chart** for the expression.
7. In the **Filter by** field, enter the name of the internal volume you are analyzing.
8. The **Label** field identifies the expression. Change the label to something meaningful like “R + W IOPS”.
9. Click **+Add** to add a line for total IOPS to the widget.
10. Change the default device to Internal volume: Click **Storage > Internal volume > IOPS Total**
11. In the **Filter by** field, enter the name of the internal volume you analyzing.



The chart displays the total IOPS as a line, with the chart showing the combination of read and write IOPS in blue. The gap between 9:30 and 9:45 shows non-read and non-write IO (overhead) operations.

Introduction to minimizing risk in thin provisioning

In today's hybrid IT data centers, administrators are pressured to stretch resource utilization beyond physical bounds by employing capacity efficiency technologies such as thin provisioning to control over allocation and leverage what was once unavailable capacities.

OnCommand Insight provides near real time capacity usage and utilization details historically across multiple thin provisioned layers within the IT service stack. Failing to properly manage oversubscription risk could result in untimely downtime to the business.

Monitoring the storage pool

Each storage pool landing page provides over-subscription ratios, identifies correlated resources, LUN and disk utilization, as well as policy breaches and violations that have occurred with the storage pool.

Use the storage pool landing page to identify any potential problems with the physical assets supporting your virtual infrastructure. You can track capacity and capacity ratios trending over 30 days or use a custom time frame. Pay attention to data in the following sections to monitor the status of the storage pool.

- **Summary**

Use this section to understand:

- Storage pool capacity information including physical capacity and the overcommitted capacity.
- Whether the aggregate is oversubscribed, and by how much.
- Any policy violations that have occurred.

- **Storage resources and Disks sections**

The storage resources section shows the LUN utilization.

The disks section shows the individual disks that make up the storage pool.

- **Resources**

Use this section to understand the VMDKs to LUNs correlation and understand the storage to VM application path.

- **Violations section**

The violations section identifies any breaches to performance policies that have been set for the storage pool.

Monitoring the Datastores

The Datastore landing page identifies over-subscription ratios, LUN and disk utilization, correlated resources, and shows policy breaches and violations that have occurred with the Datastore.

Use this landing page to identify problems with your virtual infrastructure. You can track capacity and capacity ratio trending to anticipate changes in your capacity.

- **Summary**

Use this section to understand:

- Datastore capacity information including physical capacity and the overcommitted capacity.
- The percentage of overcommitted capacity.
- Metrics for latency, IOPS, and throughput.

- **VMDKs**

The VMDKs section shows virtual disk capacity and performance.

- **Storage resources**

This section shows the capacity used and the performance metrics for the internal volume correlated to the Datastore.

- **Resources**

Use this section to understand the VMDKs to LUNs correlation, and understand the storage to VM application path.

- **Violations section**

The violations section identifies any breaches to performance policies that have been set for the Datastore.

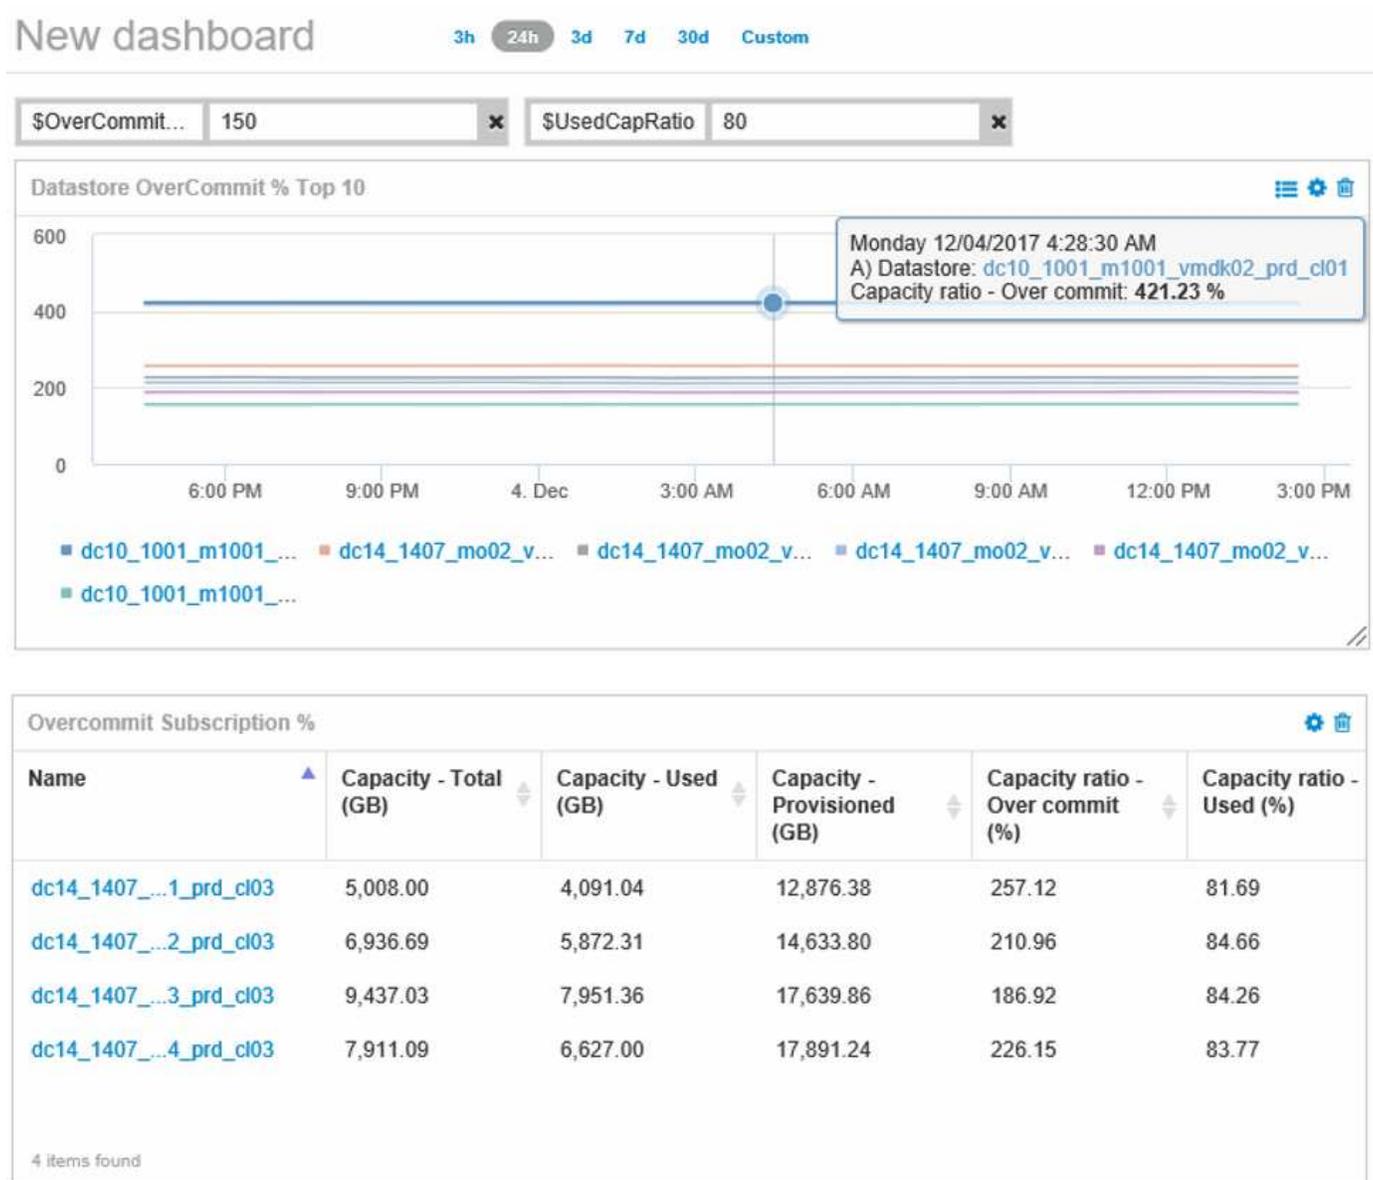
Create dashboards to monitor thin provisioned environments

OnCommand Insight's flexible dashboard widget design and display charting options allow deep analysis into capacity usage and utilization, strategic information for minimizing risks in thin provisioned data center infrastructures.

You can create dashboards that provide access to Datastore and Storage pool information that you want to monitor.

Using dashboards to access Datastore information

You might want to create dashboards that provide quick access to the data you want to monitor in your virtual infrastructure. A dashboard could include widgets similar to the following to identify the top 10 Datastores based on their overcommitted % and a widget showing the capacity data for Datastores. The dashboards use variables to highlight Datastores that are overcommitted by more than 150% and Datastores that have exceeded more than 80% used capacity.



Additional widgets that could be used to monitor your thin-provisioned environment could include some of the following information:

- VMDK capacities correlated to Datastores
- VM capacities
- Data store capacity used trending

Using dashboards to access Storage pool information

A dashboard could include widgets similar to the following, identifying the amount of physical storage capacity used, or identifying the overcommitted capacity for a Storage Pool.



Using performance policies to reduce risk in thin provisioning

You should create performance policies to raise alerts when thresholds in your virtual infrastructure have been breached. The alerts allow you to respond to changes in your environment before they cause interruptions or outages in operations.

Policies that help in monitoring the virtual infrastructure include the following:

- **Datastore**

You could use the following policies on the Datastore:

- Capacity ratio - Overcommit
- Capacity ratio - Used
- Capacity - Used
- Capacity - Total

- **Storage pool**

The following policies can protect against storage related capacity outages in thin provisioned environments:

- Capacity provisioned
- Capacity used

- Capacity ratio - Overcommit
- Capacity ratio - Used

You can expand from these policies to monitor capacity in the virtual infrastructure, including:

- Internal volumes
- LUNs
- Disks
- VMDKs
- VMs

You can configure policies using annotations. You assign the same annotation to the specific assets that support an application. For example, you can assign annotations to the Datastores and the Storage pools of a thin provisioned application. You might have annotations named Production for the production environment, Development for the development environment, and so on. You can change the thresholds and criticality of warnings depending on the type of application the assets are supporting. For example, a breach of a threshold for a production application's DataStore might raise a *critical warning*, while the same breach for a development environment might only raise a *warning*. Incorporating annotations within defined policies can help to further reduce unwanted alerting noise for non-critical assets.

Creating performance policies for Storage Pools

You can create performance policies that trigger alerts to notify you when thresholds for Storage Pool assets have been exceeded.

Before you begin

This procedure assumes that you have thin provisioned the storage pool.

About this task

You want to create policies that monitor and report changes in a storage pool that could contribute to outages. For the thin provisioned physical storage pool, you want to monitor the physical capacity and monitor the Overcommit Ratio.

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. If notifications are enabled (**Admin > Notifications**), you can configure Insight to send email when performance policies are breached.

3. Click **+Add** to create a new policy.
4. In **Policy Name** enter a policy name for the Storage Pool.
5. In **Apply to objects of type** select Storage Pool.
6. In **Apply after window of** enter First occurrence.
7. In **With severity** enter Critical

8. Configure the Email recipients that you want notified when thresholds are breached.

By default, email alerts on policy violations are 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 this policy will be sent to all recipients in the list.

9. In **Create alert if any of the following are true** enter Capacity ratio - Used > 85%

Results

This configuration results in the system sending a critical warning message when more than 85% of the physical capacity of the storage pool is used. Using 100% of the physical memory will result in application failure.

Create additional Storage Pool policies

About this task

Create an additional “Capacity ratio - Used” policy that raises a warning message when the Storage Pool capacity used exceeds 75%. If notifications are enabled (**Admin > Notifications**), you can configure Insight to send email when performance policies are breached.

Creating performance policies for Datastores

You can create performance policies with thresholds for metrics associated with the datastores that correlate to the storage pools you are monitoring. By default, performance policies apply to all devices of the specified type when you create them. You can create an annotation to include only a specific device or a set of devices in the performance policy.

Before you begin

When using an annotation in a performance policy, the annotation must exist before the policy is created.

About this task

You create a performance policy that provides notification when one or more Datastores you are monitoring exceeds a threshold you set. Your system might already contain a global policy that meets your needs or a policy using annotations might also work if you annotate your Datastores.

Steps

1. From the Insight toolbar, select **Manage > Performance Policies**

The performance policies page is displayed. Review any existing performance policies to identify existing policies that address the metrics for thresholds you want to monitor.

2. Click **+Add** to add a new policy

3. Add a “Policy Name”

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 data store. The best practice is to always use a unique name for any policy, regardless of the object type.

4. Select "Datastore" as the Object Type
5. Click "First Occurrence"

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.

6. Click "Warning"
7. For "Create alert", select **Capacity ratio - Over commit** and set the value to **> 150**

You might want to create additional capacity related alerts, such as **Capacity total** and **Capacity used**.

Collecting Host and VM file system utilization data

The Host and VM File Systems data source, combined with the Host Utilization license, enables reporting and chargeback at the file system level for known Hosts and VMs.

OnCommand Insight collects data from storage devices, most of which report their volumes as block devices. This allows Insight to report on utilization at the storage level, but not at the file system level. Storage arrays typically know which blocks have been written to, but not which blocks have been freed.

Client hosts and VMs implement file systems (ntfs, ext*...) on top of these block devices. Most file systems keep a table of contents containing directory and file metadata. When files are deleted, their entries are simply removed from the table of contents. Blocks consumed by those files are now eligible for re-use by the file system, but the storage array doesn't know this. In order for Insight to report on filesystem usage, it must be collected from the client host or VM point of view for accurate chargeback.

Insight allows this level of file system utilization data collection through the **NetApp Host and VM File System** data source, in combination with the **Host Utilization** license. VM's must be annotated with the appropriate **Compute Resource Group** name, and associated storage arrays must be annotated with appropriate **Tier** annotations with proper costs for accurate cost reporting.



The Host Utilization License is resource-based, as opposed to capacity-based as other Insight licenses.

Configure Insight for file system collection

To configure Insight for collection of file system utilization data, you must install the Host Utilization Pack license and configure the NetApp Host and VM File Systems data source.

Before you begin

If you haven't already, install the Host Utilization Pack license. You can check for the license in the **Admin > Setup** page, on the **Licenses** tab.

The Host and VM File Systems data source only reports file system utilization and file system metadata for

known **Compute Resources** (hosts and VMs) currently being collected or discovered in Insight:

- Virtual Machines are collected by hypervisor data sources such as Hyper-V and VMware.
- Hosts are discovered via device resolution.

The proper Tier annotations must be present on the appropriate storage resources.

The following connected block storage devices are supported:

- NetApp Clustered Data OnTap (cDOT)
- NetApp 7-Mode
- Clariion
- Windows: VMWare virtual disks (VMDKs) for FC, iSCSI
- Linux: VMWare VMDKs (iSCSI and FC not supported)

A **Compute Resource Group** is an annotation that allows grouping of hosts and/or virtual machines that share a common administrative credential.

Steps

1. First, annotate the hosts and/or virtual machines to be included in your **Compute Resource Group**. Go to **Queries > +New query** and search for *Virtual Machine* assets.

You will need to repeat these steps for *Host* assets.

2. Click on the column selector on the right of the table and select the **Compute Resource Group** column to display it in the query results table.
3. Select the virtual machines you wish to add to the desired compute resource group. You can use a filter to search for specific assets.
4. Click on the **Actions** button and choose **Edit annotation**.
5. Select the *Compute Resource Group* annotation, then choose the desired resource group name in the *Value* field.

The resource group annotation is added to the selected VMs. The resource group name must match the name you will configure in the Host and VM File Systems data source later.

6. To configure the Host and VM File Systems data source for a compute resource group, click on **Admin > Data sources** and **Add** the *NetApp Host and VM File Systems* data source.

Settings

*Name

Vendor

Model

Where to run

What to collect

- Clustered Data ONTAP 8.1.1+
- Clustered Data ONTAP 8.1.1+ (Unified Manager 6.0+)
- Data ONTAP 7-Mode
- E-Series (Firmware 6.x)
- E-Series (Firmware 7.x+)
- Host and VM File Systems**
- SolidFire 8.1+
- StorageGrid

Configuration

Advanced configuration

Test

Cancel Save

- In the **Configuration** section, enter a **User Name** and **Password** for an operating system user with appropriate rights to retrieve file system data. For Windows operating system users, this must include the domain prefix if your Windows environment uses it.

Note that an Insight Acquisition Unit (AU) installed on Linux can report on Linux compute resources, while an AU installed on Windows can talk to either Linux or Windows compute resources.

- Enter the name of the **Compute Resource Group** for the assets from which you will want to collect file system utilization data. This name must match the resource group name you used to annotate the assets above.

If you leave the Compute Resource Group field empty, the data source will collect data for hosts or VMs that have no Compute Resource Group annotation.

- In the **Advanced Configuration** section, enter the desired polling interval for this data source. The default of 6 hours is usually adequate.
- It is recommended to **Test** the data source connection before saving it. A successful connection result will also show you how many compute resource targets are contained in the group.
- Click **Save**. The Host and VM File Systems data source will begin collecting data on its next poll.
- Once file system data is being collected, you can view it on the host's or VM's asset page, in the File System widget:

File Systems

Name	Capacity (Used / Total GB)	Type	Storage Resource
/	9.15% (11.0 / 120.0)	xfs	vifasnane:...vm_oci_
/boot	23.79% (0.1 / 0.5)	xfs	vifasnane:...vm_oci_
/dev/dm-1	7.8	swap	vifasnane:...vm_oci_

Showing 1 to 3 of 3 entries

- Repeat these steps for each Compute Resource Group you will have. Each compute resource group must be associated with its own Host and VM File Systems data source.

Note that file system information will be collected for hosts and VM's that are already being acquired by any traditional VMware or Hyper-V data sources in your environment.

File system chargeback and reporting

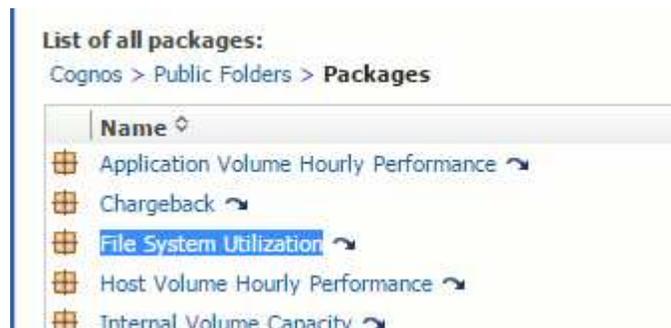
Chargeback for file systems is always performed from the storage perspective. Storage arrays associated with virtual machines annotated for a particular compute resource group will be included in chargeback reports for that resource group.

Before you begin

Any virtual machines which you wish to include in file system utilization chargeback must be annotated with the appropriate compute resource group name. Storage arrays associated with those virtual machines must be annotated with the appropriate Tier annotations. ETL to data warehouse must have occurred after these annotations are in place.

Steps

- Open a browser to your Reporting server, usually <https://<host or IP>:9300/p2pd> or <http://<host or IP>:9300/bi> (7.3.3 or later) and log in.
- Choose the **File System Utilization** package and create a new report.



- Drag and drop items from your data mart(s) to build your report.

The example below is a very simple report. You can create complex reports built around your specific business needs.

Name	Type	Allocated Capacity GB	Used Capacity GB	Tier Name	Cost	Storage Name
/	xfs	119.96	9.96	N/A		vifasnane05,vifasnane06
/	xfs	5,492.53	799.63	Tier 1	100	vifasnane
/boot	xfs	0.48	0.17	N/A		vifasnane05,vifasnane06
/boot	xfs	8.72	2.41	Tier 1	100	vifasnane
/dev/dm-1	swap	7.81	0.00	N/A		vifasnane05,vifasnane06
/dev/dm-1	swap	140.61	0.78	Tier 1	100	vifasnane
C:\	NTFS	948.27	331.98	Tier 1	100	vifasnane
PHYSICALDRIVE0: System Reserved	NTFS	1.70	1.41	Tier 1	100	vifasnane

Configuring your system to report chargeback data

Chargeback reports provide storage capacity chargeback and accountability information by hosts, application, and business entities, and include both current and historical data.

This guide describes how to configure Insight to generate a chargeback report providing accountability for service level costs and storage usage cost. The intent of the guide is to provide the steps required to create a simple chargeback report, and familiarize Insight users with the options available when configuring chargeback in their unique environment.

For each application, the example report identifies the resources provisioned and the cost of the resources. The output for the report is created by defining the following data in Insight

- Storage tiers
- Cost associated with each storage tier
- Provisioned storage capacity
- Service levels
- Cost per service level

The following sections describe the steps required to configure this data so that it can be accessed by Insight Reporting.

Defining annotations for use with chargeback

When customizing OnCommand Insight to track data for your corporate requirements, you can define specialized annotations needed to provide a complete picture of your data: for example, an annotation can define an asset's end of life, or the data center the asset resides in, or a storage tier defining the cost per GB of the storage.

About this task

The chargeback report example in this guide provides data for Service level and for Tier level. You must create annotations for each Service level and Tier level and then define costs for the Service levels and Tier levels.

Steps

1. Log in to the Insight web UI
2. Click **Manage > Annotations**

The annotations page displays.

3. Position your cursor over the Service Level, or Tier annotation, and click .

The Edit Annotation dialog box displays.

4. Click **ADD** to add new Tiers and cost.

In the report example, the Tier and Service level names use the precious metal analogy of Gold, Silver, and Bronze. You can use any naming convention chosen by your organization, for example, Tier 1, Level 2, Supreme.

5. Enter the values for the Gold-Fast, Gold, Silver, and Bronze Tiers and the costs associated with each.

The values you enter define the cost per-GB for the storage that is used by applications. The service level cost can be the cost of providing the service or the actual price to service the consumer. These costs will be reported in the Chargeback report.

6. Click **Save** when you finish.

Defining applications for use with chargeback

If you want to track cost data associated with specific applications running in your environment, you first need to define the applications.

Before you begin

If you want to associate application with a business entity, you must have already created the business entity.



This example does not associate any application with business entities.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage > Application**

After you define an application, the Applications page displays the application's name, its priority, and, if applicable, the business entity associated with the application.

3. Click **Add**

The Add Application dialog box displays.

4. Enter a unique name for the application in the Name box. Enter the Applications identified in the Report: African Tours, APAC Commercial Sales, and so on.
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. You will not use volume sharing, click to clear the **Validate** volume sharing box.
8. Click **Save**.

The applications appears in the Applications page. If you click the application's name, Insight displays the asset page for the application. After defining an application, you can go to an asset page for host, virtual machine, volume, internal volume, or hypervisor to assign an application to an asset.

Assigning applications to assets

After defining your applications, you need to associate the applications with specific assets. You can use a simple ad hoc method to apply an applications to an asset. Users looking to apply applications in bulk should use a query method to identify the assets they want to assign to an application.

Assigning applications to assets using an ad hoc method

You assign an application to an asset so that you can identify the resources of the asset that the application uses. If an asset has a cost assigned to it you can identify the cost that is incurred by the application and if the resource is measured by size, you can determine if the resource will need to be replenished.

About this task

Use the following method to assign applications 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 application by doing either of the following:

Option	Description
Navigate to the list of assets	Click Dashboard > Assets Dashboard and select the asset.
Search for 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** is displayed) and then click  (Edit application).

The list of available applications for the selected asset is displayed. 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 a 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.

Assigning applications to an asset using a query

You assign an application to an asset so that you can identify the resources of the asset that the application uses. If an asset has a cost assigned to it you can identify the cost that is incurred by the application and if the resource is measured by size, you can determine if the resource will need to be replenished.

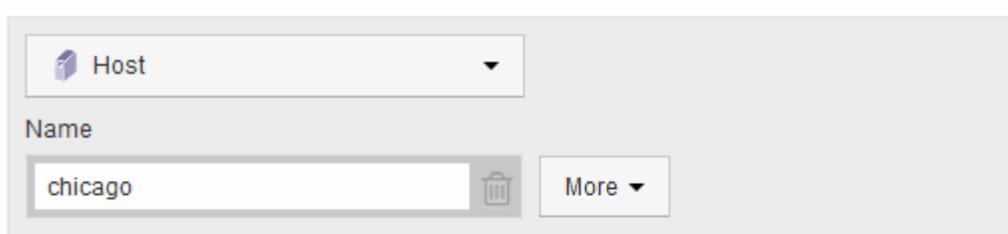
About this task

You can simplify the task of assigning multiple assets to an application by using a query.

Steps

1. Create a new query to identify the assets you want to assign an application to. For example, if you want to assign it to a host with a specific name that relates to a geographic location, Click **Queries** > **+New Query**
2. Click **Host**
3. In the **Name** field, enter `Chicago`

The system displays all of the Hosts with `Chicago` as part of their name.



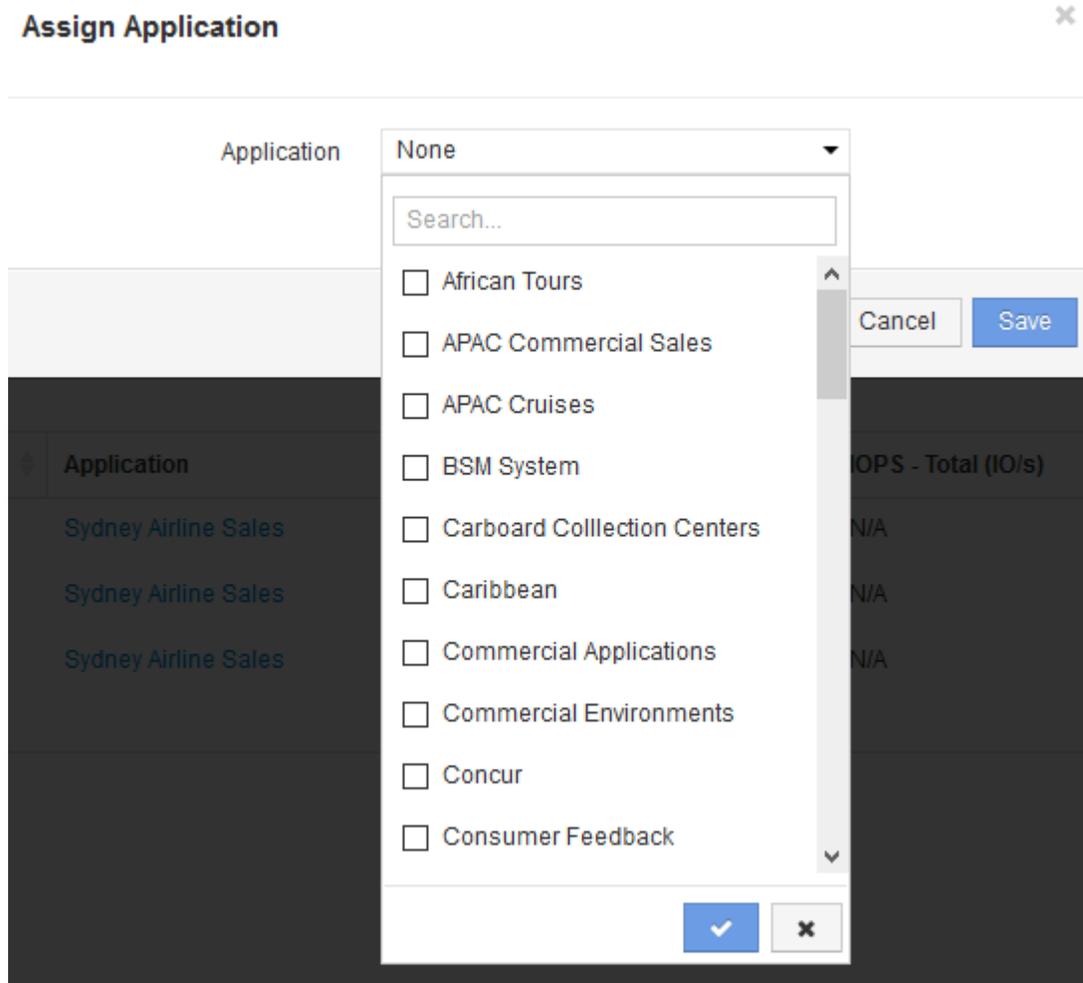
Query results

<input type="checkbox"/>	Name	IP	Application
<input type="checkbox"/>	Chicago-Host1	10.11.12.21	Sydney Airline Sales
<input type="checkbox"/>	Chicago-Host2	10.11.12.32	Sydney Airline Sales
<input type="checkbox"/>	Chicago-NAS	10.11.12.10	Sydney Airline Sales

Showing 1 to 3 of 3 entries

4. Select one or more of the Hosts identified by your query.

5. Click **Actions > Add Application**



The system displays the Assign application dialog.

6. Select the Application you want to assign to the Host and click 

7. Click **Save**

The application name appears in the User Data section.

Creating a simple Chargeback report

Chargeback reports allow administrators and managers to evaluate capacity usage by application, business entity, service level, and tier. The chargeback reports include capacity accountability, historical capacity accountability, and trending data. The data for these reports are built and scheduled from the OnCommand Insight Data Warehouse.

Before you begin

To create the sample report your system must be configured to report costs for storage tiers. The following tasks need to be completed:

- Define annotations for tiers.

- Assign costs to annotations.
- Define the applications you want to track data for.
- Assign applications to assets.

About this task

This example uses the Cognos Workspace Advanced reporting tool to create the Chargeback report. With Workspace Advanced, you create reports by dragging and dropping data elements into a report pallet.

Steps

1. In the OnCommand Insight web UI, click the reporting icon.
2. Log in to the Reporting Portal.
3. In the IBM Cognos Connection toolbar, click **Launch > Cognos Worksapce Advanced**

The Workspace Advanced package screen opens.

4. Click **Packages > Chargeback**

The IBM Workspace Advanace screen is displayed.

5. Click **New**
6. In the **New** report dialog Click **List** to specify a list report.

The report palette is displayed and the Chargeback “Simple data mart” and “Advanced data mart” are displayed under the Source heading.

7. Click the arrows next to each data mart to expand them.

The full contents of the data marts are displayed.

8. Drag “Application” from the “Simple Data Mart `” into the far left column of the report palette.

When you drag an item into the palette the column shrinks and is highlighted. Dropping the application data into the highlighted columns results in all of the applications being listed correctly in the column.

9. Drag “Tier” from the “Simple Data Mart `” into the next column of the report palette.

The storage tier associated with each application is added to the palette.

10. Drag “Tier Cost” from the “Simple Data Mart `” into the next column of the report palette.
11. Drag “Provisioned capacity” from the “Simple Data Mart `” into the next column of the report palette.
12. Hold down the **Ctrl** key and select the “Tier cost” and “Provisioned capacity” columns in the pallet.
13. Right-click the mouse in either of the selected columns.
14. Click **Calculate > Tier cost * Provisioned capacity DB**

A new column is added to the pallet with the title "Tier Cost * Provision Capacity GB".

15. Right click the **Tier Cost * Provision Capacity GB** column.
16. Click **Style > Data Type**

17. Click **Format type > Currency**

18. Click **OK**

The column data is now formatted as US currency.

19. Right click “Tier Cost * Provision Capacity GB” and select **Edit Data Item Label**

20. Replace the Name field with “Provisioned Capacity Cost”

21. To run the report, click **Run > Run report - HTML**

A report similar to the following is displayed.

Application	Service Level	Service Level Cost	Tier	Tier Cost	Provisioned Capacity GB	Provisioned Capacity Cost
APAC Commercial Sales	Gold-Fast	12	Gold-Fast	12	674.04	\$8,088.42
APAC Commercial Sales	Silver	10	Silver	7	1,903.83	\$13,326.82
APAC Cruises	Gold-Fast	12	Gold-Fast	12	730.20	\$8,762.44
African Tours	Gold	12	Gold	10	4,856.12	\$48,561.16
African Tours	Silver	10	Silver	7	1,480.85	\$10,365.93
CRM	Bronze	3	Bronze	3	5,689.08	\$17,067.23
Caribbean	Gold	12	Gold	10	4,590.41	\$45,904.08
Commercial Applications	Bronze	3	Bronze	3	14,312.88	\$42,938.64
Commercial Applications	Gold-Fast	12	Gold-Fast	12	40,308.42	\$483,701.05
Commercial Environments	Bronze	3	Bronze	3	16,812.27	\$50,436.81
Commercial Environments	Gold	12	Gold	10	9,313.51	\$93,135.13
Commercial Environments	Silver	10	Silver	7	1,480.79	\$10,365.54
Concur	Gold	12	Gold	10	247.39	\$2,473.91
Concur	Gold-Fast	12	Gold-Fast	12	575.17	\$6,902.09
Consumer Feedback	Gold	12	Gold	10	1,335.89	\$13,358.94

Ensuring IO density reports describe only internal data volumes

In NetApp storage systems the root aggregate contains the root volume. The root volume contains special directories and configuration files for managing and controlling the storage system. The management and control operations might result a large amount of activity in the root aggregate. When you query the Insight system for the top 10 internal volumes with the highest IO density, your results might include NetApp root aggregates as members of the top 10.

When monitoring your environment it is more important to determine which internal data volumes are producing high I/O density numbers. In order to accurately identify only the data volumes, you need to isolate the NetApp internal volumes from queries you use to monitor I/O density.

This guide describes how to easily identify the NetApp root aggregates, isolate them from the results of internal volume queries, and create rules that exclude any new NetApp root aggregates as they are added to the system. The following Insight features are used to insure that your I/O density reports are derived from internal data volumes.

- A query is created to identify all of the NetApp root aggregates that are monitored by Insight.

- An annotation is assigned to each of the NetApp root aggregates.
- An annotation rule is created to exclude the NetApp aggregates

Creating a query to identify NetApp root aggregates in your environment

Queries provide search at a granular level, based on user-selected criteria. Using a query allows you to search for Internal volumes in your environment that contain the NetApp root aggregate.

Steps

1. In the OnCommand Insight web UI, create a query to identify NetApp root aggregates in your environment: **Queries > New Query > Select Resource Type**
2. Click **Storage Pool**
3. Enter the name for the root aggregate

This example uses “aggr0” for the name. When creating an aggregate, only the following requirements for the name must be followed:

- It must begin with either a letter or an underscore (_).
- It can contain only letters, digits, and underscores.
- It can be 250 characters or less.

In most cases the aggregate is name aggr0, aggr_0, or something similar. It might require an iterative process to identify all of the NetApp root aggregates in your environment.

4. Click **Save** and enter a name for the new query.

As mentioned before, this might be an iterative process and require multiple queries to identify all of the NetApp root aggregates.

Create an annotation for the root volumes returned by your queries

Annotations are specialized notes that you assign to your assets, allowing you to filter assets by their annotations. The annotation you create will be used to identify the NetApp root aggregates in your environment and ensure that they are not included in a specific report.

Before you begin

You must have identified all of the root aggregates you want to exclude from the “High I/O Density” report.

Steps

1. Create an annotation to associate all of the NetApp root aggregates you identified with queries: **Manage > Annotations**
2. Click **Add**
 - a. Enter the name for the annotation: **RootAggr**
 - b. Enter a description of the annotation: **Remove root aggregate from "High I/O Density" report**

c. Enter the type of annotation: **Boolean**

3. Click **Save**

Create an annotation rule to automate excluding specific aggregates from your I/O density report

As an alternative to manually applying annotations to individual assets, you can automatically apply annotations to multiple assets using annotation rules. Annotation rules are based on queries you create and when run on the system they add new assets to existing sets of assets. When these sets of assets are excluded from a report the new assets are automatically excluded too.

Before you begin

You must have created and saved a query that identifies the NetApp root aggregates you identified in your environment.

Steps

1. Log in to the OnCommand Insight web UI.

2. Click **Manage > Annotation rules**

3. Click **Add**

The Add Rule dialog box displays.

4. Do the following:

a. In the Name box, enter a unique name that describes the rule: "RootAggrExclude"

b. Click Query and select the Query that Insight should use to apply the annotation rule to:" Aggregate0"

c. Click Annotation and select: "Root agg exclude"

d. Click Value and enter True

Collecting integration data

You can import integration data into your OnCommand insight system. Data can be imported using collectd, open source software that runs as a daemon to collect performance data, or by using the integration SNMP data source which allows you to collect generic SNMP data.

Data flow for integration data

The following applies to the total amount of integration data that is allowed to be presented to the OnCommand Insight server:

- A queue of 100 calls is maintained.

When a client waits in the queue for more than one minute, a timeout error occurs.

- The recommended ingestion rate for integration data is once per minute, per client.

- There is a limit of 300 integration object types allowed.

Accessing collectd software and documentation

You can access the output writer plugin software and documentation for collectd at NetApp's GitHub site: https://github.com/NetApp/OCI_collectd

Backup and restore of integration data

Backup and restore of integration data is modeled after OnCommand Insight performance data backup and restore policies. When a backup is configured for performance data, integration data is also included in the backup. As with performance backup, the most recent seven days of integration data is included in the backup. Any integration data that is present in a backup is restored in a restore operation.

Licenses

A Perform license is required for integration data to be reported. If a Perform license is not present an error occurs with a message "Perform license required to report integration data".

Collecting SNMP Integration data

The integration SNMP data source allows you to collect generic SNMP data in OnCommand Insight.

Integration packs

The SNMP Integration data source uses an "Integration Pack" to define what integration values are collected, and what SNMP objects provide those values.

An Integration Pack consists of:

- A JSON configuration file (integration.json) defining integration payload contents in terms of SNMP objects of a specific device type (switch, router, and so on).
- A list of MIB files that the integration pack depends on.

An integration pack can define multiple data types. For example, when integrating an RHEL host, a data type can be defined for the general system information such as uptime, number of users, and number of running processes, a second data type can be defined for data on memory and file system usage. In general, each data type must be "flat" and cannot contain nested data.

A single integration pack should not define more than 24 data types. Insight limits the amount of integration data that is collected. Attempting to ingest more than 24 reports over a period of one minute results in a rate error.

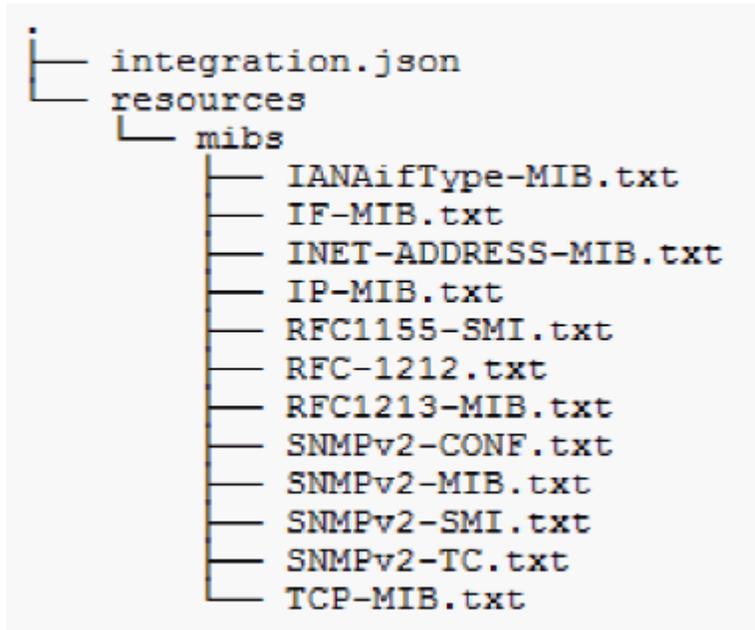
The names for integration types must adhere to the following rules:

- The name cannot start with the following characters: `_`, `-`, `or`, `+`
- The name cannot contain the following characters: `#`, `\`, `/`, `*`, `?`, `"`, `<`, `>`, `|`, `'`, ```,
- Cannot be longer than 100 UTF-8 encoded bytes
- Cannot be named `.` or `..`

Integration file format

An integration pack is a ZIP file that contains a JSON configuration file (`integration.json`) defining the integration payload contents in terms of SNMP objects. It also contains a MIBS folder that contains all of the MIB files and their MIB dependencies.

The `integration.json` file must exist at the top level of the ZIP file and the MIB files must exist in the "resources/mibs" subdirectory within the ZIP. The ZIP file may also contain files, such as a "readme.txt", if desired. An example of integration ZIP structure is:



Importing SNMP integration packs

You import SNMP integration packs into OnCommand Insight using the web UI. Integration packs are identified by the "integrationPacName" value defined in the `integration.json` configuration file contained in the ZIP file.

Before you begin

You must have created a properly formatted ZIP file that contains the integration pack you want to import to the OnCommand Insight server.

About this task

Use the following steps to import SNMP integration packs to the Insight server.

Steps

1. Click **Admin > Setup > SNMP Integration**

The system displays the Import SNMP package screen:

Import SNMP package

Select file No file selected Import

Warning: This will overwrite any conflicting package from existing database.

2. Click **Select file** to select the local file containing the SNMP package.

The file you select is displayed in the File box.



Any existing integration pack with the same name is overwritten.

3. Click **Import**

The file is imported to the Insight server.

Creating an SNMP integration data source

The Integration SNMP Data Source provides common SNMP configuration properties similar to other SNMP based data sources included with the OnCommand Insight data sources for Brocade and Cisco.

Before you begin

In order to successfully use the Integration SNMP data source to collect, the following must be true:

- You must have already imported an integration pack you will use for this SNMP data source.
- All target devices share the same credentials.
- All target devices implement the SNMP Objects referenced by the configured integration pack.

About this task

To create an SNMP Integration data source, choose vendor "Integration" and model "SNMP" in the data source creation wizard.

Steps

1. In the OnCommand Insight web UI, click **Admin > Data Sources**
2. Click **+Add**
3. Enter a name for the Data Source
4. For Vendor, select **Integration**
5. For Model, select **SNMP**

Add data source



Settings

*Name

Vendor

Model

Where to run

What to collect Integration (BETA)

[Configure](#) ↓

Configuration

Advanced configuration

Test

6. For What to collect, check **Integration**

This is the only package on this data source and is checked by default:

7. Click **Configuration**
8. Enter the IP addresses for the systems from which you will collect SNMP data
9. Select an imported SNMP Integration Pack
10. Set the integration poll interval
11. Select the SNMP version
12. Enter the SNMP community string
- For SNMP V1 and V2.
13. Add the user name and password for systems you will be collecting data from.
- For SNMP V3.
14. Click **Advanced Configuration**

The Advanced Configuration default settings are displayed. Make any changes to these settings that are required.

Integration.json file information

The integration.json file identifies the payload .

The following illustration provides a color-coded representation of a simple integration.json file. The accompanying table identifies the function of the objects in the file.

```
{
  "integrationPacName": "WindowsSnmp",
  "description": "Generic integration for mibs supported by the default
SNMP Agent for Windows 2012, including HOST-RESOURCES",
  "acquisitionType": "SNMP",
  "integrationTypes": [
    {
      "integrationType": "snmp_win2012_host",
      "name": {
        "mibModuleName": "RFC1213-MIB",
        "objectName": "sysName"
      },
      "identifiers": {
        "hostname": {
          "mibModuleName": "RFC1213-MIB",
        }
      },
      "attributes": {
        "description": {
          "mibModuleName": "RFC1213-MIB",
          "objectName": "sysDescr"
        },
        "snmp_sys_obj_id": {
          "mibModuleName": "RFC1213-MIB",
          "objectName": "sysObjectID"
        }
      },
      "dataPoints": {
        "uptime": {
          "num": {
            "mibModuleName": "RFC1213-MIB",
            "objectName": "sysUpTime"
          }
        }
      }
    }
  ]
}
```

Blue	Reserved
Red	User customizable strings and IDs
Green	MIB names
Purple	MIB object
Black	JSON structure

About integration.json files

Each field has the following characteristics:

- The "identifiers" section forms a unique compound key to create a new "object" in Insight
- The "attributes" provide supporting meta-data about the object.

In both of these cases, only the value of the latest report for that object (identified by the identifiers) is preserved.

- The "dataPoints" are time-series data and must be numeric values. Insight keeps each and every value reported here for 90 days (by default) and links them time-series to the object identified.

Numeric Expressions

By default, all value expressions are reported as strings in the integration payload. "identifiers" and "attributes" may only define string values. "dataPoints" may define string or numeric values. Numeric values are defined using one of the following modifier keys:

- num - the total number of bytes received since the counter was last initialized
- delta - the number of bytes received during the poll interval
- rate - the average receive rate during the poll interval in bytes per second

An average receive rate during the poll interval in megabytes per second can be accomplished using a combination of rate and math operations

Math operations

The `integration.json` file supports the following math operations: add, subtract, multiply, divide. The following example shows multiplication, division, and sum operations in a JSON file.

```

"network_utilization":
{
  "mult": [
    {
      "div": [
        {
          "sum": [
            "rate": {
              "mibModuleName": "IF-MIB",
              "objectName": "ifHCOutOctets",
              "comment": "bytes per second out"
            },
            "rate": {
              "mibModuleName": "IF-MIB",
              "objectName": "ifHCInOctets",
              "comment": "bytes per second in"
            }
          ]
        },
        {
          "num": {
            "mibModuleName": "IF-MIB",
            "objectName": "ifSpeed",
            "comment": "1,000,000 bits per second"
          }
        }
      ]
    },
    {
      "const": 0.0008,
      "comment": "normalize to ratio of bits and convert to percent:
8 * 100 / 1,000,000 = 0.0008"
    }
  ]
}

```

Keywords

An integration pack keyword, string, is implemented to force OCTET STRINGS or proprietary types derived from OCTET STRING that would normally be rendered in hexadecimal format to instead be rendered as ASCII characters.

Often OCTET STRINGS contain binary data, for example MAC addresses and WWNs:

```

"interface_mac": {
  "mibModuleName": "IF-MIB",
  "objectName": "ifPhysAddress"
}

```

ifPhysAddress is type PhysAddress, which is just an OCTET STRING:

```

PhysAddress ::= TEXTUAL-CONVENTION
                DISPLAY-HINT "1x:"
                STATUS        current
                DESCRIPTION
                    "Represents media- or physical-level
addresses."
                SYNTAX        OCTET STRING

```

When ifPhysAddress is rendered as hex by default, the result is:

```
"interface_mac": "00:50:56:A2:07:E7"
```

However if you have an OCTET STRING or proprietary type derived from OCTET STRING that you want to interpret as ASCII, you can use the "string" keyword:

```

"string_test_1": {
  "string": {
    "mibModuleName":      "IF-MIB",
    "objectName":         "ifPhysAddress"
  }
},

"string_test_2": {
  "string": [
    {
      "mibModuleName":      "IF-MIB",
      "objectName":         "ifPhysAddress"
    },
    {
      "const": "JSD"
    },
    {
      "mibModuleName":      "IF-MIB",
      "objectName":         "ifPhysAddress"
    }
  ]
}

```

The keyword follows the existing string concatenation rules, inserting a single space between terms in the following example:

```
"string_test_1": "PVçç",  
  "string_test_2": "PVçç JSD PVçç"
```

The "string" keyword acts on a single term or a list of terms, but not nested expressions. Nested expressions are only supported for dataPoint expressions. Attempting to use a "string" expression in a dataPoint expression will result in an error similar to the following:

```
java.lang.IllegalArgumentException: Integration pack 'GenericSwitch32' index 'snmp_generic_interface_32'  
section 'dataPoints' key 'string_test_3' unsupported JSON numeric expression  
'{"string":{"mibName":"IF-MIB","objectName":"ifPhysAddress"}}'
```

Some derived OCTET STRING types such as DisplayString, SnmpAdminString have hard-coded precedence over the "string" keyword. This is because SnmpAdminString is specifically UTF-8 encoded, and we want to handle it correctly, whereas the "string" keyword forces the default string representation returned by the snmp_framework, which assumes single byte ascii code points per character.

Analyzing an application performance problem

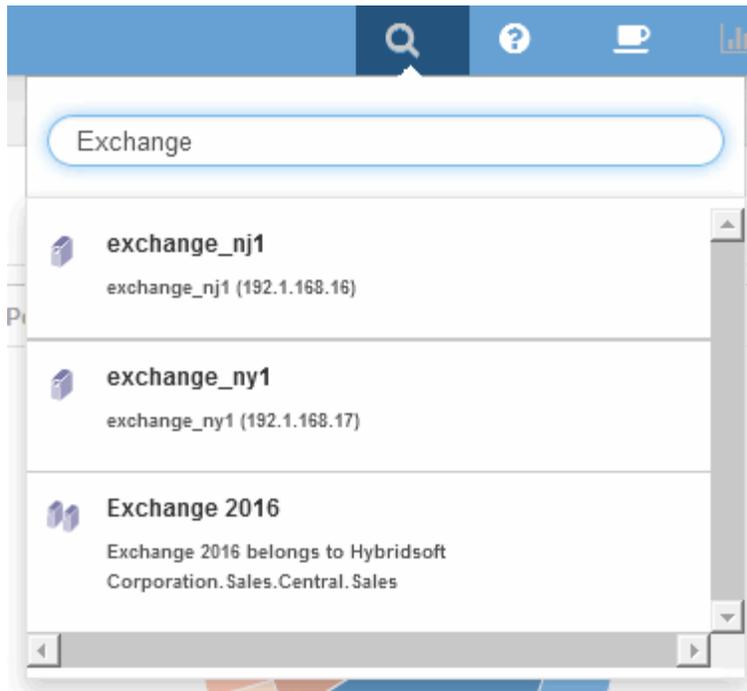
This document describes steps you might take to address reports of performance problems for an application that are impacting users or administrators. For example, users are complaining that their Exchange application is experiencing periods of slowness throughout the day.

About this task

In OnCommand Insight, an application is a configured entity. You assign a name and business entity to the application and you assign compute and storage resources to the application. This allows a better end-to-end view of infrastructure health and more pro-active management of infrastructure asset management.

Steps

1. To begin investigating the issue, use the Insight toolbar to perform a global search for the Exchange application.



When performing a search, you can add an object descriptor before the object name to refine the search results.

- When you select "Exchange 2016" from the search results, the system displays the Application landing page.

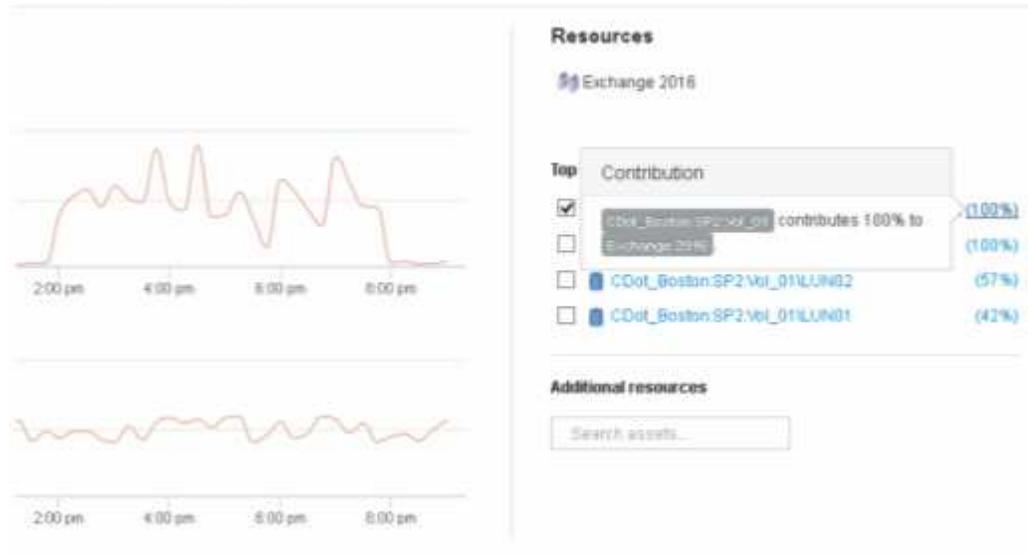


In the Application landing page, the following information is of interest:

- In the 24-hour time period selected, an increase in latency is shown on the right of the latency graph.
- During the period of increased latency there is no significant change in the level of IOPS. It appears the latency increase is not caused by a heavier application usage. We are not really seeing a high IOPS

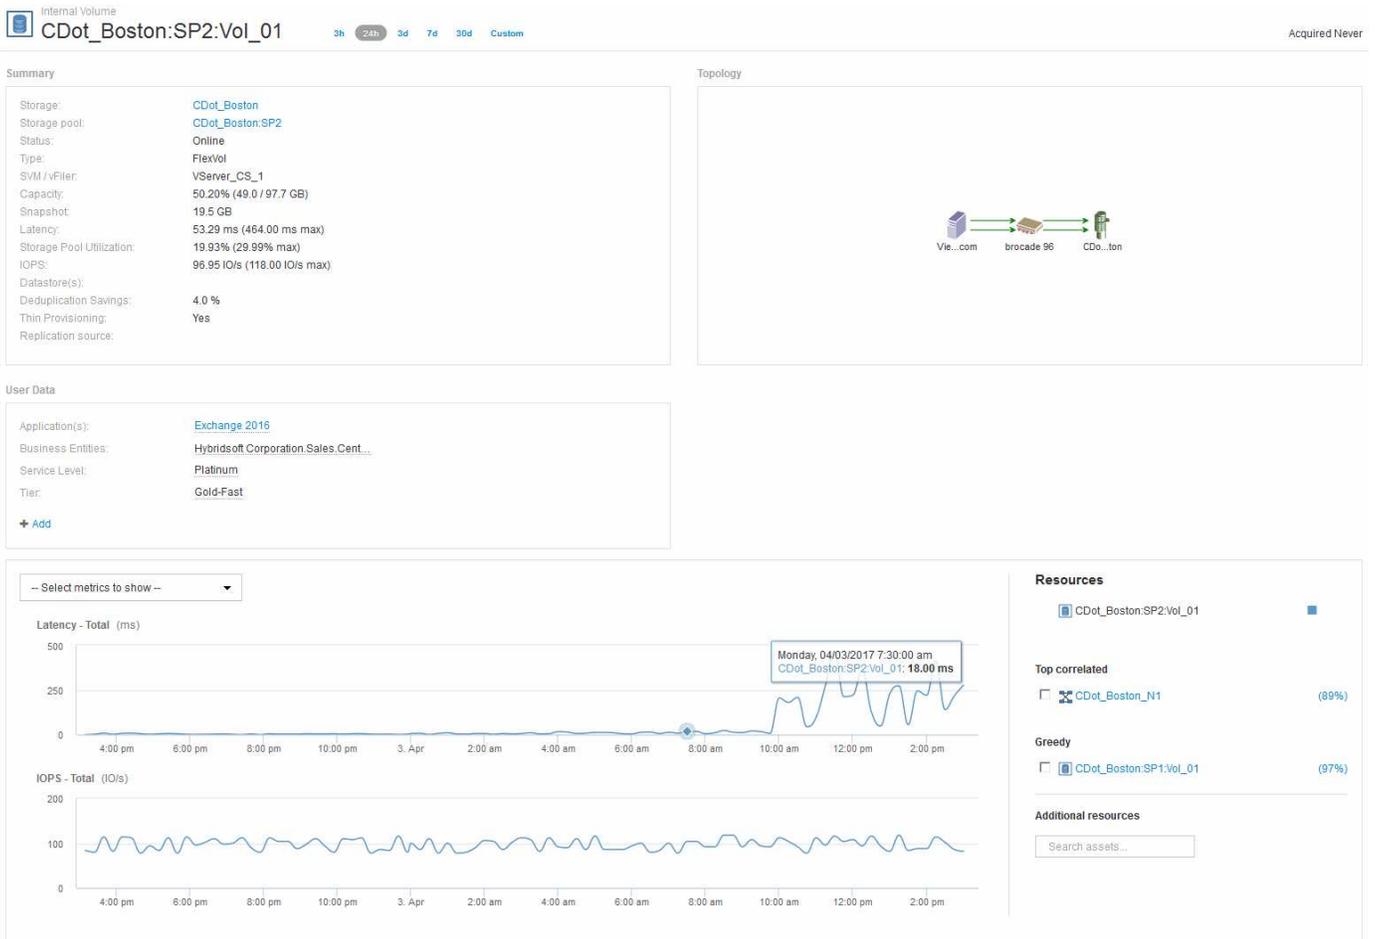
demand on the storage that could account for the latency spike. The increase in latency could be due to an external factor.

- On the right of the charts in the Top contributors section, click on the 100% for the selected internal volume (CDot_Boston:SP2:Vol_01). The system shows this resource is contributing 100% to the Exchange 2016 application.



- Click on the navigation link for this internal volume (CDot_Boston:SP2:Vol_01) to access the internal volume landing page. Analysis of the internal volume might provide information pertaining to the latency spike.

Examining the internal volume



In the Internal Volume landing page, you see:

- The performance charts for the internal volume match what was previously seen for the application performance for both latency and IOPS.
- In the Resources section, where the correlated assets are displayed, a “Greedy” resource is identified (CDot_Boston:SP1:Vol_01).

A greedy resource is identified by insight correlation analytics. Greedy/degraded resources are “peers” that utilize the same shared resource. The greedy resource has IOPS or utilization rates that negatively impact the degraded resource’s IOPS or latency.

Greedy and Degraded resources can be identified on Virtual Machine, Volume, and Internal Volume landing pages. A maximum of two greedy resources will be displayed on each landing page.

Selecting the correlation ranking (%) provides the Greedy resource analysis findings. For example, clicking a greedy percentage value identifies the operation on an asset that impacts the operation on the Degraded asset, similar to what is shown in the following example.

The screenshot shows a resource page for **CDot_Bosto...I_01\LUN01**. Under the **Resources** section, the resource is listed with a blue square icon. Below it, the **Top correlated** section lists **VM_Exchange_1** (98%) and **CDot_Boston_N1** (85%). A **Greedy** section is highlighted with a yellow border, showing **CDot_Boston:SP1:Vol...** (98%). A callout box points to this entry, stating: "Greedy: IOPS of **CDot_Bosto...I_01\LUN01** impacts Latency of **CDot_Bosto...I_01\LUN01** by 98%." (98%).

When a degraded resource is identified, you can select the degraded (%) score to identify the operation and the resource that is impacting the degraded resource.

The screenshot shows a resource page for **CDot_Bosto...I_01\LUN01**. Under the **Resources** section, the resource is listed with a blue square icon. Below it, the **Top correlated** section lists **VM_Cs_travBook** (99%) and **CDot_Boston:SP1** (56%). A **Degraded** section is highlighted with a yellow border, showing **CDot_Boston:SP2:Vol...** (98%). A callout box points to this entry, stating: "Degraded: IOPS of **hionpcmsac...p13_splunk** impacts Latency of **hionpcmsac...4_prd_cl03** by 88%." (88%) and "(40%)".

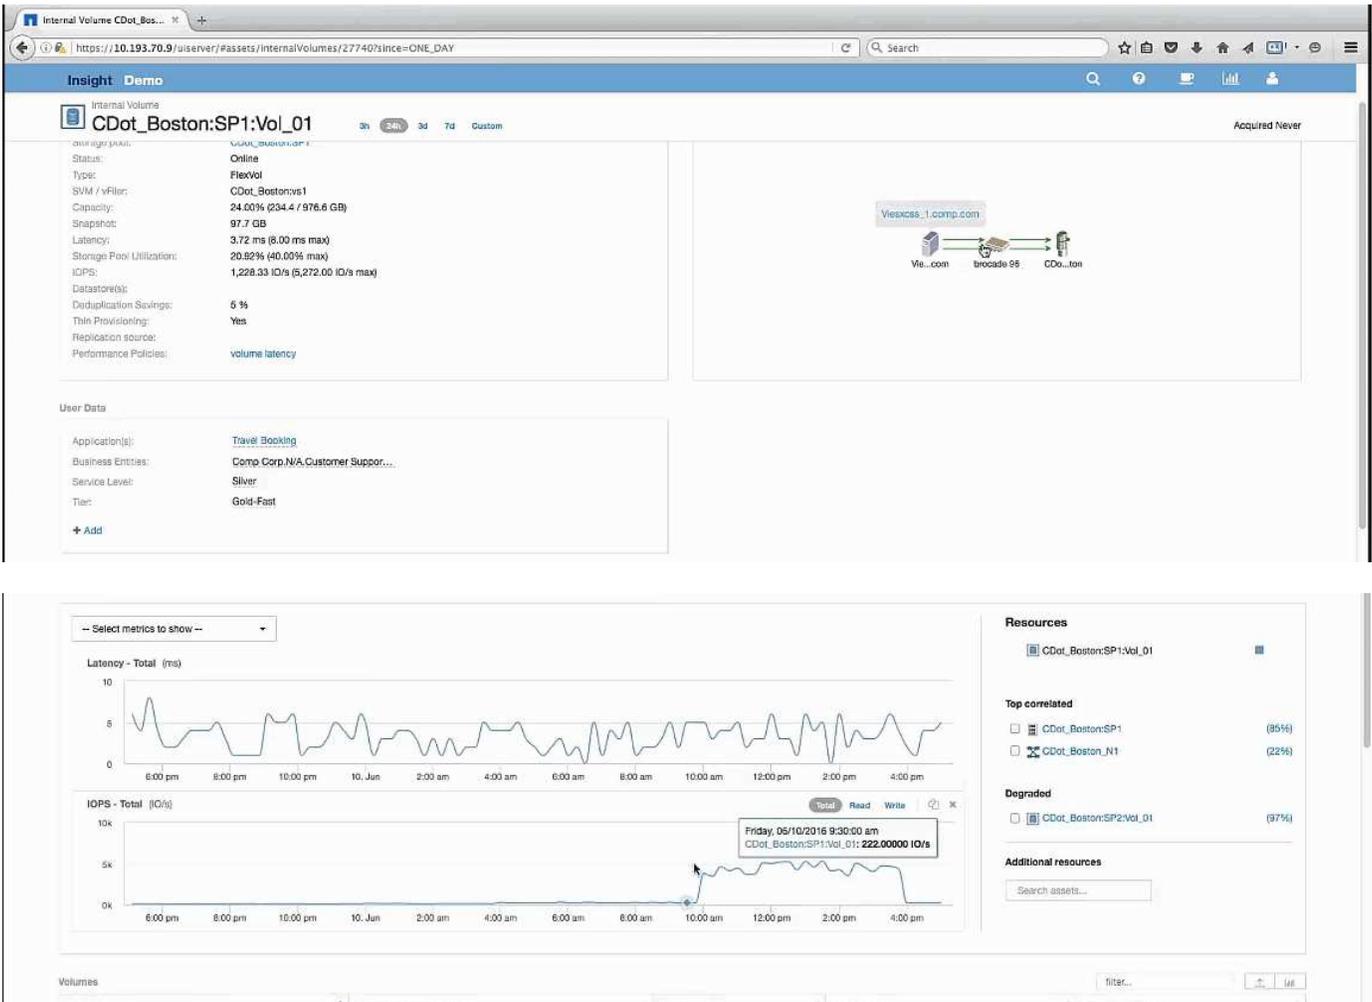
Below the degraded section, the **Additional resources** section contains a search box labeled "Search assets...".

The bottom part of the screenshot shows a resource page for **hionpcmsac...p13_splunk**. Under the **Resources** section, the resource is listed with a blue square icon. Below it, the **Top correlated** section lists **hionpcmsaclu01n01b:...saciu01n01b_ex...**. A **Degraded** section is highlighted with a yellow border, showing **hionpcmsaclu01:svmn...170_vmdk04_p...** (88%) and **hionpcmsaclu01:svmn...180_vmdk04_p...** (40%).

Examining the greedy resource

Clicking on the internal volume identified as the greedy resource opens the landing page for the volume **CDot_Boston:SP1:Vol_01**.

Note in the summary details this internal volume is a resource for a different application (Travel Booking) and although contained in a different storage pool is on the same node as the internal volume for Exchange 2016 (**CDot_Boston_N1**)



The landing page shows:

- The internal volume associated with a Travel Booking application.
- A new storage pool is identified in the correlated resources.
- The original internal volume you were examining (CDot_Boston:SP2:Vol_01) is identified as “Degraded”.
- In the performance graph, the application has a steady latency profile and does have an IOPS spike roughly at the same time we see the latency spike on the Exchange application.

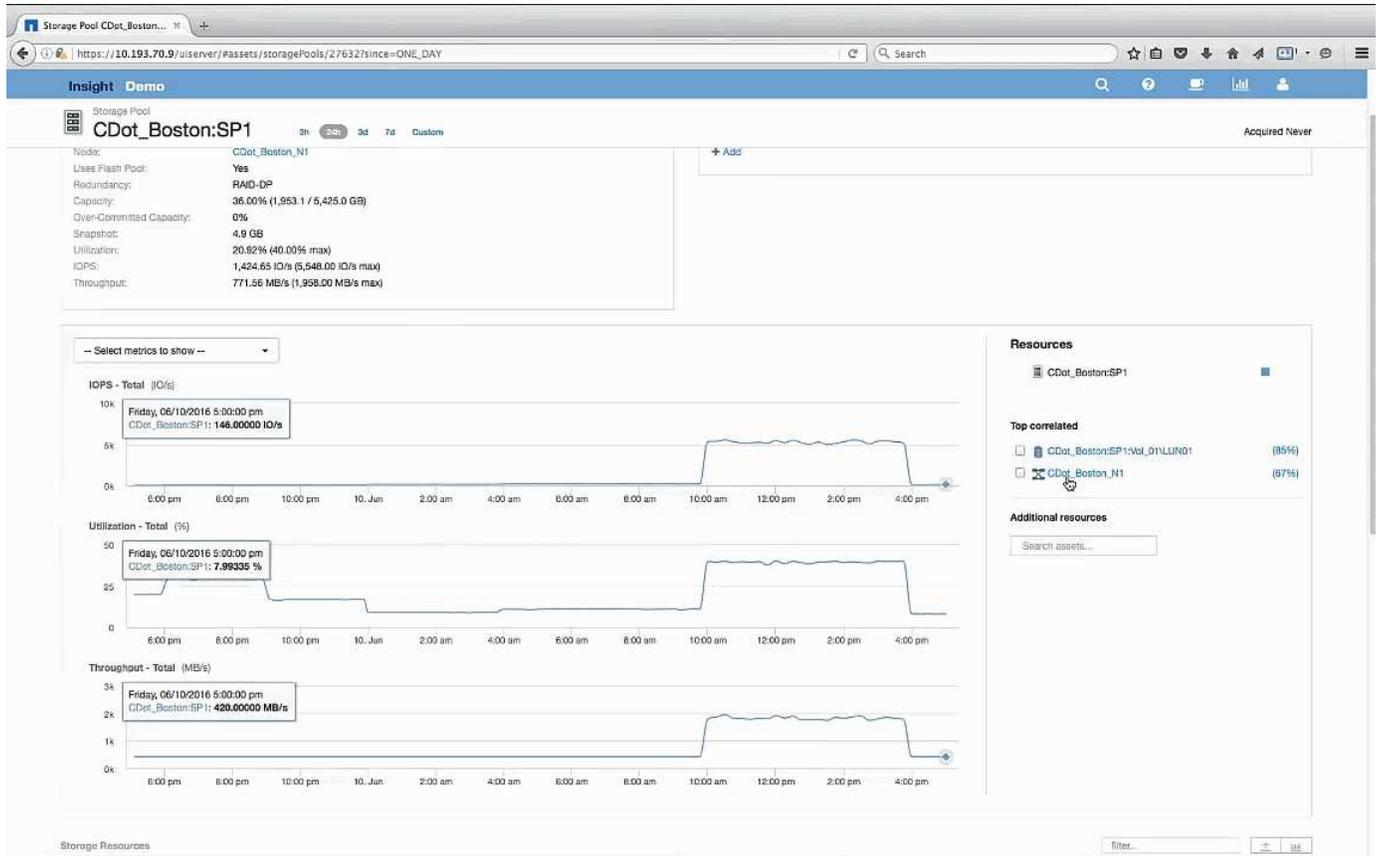
This might indicate that the latency spike on the Exchange application is likely caused by the IOPS spike on this volume.

To the right of the charts in the Resource section notice the correlated Degraded resource which is the Exchange 2016 internal volume (CDot_Boston:SP2:Vol_01). Click on the check box to include the degraded internal volume in the in the performance graphs. Aligning the two performance graphs shows that the latency and IOPS spikes occur at nearly the exact same time. This tells us that we want to get a better understanding of the Travel Booking application. We need to understand why the application is experiencing such a prolonged IOPS spike.

Examining the Storage pool associated with the Travel Booking application might identify why the application is experiencing the IOPS spike. Click CDot_Boston:SP1 to view the Storage Pool landing page.

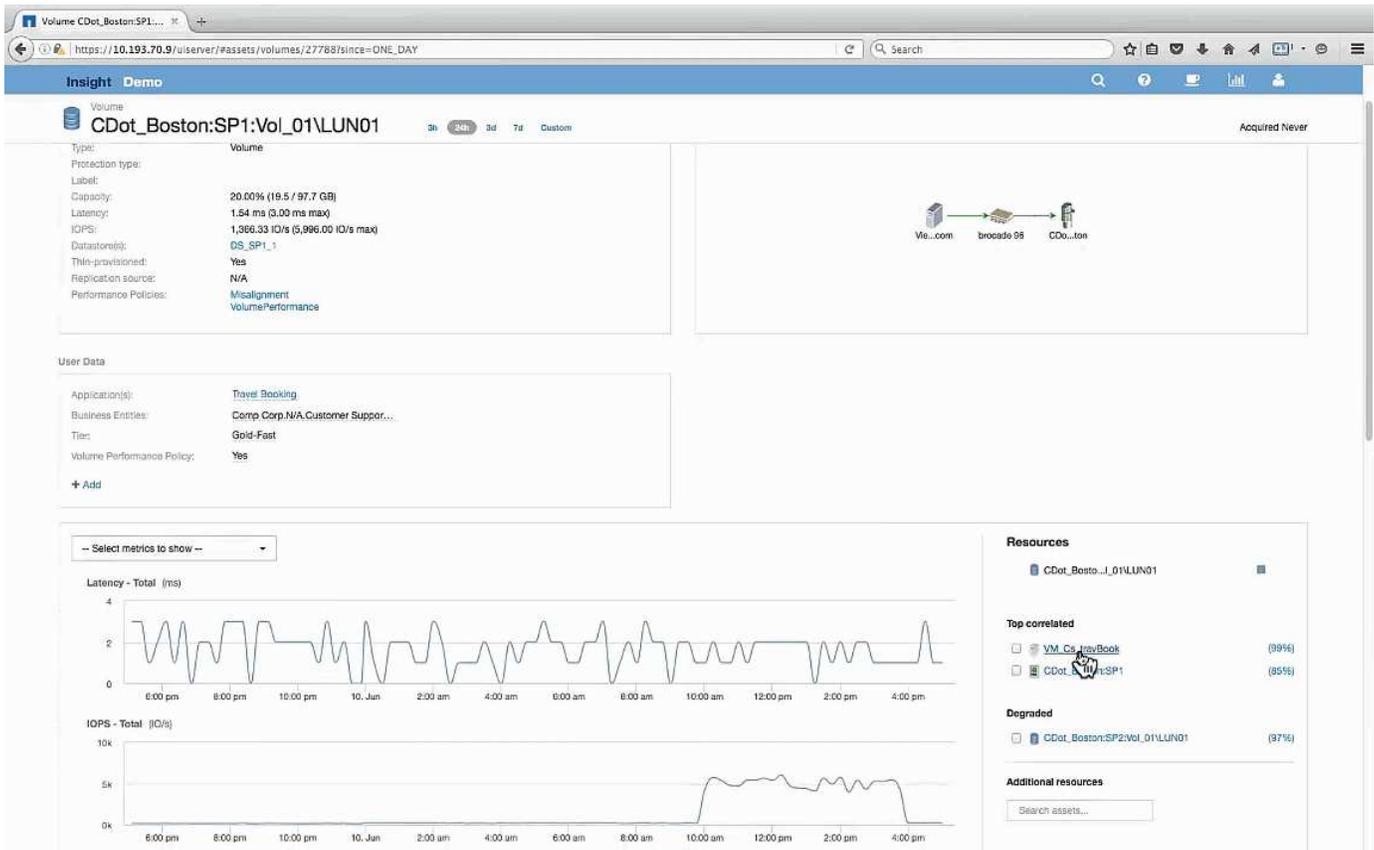
Examine the storage pool

Examining the storage pool landing page shows the same IOPS spike seen in its correlated assets. In the Resources section you can see that this storage pool landing page links to the volume of the travel application. Click on the volume to open the volume landing page.



Examining the volume

The volume landing page shows the same familiar IOPS spike seen in its correlated assets.



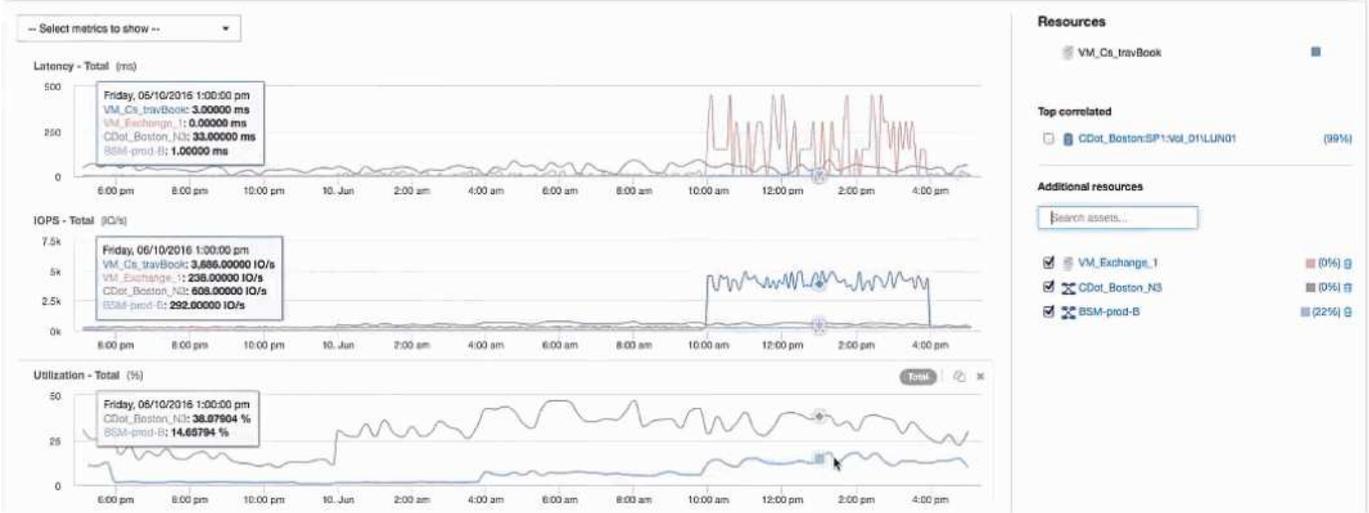
In the resources section the VM for the Travel Booking application is identified. Click on the VM link to view the VM landing page.

Examining the VM

In the VM landing page, select additional metrics to display and include CPU utilization and Memory utilization. The graphs for CPU and Memory utilization show that both are operating at nearly 100% of their capacity. This tells us that the problem with the Exchange server is not a storage problem, but instead is the result of the high VM CPU and memory utilization and the consequential memory swapping of I/O to disk.



To solve this problem, you can look for additional similar resources. Enter “Node” in the Additional resources input dialog to show metrics for assets similar to the Exchange VM. The comparison can help identify a node that might be a better fit for hosting the workload should a change be necessary.



Collecting and reporting AWS billing data

The Amazon AWS Cloud Cost data source imports billing data generated by Amazon into Insight as integration data, making it available to the data warehouse for reporting.

There are three parts to making cloud billing data available to Insight:

Verifying your AWS account information.

Configuring the AWS Cloud Cost data source in Insight to collect the data.

Sending the data to Data Warehouse via ETL for use in reports.

Preparing AWS for Insight data collection

Your AWS account must be properly configured to allow Insight to collect cloud cost data.

About this task

The following steps are done through your AWS account. See the Amazon documentation for more information: <http://docs.aws.amazon.com>. If you are unfamiliar with setting up an AWS cloud account, contact your cloud provider for assistance.



These steps are provided here as a courtesy and are believed correct as of the time of publication. NetApp makes no guarantee of the correctness of these steps. Contact your cloud provider or AWS account holder for information or assistance on configuring your AWS account.

Best practice: Insight recommends that you create a primary IAM user on the same account that owns the S3 bucket where the billing reports are uploaded, and use this user to configure and collect AWS billing data.

To configure your AWS account to allow Insight to collect data, perform the following steps:

Steps

1. Log in to your AWS account as an Identity Access Management (IAM) user. For proper collection, log in to the primary IAM account, as opposed to a group IAM account.
2. Go to **Amazon S3** to create your bucket. Enter a unique bucket name and verify the correct Region.
3. Turn on your Amazon Cost and Usage Report. See <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-reports-gettingstarted-turnonreports.html> for information.
 - a. Go to the **AWS Billing and Cost Management Dashboard** and choose **Reports**.
 - b. Click on **Create report** and enter in the Report Name. For **Time unit**, choose Daily. Check the box to include **Resource IDs**, and click **Next**.
 - c. Click on the **Sample Policy** link in the Select delivery options page. Copy the Sample Policy text in the box to the clipboard. Click **Close**.
 - d. Go back to the S3 Bucket that was created, click on the **Permissions** tab and select the **Bucket Policy** button.
 - e. Paste the text from the Sample Policy, and replace <bucketname> with your actual bucket name in each line that contains the following: "Resource": "arn:aws:s3:: <bucketname>". **Save** the policy.
 - f. Go back to your Create Report screen, enter in your S3 Bucket and click the **Verify** button. Click **Next**.
 - g. Verify your information and click **Review and Complete**.
4. You must grant permissions in order for Insight to collect data from AWS. The following link provides details on how to grant permissions to **List All Buckets** (Step 4.1) and set permissions on the objects in the folder (Step 5.2): <https://docs.aws.amazon.com/AmazonS3/latest/dev/walkthrough1.html>.

5. In the IAM console, go to **Policies** and click **Create policy**.
6. Enter a name in the **Policy Name** field, and click **Create policy** at the bottom.
7. In the IAM console, select your user, then select **Add Inline Policy** at the bottom of the screen.
8. Click on **Choose a service** and select S3.
9. Go to the **JSON** tab. Copy the JSON sample text from step 5.1.2.g of the AWS walkthrough into the JSON box.
10. Replace the *companybucket* and *Development* fields in the JSON with your S3 information.
11. Click **Review Policy** to review your policy settings.

Configuring the AWS Cloud Cost data source

You configure the AWS Cloud Cost data source as you would for any Insight data source.

Before you begin

You must have your Amazon AWS account already set up and prepared for Insight data collection, and have the following pieces of information at hand.

- Report Name
- S3 Bucket Name
- AWS Region where your S3 bucket resides.
- Report path prefix

About this task

Once your AWS account is ready and has the proper permissions set, you are ready to configure OnCommand Insight to collect billing report data.



You will need to add a separate AWS Cloud Cost data source for each billable user/account from which you wish to retrieve billing data.

Steps

1. Log in to OnCommand Insight as an administrator.
2. Click on **Admin > Data sources** to open the Insight Data Source page.
3. Click **+Add** to add a new data source. Choose **Amazon** and select **AWS Cloud Cost**.
4. In the **Configuration** section, fill in the *Report name*, *S3 Bucket name*, *S3 Region* (must be the region where your S3 bucket resides), *Report path prefix*, *AWS IAM Access Key ID*, and *AWS IAM Secret Access Key*. If you are unsure of any of these, check with your cloud provider or AWS account holder.
5. Click the checkbox to verify your understanding that AWS will bill you for API requests and data transfers made by the Insight data source.
6. In **Advanced Configuration**, enter the HTTP connection and socket timeout. The default is 300 seconds.
7. Click **Save**.

Processing AWS Cloud Cost data in Insight

Insight collects data from your AWS billing report once a month for the previous month, and reflects the finalized cloud cost for that month.

After you set up your AWS Cloud Cost data source(s), if you already had billing reports generated to S3, you will get up to three months of past data immediately after the first data source poll.

Insight collects AWS “final” data once a month. This collection occurs a few days after the close of the previous month, allowing AWS time to finalize the actual data.

AWS billing data is sent to Insight’s Data Warehouse for use in reporting.

Keep in mind that each data source must be configured for a single billable account/user.

Reporting on Cloud Cost data in Insight

Cloud cost monthly data collected in Insight is sent to the data warehouse and is available in the Cloud Cost datamart for use in reports.

Before you begin

You must have data sources configured to collect cloud cost data from AWS. Each billable user/account must have a separate data source.

Allow Insight at least 36 hours to begin collecting data.

Allow ETL to run at least once after that time, to send the data to the data warehouse.

About this task

After your data has been collected and sent to the data warehouse, you can view it in any of several pre-configured reports, or create custom reports. Insight stores the data in its own Cloud Cost datamart.

To view your cloud cost data in one of the pre-configured reports:

Steps

1. Open Insight Reporting by one of these methods:
 - Click on the Reporting Portal icon  in the Insight server web UI or in the Data Warehouse UI.
 - Launch Reporting directly by entering the following URL: https://<dw_warehouse_name>:9300/p2pd/servlet/dispatch or https://<dw_warehouse_name>:9300/bi (7.3.3 and later)
2. Once you are logged in to Reporting, click on **Public Folders** and select **Cloud Cost**.
3. You can view your AWS billing data in the available reports located in the **Cloud Cost** folder, or create your own custom report using the **Cloud Cost datamart** available from the **Packages** folder.

Integrating with ServiceNow

OnCommand Insight integrates with ServiceNow management software to provide greater value than the products have separately.

Using a Python script, Insight can integrate data with ServiceNow, synchronizing the following information:

- Storage asset data for ServiceNow servers
- Host and VM URLs for ServiceNow servers
- Relationships between Hosts/VMs and Storage

Preparation and prerequisites for Service Now integration

The necessary preparations and prerequisites must be satisfied for ServiceNow, Insight, and the Python middleware connector prior to integration.

Recommended workflow

The following workflow is strongly recommended when integrating ServiceNow with Insight:

1. Deploy the Python middleware connector in your development instance first.
2. Once you have confirmed all faults have been identified and corrected in your development instance, deploy the connector in your test/stage instance.
3. Once you have confirmed correct operation in your staging instance, deploy the connector in your production instance.

If problems are found during any of these stages, follow your rollback steps and disable the connector, then troubleshoot the problem and re-deploy.

General prerequisites:

- You can use either a standalone host or VM (recommended) or the Insight server host/VM to host the python middleware connector.
- It is highly recommended to backup the production Insight server and deploy it on a development instance.
- ServiceNow must be accurately discovering servers in the CMDB.
- Insight must be accurately discovering your storage and compute environments.
- Port 443 and 80 to the Insight Server and ServiceNow Instance.

ServiceNow prerequisites:

- It is highly recommended to use a development/test instance.
- Permission to load ServiceNow update sets.
- Permission to create users.
- ServiceNow version Jakarta or later

Insight prerequisites:

- It is highly recommended to use a development/test instance.
- Permission to create users (Admin permissions).
- Insight version 7.3.1 or later is supported, but to get the most out of Insight, use the latest version.

Python middleware connector prerequisites:

- Python version 3.6 or greater installed.
- When installing Python, check the box to enable all users. This sets Python for standard application install locations.
- When installing Python, check the box to enable the installer to update the path. Otherwise, you will have to update the path manually.
- Download the Python **pysnow** and **requests** libraries.

Downloading the ServiceNow Python connector

You must download the Python connector for ServiceNow integration and extract it to a location of your choosing.

Steps

1. Download the **ServiceNow Integration connector** from the [NetApp Storefront](#).
2. Extract the .zip file to a folder, for example `c:\OCI2SNOW`.

The integration connector script is named `oci_snow_sync.pyz`.

Configuring ServiceNow for integration

Integrating ServiceNow with Insight requires several setup tasks.

About this task

The following tasks must be performed when integrating ServiceNow with Insight:

On the ServiceNow side:

- Elevate Role
- Install Update Sets
- Set up users

On the Insight side:

- Add the ServiceNow user

On the Python connector side:

- Install Python
- Install additional libraries
- Initialize the connector
- Edit the config.ini file
- Test the connector
- Synchronize the connector
- Schedule daily task execution

Each of these is explained in greater detail in the following sections.

Elevate role

You must elevate your ServiceNow role to `security_admin` before you can integrate with insight.

Steps

1. Log into your ServiceNow instance with administrator permissions.
2. Under the **System Administrator** drop-down, choose **Elevate Roles** and elevate your role to `security_admin`. Click OK.

Install update set

As part of the integration between ServiceNow and OnCommand Insight you must install an Update Set, which loads pre-configured data into ServiceNow in order to provide the connector with specific fields and tables for extracting and loading data.

Steps

1. Navigate to the remote update sets table in ServiceNow by searching for “Retrieved update sets”.
2. Click on **Import Update Set from XML**.
3. The update set is in the Python connector .zip file previously downloaded to your local drive (in our example, the `C:\OCI2SNOW` folder) in the `\update_sets` sub-folder. Click on **Choose File** and select the .xml file in this folder. Click **Upload**.
4. Once the Update Set is loaded, open it and click on **Preview Update Set**.

If errors are detected, you must correct them before you can commit the Update Set.

5. If there are no errors, click **Commit Update Set**.

Once the Update Set has been committed it will show on the **System Update Sets > Update Sources** page.

ServiceNow integration - Set up user

You must set up a ServiceNow user for Insight to connect with and synchronize data.

About this task

Steps

1. Create a services account in ServiceNow. Login to ServiceNow and navigate to **system security > users and groups > users**. Click on **New**.
2. Enter a user name. In this example, we will use “OCI2SNOW” as our integration user. Enter a password for this user.



In this How-to we use a services account user named “OCI2SNOW” across the documentation. You may use a different services account, but be sure it is consistent across your environment.

3. Right-click on the menu bar and click **Save**. This will allow you to stay on this user in order to add roles.
4. Click **Edit** and add the following roles to this user:
 - asset
 - import_transformer
 - rest_service
5. Click **Save**.
6. This same user must be added to OnCommand Insight. Log in to Insight as a user with Administrator permissions.
7. Navigate to **Admin > Setup** and click on the **Users** tab.
8. Click the **Actions** button and select **Add user**.
9. For name, enter “OCI2SNOW”. If you used a different user name above, enter that name here. Enter the same password you used for the ServiceNow user above. You may leave the email field blank.
10. Assign this user the **User** role. Click **Save**.

Install Python and libraries

Python can be installed on the Insight server or on a standalone host or VM.

Steps

1. On your VM or host, download Python 3.6 or later.
2. Choose custom installation and choose the following options. These are either necessary for proper connector script operation or are highly recommended.
 - Install launcher for all users
 - Add Python to the PATH
 - Install pip (which allows Python to install other packages)
 - Install tk/tcl and IDLE
 - Install the Python test suite
 - Install py launcher for all users
 - Associate files with Python
 - Create shortcuts for installed applications
 - Add python to environment variables
 - Precompile standard library
3. After Python is installed, install the “requests” and “psnow” Python libraries. Run the following command:

```
python -m pip install requests pysnow
```

NOTE: This command might fail when you are operating in a proxy environment. To work around this issue, you need to manually download each one of the Python Libraries and run the install requests one by one and in the correct order.

The command will install several files.

4. Verify the Python libraries are installed correctly. Start Python using one of the following methods:
 - Open a cmd prompt and type `python`
 - On Windows, open **Start** and choose **Python > python-<version>.exe**
5. At the Python prompt, type `modules`

Python will ask you to wait a moment while it gathers a list of modules, which it will then display.

Setup Python middleware

Now that Python and the necessary libraries are installed, you can configure the middleware connector to communicate with OnCommand Insight and ServiceNow.

Steps

1. On the host or VM where you downloaded the connector software, open a cmd window as administrator and change to the `\OCI2SNOW\` folder.
2. You must initialize the script to generate an empty **config.ini** file. Run the following command:
`oci_snow_sync.pyz init`
3. Open the **config.ini** file in a text editor and make the following changes in the [OCI] section:
 - Set **url** to `https://&name.domain&name.domain` or `https://&ip` address for the Insight instance.
 - Set **user** and **password** to the Insight user created, for example, OCI2SNOW.
 - Set **include_off_vms** to **false**
4. In the [SNOW] section, make the following changes:
 - Set **Instance** to the FQDN or ip address for your ServiceNow instance
 - Set **User** and **Password** to the ServiceNow service account user, for example, the OCI2SNOW.
 - Under **Field for the OCI URL**, set the **url** field to "u_oci_url". This field is created as part of the connector OCI update set. You can change this in the customer environment, but if you do so, you need to modify it here and in ServiceNow. Best practice is to leave this field as is.
 - Set the **filter_status** field to "Installed, In Stock". If you have a status that is different, you must set that status here in order to get all the records to match with Insight records prior to upload of new records. In most cases this field should remain unchanged.
 - Set **stale_status** to "Retired".
5. The [Proxy] section is only required if you use a proxy server. If you need to use this section, ensure the following settings:
 - `;https = http://&host&host:&port&port`;
 - `;http = http://&host&host:&port&port`;
 - `;include_oci = True`
 - `;include_snow = True`
6. Edit the [Log] section only if you need deeper debug information.

7. To test the connector, open a cmd prompt as administrator and change to the \OCI2SNOW folder. Run the following command: `oci_snow_sync.pyz test`

Details can be seen in the `logs\` folder.

Syncing the connector

Once ServiceNow, Insight and the connector are properly configured, you can synchronize the connector.

Steps

1. Open a cmd prompt and change to the \OCI2SNOW folder.
2. Run the following command twice. The first sync updates the items, the second sync updates the relationships: `oci_snow_sync.pyz sync`
3. Verify that the Storage Server table in your ServiceNow instance is populated. Open a storage server and verify that resources related to that storage are listed.

Scheduling synchronization to occur daily

You can use the Windows task scheduler to automatically sync the ServiceNow connector.

About this task

Automatic synchronization ensures that Insight data is regularly moved to ServiceNow. You can use any method for scheduling. The following steps use the Windows task scheduler to accomplish automatic syncing.

Steps

1. On the Windows screen, click **Start** and enter **run > task scheduler**.
2. Click **Create Basic Task...**
3. Enter a meaningful name, such as "OCI2SNOW Connector Sync". Enter a description of the task. Click **Next**.
4. Select to run the task **Daily**. Click **Next**.
5. Choose a time of day to run the task. Click **Next**.
6. For Action, select **Start a program**. Click **Next**.
7. In the **Program/script** field, enter `C:\OCI2SNOW\oci_snow_sync_pyz`. In the **Arguments** field, enter `sync`. In the **Start in** field, enter `C:\OCI2SNOW`. Click **next**.
8. Review the Summary details, and click **Finish**.

The synchronization is now scheduled to run daily.

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

<http://www.netapp.com/us/legal/copyright.aspx>

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/us/media/patents-page.pdf>

Privacy policy

<https://www.netapp.com/us/legal/privacypolicy/index.aspx>

Notice

Notice files provide information about third-party copyright and licenses used in NetApp software.

[Notice for OnCommand Insight 7.3.15](#)

[Notice for OnCommand Insight 7.3.14](#)

[Notice for OnCommand Insight 7.3.13](#)

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.