


Isilon InsightIQ

Administration Guide

4.1.4.x

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Chapter 1: Introduction to this guide.....	5
About this guide.....	5
Where to go for support.....	5
Chapter 2: Configuration.....	6
Configuration and setup overview.....	6
Configuring the InsightIQ system.....	6
Specify the InsightIQ port.....	6
Specify an SSL certificate.....	7
Connect to InsightIQ over IPv6.....	7
Configure outbound email support.....	8
Monitoring the InsightIQ system.....	8
View InsightIQ status.....	8
Disable status alert email messages.....	8
Enable InsightIQ status alerts.....	9
yum update commands.....	9
Configuring File System Analytics job settings.....	10
Viewing FSA job settings.....	10
Enable and disable FSA.....	10
Configure File System Analytics (FSA) job settings.....	11
Managing File System Analytics result sets.....	12
Configuring cluster monitoring.....	14
Monitor a cluster.....	14
View monitored cluster settings.....	14
Modify cluster login credentials.....	15
Suspend and resume InsightIQ monitoring.....	15
Delete an InsightIQ datastore.....	15
View the connection status of a monitored cluster.....	16
Configuring datastores.....	16
Datastore requirements.....	17
View datastore settings.....	18
Export a datastore.....	18
Import a datastore.....	19
Move a datastore.....	20
Automatic data pruning.....	20
Configuring LDAP for authentication.....	21
Connect InsightIQ to an LDAP server.....	21
Add an LDAP group.....	21
Setting up data sets.....	22
Data set configuration table.....	22
Data sets, modules, and breakouts.....	22
View data set sizes.....	24
Enable or disable retrieval of data sets.....	24
Setting up user accounts.....	25

Create a user account.....	25
Modify a user account password.....	25
Delete a user account.....	26
Chapter 3: Troubleshooting InsightIQ.....	27
Troubleshooting overview.....	27
Troubleshooting configuration issues.....	27
Troubleshooting datastore issues.....	28
Troubleshooting data retrieval issues.....	28
Troubleshooting datastore move issues.....	28
Troubleshooting memory issues.....	29
Troubleshooting InsightIQ web application issues	29
Chapter 4: InsightIQ commands.....	30
InsightIQ commands overview.....	30
iiq_data_export fsa describe	30
iiq_data_export fsa export.....	31
iiq_data_export fsa list.....	35
iiq_data_export perf describe.....	36
iiq_data_export perf export.....	38
iiq_data_export perf list.....	43
iiq_restart.....	44
iiq_stop.....	44
iiq_start.....	44

Introduction to this guide

Topics:

- [About this guide](#)
- [Where to go for support](#)

About this guide

This guide describes how to configure and manage InsightIQ.

Your suggestions help us to improve the accuracy, organization, and overall quality of the documentation. Send your feedback to <http://bit.ly/isilon-docfeedback>. If you cannot provide feedback through the URL, send an email message to docfeedback@isilon.com.

Where to go for support

This topic contains support resources for Dell EMC Isilon InsightIQ.

Telephone support	<ul style="list-style-type: none">• United States: 1-800-SVC-4EMC (1-800-782-4362)• Canada: 1-800-543-4782• Worldwide: 1-508-497-7901• Local phone numbers for a specific country or region are available at Dell EMC Customer Support Centers.
Info Hubs	<ul style="list-style-type: none">• InsightIQ Info Hub: https://www.dell.com/support/kbdoc/000129563/• OneFS Info Hubs: https://www.dell.com/support/article/sln318794
Community support	<ul style="list-style-type: none">• Dell Community

Configuration

This section contains the following topics:


Topics:

- [Configuration and setup overview](#)
- [Configuring the InsightIQ system](#)
- [Monitoring the InsightIQ system](#)
- [Configuring File System Analytics job settings](#)
- [Configuring cluster monitoring](#)
- [Configuring datastores](#)
- [Configuring LDAP for authentication](#)
- [Setting up data sets](#)
- [Setting up user accounts](#)

Configuration and setup overview

InsightIQ system settings are managed in the InsightIQ web application.

The InsightIQ Administrator account can manage InsightIQ monitored cluster, data store, email, and file system settings by using the InsightIQ web application. These settings are not available through the command-line interface. You can also manage network and authentication settings for the InsightIQ virtual machine.

 **NOTE:** Depending on which version of the OneFS operating system the monitored cluster is running, certain InsightIQ features might not be available.

Configuring the InsightIQ system

Configure InsightIQ.

InsightIQ configuration tasks include installation of SSL certificates, modifying port access, and configuring email settings.

Specify the InsightIQ port

Specify the port through which InsightIQ connects.

About this task

By default, InsightIQ runs on the standard HTTPS port of 443.

InsightIQ also listens on the standard HTTP port of 80, but this port is only used to redirect users to the HTTPS port, allowing users to type in a URL in a browser without explicitly adding the port number.

To run InsightIQ on a different TCP port you must modify two different configuration files, make the necessary firewall updates, and then restart InsightIQ.

Steps

1. Open an SSH connection to the InsightIQ virtual machine and log in.
2. In a text editor, open the `/etc/isilon/insightiq.ini` file, and specify the values for the following options.

Option	Description
<i>redirect_from_port</i>	The standard non-secure HTTP port. This value should be set to 80.
<i>redirect_to_port</i>	The secure HTTPS port used to connect to InsightIQ. The default value is 443.

You must have root permissions to modify the `/etc/isilon/insightiq.ini` file. If you are logged in the InsightIQ administrator account, you can gain root access by beginning a command with `sudo`.

Save and close the `/etc/isilon/insightiq.ini` file.

- Update the port number used in the SSL certificate: In a text editor, open the `/etc/isilon/uwsgi.ini` file, and update the port number:

```
https = [::]:<port_number>,/etc/ssl/certs/server.crt,/etc/ssl/certs/server.key,HIGH
```

Save and close the `/etc/isilon/uwsgi.ini` file.

- Restart InsightIQ by running the command `iiq_restart`.

Specify an SSL certificate

Specify a custom SSL certificate.

About this task

Although InsightIQ includes a default SSL certificate, you can specify a custom SSL certificate.

Steps

- Open an SSH connection to the InsightIQ virtual machine, and log in.
- On the InsightIQ virtual machine, save a copy of the SSL certificate files that you want to specify.
The certificate files must be of the `.crt` and `.key` file types.
- Open the `/etc/isilon/uwsgi.ini` file in a text editor and update the value of the `https` setting with the path of the SSL certificate files.

For example, the following text specifies the SSL certificate files that are in the `/etc/ssl/certs/` directory:

```
https = [::]:443,/etc/ssl/certs/server.crt,/etc/ssl/certs/server.key,HIGH
```

- Restart InsightIQ by running the following command:

```
iiq_restart
```

Connect to InsightIQ over IPv6

You can connect to InsightIQ, monitor clusters, and connect to NFS data stores over IPv6.

You can monitor an Isilon cluster (running OneFS 7.2.1 or later) over IPv6, only if IPv6 addresses are configured for both the monitored cluster and InsightIQ. The browser you use determines what IPv6 connection types are available.

Table 1. Browser support with IPv6

Web browser	Supported connection type
Apple Safari	DNS hostname or an SSL connection
Google Chrome	DNS hostname
Mozilla Firefox	DNS hostname or IP address

Configure outbound email support

To send reports and InsightIQ status alerts by email, configure InsightIQ to use an email server.

Steps

1. In the InsightIQ web application, click the **Settings** tab and then **Email** on the **Settings** ribbon. The **Configure Email Settings (SMTP)** view appears.
2. In the **SMTP server** field, type the hostname or IP address of an SMTP server that handles email for the organization.
3. In the **SMTP port** field, type the number of the port that is used to connect to the SMTP server that you specified.
4. If the SMTP server that you specified requires a username and password for authentication, in the **Username** and **Password** fields, specify a valid username and password.
5. If the SMTP server you specified accepts email only from valid email addresses, type a valid email address in the **From Email** field.

The address that you type appears in the **From** field of email messages.

6. If either the Transport Layer Security (TLS) or Secure Sockets Layer (SSL) protocol is required to connect to the SMTP server that you specified, select the **TLS Connection** checkbox.
7. Click **Submit**.

Testing the email configuration

Validate the InsightIQ email configuration by sending a test email.

Steps

1. In the InsightIQ web application, click the **Settings** tab and then **Email** on the **Settings** ribbon.
2. In the **Send a test email** field, type the name of an email address.
3. Click **Send**.
4. Check the recipient email inbox.
If you have configured InsightIQ correctly, a test email arrives.

Monitoring the InsightIQ system

View the status of InsightIQ, configure InsightIQ to send alert email messages, and confirm virtual machine updates.

View InsightIQ status

View the status of InsightIQ, including the clusters InsightIQ monitors.

Steps

1. In the InsightIQ web application, click the **Settings** tab.
2. Click **Status** on the **Settings** ribbon.
3. On the **InsightIQ Status** page, view the status of InsightIQ.

Disable status alert email messages

Disable InsightIQ status alert email messages.

Steps

1. In the InsightIQ web application, click the **Settings** tab.
2. Click **Status** on the **Settings** ribbon.

3. Clear the checkbox **Enable InsightIQ status email alerts**.

Enable InsightIQ status alerts

Configure InsightIQ to send an alert email message on the status of InsightIQ.

Steps

1. In the InsightIQ web application, click the **Settings** tab.
2. Click **Status** on the **Settings** ribbon.
3. Select **Enable InsightIQ status email alerts**.
4. In the **Send InsightIQ status email alerts to** field, type the name of an email address that you want to send alerts to.
5. Optional: Optionally, to limit the number of alert email messages of the same type within a specified period, perform the following steps.
 - a. Select **Limit alert email repetition**.
 - b. In the **Do not send multiple status alerts of the same type within (hours)** box, type the number of hours to limit alerts for a single type of alert.

For example, if you configure InsightIQ to **Limit alert repetition** to 2 hours, and InsightIQ encountered three Connection Errors over 2 hours, InsightIQ sends an alert only for the first error.
6. Click **Submit**.

yum update commands

Information about the yum update commands for InsightIQ.

NOTE: It is a best practice to periodically keep the operating system on the OVA updated. This section provides some commands to do so.

NOTE: If your environment prevents direct yum updates, it is advised that you review the [Red Hat document](#) regarding offline updates.

yum update

To keep your InsightIQ Linux virtual machine up to date.

Syntax

```
yum update
```

Options

To see all the options, run the following command:

```
man yum
```

yum update [package-name]

To update specific packages, where [package-name] is the name of the package you are updating.

Syntax

```
yum update [package-name]
```

yum check-update

To see what packages are available for updating.

Syntax

```
yum check-update
```

yum-cron

For more information about enabling automatic updates, see the [Red Hat System Administrator's Guide](#).

Configuring File System Analytics job settings

Setup File System Analytics (FSA).

You can configure File System Analytics (FSA) job settings to enable and disable the FSA job, and how the FSA job collects data on the cluster. FSA job settings also control how long result sets are retained on the cluster before automatic deletion.

Viewing FSA job settings

View the settings for the InsightIQ FSA job.

Steps

1. In the InsightIQ web application, click the **Settings** tab.
2. Click **Monitored Clusters** on the **Settings** ribbon.
The **Monitored Clusters** view appears and shows a list of all clusters that InsightIQ is configured to monitor.
3. In the **Actions** column of the cluster that you want configure, click **Configure**.
The **Cluster Credentials** view appears.
4. Click the **Enable FSA** tab.
The **Job Control** view appears.
5. Review the FSA job settings.
The FSA job settings indicate whether you can view File System Analytics reports from the InsightIQ web application and whether the cluster generates new reports.
6. Click the **FSA Configuration** tab.
The **Job Settings** view appears and shows the current File System Analytics settings.
7. Review the FSA configuration settings.

Enable and disable FSA

Enables and disables the FSA job and reports.

About this task

To view and analyze File System reports with InsightIQ, enable the FSA job.

Steps

1. In the InsightIQ web application, click the **Settings** tab.
2. Click **Monitored Clusters** on the **Settings** ribbon.
The **Monitored Clusters** view appears and shows a list of all clusters that InsightIQ is configured to monitor.
3. In the **Actions** column of the cluster that you want configure, click **Configure**.

The **Cluster Credentials** view appears.

4. Click the **Enable FSA** tab.
The **Job Control** view appears.
5. To enable or disable whether you can view File System Analytics reports in InsightIQ, and whether the cluster generates new reports, select or clear the checkboxes.
 - a. To configure the FSA job, select or clear the **Generate FSA reports on the monitored cluster** checkbox.
 - b. To configure InsightIQ for File System Analytics reports, select or clear the **View FSA reports in InsightIQ** checkbox.
6. Click **Submit**.

Configure File System Analytics (FSA) job settings

Configure settings for the InsightIQ File System Analytics (FSA) job.

About this task

You can configure settings for the InsightIQ File System Analytics (FSA) job, including the generations of FSA result sets from snapshots on the monitored cluster.

Steps

1. In the InsightIQ web application, click the **Settings** tab.
2. Click **Monitored Clusters** on the **Settings** ribbon.
The **Monitored Clusters** view appears and shows a list of all clusters that InsightIQ is configured to monitor.
3. In the **Actions** column of the cluster that you want configure, click **Configure**.
The **Cluster Credentials** view appears.
4. Click the **FSA Configuration** tab.
The **Job Settings** view appears and shows the current File System Analytics settings for the monitored cluster.
5. To configure the **Result Set Options** settings, perform the following steps.
 - a. In the **Directory chart maximum depth** field, type an integer that represents the maximum depth of the directories that you want InsightIQ to analyze.
To specify an unlimited depth, type **-1**.
 - b. In the **File/directory list size** field, type an integer that represents the maximum number of top-contributing files and directories to include in File System Analytics reports.
 - c. In the **Path squash depth** field, type an integer that represents the maximum number of directory levels to include in a path in the `/ifs` directory.
For example, if the **Path squash depth** value is **3**, the path `/ifs/corp/marketing/archive` is represented as `/ifs/corp/marketing`, and all sub-directories of `marketing` are treated as part of `marketing`.
 - d. To generate FSA results sets from snapshots on the monitored cluster, select the **Take snapshot** checkbox.
The default setting is unselected, which generates FSA result sets directly from the cluster.
6. To configure the **Result Set Retention** settings, perform the following steps.
 - a. In the **Maximum result set age in days** field, type an integer that represents, in days, how long to retain each result set before automatically deleting it.
This setting prevents result sets older than the specified number of days from being retained. It also removes result sets that are older than the specified age. Pinned result sets are not deleted, even if they are older than the specified setting.
To retain result sets indefinitely, type **0**. You can manually delete a result set at any time.
 - b. In the **Maximum result set count** field, type the maximum number of result sets to retain.
To retain unlimited result sets, type **0**. You can manually delete a result set at any time.
7. Click **Submit**.

Managing File System Analytics result sets

The File System Analytics feature allows you to view File System reports.

When File System Analytics (FSA) is enabled on a monitored cluster, a File System Analytics job runs on the cluster and collects data that InsightIQ uses to populate file system reports. You can modify how much information is collected by the FSA job through OneFS. You can also configure the level of detail displayed in file system reports through InsightIQ.

NOTE: When enabled, the FSA job consumes computing resources on the monitored cluster and can affect cluster performance. If cluster performance is negatively affected, you can disable the FSA feature. Disabling the FSA job prevents the job from running.

File System Analytics (FSA) result sets overview

InsightIQ generates all File System reports from File System Analytics (FSA) result sets.

If File System Analytics (FSA) is enabled for a monitored cluster, InsightIQ generates File System reports from FSA result sets. Each File System report is generated from a single FSA result set.

Unlike InsightIQ data sets, which are stored in the InsightIQ data store, result sets are stored on the monitored cluster in the `/ifs/.ifsvar/modules/fsa` directory. The monitored cluster routinely deletes result sets to save storage capacity. You can manage result sets by specifying the maximum number of result sets that are retained.

You can configure whether File System reports are visible on the InsightIQ web application and whether new File System reports are generated on the monitored cluster. You can view any existing File System reports, even when the cluster does not generate new reports.

Disabling the FSA job does not affect any of the existing result sets. Result set expiration occurs based on the maximum number of result sets configured. If a result set is marked to be saved, or pinned, it is not deleted. By default, the monitored cluster generates one result set per day.

NOTE: This note applies to versions of OneFS earlier than 8.0.
If NFS is disabled on the monitored cluster, and File System Analytics reports are viewable in the InsightIQ web application, the InsightIQ installation continues to generate FSA connection errors. These errors are generated even if new reports are not created on the monitored cluster.

File System Analytics result set attributes

The **Manage FSA Result Sets** section on the **Manage FSA Result Sets** tab of the **Configuration** view displays information about all retained File System Analytics result sets.

The following list describes each column in the **Manage FSA Result Sets** table.

ID	Displays the unique identifier for the result set, as assigned by the monitored cluster.
Start Time	Displays the date and time at which the data-collection process started.
End Time	Displays the date and time at which the data-collection process ended.
Status	Indicates the status of the result set or, if currently running, the data-collection job.
Size	Indicates the size of the result set.
Pinned	Indicates whether the result set is pinned or unpinned. A pinned result set is never automatically deleted, even if it has expired or exceeds the maximum result-set count.
Actions	Displays links for any actions that you can perform.

View File System Analytics result sets

You can view a list of all stored File System Analytics result sets.

About this task

InsightIQ does not display the information contained in a File System Analytics result set when you view it. To view the data contained in a File System Analytics result set, see [File System Analytics \(FSA\) result sets overview](#) on page 12 .

Steps

1. Click **Settings > Monitored Clusters**.
The **Monitored Clusters** view appears.
2. In the **Actions** column for the cluster whose File System Analytics result sets you want to view, click **Configure**.
The **Configuration** view appears.
3. Click the **Manage FSA Result Sets** tab.
The **Manage FSA Result Sets** tab appears and displays a list of all retained File System Analytics result sets.

Pin or unpin a File System Analytics result set

Pin a File System Analytics result set so that it will not be automatically deleted.

About this task

When a File System Analytics result set is pinned, it will not be automatically deleted, regardless of the expiration schedule. Pinned File System Analytics result sets can only be deleted manually.

If you unpin a result set that would have otherwise been deleted according to an expiration schedule, that result set will be deleted the next time InsightIQ refreshes the result sets, typically within a day.

Steps

1. Click **Settings > Monitored Clusters**.
The Monitored Clusters page appears.
2. In the **Actions** column of the cluster whose File System Analytics result set you want to pin or unpin, click **Configure**.
The **Configuration** view appears.
3. Click the **Manage FSA Result Sets** tab.
The **Manage FSA Result Sets** tab appears and displays a list of all retained File System Analytics result sets.
4. In the **Pinned** column of the result set that you want to pin or unpin, specify an action.

Pin a result set	Select the check box.
Unpin a result set	Clear the check box.

Delete a File System Analytics result set

Delete a File System Analytics result set.

About this task

Any result set can be deleted manually, even if the result set is pinned.

Steps

1. Click **Settings > Monitored Clusters**.
The **Monitored Clusters** view appears.
2. In the **Actions** column for the cluster whose File System Analytics result set you want to delete, click **Configure**.

The **Configuration** view appears.

3. Click the **Manage FSA Result Sets** tab.

The **Manage FSA Result Sets** tab appears and shows a list of all retained File System Analytics result sets.


4. In the **Action** column for the result set that you want to delete, click **delete**.

Configuring cluster monitoring

About this task

You can configure InsightIQ to monitor multiple OneFS clusters. The maximum number of clusters that you can simultaneously monitor varies depending on the resources available to InsightIQ and the following recommendations:

- It is recommended that each instance of InsightIQ monitor no more than 80 nodes.
- It is recommended that each instance of InsightIQ monitor no more than 8 separate clusters.

 **NOTE:** If you want to monitor additional nodes or clusters, it is recommended that you deploy additional instances of InsightIQ.

Monitor a cluster

Configure InsightIQ to monitor an initial cluster.

Prerequisites

On the OneFS cluster that you want to monitor, a local user account for **insightiq** must be enabled, configured with a password, and within the System Zone


Steps


1. In the InsightIQ web application, click the **Settings** tab and then **Monitored Clusters** on the **Settings** ribbon, and then click **Add Cluster**.

 **NOTE:** The first time that you configure InsightIQ, the **Add Cluster** dialog box might already be displayed by default.

2. In the **Add Cluster** dialog box, in the **Isilon cluster address** field, type the hostname, SmartConnect zone name, or IP address of any node in the cluster that you want to monitor.

In general, it is recommended that you specify a cluster by a SmartConnect zone that includes CPU load-balancing policies. Alternatively, for clusters with high usage and file-heat data, you might want to specify the cluster by the IP address or by the hostname of a specific node.

 **NOTE:** It is recommended that you monitor clusters over a LAN connection. Monitoring clusters over a WAN connection can significantly degrade the performance of InsightIQ.

 **NOTE:** Specify IPv6 addresses without surrounding brackets. You can connect over IPv6 only if IPv6 addresses are configured for both the InsightIQ instance and the monitored cluster.

3. In the **InsightIQ user name** field, type **insightiq**.
4. In the **InsightIQ user password** field, type the password of the OneFS cluster's **insightiq** administrator account.
5. Click **OK**.

View monitored cluster settings

View InsightIQ connection and configuration settings.

About this task

You can view InsightIQ configuration settings from the InsightIQ web application. You cannot view these settings using the command-line interface.


Steps

1. In the InsightIQ web application, click the **Settings** tab.
2. Click **Monitored Clusters** on the **Settings** ribbon.
The **Monitored Clusters** view appears and shows a list of all clusters that InsightIQ is configured to monitor.
3. To view the settings for a specific cluster, in the **Actions** column, click **Configure**.
The **Cluster Credentials** view appears, and displays information about the specified cluster.
4. To view specific settings for the selected cluster, click any of the **Configuration** tabs.

Modify cluster login credentials

Modify the InsightIQ user account credentials of a monitored cluster.

About this task

 **NOTE:** If you modify InsightIQ user account credentials in InsightIQ, modify the user account credentials on the cluster.

Steps

1. In the InsightIQ web application, click the **Settings** tab and then **Monitored Clusters** on the **Settings** ribbon.
The **Monitored Clusters** view appears and shows a list of all clusters that InsightIQ is configured to monitor.
2. To modify the settings for a specific cluster, in the **Actions** column, click **Configure**.
3. In the **Configuration** view, click the **Cluster Credentials** tab.
4. Modify the account credentials.
5. Click **Update**.

Suspend and resume InsightIQ monitoring

Suspend and resume InsightIQ monitoring of a cluster.

About this task

You can temporarily suspend InsightIQ monitoring of the cluster and then later resume monitoring.

If you suspend monitoring for a cluster, InsightIQ completes any data collection queries that are in process. No new queries are run until you resume monitoring. The data store remains intact, but InsightIQ does not collect or store any new data. Based on how long monitoring is suspended, reports might display periods of missing data.

Steps

1. In the InsightIQ web application, click the **Settings** tab and then **Monitored Clusters** on the **Settings** ribbon.
The **Monitored Clusters** view appears and shows a list of all clusters that InsightIQ is configured to monitor.

The current monitoring state of each cluster configuration is indicated in the **Actions** column. If the **Suspend** link is visible, InsightIQ is monitoring the cluster. If the **Resume** link is visible, InsightIQ is not monitoring the cluster.
2. For the cluster monitoring state that you want to modify, click **Suspend** or **Resume**.
The selected link changes based on the new monitoring state.

Delete an InsightIQ datastore

Permanently stop InsightIQ monitoring on a cluster and delete all historical data about the cluster's performance.

About this task

If you no longer want to collect new data or view historical data for a monitored cluster, you can permanently stop monitoring on the cluster and delete all historical data.

 **CAUTION:** If you delete the InsightIQ monitoring configuration for a cluster, all historical data for that cluster is deleted. You cannot undo a deletion.

As an alternative to deleting the InsightIQ monitoring configuration and all associated data, you can temporarily suspend monitoring of the cluster and then resume monitoring at a later time.

Steps

1. In the InsightIQ web application, click the **Settings** tab and then **Monitored Clusters** on the **Settings** ribbon.
The **Monitored Clusters** view appears and shows a list of all clusters that InsightIQ is configured to monitor.
2. In the **Actions** column of the cluster configuration that you want to delete, click **Delete**.
A **Remove Cluster** dialog box appears.
3. Click **OK**.

View the connection status of a monitored cluster

View a monitored cluster's connection status.

About this task

You can view each monitored cluster's connection status. This information can be helpful if you suspect that communication has been interrupted between InsightIQ and a monitored cluster.

Steps

1. In the InsightIQ web application, click the **Settings** tab and then **Monitored Clusters** on the **Settings** ribbon.
The **Monitored Clusters** page appears and displays a list of all clusters that InsightIQ is configured to monitor.
2. In the **Monitoring Status** column, review each monitored cluster's status as indicated by its colored icon and message:

Green	A green icon indicates that communication between the monitored cluster and InsightIQ is normal with no errors.
Yellow	A yellow icon indicates that communication between the monitored cluster and InsightIQ has been temporarily interrupted. This error might be due to a brief timeout. This error condition typically resolves automatically.
Red	A red icon indicates that communication between the monitored cluster and InsightIQ has been interrupted indefinitely. This error might be due to an authorization issue, an unconfigured license, or a prolonged timeout. An administrator must resolve this error condition.
Gray	A gray icon indicates that InsightIQ monitoring is suspended and that there is no communication between the monitored cluster and InsightIQ. To resume InsightIQ monitoring, click Resume in the Actions column.

Configuring datastores

Manage the InsightIQ datastore.

InsightIQ data is stored in either the local datastore that is part of the InsightIQ installation or on any NFS-based server, including an Isilon cluster. You can move the InsightIQ datastore to a new location and migrate all information from the current datastore to the new location. If InsightIQ monitors the cluster that contains the datastore, the cluster appears as a client of itself in the InsightIQ web application.

Datastore requirements

InsightIQ collects data and stores that information in the **datastore**. Below are the datastore requirements based on if the datastore is being stored locally, on a OneFS cluster, or on other NFS servers.

Local datastore

If InsightIQ is installed on a virtual machine, InsightIQ stores the datastore locally on a virtual hard drive that is included on the InsightIQ virtual machine. If InsightIQ is installed on a Linux system, InsightIQ stores the datastore on the local hard drive.

If you want to store the datastore on a local drive, the drive must meet the following requirements:

- Ensure that the local drive contains **at least** 70 GB of free disk space.
- The datastore location cannot contain an existing InsightIQ datastore.

Isilon OneFS datastore

If you want to store the datastore on an Isilon OneFS system, the cluster must meet the following requirements:

- Ensure that the directory contains **at least** 70 GB of free disk space.
- The datastore location cannot contain an existing InsightIQ datastore.
- Verify that the Isilon OneFS operating system supports the version of InsightIQ to be installed.
- Verify that the server includes a correctly configured NFS export rule. The export rule must export the datastore path, and map the root user on the InsightIQ server to the user account that owns the export on the NFS server.
- Enable read and write access to the export. Isilon OneFS ships with a default NFS export rule for the `/ifs` directory that you can use for InsightIQ. If that default NFS export has been modified or deleted, create NFS export rule that allows write access for InsightIQ.

NOTE: Do not place the InsightIQ datastore in a location included in snapshots, as the datastore has a high chance of becoming corrupted.

NOTE: Do not apply a quota to the InsightIQ datastore through the SmartQuotas module. If you limit the size of the InsightIQ datastore through a quota, InsightIQ cannot detect the available space. The datastore might become full before InsightIQ can delete older data to make space available for newer data.

Other NFS datastore

If you want to store the datastore on a different NFS-based server the NFS server must meet the following requirements:

- Ensure that the directory contains **at least** 70 GB of free disk space.
- The datastore location cannot contain an existing InsightIQ datastore.
- Verify that the server includes a correctly configured NFS export rule. The export rule must export the datastore path, and map the root user on the InsightIQ server to the user account that owns the export on the NFS server.
- Enable read and write access to the export rule. This configuration allows InsightIQ to mount the server and create the necessary directories and files on the server.
- If you want to connect to an NFS datastore over IPv6, confirm that the IPv6 addresses are configured for both InsightIQ and the NFS server.

NOTE: To avoid datastore connection issues, ensure the NFS server is in the same subnet and has good network connectivity.

Additional information on datastore disk space

An adequate amount of disk space must be available for InsightIQ to store data. It is recommended that you have at least 70 GB of free disk space, but the amount of data that your specific InsightIQ datastore requires really depends on the number of nodes that are monitored and the length of time that you want InsightIQ to retain data.

It is recommended that you increase the size of the InsightIQ datastore by 2 GB per node for each month of retained data.

Use the following equation to calculate the minimum size of your specific datastore:

$$\text{<minimum_datastore_size_in_GB>} = 2 * \text{<number_of_nodes>} * \text{<number_of_months>} + 10$$

NOTE: The +10 at the end of the equation is to cover the recommended 10 GB of free disk space. If InsightIQ is unable to free at least 5 GB of datastore disk space, InsightIQ stops monitoring the cluster.

NOTE: There is no maximum size for an InsightIQ datastore. The size of the datastore does not significantly affect performance.

View datastore settings

View InsightIQ datastore settings.

About this task

You can view the configured settings for the location where InsightIQ stores collected performance data of all monitored clusters.

Steps

1. In the InsightIQ web application, click the **Settings** tab and then **Datastore** on the **Settings** ribbon.
2. Review the datastore settings on the **Datastore configuration** view:

Datastore location	The path of the InsightIQ datastore. If the InsightIQ datastore is stored on an NFS server, the hostname or IP address of that server is displayed before the path.
Datastore total capacity	The total amount of disk space in the datastore.
InsightIQ data	The total amount of InsightIQ data that is stored in the datastore.
Percentage used	The percentage of used space in the datastore.
Datastore free disk space	The total amount of free disk space in the datastore.

Export a datastore

Export an InsightIQ datastore. An export can be useful if you want to create a backup of an InsightIQ datastore before performing maintenance, such as moving the datastore to another location.

About this task

InsightIQ datastores are exported as compressed zip files. After an export file has been created, you can import data from that zip file into instances of InsightIQ that are running InsightIQ 4.1.x or later.

NOTE: The export process might take several minutes or longer and you cannot perform the following actions while a datastore export is in progress:

- Starting an additional datastore export.
- Importing data from a datastore.
- Adding a cluster to monitor.
- Deleting data that is being exported.
- Modifying the InsightIQ configuration settings of any referenced cluster.

Steps

1. In the InsightIQ web application, click the **Settings** tab and then **Monitored Clusters**.
2. Click **Export Datastore**.
3. In the **Export Datastore** dialog box, select the location for the export data.

- To export data to a local directory on the InsightIQ Linux computer or virtual machine, select **Export to local directory** and specify a path to the export data target directory.
- To export data to a directory on an NFS server, select **Export to NFS server** and specify the address of the NFS server and a path to the export data target directory.

NOTE: You can specify either an IPv4 or an IPv6 address. Specify IPv6 addresses without surrounding brackets.

4. In the **Export** table, select the cluster data that you want to export.
The export file contains only data that is related to the clusters that you specify.
5. Ensure that the **Estimated compressed export size** is smaller than the amount of free space currently available on the specified target directory.
If the target directory does not contain enough free space, the export fails.
6. Click **Export**.
7. Click **Yes**.

Results

Once the export is complete, InsightIQ displays the location of the export zip file on the bottom of the **Monitored Clusters** view.

Import a datastore

You can import data from a datastore export zip file that was created with InsightIQ. You can import datastores to restore an InsightIQ datastore that you backed up or to import data from another InsightIQ server.

About this task

InsightIQ import requires an exported datastore from InsightIQ 4.1.x or later.

NOTE: The import process might take several minutes or longer and you cannot perform the following actions while a datastore export is in progress:

- Starting an additional datastore import.
- Exporting data from a datastore.
- Adding a cluster for InsightIQ to monitor.

Steps

1. In the InsightIQ web application, click the **Settings** tab and then **Monitored Clusters**.
2. Click **Import Datastore**.
3. In the **Import Datastore** dialog box, select where to import data from.
 - To import data from data on a local directory on the InsightIQ Linux computer or virtual machine, select **Import from local directory** and specify a path that contains the export data.
 - To import data from a directory on an NFS server, select **Import from NFS server** and specify the address of the NFS server and a path that contains the export data.

NOTE: You can specify either an IPv4 or an IPv6 address. Specify IPv6 addresses without surrounding brackets.

InsightIQ displays a list of export files that are contained in the specified directory.

NOTE: To avoid datastore connection issues, confirm your NFS share is on the same subnet and has good network connectivity.


4. Select the export file that you want to import and then click **Next**.
InsightIQ displays the list of clusters that are contained in the file.
5. In the **Import** column, select the clusters that you want to import.
InsightIQ imports only the data that is related to clusters that you specify.

NOTE: If InsightIQ has existing data for a cluster in the import file, you cannot select that cluster. To import data from a cluster that has monitored data, delete the datastore for that cluster. You can then import data for that cluster.
6. Optional: To begin monitoring a cluster after the import process completes, select the checkbox in the **Begin Monitoring** column for that cluster, and then type the password of the InsightIQ user account in the **InsightIQ User Password** column.
7. Click **Import**.

8. Click **Yes**.

Move a datastore

Prerequisites



 **NOTE:** It is recommended that you back up the datastore before moving InsightIQ. Back up the datastore by exporting the datastore to a .tar file.

About this task

To move an InsightIQ datastore, follow these steps:

Steps

1. In the InsightIQ web application, click the **Settings** tab and then click **Datastore**.
 2. Specify the datastore location.
 - To migrate data to a local datastore, click **Local Datastore** and then select the location for the export data.
 - To migrate data to an NFS datastore, click **NFS Mounted Datastore**, and then specify the address of the NFS server and select the location for the export data.

 **NOTE:** Verify that the new datastore location meets the requirements for an NFS server. Refer to the topic, [Datastore requirements](#) on page 17, for more information.
 3. Click **OK**.
-  **NOTE:** During datastore movement, datastore configuration sizes will be shown as NA. The correct values will populate once the move is complete.

Automatic data pruning

InsightIQ deletes older data when a datastore runs out of disk space.

When the InsightIQ datastore fills the storage capacity of the disk, InsightIQ deletes the oldest data until storage capacity has increased.

If multiple clusters are monitored, InsightIQ favors deletion of the oldest data among all the clusters. For example, if ClusterA has data from January 2015, and ClusterB has data from January 2017, the data from ClusterA is deleted until the oldest data has the same date as the data on ClusterB.


Table 2. InsightIQ Actions at Datastore Disk Capacity Thresholds

Remaining available disk capacity	InsightIQ Action
> 5 GB	InsightIQ has the required datastore disk capacity. InsightIQ cancels any preexisting alerts about datastore disk capacity limits.
1–5 GB	InsightIQ continues monitoring the cluster. InsightIQ alerts the user that the datastore disk capacity limit is at risk. These alerts continue until datastore disk capacity changes.
< 1 GB	InsightIQ suspends monitoring the cluster. InsightIQ tries to delete the oldest datastore summary data. If InsightIQ cannot delete datastore data, it suspends monitoring.

Configuring LDAP for authentication

Enable InsightIQ to use an LDAP server for authentication.

You can configure Lightweight Directory Access Protocol (LDAP) as the authentication method for accessing InsightIQ. Once LDAP is enabled, InsightIQ checks the configured LDAP server's users and groups for authentication. The LDAP server, as well as its users and groups, must be configured on the LDAP server before InsightIQ can use the information for authentication.

 **NOTE:** InsightIQ is RFC 2307-compliant.

Connect InsightIQ to an LDAP server

Connect InsightIQ to an LDAP server for authentication.

Prerequisites

The LDAP server, users, and groups must be configured on the LDAP server before InsightIQ can access the information.

About this task

You can enable InsightIQ to use an LDAP server for user authentication.

Steps

1. In the InsightIQ web application, click the **Settings** tab.
2. Click **Users** on the **Settings** ribbon.
3. Click **Configure LDAP**.
4. Check **Enable LDAP**.
Enabling LDAP allows you to edit the remaining fields on this page.
5. Type the LDAP server URI into the **LDAP server** field.
6. Type the **Base Search Entry**.
7. Type LDAP server credentials in the **Bind entry** and **Bind password** fields.
InsightIQ uses these credentials when connecting with the LDAP server.
8. Click **Submit**.
The system connects with the LDAP server and, if successful, saves the settings.

Results

Once a connection is made between InsightIQ and the LDAP server, you can add LDAP groups and users to InsightIQ.

Add an LDAP group

Configure an LDAP group to access InsightIQ.

Prerequisites

The LDAP server, users, and groups must be configured on the LDAP server before InsightIQ can access the information. Also, InsightIQ must have LDAP connectivity that is enabled and must be connected to the LDAP server. LDAP users must be direct members of a group in order to access InsightIQ, and not a member of a group nested in a group.

About this task

Configure an LDAP group to access InsightIQ.

Steps

1. Click **Add Group**.
A separate window opens to search for a Group in the connected LDAP server.
2. Type the group name, and click **Search**.

A list of groups that match the search is displayed.

3. Select the group that you want to give access to InsightIQ.

4. Select the privilege level for the group.

Read-only permits a group to view InsightIQ reports.

Administrator permits a group to configure InsightIQ and view reports.

5. Click **Add Group**.

Results

The group that you added is now listed in the **Add Privileged Groups** table under the **Configure LDAP** tab. The users within that group now have access to InsightIQ.



Setting up data sets

The information in the InsightIQ datastore is divided into data sets.

Each Performance module references data from one or more data sets. You can enable or disable the retrieval of an individual data set. However, disabling a data set causes the Performance modules that reference the data set to be incomplete.

Data set configuration table

Detailed data set information is available in the **Data Set Configuration** table.

Data set	Identifies the name of the data set.	
Raw data	 NOTE: This information is available only if you click Show data set size details .	
	Identifies the amount of data that is retrieved for the data set that is not downsampled. This field also shows the period that the data represents.	
Summary data	 NOTE: This information is available only if you click Show data set size details .	
	Identifies the amount of data that is downsampled for the data set. This field also shows the period for the downsampled data. This data is used to generate the minimum and maximum values that are displayed on a Performance module.	
Total data	Identifies the combined total of downsampled and non-downsampled data that is stored for the data set.	
Latest data that is retrieved	Identifies the most recent time that data was retrieved for the data set. The <time_zone> value indicates the time zone of the reported time.	
Status	Identifies the retrieval status for the data set.	
	Enabled	InsightIQ is configured to retrieve data for the data set.
	Disabled	InsightIQ does not retrieve data for the data set.
	Unsupported	The cluster is running a version of OneFS that does not support the InsightIQ features for this data set. InsightIQ does not retrieve data from clusters running unsupported versions of OneFS.
	Delayed	InsightIQ is configured to retrieve data for the data set, but InsightIQ was not able to retrieve data for at least 15 minutes.
Actions	Lists the actions that you can perform on the data set.	

Data sets, modules, and breakouts

Data sets are associated with specific report modules and breakouts.

Data that InsightIQ retrieves is grouped into data sets. Each data set is associated with specific modules and breakouts. When you disable a data set, the disabled module or breakout displays no information for the period when data retrieval was disabled.

Table 3. Data set, module, and breakout relationship

Data set	Affected modules	Affected breakouts
Active Client Count	Active Clients	--
Aggregate External Network Counters	<ul style="list-style-type: none"> External Network Errors External Network Packets Rate External Network Throughput Rate Node Summary 	--
Connected Client Count	Connected Clients	--
CPU Usage	CPU % Use	--
Deduplication	<ul style="list-style-type: none"> Deduplication Summary (Logical) Deduplication Summary (Physical) 	--
Disk Performance	<ul style="list-style-type: none"> Average Disk Hardware Latency Average Disk Operation Size Average Pending Disk Operations Count Disk Activity Disk Operations Rate Disk Throughput Rate Pending Disk Operations Latency Slow Disk Access Rate 	--
Disk Storage	Cluster Capacity	Disk breakouts for the following modules: <ul style="list-style-type: none"> Total Capacity Protection Performance modules
Events	Event Summary	--
External NIC Counters	--	Interface breakouts for the following Performance modules: <ul style="list-style-type: none"> External Network Errors External Network Packets Rate External Network Throughput Rate
IFS Cache Performance	<ul style="list-style-type: none"> Cache Data Age L1 and L2 Cache Prefetch Throughput Rate L1 Cache Throughput Rate L2 Cache Throughput Rate L3 Cache Throughput Rate Overall Cache Hit Rate Overall Cache Throughput Rate 	--
IFS Operation Counters	<ul style="list-style-type: none"> Blocking File System Events Rate Deadlocked File System Events Rate File System Events Rate Locked File System Events Rate 	--
IFS Throughput	File System Throughput Rate	--
IFS Usable Capacity	This data set affects only the following entries in the Cluster Capacity Performance module: <ul style="list-style-type: none"> Allocated Capacity Writable Capacity 	--
IFS Usage	This data set affects only the following entries in the Cluster Capacity Performance module:	--

Table 3. Data set, module, and breakout relationship (continued)

Data set	Affected modules	Affected breakouts
	<ul style="list-style-type: none"> Total Capacity User Data Including Protection 	
Job Engine Statistics	<ul style="list-style-type: none"> Job Workers Jobs 	--
Jobs and Services Statistics	<ul style="list-style-type: none"> Cache Hits CPU Usage Rate Disk IOPS Jobs and Services Summary 	--
Per Client Statistics	Client Summary	--
Per File IFS Operation Counters	--	File breakouts for the following modules: <ul style="list-style-type: none"> Blocking File System Events Rate Deadlocked File System Events Rate File System Events Rate Locked File System Events Rate
Per Protocol Performance	<ul style="list-style-type: none"> Client Summary External Network Throughput Rate Protocol Operations Average Latency Protocol Operations Rate 	--
Quotas	<ul style="list-style-type: none"> Quota Browser Quota Limit Usage 	--

View data set sizes

View the size of each data set of a monitored cluster.

Steps


1. In the InsightIQ web application, click the **Settings** tab and then click **Monitored Clusters** on the **Settings** ribbon.
The **Monitored Clusters** view appears and shows a list of all clusters that InsightIQ is configured to monitor.
2. In the **Actions** column for the cluster, click **Configure**.
The **Cluster Credentials** view appears, and displays information about the specified monitored cluster.
3. In the **Configuration** view, click the **Data Set Configuration** tab.
4. Click **View data set size details** above the **Data Set Configuration** table.

Enable or disable retrieval of data sets

Enable or disable the retrieval of data in InsightIQ.

About this task

You can enable or disable the retrieval of an InsightIQ data set. If you are not interested in a specific type of information, you can disable a data set. Disabling a data set reduces InsightIQ resource consumption.

 **NOTE:** After you enable retrieval of a data set, InsightIQ begins retrieving the most recent data for that data set. InsightIQ does not recover data for the time that data retrieval was disabled. Affected modules display zero values for the period that data set retrieval was disabled.

Steps

1. In the InsightIQ web application, click the **Settings** tab and then click **Monitored clusters** on the **Settings** ribbon.
The **Monitored Clusters** view appears and shows a list of all clusters that InsightIQ is configured to monitor.

2. In the **Actions** column for the cluster that you want to modify, click **Configure**.
3. Click **Data Set Configuration**.
4. In the **Actions** column for the data set that you want to enable or disable, click **enable** or **disable**.


Results

The action takes place immediately and the InsightIQ web application updates the enable or disable status of the data set.

Setting up user accounts

Create and delete user accounts.

A user account allows anyone to log in to the InsightIQ web application and monitor cluster activity. A read-only user account cannot modify InsightIQ configuration settings, create performance reports, or add user accounts.

 **NOTE:** An LDAP user may be granted administrator access based on the LDAP group configuration.

During the InsightIQ setup process, the administrator account is configured. The administrator account can create, modify, and delete user accounts. There is only one administrator account per installation of InsightIQ. There is no limit to the number of user accounts.

Create a user account

Create a user account.

Prerequisites

You must be logged in as the administrator.

About this task

During the user account creation process, specify a username and password for the user. The user must use these credentials to log in to the InsightIQ application.

Steps

1. In the InsightIQ web application, click the **Settings** tab.
2. Click **Users** on the **Settings** ribbon.
The **Read-Only Users** view appears, and displays a list of all configured user accounts.
3. Click **Add User**.

The **Add a Read-Only User** dialog box appears.

4. In the **Username** field, type a name for the user.
5. In both the **Password** and **Confirm Password** fields, type a password for the user.

Make a note of the password that you configured. InsightIQ does not allow the administrator or users to retrieve a lost password.

6. Click **Submit**.


The user account is added to the **Read-Only Users** list.

Modify a user account password

Modify the password of a user account.

About this task

The administrator can modify the password of a user account in the InsightIQ web application.

 **NOTE:** Users cannot modify passwords for their accounts.

Steps

1. In the InsightIQ web application, click the **Settings** tab.
2. Click **Users** on the **Settings** ribbon.
The **Read-Only Users** view appears and shows a list of all configured users.
3. In the **Actions** column, for the user whose password you want to modify, click **Change Password**.
The **Change Password** dialog box appears.
4. In the **New password** field, type a new password for the user.
5. In the **Confirm password** field, retype the new password for the user.
Make a note of the password you configured. InsightIQ does not enable the administrator or user accounts to retrieve a lost password.
6. Click **Submit**.

Delete a user account

Delete a user account.

Steps

1. In the InsightIQ web application, click the **Settings** tab.
2. Click **Users** on the **Settings** ribbon.
The **Read-Only Users** view appears, and displays a list of all configured read-only users.
3. In the **Actions** column for the user account that you want to delete, click **Delete**.
A **Delete Read-Only User** dialog box appears.
4. Click **Yes**.
The user account is removed, and the user can no longer log in to the InsightIQ web application.

Troubleshooting InsightIQ

This section contains the following topics:

Topics:

- [Troubleshooting overview](#)
- [Troubleshooting configuration issues](#)
- [Troubleshooting datastore issues](#)
- [Troubleshooting data retrieval issues](#)
- [Troubleshooting datastore move issues](#)
- [Troubleshooting memory issues](#)
- [Troubleshooting InsightIQ web application issues](#)

Troubleshooting overview

Overview of troubleshooting indicators.

In the upper-right corner of the InsightIQ web interface, the **InsightIQ Status** area displays at-a-glance information about the status of InsightIQ.

The status of InsightIQ is indicated by color:

Green	InsightIQ is operating normally.
Yellow	At least one transient, nonfatal error has occurred.
Red	InsightIQ could not save data to the data store. For example, InsightIQ could not save data to the data store because the data store is full or because InsightIQ could not contact the server hosting the data store. InsightIQ does not collect additional data until the issue is resolved.

Troubleshooting configuration issues

Troubleshoot InsightIQ configuration issues.

Troubleshooting InsightIQ configuration issues typically fall into one of two general categories: InsightIQ virtual machine configuration issues, and monitored cluster configuration issues.

Virtual machine configuration issues

If InsightIQ does not detect a virtual machine network adapter during the boot process, shut down the InsightIQ virtual machine. Then, check the settings for the network adapter, and ensure that the network type is correct for the virtualization environment. For example, if you reconfigure the network adapter to run in bridged mode. If this issue persists, delete the network interface in the virtualization environment, add network interface, and then restart the InsightIQ virtual machine.

If you cannot log in to the InsightIQ web application, open the InsightIQ virtual machine command-line interface and configure InsightIQ networking to run in DHCP networking mode. Then try accessing the InsightIQ web application through the DHCP-generated IP address. The results can indicate whether the issue is related to the IP address.

Monitored-cluster configuration issues

If InsightIQ cannot connect to the monitored cluster, try to ping the IP address of any node in the cluster. If you cannot ping a cluster node, verify that the node at the specified IP address is operating correctly and that the IP address is valid. If the issue persists, try connecting to a different node in the monitored cluster by the node's hostname or IP address. If the issue

persists, configure a SmartConnect zone on the monitored cluster, and then try to connect to that SmartConnect zone instead of a hostname or IP address.

If an `Unlicensed` error message appears, log in to the monitored cluster and verify that a valid InsightIQ license has been activated on the cluster.

If an `Unauthorized` error message appears, verify that the local InsightIQ or OneFS Platform API user on the monitored cluster is enabled and is configured with a valid password. Verify that the corresponding InsightIQ user settings in the InsightIQ application match the settings that are configured on the monitored cluster.

 **NOTE:** The local user on the monitored cluster is the InsightIQ user named InsightIQ.

Troubleshooting datastore issues

If InsightIQ cannot write to the datastore, check the following settings.

If the InsightIQ datastore is on an Isilon cluster and an NFS datastore permissions error message appears, verify that a valid NFS export is configured on that cluster. The NFS export must be configured to grant write access to the root user. This configuration enables InsightIQ to mount the cluster or server and create the necessary directories and files on the cluster or server. InsightIQ connects to the NFS host as the root user. The configured NFS export must grant the root user write access for the specified InsightIQ virtual machine IP address.


If InsightIQ is configured to use a local datastore and a permissions error message appears, connect to the virtual machine command-line interface. Then, verify that the parent directory of the datastore is configured with a permission setting of 755 or higher.

If InsightIQ cannot write to the datastore, review the permissions settings for the datastore directory and for all the files contained in the directory. All the files in the datastore directory must be configured with a permission setting of 744 or higher. If the issue persists, verify that the directory's owner and group settings are correctly configured. For an NFS datastore, the **owner:group** setting must be **nobody:nobody**. For a local datastore, the **owner:group** setting must be **root:root**.

Troubleshooting data retrieval issues

Guidance on data retrieval problems.

If you experience Data Retrieval Delayed errors, consider disabling the retrieval of a specific data set. If InsightIQ is unable to retrieve a data set, disabling retrieval for delayed data sets removes the Data Retrieval Delayed errors caused by those data sets. Disabling retrieval can be useful if you are receiving Data Retrieval Delayed errors for data sets that you do not care about and want to reserve those errors for data sets that you do care about.

 **NOTE:** The following data sets do not contribute to Data Retrieval Delayed errors. If retrieval for any of the following data sets is delayed, it does not cause a Data Retrieval Delayed error:

- Disk Performance
- IFS Operation Counters
- IFS Usable Capacity
- Job Engine Statistics
- Per File IFS Operation Counters

Troubleshooting datastore move issues

If there is not enough free space on the target datastore or if an NFS connection gets interrupted, the datastore move operations can fail.

If InsightIQ is migrating data from or to an NFS datastore, and the connection to the NFS datastore is interrupted, InsightIQ stops migrating data until the connection is restored. If the connection is permanently severed, you can recover the data only if you have created a backup of the datastore by exporting it to a .tar file. If you have created a backup, you can import the backup datastore to a new instance of InsightIQ to recover the data. For this reason, The recommendation is to back up the datastore before moving an InsightIQ datastore.

InsightIQ tries to calculate whether there is enough free space on the target location before migrating data. However, if a quota is applied to the target location and the quota is configured to report the size of the entire file system instead of the quota

limit, there is less space on the target location than InsightIQ requires. The migration might fail. If this failure occurs, InsightIQ automatically transfers the datastore back to the original datastore.

Troubleshooting memory issues

Troubleshoot issues with memory and timeouts.

If many breakouts or data modules are applied to a single report, InsightIQ might run out of memory, causing data modules to timeout.

If a timeout occurs, restart InsightIQ by running the following command:

```
iiq_restart
```

Troubleshooting InsightIQ web application issues

Troubleshoot issues with the InsightIQ web application.

If you are unable to access the InsightIQ web application through a web browser, the browser might be trying to access InsightIQ through an incorrect port or the InsightIQ web application might be unable to communicate with InsightIQ.

Some web browsers automatically store the port number through which a site is accessed the first time that the browser goes to the site. If you access InsightIQ before modifying the InsightIQ port, the web browser might later try to connect to InsightIQ through the original port. To resolve this issue, clear the cache of the web browser and then connect to InsightIQ.

If the issue persists, verify that InsightIQ is running, and then try to ping InsightIQ or log in through another supported browser. If you are still unable to access the InsightIQ web application, log in to the InsightIQ command-line interface and restart InsightIQ by running the following command:

```
iiq_restart
```

InsightIQ commands

This section contains the following topics:

Topics:

- [InsightIQ commands overview](#)
- [iiq_data_export fsa describe](#)
- [iiq_data_export fsa export](#)
- [iiq_data_export fsa list](#)
- [iiq_data_export perf describe](#)
- [iiq_data_export perf export](#)
- [iiq_data_export perf list](#)
- [iiq_restart](#)
- [iiq_stop](#)
- [iiq_start](#)

InsightIQ commands overview

Information about the InsightIQ commands

You can manage InsightIQ from a command-line interface.

iiq_data_export fsa describe

Displays descriptions of File System Analytics modules.

Syntax

```
iiq_data_export fsa describe --cluster <cluster>  
--data-module <module_value>
```

Options

{--cluster | -c} <cluster>

Displays descriptions of File System Analytics modules on the specified cluster.

{--data-module | -d} <module>

Displays a description of the specified File System Analytics modules. The following values are valid:

Table 4. Module values

Module	Value
Directories	directories
File Count by Logical Size	file_count_by_logical_size
File Count by Last Changed	file_count_by_changed_time
File Count by Physical Size	file_count_by_physical_size

Table 4. Module values (continued)

Module	Value
Top Directories	top_directories
Top Files	top_files

Examples

The following command displays the description of the Top Files module while monitoring cluster1.

```
iiq_data_export fsa describe --cluster cluster1 --data-module top_files
```

iiq_data_export fsa export

Export InsightIQ File System Analytics data to a comma-separated value (CSV) file.

Syntax

```
iiq_data_export fsa export --cluster <cluster-name>
--data-module <data-module> --report <id>
[--comp-report <id>]
[--breakout-by <breakout>]
[--filter-rule <breakout>:<filter>...]
[--path <path>]
[--name <name>]
[--number-breakouts <integer>]
```

{--cluster | -c} <cluster>

Exports data about the specified cluster.

{--data-module | -d} <data-module>

Exports the specified data module. The following values are valid:

Table 5. Module descriptions

Module description	Module Value
Directories	directories
File Count by Logical Size	file_count_by_logical_size
File Count by Last Changed	file_count_by_changed_time
File Count by Physical Size	file_count_by_physical_size
Top Directories	top_directories
Top Files	top_files

{--report | -o} <report-id>

Exports data contained in the specified report.

{--comp-report | -z} <id>

Instead of exporting the data contained in the report specified by the --report option, exports the comparison between the older report and the newer report specified by the --report and --comp-report options.

{--breakout-by | -b} <breakout>

Applies the specified breakout to the exported data.

Table 6. Breakout values

Breakout description	Breakout value
Accessed Time	atime
Directory	directory
File Extension	path_ext
Logical Size	log_size
Changed Time	ctime
Node Pool	node_pool
Physical Size	phys_size
Tier	tier
User Attribute	attribute

`{--filter-rule | -r} <breakout>:<filter>...`

Applies the specified filter rule to the exported data. To specify a filter, you must specify the breakout that the filter applies to.

Table 7. Filter values

Breakout description	Breakout value	Filter values
Accessed Time	atime	The following values are valid:
		-60 Specifies 0:00:00 - 0:01:00
		-3600 Specifies 0:01:00 - 1:00:00
		-86400 Specifies 1:00:00 - 1 day
		-604800 Specifies 1 day - 7 days
		-2592000 Specifies 7 days - 30 days
		-5184000 Specifies 30 days - 60 days
		-7776000 Specifies 60 days - 90 days
		-15552000 Specifies 90 days - 180 days
		-31536000 Specifies 180 days - 365 days
		-63072000 Specifies 365 days - 730 days
		Infinity Specifies 730 days - Infinity
For example:		
<code>--filter-rule atime:-3600</code>		

Table 7. Filter values (continued)

Breakout description	Breakout value	Filter values																				
Directory	directory	<div>The path of a directory starting with / ifs. For example:</div> <div><pre>--filter-rule directory:/ifs/data/media</pre></div>																				
File Extension	path_ext	<div>The name of a file extension. For example:</div> <div><pre>--filter-rule path_ext:txt</pre></div>																				
Logical Size	log_size	<div>A range of sizes in the following format:</div> <div><pre><integer><unit>-<integer><unit></pre></div> <div><table><tr><td>0</td><td>Specifies 0 bytes - 8 KB</td></tr><tr><td>8192</td><td>Specifies 8 KB - 128 KB</td></tr><tr><td>131072</td><td>Specifies 128 KB - 1 MB</td></tr><tr><td>1048576</td><td>Specifies 1 MB - 10 MB</td></tr><tr><td>10485760</td><td>Specifies 10 MB - 100 MB</td></tr><tr><td>104857600</td><td>Specifies 100 MB - 1 GB</td></tr><tr><td>1073741824</td><td>Specifies 1 GB - 10 GB</td></tr><tr><td>10737418240</td><td>Specifies 10 GB - 100 GB</td></tr><tr><td>107374182400</td><td>Specifies 100 GB - 1 TB</td></tr><tr><td>1099511627776</td><td>Specifies 1 TB - Infinity</td></tr></table></div> <div>For example:</div> <div><pre>--filter-rule log_size:1GB-10GB</pre></div>	0	Specifies 0 bytes - 8 KB	8192	Specifies 8 KB - 128 KB	131072	Specifies 128 KB - 1 MB	1048576	Specifies 1 MB - 10 MB	10485760	Specifies 10 MB - 100 MB	104857600	Specifies 100 MB - 1 GB	1073741824	Specifies 1 GB - 10 GB	10737418240	Specifies 10 GB - 100 GB	107374182400	Specifies 100 GB - 1 TB	1099511627776	Specifies 1 TB - Infinity
0	Specifies 0 bytes - 8 KB																					
8192	Specifies 8 KB - 128 KB																					
131072	Specifies 128 KB - 1 MB																					
1048576	Specifies 1 MB - 10 MB																					
10485760	Specifies 10 MB - 100 MB																					
104857600	Specifies 100 MB - 1 GB																					
1073741824	Specifies 1 GB - 10 GB																					
10737418240	Specifies 10 GB - 100 GB																					
107374182400	Specifies 100 GB - 1 TB																					
1099511627776	Specifies 1 TB - Infinity																					
Changed Time	ctime	<div>The following values are valid:</div> <div><table><tr><td>-60</td><td>Specifies 0:00:00 - 0:01:00</td></tr><tr><td>-3600</td><td>Specifies 0:01:00 - 1:00:00</td></tr><tr><td>-86400</td><td>Specifies 1:00:00 - 1 day</td></tr><tr><td>-604800</td><td>Specifies 1 day - 7 days</td></tr></table></div>	-60	Specifies 0:00:00 - 0:01:00	-3600	Specifies 0:01:00 - 1:00:00	-86400	Specifies 1:00:00 - 1 day	-604800	Specifies 1 day - 7 days												
-60	Specifies 0:00:00 - 0:01:00																					
-3600	Specifies 0:01:00 - 1:00:00																					
-86400	Specifies 1:00:00 - 1 day																					
-604800	Specifies 1 day - 7 days																					

Table 7. Filter values (continued)

Breakout description	Breakout value	Filter values
		<p>-2592000 Specifies 7 days - 30 days</p> <p>-5184000 Specifies 30 days - 60 days</p> <p>-7776000 Specifies 60 days - 90 days</p> <p>-15552000 Specifies 90 days - 180 days</p> <p>-31536000 Specifies 180 days - 365 days</p> <p>-63072000 Specifies 365 days - 730 days</p> <p>Infinity Specifies 730 days - Infinity</p> <p>For example:</p> <pre>--filter-rule ctime:-3600</pre>
Node Pool	node_pool	<p>The name of a node pool. For example:</p> <pre>--filter-rule node_pool:pool1</pre> <p>i NOTE: This filter option is available only if you monitor a cluster that is running OneFS 8.0 or later</p>
Physical Size	phys_size	<p>A range of sizes in the following format:</p> <pre><integer><unit>-<integer><unit></pre> <p>0 Specifies 0 bytes - 8 KB</p> <p>8192 Specifies 8 KB - 128 KB</p> <p>131072 Specifies 128 KB - 1 MB</p> <p>1048576 Specifies 1 MB - 10 MB</p> <p>10485760 Specifies 10 MB - 100 MB</p> <p>104857600 Specifies 100 MB - 1 GB</p> <p>1073741824 Specifies 1 GB - 10 GB</p> <p>10737418240 Specifies 10 GB - 100 GB</p> <p>107374182400 Specifies 100 GB - 1 TB</p>

Table 7. Filter values (continued)

Breakout description	Breakout value	Filter values
		1099511627776 Specifies 1 TB - Infinity For example: <pre>--filter-rule phys_size:1GB-10GB</pre>
Tier	tier	The name of a tier. For example: <pre>--filter-rule tier:tier_name</pre> <p>NOTE: This filter option is available only if you monitor a cluster that is running OneFS 8.0 or later</p>
User Attribute	attribute	The name of a user attribute defined on the cluster. All characters are valid.

{--path | -p} <path>

Specifies where to create the .csv file. Specify a full directory path. The default value is the full path of the working directory.

{--name | -n} <name>

Specifies a name for the .csv file. A time stamp is appended to the specified name. If no name is specified, the file is named according to the following pattern:

```
<data_module>_[<breakout_by>_]<report_id>[-
<comp_report_id>]_<time_generated>.csv
```

<breakout_by> and <comp_report_id> are included only if those options are specified.

{--number-breakouts | -m} <integer>

Specifies the number of breakout components to include for each breakout. For example, if you broke out a data module by directory, specifying `--number-breakouts 5` would export data about the top 5 directories. The default number is 12.

Examples

The following command exports data about the number of files on the cluster organized by logical size:

```
iiq_data_export fsa export --cluster cluster1 --data-module file_count_by_logical_size
--report 411
```

iiq_data_export fsa list

Displays the names of monitored clusters, File System Analytics data modules, breakouts, filters, and File-System Analytics reports.

Syntax

```
iiq_data_export fsa list [--clusters | --data-modules
| --breakouts <cluster> | --filters <cluster> | --reports <cluster>]
```

Options

--clusters

Displays the names of all clusters that InsightIQ is monitoring.

--data-modules

Displays the names of all data modules that InsightIQ supports for File-System Analytics reports. Each data module supports a subset of breakouts.

--breakouts <cluster>

Displays the names of all breakouts that InsightIQ supports for File-System Analytics data modules for the specified cluster. Each data module supports a subset of breakouts.

--filters <cluster>

Displays the names of all filters that have been created for the specified cluster.

--reports <cluster>

Displays the names of all available File-System Analytics reports that have been created for the specified cluster.

Examples

The following command displays the names of all file-system reports created for cluster1:

```
iiq_data_export fsa list --reports cluster1
```

iiq_data_export perf describe

Displays descriptions of performance data modules.

Syntax

```
iiq_data_export perf describe --data-module <module-value>
```

Options

--data-module <data-module>

Displays a description of the specified performance data module. The following values are valid:

Table 8. Performance report module values

Module description	Module value
Active Clients	client_active
Average Cached Data Age	cache_oldest_page_age
Average Disk Hardware Latency	disk_adv_access_latency
Average Disk Operation Size	disk_adv_op_size
Average Pending Disk Operations Count	disk_adv_io_queue
Blocking File System Events Rate	ifs_blocked
Cache Hits	cache_hits
Cluster Capacity	ifs_cluster_capacity

Table 8. Performance report module values (continued)

Module description	Module value
Connected Clients	client_connected
Contented File System Events Rate	ifs_contended
CPU %Use	cpu_use
CPU Usage Rate	cpu_usage_rate
Deadlocked File System Events Rate	ifs_deadlocked
Deduplication Summary (Logical)	dedupe_logical
Deduplication Summary (Physical)	dedupe_physical
Disk Activity	disk_adv_busy
Disk IOPS	disk_iops
Disk Operations Rate	disk_adv_op_rate
Disk Throughput Rate	disk_adv_bytes
External Network Errors	ext_error
External Network Packets Rate	ext_packet
External Network Throughput Rate	ext_net_bytes
File System Events Rate	ifs_heat
File System Throughput Rate	ifs_total_rate
Job Workers	worker
Jobs	job
L1 and L2 Cache Prefetch Throughput Rate	cache_all_prefetch
L1 Cache Throughput Rate	cache_l1_read
L2 Cache Throughput Rate	cache_l2_read
L3 Cache Throughput Rate	cache_l3_read
Locked File System Events Rate	ifs_lock
Overall Cache Hit Rate	cache_all_read_hitrates
Overall Cache Throughput Rate	cache_all_read
Pending Disk Operations Latency	disk_adv_io_latency
Protocol Operations Average Latency	proto_latency
Protocol Operations Rate	proto_op_rate
Slow Disk Access Rate	disk_adv_access_slow

Examples

To view a description of the Disk Activity data module, run the following command:

```
iiq_data_export perf describe --data-module ifs_contended
```

iiq_data_export perf export

Exports InsightIQ performance data to a comma-separated value (CSV) file.

Syntax

```
iiq_data_export perf export --data-module <data-module>
--cluster <cluster-name>
[--path <path>]
[--name <name>]
[--breakout-by <breakout>...]
[--filter-rule <breakout>:<filter>...]
[--number-breakouts <integer>]
[--end {<timestamp> | now}]
[--interval <integer> <units>]
[--fmt-time]
[--min-max]
```

Options

{--data-module | -d} <data-module>

Exports the specified performance data module.

Table 9. Performance data module values

Module description	Module value
Active Clients	client_active
Average Cached Data Age	cache_oldest_page_age
Average Disk Hardware Latency	disk_adv_access_latency
Average Disk Operation Size	disk_adv_op_size
Average Pending Disk Operations Count	disk_adv_io_queue
Blocking File System Events Rate	ifs_blocked
Cache Hits	cache_hits
Cluster Capacity	ifs_cluster_capacity
Connected Clients	client_connected
Contented File System Events Rate	ifs_contented
CPU %Use	cpu_use
CPU Usage Rate	cpu_usage_rate
Deadlocked File System Events Rate	ifs_deadlocked
Deduplication Summary (Logical)	dedupe_logical
Deduplication Summary (Physical)	dedupe_physical
Disk Activity	disk_adv_busy
Disk IOPS	disk_iops
Disk Operations Rate	disk_adv_op_rate
Disk Throughput Rate	disk_adv_bytes
External Network Errors	ext_error
External Network Packets Rate	ext_packet

Table 9. Performance data module values (continued)

Module description	Module value
External Network Throughput Rate	ext_net_bytes
File System Events Rate	ifs_heat
File System Throughput Rate	ifs_total_rate
Job Workers	worker
Jobs	job
L1 and L2 Cache Prefetch Throughput Rate	cache_all_prefetch
L1 Cache Throughput Rate	cache_l1_read
L2 Cache Throughput Rate	cache_l2_read
L3 Cache Throughput Rate	cache_l3_read
Locked File System Events Rate	ifs_lock
Overall Cache Hit Rate	cache_all_read_hitrate
Overall Cache Throughput Rate	cache_all_read
Pending Disk Operations Latency	disk_adv_io_latency
Protocol Operations Average Latency	proto_latency
Protocol Operations Rate	proto_op_rate
Slow Disk Access Rate	disk_adv_access_slow

{--cluster | -c} <cluster-name>

Exports performance data about the specified cluster.

{--path | -p} <path>

Specifies where to create the .csv file. Specify a full directory path. If no name is specified, the file is named according to the following pattern:

```
<data_module>_[<breakout_by>_]<report_id>_<time_generated>.csv
```

<breakout_by> is included only if the option is specified.

{--name | -n} <name>

Specifies the name of the .csv file. A timestamp is appended to the specified name.

{--breakout-by | -b} <breakout>...

Applies the specified breakouts to the exported data.

Table 10. Breakout values

Breakout description	Breakout value
Client	remote_addr
Direction	direction
Disk	disk
Disk Pool	disk_pool
Event	event_name
Interface	interface
Job ID	job_id
Job Type	job_name

Table 10. Breakout values (continued)

Breakout description	Breakout value
Node (logical node number)	node
Node Pool	nodepool
Operation Class	op_class
Path	path
Protocol	protocol
Service	service
Tier	tier

`{--filter-rule | -r} <breakout>:<filter>...`

Applies the specified filter rule to the exported data. To specify a filter, specify the breakout that the filter applies to.

Table 11. Filter values

Breakout description	Breakout value	Filter values
Client	remote_addr	<p>The hostname or IP address of a client. For example:</p> <pre>--filter-rule remote_addr:client.ip.com</pre>
Direction	direction	<ul style="list-style-type: none"> • out • in <p>For example:</p> <pre>--filter-rule direction:out</pre>
Disk	disk	<p>The disk number in the following format:</p> <pre><devid>/bay <number></pre> <p>For this filter, enclose both the breakout value and the filter value in quotation marks. For example:</p> <pre>--filter-rule "disk:l/bay 1"</pre>
Disk Pool	disk_pool	<p>The name of a disk pool. For example:</p> <pre>--filter-rule disk_pool:pool1</pre>
Event	event_name	<ul style="list-style-type: none"> • lookup • getattr • read • write • setattr • rename • link • unlink

Table 11. Filter values (continued)

Breakout description	Breakout value	Filter values
		<p>For example:</p> <pre>--filter-rule event_name:write</pre>
Interface	interface	<p>The interface number and type in the following format:</p> <pre><devid>/<type></pre> <p>The format mirrors the output format of the <code>isi networks list interfaces</code> command, except that the ":" is replaced with a "/".</p> <p>For example, if the following interface name appears in the output of the <code>isi networks list interfaces</code> command:</p> <pre>1:ext-1</pre> <p>The following filter value would be valid:</p> <pre>--filter-rule interface:1/ ext-1</pre>
Job ID	job_id	<p>The numeric ID of a job. For example:</p> <pre>--filter-rule job_id:4</pre>
Job Type	job_name	<ul style="list-style-type: none"> • AutoBalance • AutoBalanceLin • AVScan • Collect • Dedupe • DedupeAssessment • DomainMark • FlexProtect • FlexProtectLin • FSAnalyze • IntegrityScan • MediaScan • MultiScan • PermissionRepair • QuotaScan • SetProtectPlus • ShadowStoreDelete • SmartPools • SnapRevert • SnapshotDelete • TreeDelete <p>For example:</p> <pre>--filter-rule job_name:AVScan</pre>

Table 11. Filter values (continued)

Breakout description	Breakout value	Filter values
Node	node	The logical node number (LNN) of a node. For example: <code>--filter-rule node:1</code>
Node Pool	nodepool	The name of the node pool. For example: <code>--filter-rule nodepool:pool1</code>
Operation Class	op_class	<ul style="list-style-type: none"> • read • write • other • namespace_read • file_state • create • namespace_write • delete • session_state For example: <code>--filter-rule op_class:delete</code>
Path	path	The path of a directory beginning with /ifs. For example: <code>--filter-rule path:/ifs/data/media</code>
Protocol	proto_name	<ul style="list-style-type: none"> • nfs3 • nfs4 • smb2 • http • papi • siq For example: <code>--filter-rule proto_name:siq</code>
Service	service	
Tier	tier	The name of the tier. For example: <code>--filter-rule tier:archive</code>

`{--number-breakouts | -m} <integer>`

Specifies the number of breakout components to include for each breakout. For example, if you breakout a performance data module by client, specifying `--number-breakouts 5` exports data about the top five clients. The default number is 12.

`{--end | -e} {<timestamp> | now}`

Specifies the end of the data collection period. The default is `now`, which specifies the current time.

Specify `<timestamp>` in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

{--interval | -i} <integer> <units>

Specifies the length of the data collection period.

The default interval is one hour, or 1H. The following units are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

{--fmt-time | -f}

Creates an easily readable timestamp in the CSV file. If this option is not specified, the timestamp is in POSIX format.

{--min-max | -x}

Includes the maximum and minimum values for each 10 minute interval.

Examples

The following command exports data about active clients over the past hour, broken out by node:

```
iiq_data_export perf export --cluster cluster1 \  
--data-module client_active --path /home/exports --name disk-act \  
--breakout-by node
```

The following command exports data about active clients during January of 2015:

```
iiq_data_export perf export --cluster cluster1 \  
--data-module client_active --path /home/exports --name used \  
--end 2015-02-01 --interval 1M
```

iiq_data_export perf list

Displays the names of monitored clusters, performance data modules, and breakouts.

Syntax

```
iiq_data_export perf list {--clusters | --all-breakouts  
| --list-breakouts <data-module> | --all-modules}
```

Options

--clusters

Displays the names of all clusters that InsightIQ is monitoring.

--all-breakouts

Displays the names of all breakouts that InsightIQ supports for performance data modules. Each data module supports a subset of breakouts.

--data-modules

Displays the names of all available performance data modules.

Examples

To view the names of all clusters that are currently being monitored by InsightIQ, run the following command:

```
iiq_data_export perf list --clusters
```

iiq_restart

Stops and restarts InsightIQ.

Syntax

```
iiq_restart
```

Stop and restart InsightIQ.

This command can be used reset the InsightIQ server when the InsightIQ web application is unable to communicate with the server.

Options

There are no options for this command.

iiq_stop

Stops InsightIQ.

Syntax

```
iiq_stop
```

Options

There are no options for this command.

iiq_start

Starts InsightIQ after it has been stopped.

Syntax

```
iiq_start
```

Options

There are no options for this command.