



# Isilon InsightIQ

Version 3.0

## User Guide

Copyright © 2010-2014 EMC Corporation. All rights reserved. Published in USA.

Published January, 2014

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC<sup>2</sup>, EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>). For documentation on EMC Data Domain products, go to the EMC Data Domain Support Portal (<https://my.datadomain.com>).

EMC Corporation  
Hopkinton, Massachusetts 01748-9103  
1-508-435-1000 In North America 1-866-464-7381  
[www.EMC.com](http://www.EMC.com)

# CONTENTS

<b>Chapter 1</b>	<b>Introduction to InsightIQ</b>	<b>7</b>
	InsightIQ downsampling.....	9
	Maximum and minimum values.....	9
	Cluster monitoring through reports.....	9
	Performance reports.....	10
	File-system reports.....	11
	Report configuration components.....	12
	Data module information.....	12
	Filter rules and filters.....	17
	Breakouts.....	18
	Breakout definitions.....	18
	File System Analytics.....	20
	File System Analytics snapshots.....	20
	File System Analytics result sets.....	20
	Cluster-monitoring summary.....	21
	Read-only users.....	21
	Permalinks.....	21
<b>Chapter 2</b>	<b>Configuring InsightIQ</b>	<b>23</b>
	Managing monitored clusters .....	24
	Add a new cluster to monitor.....	24
	Import a cluster to monitor.....	25
	Modify monitored-cluster settings.....	26
	Suspend or resume cluster monitoring.....	26
	Stop monitoring a cluster and delete cluster data.....	26
	View monitored-cluster settings.....	27
	View cluster-monitoring status.....	27
	Managing datastores.....	28
	Specify a datastore.....	28
	View datastore settings.....	31
	View data set sizes.....	31
	Managing File System Analytics result sets.....	31
	Pin or unpin a File System Analytics result set.....	32
	Delete a File System Analytics result set.....	32
	View File System Analytics result sets.....	32
	File System Analytics result set information.....	33
	Configuring File System Analytics settings.....	33
	Configure File System Analytics settings.....	34
	Enable or disable File System Analytics.....	35
	View File System Analytics settings.....	35
	Managing InsightIQ virtual appliance settings.....	36
	Modify the InsightIQ virtual appliance time zone.....	36
	Modify the InsightIQ virtual appliance administrator password.....	36
	Modify InsightIQ virtual appliance network settings.....	37
	View InsightIQ virtual appliance network settings.....	37
	Configure the InsightIQ virtual appliance to monitor clusters through a proxy server.....	37
	Monitoring the status of InsightIQ.....	38
	View InsightIQ status.....	38

	Enable and configure InsightIQ status alert emails.....	38
	Disable InsightIQ status alert emails.....	38
	InsightIQ cluster monitoring errors.....	38
	Managing read-only users.....	39
	Add a read-only user account.....	39
	Modify the password of a read-only user account.....	40
	Delete a read-only user account.....	40
	Disable gathering event data.....	40
	Configure LDAP authentication.....	41
	Specify a custom SSL certificate.....	42
	Specify the InsightIQ port.....	42
	Configure email settings.....	43
	Test email configuration.....	43
	Create a permalink.....	43
<b>Chapter 3</b>	<b>Monitoring cluster status summaries</b>	<b>45</b>
	Monitor aggregated-cluster status.....	46
	Monitor the status of an individual cluster.....	46
<b>Chapter 4</b>	<b>Monitoring cluster status through reports</b>	<b>47</b>
	Managing live performance reports.....	48
	Create a live performance report.....	48
	View a live performance report.....	49
	Schedule a live performance report.....	50
	Modify a live performance report.....	50
	Delete a live performance report.....	50
	Managing performance report schedules.....	51
	Create a performance report schedule.....	51
	Disable a performance report schedule.....	53
	Enable a performance report schedule.....	54
	Modify a performance report schedule.....	54
	Delete a performance report schedule.....	54
	Generating performance reports.....	55
	Generate a performance report on demand.....	55
	Send a generated performance report.....	56
	Delete a generated performance report.....	56
	Modify the schedule of a generated performance report.....	57
	View a generated performance report.....	57
	View a live file-system report.....	58
	Managing rules and filters.....	59
	Create a filter.....	59
	Modify a filter.....	59
	Delete a filter.....	60
	Filter values.....	61
	Exporting InsightIQ data.....	61
	Export InsightIQ data to a CSV file.....	62
<b>Chapter 5</b>	<b>InsightIQ commands</b>	<b>63</b>
	iiq_data_export export.....	64
	iiq_data_export list.....	69
	iiq_data_export describe.....	71
	iiq_restart.....	73
	iiq_stop.....	73

	iiq_start.....	73
	iiq_network.....	73
<b>Chapter 6</b>	<b>InsightIQ troubleshooting</b>	<b>75</b>
	Configuration issues.....	76
	Datastore issues.....	76
	OneFS Platform API issues.....	77
	Web application issues .....	77
	Report schedule issues.....	78

## CONTENTS

# CHAPTER 1

## Introduction to InsightIQ

InsightIQ enables you to monitor and analyze the performance and file systems of Isilon clusters.

With InsightIQ, you can perform the following tasks:

- ◆ Determine whether a cluster is performing optimally.
- ◆ Compare changes in performance across multiple metrics, such as CPU usage, network traffic, protocol operations, and client activity.
- ◆ Correlate critical cluster events with performance changes.
- ◆ Determine the effect of workflows, software, and systems on cluster performance over time.
- ◆ View and compare properties of the data on the file system.

InsightIQ enables you to monitor and analyze Isilon cluster activity through flexible, customizable reports in the InsightIQ web-based application; you can customize these reports to provide detailed information about cluster hardware, software, and protocol operations. InsightIQ transforms data into visual information that emphasizes any performance outliers, enabling you to quickly and easily diagnose bottlenecks or optimize workflows. By transforming complicated data into relevant, visual information, you can perform the following tasks:

### **Determine a cause or validate a theory**

You can create and view specific reports to identify or confirm the cause of a performance issue. For example, if end users were reporting client connectivity problems during certain time periods, you could create an InsightIQ report that indicates whether the issue occurred at the cluster level, the node level, the client level, the disk level, or the network level. InsightIQ enables you to easily correlate seemingly unrelated data across both present and historical conditions.

### **Measure the effects of configuration changes**

If you modify your cluster environment, you can determine the effect of those changes by creating a report that compares the past, baseline performance with the performance since the changes were made. For example, if you added 20 new clients, you could create a report that illustrates the before-and-after performance; this would enable you to determine whether the added clients had an effect on a specific aspect of system performance, such as network throughput or CPU usage.

### **Optimize application workflows or performance**

You can create InsightIQ reports that help you identify bottlenecks or inefficiencies in your systems or workflows. For example, if you wanted to ensure that all clients are able to access the cluster quickly and efficiently, you could create a report that indicates whether some client connections are significantly faster or slower than others. You could then modify the report by adding breakouts and filters as needed to identify the root causes, which could be related to specific files that the clients are working with or the operations that the clients are performing.

### Analyze detailed operations data

You can customize InsightIQ reports to provide specific information about general cluster operations. For example, if you recently deployed an Isilon cluster, you might want to view a customized report that illustrated how the cluster and its individual components are performing.

### Forecast future needs

Reviewing past performance trends can help you predict future trends and needs. For example, if you deployed an Isilon cluster for a data-archival project six months ago, you might want to estimate when the cluster will reach its maximum storage capacity. You could customize a report to illustrate storage capacity usage by day, week, or month in order to determine approximately when the cluster will reach capacity if current trends were to continue.

You can install InsightIQ through Red Hat Package Manager (RPM) on a Linux machine or as a virtual appliance. If you install InsightIQ as a virtual appliance, you can configure virtual appliance settings, such as the IP address of InsightIQ and the administrator password, through InsightIQ. However, if InsightIQ is installed through RPM, you must configure the equivalent settings through the Linux interface.

- ◆ [InsightIQ downsampling](#).....9
- ◆ [Maximum and minimum values](#).....9
- ◆ [Cluster monitoring through reports](#).....9
- ◆ [Report configuration components](#).....12
- ◆ [File System Analytics](#).....20
- ◆ [Cluster-monitoring summary](#)..... 21
- ◆ [Read-only users](#)..... 21
- ◆ [Permalinks](#)..... 21



## InsightIQ downsampling

InsightIQ downsamples the data that it collects, meaning that InsightIQ converts higher-resolution data to lower-resolution data by adding or averaging several higher-resolution samples.

InsightIQ collects a data sample every five seconds. After a sample is more than 24 hours old, InsightIQ averages or adds the sample together with other samples to create a single, low-resolution sample that represents a 10 minute time period. InsightIQ then deletes the averaged samples to limit the size of the datastore. After 24 hours of monitoring, the growth rate of the InsightIQ datastore stabilizes and remains relatively constant over time.

## Maximum and minimum values

If you are viewing downsampled data in InsightIQ, you can view the maximum and minimum values for the downsampled time period.

Maximum and minimum values are shown as highlighted areas surrounding the line that represents the averaged value. You can use maximum and minimum values to see how cluster activity varied over a period of time. For example, if the maximum value for a time period greatly exceeds the averaged value, and the minimum time value is very close to the averaged value, it is likely that the value increased and decreased significantly over a very short period of time.

Maximum and minimum values are not displayed for the following data modules:

Data module	Notes
Protocol Operations Rate	You can view maximum and minimum values if you filter by protocol.
Protocol Operations Average Latency	
Active Clients	
Connected Clients	
File System Events Rate	You can view maximum and minimum values if you filter by file system (FS) event.
Jobs	
Job Workers	
L2 Cache Throughput Rate	
Storage Capacity	
HDD Storage Capacity	

## Cluster monitoring through reports

You can monitor clusters through customizable reports that display detailed data about clusters over specific periods of time.

InsightIQ enables you to view two general types of reports: performance reports and file-system reports. Performance reports include information about cluster activity and

capacity. Performance reports can be useful if, for example, you want to determine whether clusters are performing as expected or you want to investigate the specific cause of a performance issue. File-system reports include data about the files that are stored on a cluster. File-system reports can be useful if, for example, you want to identify the types of data being stored and where on a cluster that data is stored. Before you can apply a file-system report to a cluster, you must enable the InsightIQ File System Analytics feature for that cluster.

InsightIQ supports live versions of reports that are available only through the InsightIQ web application. You can create live versions of both performance and file-system reports. You can modify certain attributes of live reports as you are viewing the reports, including the time period, breakouts, and filters.

In addition to viewing performance reports live, you can also access them through static PDF reports. PDF reports are generated based on a report schedule that you configure, and cannot be modified to show different data after they are generated. However, you can configure InsightIQ to send generated reports to email addresses when they are generated, making the same data available to multiple people instantly. You can use generated reports to verify cluster health periodically and distribute InsightIQ information to people who do not have access to the InsightIQ web application.

## Performance reports

Performance reports enable you to view information about the activity of the nodes, networks, clients, disks, and IFS cache of a cluster. You can also view information about CPU usage, available capacity and recent cluster events.

You can use performance reports to verify whether a cluster is performing as expected, or identify the cause of a performance issue. For example, you can view performance reports that display information about network throughput during a period of time. If you find that throughput is not as high as expected, you can search for a cause of the issue by viewing information about any network throughput errors.

Performance reports can also be useful when you want to monitor cluster capacity. For example, you might create a performance report that shows the changes in available capacity over the past year. You might also create a report schedule that sends a weekly report that illustrates the cluster capacity over the past week.

## Live performance reports

Live performance reports enable you to dynamically modify a report as you view it. You can only view live performance reports through the InsightIQ web application.

While viewing a live performance report, you can modify the range of data that is displayed, as well as breakouts and filters. Live performance reports can be useful when you want to investigate a performance issue that is currently affecting a cluster. For example, you might notice through the InsightIQ dashboard that throughput is higher than expected. You could then view a live performance report to see the individual contributions of specific components, such as the direction in which the throughput is traveling.

In a live performance report, you can view the current status the monitored cluster, including the current date and time as reported by the cluster, the number of nodes in the cluster, a client-activity overview, a network throughput overview, the capacity usage of the cluster, and the CPU usage of the cluster.

## Scheduled performance reports

You can configure a report schedule that InsightIQ will generate static PDF reports from. Report schedules specify what is included in the reports, when the reports will be generated, and, optionally, one or more email addresses to send the reports to.

For each schedule, you specify when InsightIQ will generate the reports. You can also manually generate a report based on a report schedule at any time. All generated reports are archived and can be viewed and sent to email addresses. Generated reports are not automatically deleted by the system but can be manually deleted by a user.

---

### Note

InsightIQ generates PDF reports based on the time settings on the VM, not the time settings on the monitored cluster.

---

You can use performance report schedules to keep track of cluster performance over a long period of time. For example, if you wanted to keep track of the remaining unused capacity on a cluster, you could create a report schedule that InsightIQ would use to send you data about the capacity of the cluster. You could also configure the report to send copies of the report to other people who might not have InsightIQ user accounts.

## Generated performance reports

InsightIQ generates static PDFs, otherwise referred to as generated performance reports, at intervals specified in report schedules. You can configure InsightIQ to send the reports to up to 10 specified email addresses, or you can view the generated reports through the InsightIQ web application.

## File-system reports

File-system reports include information about the files stored on a cluster, such as the number of files stored on the cluster, the types of files stored on the cluster and where on the cluster the files are located.

InsightIQ supports two kinds of file-system reports: data-usage reports and data-property reports. Data-usage reports include data about the individual files on a cluster, such as which files have remained unmodified for the longest amount of time. Data properties reports include data about the overall cluster file system, such as how many files were last modified during a specific period of time. Data properties reports can also be compared between two different points in time; for example, you can view the relative changes in file count by physical size.

File-system reports can be useful if you are trying to identify the source of a performance issue on the file system. For example, if you viewed a performance report that showed that a cluster was almost out of storage space, you could view the data usage report and find the largest files or directories on the cluster.

File-system reports can also be useful if you are managing storage space among multiple clusters. For example, if you have two clusters, one that is used for development and one that is used for archiving, you could view the data properties report to verify that the development-cluster files were being modified often and the archive-cluster files were only being modified periodically. Furthermore, if the data properties report revealed that a significant number of files on the development cluster had not been accessed in a very long time, you could view the data usage report to identify the inactive files and consider transferring those inactive files to the archive cluster.

## Live file-system reports

You can view file-system reports in a live mode that enables you to adjust the data displayed as you view it.

You can view live file-system reports through the InsightIQ web application, and can change the range of data that is displayed, as well as apply breakouts and filters.

Live file-system reports can be useful when you need to inspect the changes in file layout of your system. For example, you might want to view how the number of files being stored in different directories has changed recently. With the data properties file-system report, you can compare the file system at a previous date with the most recent time that data was collected on the cluster.

You can use live file-system reports to verify that the file system of a cluster is behaving as expected. For example, if you had a cluster that designated for archiving data, you could use the data properties file-system report to view the file count by last modified file-system module. This would show you the numbers of files on your cluster organized by how recently they had last been modified. If the count of files that had been modified was particularly high, you could consider moving those files to a cluster intended for development.

You can also use live file-system reports to investigate files. For instance, if you discovered from a performance report or InsightIQ Dashboard that you had less capacity remaining than you expected to, you could view the data usage file-system report for that cluster. The data usage report includes a file-system module that can be customized to show the top 1000 largest files on the cluster. You could use this to verify which files were taking up the most space on the cluster.

At the top of each file-system report, you can view a **Current Summary** section that displays at-a-glance information that includes the name of the monitored cluster; the current date and time as reported by the cluster; an overview of cluster capacity versus usage, the number of files, directories, and other objects; the average number of files per directory; and the ratio of physical versus logical file size.

---

### Note

The information displayed in the **Current Summary** area represents the status of the cluster at the time displayed in the **Current Summary** area; the information does not reflect the cluster status during the time period that the report is displaying data about.

---

## Report configuration components

The data that reports contain is determined by three configuration components: data modules, breakouts, and filter rules.

A data module is a section of a report that displays data about the performance or file system of a cluster. You can apply a breakout to a data module to view the individual contributions of separate components. A filter rule can be applied across an entire report, affecting the data displayed in all data modules. Filter rules can be combined into collections called filters, which can be applied to reports.

## Data module information

Data modules display information about the performance or file system of a cluster, such as the number of protocol operations per second or the changes in file count organized by physical size. A data module that appears in a performance report is referred to as a

performance module. A data module that appears in a file-system report is referred to as a file-system module.

### Performance modules

---

#### Note

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

---

#### Active Clients

Displays the number of unique client addresses generating protocol traffic on the monitored cluster. Clients that are connected but not generating any traffic are not counted. You can optionally break out this data by protocol or node.

---

#### Note

Some protocols might issue a type of ping message, which can cause clients to appear active even if they are only sending these ping messages.

---

#### Average Disk Hardware Latency

Displays the average amount of time it takes for the physical disk hardware to service an operation or transfer. You can optionally break out this data by node or disk.

#### Average Disk Operation Size

Displays the average size of the operations or transfers that the disks in the cluster are servicing. You can optionally break out this data by direction, node or disk.

#### Average Pending Disk Operations Count

Displays the average number of operations or transfers that are in the processing queue for each disk in the cluster. You can optionally break out this data by node or disk.

#### Blocking File System Events Rate

Displays the number of file blocking events occurring in the file system per second. You can optionally break out this data by path or node.

#### CPU % Use

Displays the average CPU usage for all nodes in the monitored cluster. As some nodes may consume significantly more or less CPU resources than others, the average reflects the sum of the individual CPU-usage averages for each node. You can optionally break out this data by node; this breakout indicates the average CPU usage of each node. For example, at 10:52:22 AM on 2010-06-15, the specified node was using 14.35% of the total available node CPU capacity.

#### Cached Data Age

Indicates how long the oldest data has been in the cache. The shorter the time, the faster data is moving through the cache and old data in the cache is being replaced with newer data. You can optionally break out this data by node.

### **Clients Summary**

Displays clients that are currently consuming the most bandwidth. The list shows up to 64 clients, ordered from most to least active as determined by their combined input and output throughput rates during the specified time range.

---

#### **Note**

If the monitored cluster is running the Isilon SyncIQ module, the InsightIQ client list might include a client with an IP address of 0.0.0.0. This is expected behavior.

---

### **Connected Client**

Displays the number of unique client addresses with established TCP connections to the cluster on known ports. You can optionally break out this data by protocol or node.

---

#### **Note**

UDP connections do not appear as connected. In addition, some short-lived TCP connections might not appear as connected even though they are active.

---

### **Contended File System Events Rate**

Displays the number of file contention events, such as lock contention or read/write contention, occurring in the file system per second. You can optionally break out this data by path or node.

### **Deadlocked File System Events Rate**

Displays the number of file system deadlock events the file system is processing per second. This information can be useful if you want to identify a specific file state that might be contributing to performance issues. Deadlocked events occur regularly during normal cluster operation, and the file system is designed to automatically detect and break them. You can optionally break out this data by path or node.

### **Disk Activity**

Displays the average percentage of time that disks in the cluster spend performing operations instead of sitting idle. You can optionally break out this data by node or disk.

### **Disk Operations Rate**

Displays the average rate at which the disks in the cluster are servicing data read/write/change requests, also referred to as operations or disk transfers. You can optionally break out this data by disk, direction, or node.

### **Disk Throughput Rate**

Displays the total amount of data being read from and written to the disks in the cluster. You can optionally break out this data by disk, direction, or node.

**Event Summary**

Displays events generated by the monitored cluster during the specified time range. Also includes events that were started prior to but ended during the specified time range. This data module is available only for clusters running OneFS 7.0 or later. The following list describes the specific columns of the **Events** table.

**Event Timeline**

The horizontal red event timeline represents the same period of time that is currently represented in the performance charts. A vertical red line on the event timeline indicates the start of an event. You can rest the pointer over specific sections of the horizontal timeline to view detailed information about any time-specific events or alerts.

**Severity**

Indicates the severity level of the event.

**Start**

Indicates when the event started.

**End**

If applicable, indicates when the event ended.

**Instance**

Displays the unique identifier of the event.

**Message**

Displays a description of the event.

**External Network Errors**

Displays the number of errors generated for the external network interfaces. You can optionally break out this data by direction, interface, or node.

**Note**

During normal operations, this chart indicates an error count of 0. Errors reported in this chart are often the result of network infrastructure issues (for example, malformed frames or headers, handoff errors, or queuing errors) rather than OneFS issues. When investigating the cause of these errors, first review the logs and reports for your network switch and other network infrastructure components.

**External Network Packets Rate**

Displays the total number of packets that passed through the external network interfaces in the monitored cluster. You can optionally break out this data by direction, interface or node.

**External Network Throughput Rate**

Displays the total amount of data that passed through the external network interfaces in the monitored cluster. You can optionally break out this data by interface, direction, client, operation class, protocol, or node.

**File System Events Rate**

Displays the number of file system events, or operations, (such as read, write, lookup, or rename) that the file system is servicing per second. You can optionally break out this data by direction, operation class, path, node, or event.

**File System Output Rate**

Displays the rate at which data is being read from the file system. You can optionally break out this data by node.

### **HDD Storage Capacity**

Displays the used and total storage capacity for all hard-disk drives in the monitored cluster. You can optionally break out this data by node.

### **Jobs**

Displays the number of active and inactive jobs on the cluster. An active job is a system job that is currently being performed by workers. An inactive job is a system job that has been assigned workers, but the workers are not currently performing the job.

### **Job Workers**

Displays the number of active and assigned workers on the cluster. An active worker is a worker that is currently performing a system job. An assigned worker is a worker that has been assigned to a system job but is not currently performing the job.

### **L2 Cache Throughput Rate**

Indicates how much data was requested from the level-two (L2) cache and how much data matched the requests. On the chart, the **starts** line indicates how much data was requested from the file system's L2 cache. The **hits** line indicates the amount of data in the L2 cache that matched those requests. You can optionally break out this data by node.

---

#### **Note**

The level-two cache is the cache on the storage node that contains the cached disk. This is in contrast to the level-one cache, which is the cache on the node that processes protocol requests.

---

### **Locked File System Events Rate**

Displays the number of file lock operations occurring in the file system per second. You can optionally break out this data by path or node.

### **Node Summary**

Displays status and performance metrics for all nodes in the monitored cluster.

### **Pending Disk Operations Latency**

Displays the average amount of time that disk operations spend in the input/output scheduler. You can optionally break out this data by node or disk.

### **Protocol Operation Average Latency**

Displays the average amount of time required for protocols to process incoming operations. These values are typically represented in fractions of seconds. You can optionally break out this data by client, operation class, protocol, or node.

### **Protocol Operations Rate**

Displays the total number of requests that were originated by clients for all file data access protocols. Combined with data from the disk throughput and disk operations rate data elements, this information can help you identify specific clients that might be significantly contributing to cluster load. You can optionally break out this data by client, operation class, protocol, or node.

### **SSD Storage Capacity**

Displays the used and total storage capacity for all solid-state drives in the monitored cluster. You can optionally break out this data by node.

### **Slow Disk Access Rate**

Displays the rate at which slow (long-latency) disk operations occur. You can optionally break out this data by node or disk.



**Storage Capacity**

Displays the used and total storage capacity for the monitored cluster. You can optionally break out this data by node.

**File-system modules****Directories**

Displays disk-usage data and file counts per directory, recursively. You can sort the information by clicking a column heading in the table view; this causes the chart on the left to reload and display a visual representation of the specified information. You can click path names to view more specific information about any subdirectories contained in the directory. You can also create a filter for a specific directory that you can use to filter data in other InsightIQ views.

**File Count by Physical Size**

Displays the total number of files in a specified physical-size range. For example, the monitored cluster might contain 23,724 files that are between 100 GB and 1 TB in physical size. You can optionally break out this data by modified time, disk pool, extension, or directory.

**File Count by Logical Size**

Displays the total number of files in a specified logical-size range. For example, the monitored cluster might contain 3,329,210 files that are between 1 GB and 10 GB in logical size. You can optionally break out this data by modified time, disk pool extension, or directory.

**File Count by Last Modified**

Displays the total number of files that were last modified within a specified time range. For example, the monitored cluster might contain 38,734 files that were last modified between 14 days and 21 days ago. You can optionally break out this data by disk pool, logical size, extension, directory, or physical size.

**Top <number> Files by <element>**

Displays the top files sorted by the specified element. You can configure the maximum number of top-contributing files to display.

**Top <number> Directories by <element>**

Displays the top directories sorted by the specified element. You can configure the maximum number of top-contributing directories to display.

## Filter rules and filters

You can apply filter rules across live reports to isolate data relating to specific factors. While breakouts specify a category of contributors, filter rules specify an individual contributor belonging to a category. For example, you can breakout a data module by protocol, and you can apply a filter rule for iSCSI. Breakouts appear beneath data modules without modifying the data module; filter rules modify what data is displayed in data modules. Filters are customized collections of filter rules that you can create, save, and apply to various reports.

If you apply a filter rule to a report, all data modules in the report will only display data about the filter rule. For example, you could apply a filter rule for a specific node in a cluster. The data modules in the report that the filter rule was applied to would display data about that node.

You can create and apply filters, which contain one or more filter rules. For example, a filter could contain both a filter rule for a specific node in a cluster and a filter rule for a specific client accessing that cluster. Applying this filter would cause data modules to display information about the interactions between only that node and client. You can save filter rules and then apply them to specific reports.

## Breakouts

You can apply breakouts to data modules in order to view the individual contributions of various factors. You can apply only one breakout to a data module at a time.

Breakouts are heat maps that display variations of color that visually represent each component's contribution to overall performance. The darker the color on a heat map, the greater the associated component's contribution to performance. Heat maps enable you to track performance trends and to quickly identify hot spots where performance might be constrained by specific components. If you rest the mouse pointer over any point in a heat map, InsightIQ displays performance data for the specified component for that moment in time. Breakouts are sorted in order of significance, with the most significant elements at the top of the list.

Breakouts can be useful when trying to determine the cause of a cluster issue. For example, if you broke out a **CPU usage** data module by node, you could then view the individual CPU usage of each node.

---

### Note

The sum of individual breakouts might not always match the total aggregate count for a specified element. For example, not all network traffic is associated with a specific protocol. Therefore, the sum of the individual protocol-related breakouts might not match the total reported protocol-related counts.

---

## Breakout definitions

Each breakout displays the individual contributions of specific factors. Each data module supports specific breakouts.

The following list describes the factors that each breakout includes.

### Accessed Time

Breaks out data by when a file was last accessed.

### Client

Breaks out data by individual client. Only the most active clients are represented in the data that is displayed.

### Disk

Breaks out data by individual disk.

### Disk Pool

Breaks out data by the size of files in each disk pool.

### Direction

Breaks out data by the direction of cluster throughput.

### Event

Breaks out data by individual network interface.

### File Extension

Breaks out data by file extension.

### Direction

Breaks out data by inbound traffic.

### Interface

Breaks out data by individual network interface.

### Job ID

The numerical ID of a job.

**Job Name**

The name of a job type.

**Logical Size**

Breaks out data by logical file size. Logical file size calculations include only data, and do not include data-protection overhead.

**Modified Time**

Breaks out data by when a file was last modified.

**Node**

Breaks out data by individual node. Nodes are identified by their logical node number (LNN).

**Node (devid)**

Breaks out data by individual node. Nodes are identified by their device ID.

**Operation Class**

Breaks out data by the type of operation being performed. The following operation types are supported:

**Read**

File and stream reading.

**Write**

File and stream writing.

**Create**

File, link, node, stream, and directory creation.

**Delete**

File, link, node, stream, and directory deletion.

**Namespace Read**

Attribute, statistic, and ACL reading; lookups; and directory reading.

**Namespace Write**

Renaming; attribute setting; and permission, time, and ACL writing.

**File State**

Opening, closing, locking, acquiring, releasing, breaking, checking, and notifying.

**Session State**

Negotiating, inquiring, and manipulating protocol connections or sessions.

**Other**

File-system information and other operations that cannot be categorized.

**Path**

Breaks out data by individual file name or directory name.

**Physical Size**

Breaks out data by physical size. Physical size calculations include data-protection overhead.

**Protocol**

Breaks out data by protocol (for example, NFS or iSCSI).

**User Attribute**

Breaks out data by a user-defined attribute, if an attribute is defined. You can define attributes through the OneFS command-line interface.

## File System Analytics

The InsightIQ File System Analytics feature enables you to create customizable reports that analyze data usage, data properties, and other performance data at the file-system level.

For example, you can create reports that show the distribution of file sizes, ages, and locations. This type of data can be particularly useful if you want to analyze capacity trends in order to plan for future storage needs.

If you want to enable the File System Analytics functionality, the monitored cluster must be running OneFS 6.0 or later.

### File System Analytics snapshots

If File System Analytics is enabled, you can configure InsightIQ to generate File System Analytics result sets from snapshots taken on a monitored cluster.

If you enable snapshots for the File System Analytics job, OneFS takes snapshots initiated by InsightIQ on the monitored cluster. InsightIQ then generates File System Analytics result sets from those snapshots. Generating result sets from snapshots enables you to view data that represents an exact point in time. However, the process of taking a snapshot can negatively affect cluster performance.

---

#### Note

The snapshots initiated by InsightIQ are used only by InsightIQ and are not part of the Isilon SnapshotIQ module.

---

If you disable snapshots for the File System Analytics job, InsightIQ generates File System Analytics result sets based on the current state of the monitored cluster. Because the state of the cluster might change during the data-collection process, generating result sets based on the current state of the cluster prevents you from viewing exact point-in-time data. However, generating result sets based on the current state of the cluster does not result in the potential cluster-performance issues that can occur from the snapshot method.

### File System Analytics result sets

If File System Analytics is enabled on a monitored cluster, the cluster generates File System Analytics result sets. InsightIQ generates file system reports from the File System Analytics result sets that it collects.

One file-system report is generated from one result set. Result sets are named for the day and time that they were collected. By default, the monitored cluster generates a result set once per day.

Unlike InsightIQ data sets, which are stored on the InsightIQ datastore, File System Analytics result sets are stored on the monitored cluster. The monitored cluster routinely deletes result sets to avoid any negative effects on system performance. You can specify the maximum number of result sets that are retained. If you do not want a specific result set to be deleted, you can pin a result set so that it cannot be deleted by the system.

## Cluster-monitoring summary

InsightIQ includes a cluster-monitoring summary page that enables you to view overview information about all monitored clusters.

You can view the cluster-monitoring summary to verify the overall status of all monitored clusters. The cluster-monitoring summary includes overview information about cluster capacity, clients, throughput, and CPU usage. Current information about clusters appears alongside graphs that show the relative changes in these statistics over the past 12 hours.

The cluster-monitoring summary includes an aggregated-cluster view that displays the total or average values of monitored-cluster status information. The aggregated-cluster view can be useful when deciding what to include in a performance report. For example, if you see that the total network throughput of all monitored clusters is higher than anticipated for the current time and day, you can create a customized performance report that includes data about network throughput. You can then break out network throughput by direction in the report to determine whether one direction of throughput is contributing to the total more than the other.

The cluster-monitoring summary also includes separate individual-cluster views that displays information on a specific monitored cluster. Individual-cluster views can be useful when deciding which clusters to apply a performance report to. For example, if you identify a throughput issue from the aggregated-cluster view, you can quickly review the individual-cluster views to determine which cluster is the source of the issue. You can then create a performance report in order to view more detailed information about the problem. You might determine that the issue is isolated on one specific cluster, and that you therefore need to apply the report to only one cluster. The individual-cluster views could also reveal that the problem is occurring on all clusters, and you need to apply the report to all clusters individually in order to identify the source of the issue.

## Read-only users

A read-only user can log in to the InsightIQ application and monitor cluster activity but cannot modify any configuration settings. A read-only user cannot create performance reports or add read-only user accounts.

During the initial InsightIQ setup process, you configured the default administrator user account; the administrator user can configure InsightIQ settings in addition to monitoring cluster activity. The administrator can also create, modify, and delete accounts for read-only users. There is exactly one administrator user account. However, the administrator can configure an unlimited number of read-only users.

## Permalinks

A permalink is a static link that enables you to save or share an InsightIQ report in a specific configuration. Permalinks can save time if you want to frequently view the same data for a certain period of time, or want to share that data with other InsightIQ users.

If you create a permalink, InsightIQ generates and displays a URL that represents a shortcut to a configured report. You can then go to this URL at any time to view the report in that specific configuration.

For example, you might want to share a report configured to display information about the external network throughput of Node 5 from 11:20:00 to 12:20:00 on 2011-02-15.

You could view, configure, and create a permalink of the report, and then send the generated URL to other InsightIQ users to enable them to view that information.

---

**Note**

The high-resolution data displayed in permalinks will degrade over time due to the InsightIQ data-retention and downsampling policies.

---

# CHAPTER 2

## Configuring InsightIQ

You can manage InsightIQ monitored-cluster, datastore, email and File System Analytics settings through the InsightIQ web application; these settings are not available in the command-line interface. You can also manage network and authentication settings of the InsightIQ virtual appliance.

---

### Note

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

---

◆ <a href="#">Managing monitored clusters</a> .....	24
◆ <a href="#">Managing datastores</a> .....	28
◆ <a href="#">Managing File System Analytics result sets</a> .....	31
◆ <a href="#">Configuring File System Analytics settings</a> .....	33
◆ <a href="#">Managing InsightIQ virtual appliance settings</a> .....	36
◆ <a href="#">Monitoring the status of InsightIQ</a> .....	38
◆ <a href="#">Managing read-only users</a> .....	39
◆ <a href="#">Disable gathering event data</a> .....	40
◆ <a href="#">Configure LDAP authentication</a> .....	41
◆ <a href="#">Specify a custom SSL certificate</a> .....	42
◆ <a href="#">Specify the InsightIQ port</a> .....	42
◆ <a href="#">Configure email settings</a> .....	43
◆ <a href="#">Test email configuration</a> .....	43
◆ <a href="#">Create a permalink</a> .....	43

## Managing monitored clusters

You can manage settings that control how InsightIQ accesses a monitored cluster, and you can add or remove monitored clusters. You can also import existing monitored clusters from other instances of InsightIQ, and you can temporarily suspend cluster monitoring.

You can configure InsightIQ to monitor more than one Isilon cluster simultaneously. The maximum number of clusters that you can simultaneously monitor varies depending on the resources available to the virtual machine. It is recommended that you monitor no more than eight clusters or 150 nodes at one time. If you want to monitor more clusters or nodes than this, it is recommended that you deploy an additional instance of InsightIQ.

To enable InsightIQ to connect to a monitored cluster, you must specify the host name or IP address of any node on the monitored cluster, and the authentication credentials for the local InsightIQ or API user as they are configured on the monitored cluster. If the monitored cluster is running OneFS 6.0 or later, the corresponding local user on the monitored cluster is the InsightIQ (**insightiq**) user. If the monitored cluster is running a version of OneFS earlier than 6.0, the corresponding local user on the monitored cluster is the OneFS Platform API (**api**) user.

You can import a monitored cluster from another instance of the InsightIQ. This can be useful if, for example, you have upgraded the InsightIQ software and you want to continue monitoring a cluster that the older version of InsightIQ was monitoring. If you import a monitored cluster, InsightIQ collects the configuration settings for the monitored cluster, as well as all data previously collected for the cluster, and adds them to the datastore associated with the current instance of InsightIQ.

### Add a new cluster to monitor

Before you can view data about a cluster through InsightIQ, you must configure InsightIQ to monitor that cluster.

#### Before you begin

- ◆ Verify that a valid InsightIQ license is enabled on the monitored cluster; for more information, contact your Isilon representative.
- ◆ Verify that the local InsightIQ user is enabled and configured with a password on the monitored cluster.

#### Procedure

1. Click **Settings** > **Monitored Clusters**, and then, on the **Monitored Clusters** page, click **Add Cluster**.

The **Add Cluster** dialog box appears.

---

#### Note

If you are installing InsightIQ for the first time, the **Add Cluster** dialog box might already be displayed.

---

2. In the **Add Cluster** dialog box, click **I want to monitor a new cluster**.
3. In the **Isilon cluster** field, type the host name or IP address of any node in the cluster that you want to monitor. Alternatively, you can type the name of an Isilon SmartConnect zone.



**Note**

In general, it is recommended that you specify a monitored cluster by a SmartConnect zone name. However, if the monitored cluster is heavily loaded and you rely on InsightIQ file-heat data, it is recommended that you specify the monitored cluster by the IP address or host name of a specific node in the cluster; avoid specifying an IP address that can be transferred from node to node. If you choose to identify the monitored cluster by a SmartConnect zone, specify a SmartConnect zone that includes a CPU load-balancing policy. By balancing connections to nodes with lower CPU usage, the monitored cluster can respond to InsightIQ data-collection queries more quickly and efficiently.

4. In the **Username** field, type the user name of the monitored cluster's InsightIQ user:
  - If the cluster is running OneFS 6.0 or later, type `insightiq`
  - If the cluster is running a version of OneFS earlier than 6.0, type `api`
5. In the **Password** field, type the password of the monitored cluster's InsightIQ user.
6. Click **OK**.  
InsightIQ begins monitoring the cluster.

## Import a cluster to monitor

If you import a monitored cluster from another instance of InsightIQ, InsightIQ immediately begins monitoring that cluster. In addition, InsightIQ initiates a one-time process that imports data from the other instance of InsightIQ. Depending on the amount of data being imported, this import process can take a long time to complete.

InsightIQ merges the imported data with the existing monitored-cluster data and stores it in the datastore associated with the new instance of InsightIQ. During the import process, in order to ensure data integrity, the remote instance of InsightIQ that is being imported from might suspend service, shut down, and then restart.

You can import data only from instances of InsightIQ running InsightIQ 3.0

### Procedure

1. Ensure that the InsightIQ port configurations match.
  - a. Open an SSH connection to both InsightIQ appliances and log in.
  - b. View the contents of the `/etc/isilon/insightiq.ini` files on both InsightIQ instances. The `redirect_to_port` and `port` values must match for both InsightIQ instances. If the ports do not match, modify the value on the InsightIQ instance you are importing to.
  - c. If you modified the `/etc/isilon/insightiq.ini` file, save the file and then restart InsightIQ by running the following command:
 

```
iiq_restart
```
2. Click **Settings** > **Monitored Clusters**, and then, on the **Monitored Clusters** page, click **Add Cluster**.
3. Click **Add Cluster**.
4. Click **I want to monitor a cluster already being monitored**.
5. In the **InsightIQ VM host** field, type the host name or IP address of InsightIQ.
6. In the **Username** and **Password** fields, type a valid username and password.

- If you are importing from an InsightIQ virtual appliance, specify the username and password of the administrator account.
  - If you are importing from a Linux machine, specify the username and password of any user account other than root.
7. Click **OK**.
  8. In the **Password** field for the cluster that you want to import, type the password of the cluster's InsightIQ user and then press ENTER.
  9. Select the check box for the cluster that you want to monitor, and then click **Import**.

## Modify monitored-cluster settings

If you modify the host name or IP address of a monitored cluster, or its local InsightIQ user credentials, you must also modify those settings accordingly in the InsightIQ application. InsightIQ uses these values to connect to the monitored cluster.

### Procedure

1. Click **Settings > Monitored Clusters**.

The **Monitored Clusters** page appears and displays a list of all clusters that InsightIQ is configured to monitor.

2. For the cluster whose settings you want to view, in the **Actions** column, click **Configure**.

The **Configuration** page appears.

3. Modify the monitored-cluster settings as needed, submitting changes on each tab.

## Suspend or resume cluster monitoring

As an alternative to permanently deleting a monitored-cluster entry, you can instead temporarily suspend monitoring for the cluster and then resume monitoring later.

If you suspend monitoring for a cluster, InsightIQ completes any data collection queries that are currently in process but does not start any new queries until you resume monitoring. The datastore remains intact, but InsightIQ does not collect or store any new data. Depending on how long monitoring is suspended, reports for the affected cluster might display periods of missing data.

### Procedure

1. Click **Settings > Monitored Clusters**.

The **Monitored Clusters** page appears and displays a list of all clusters that InsightIQ is configured to monitor. The current monitoring state of each cluster is indicated in the **Actions** column. If the **Suspend** link is displayed, InsightIQ is currently monitoring the cluster. If the **Resume** link is displayed, InsightIQ is not monitoring the cluster.

2. For the cluster whose monitoring state you want to modify, click **Suspend** or **Resume**.

## Stop monitoring a cluster and delete cluster data

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



**If you stop monitoring a cluster, InsightIQ deletes all data that was previously collected from that cluster; you cannot undo the deletion.**

As an alternative to this procedure, if you want to retain previously collected data for a cluster but want to temporarily stop monitoring the cluster, you can temporarily suspend monitoring and then resume monitoring later.

#### Procedure

1. Click **Settings** > **Monitored Clusters**.

The **Monitored Clusters** page appears and displays a list of all clusters that InsightIQ is configured to monitor.

2. In the **Actions** column, click the **Delete** link for the cluster that you want to permanently stop monitoring.

A confirmation dialog box appears.

3. Click **OK**.

## View monitored-cluster settings

In the InsightIQ web application, you can view the configuration settings that InsightIQ uses to connect to the monitored cluster. You cannot view these settings in the command-line interface.

#### Procedure

1. Click **Settings** > **Monitored Clusters**.

The **Monitored Clusters** page appears and displays a list of all clusters that InsightIQ is configured to monitor.

2. In the row of the cluster whose settings you want to view, click **Configure**.

The **Configuration** page appears, and displays information about the specified monitored cluster.

## View cluster-monitoring status

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

#### Procedure

1. Click **Settings** > **Monitored Clusters**.

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 cluster's status as indicated by its icon color:

#### Green

A green icon indicates that communication between cluster and the InsightIQ application is normal with no errors.

#### Yellow

A yellow icon indicates that communication between cluster and the InsightIQ application has been interrupted by a transient error. This might be due to a brief timeout. This condition will most likely resolve itself without need for user intervention.

**Red**

A red icon indicates that communication between the cluster and the InsightIQ application has been interrupted indefinitely. This might be due to an authorization issue, unconfigured license, or prolonged timeout. This condition cannot be resolved without user intervention.

## Managing datastores

InsightIQ collects data from all monitored clusters and stores it on the local datastore that is included with InsightIQ, or on an NFS-based server or Isilon cluster that you specify.

If you do not want to use the local datastore that is included with InsightIQ to store monitored-cluster data, you can instead store the data on any NFS-based server or on an Isilon cluster—either a separate Isilon cluster or the same Isilon cluster that InsightIQ is monitoring. If the InsightIQ datastore is on the monitored cluster, the cluster appears as a client of itself in the InsightIQ application.

InsightIQ does not store high-resolution data indefinitely. The amount of stored data and the length of time that data is retained depend on many variables. In general, as the InsightIQ data samples grow older, InsightIQ purges the higher-resolution data sets as needed to limit the size of the datastore. You can also manually delete specific data sets. For example, you might want to delete all CPU-usage data for a specific monitored cluster if you no longer need that data and want to free up some space on the datastore.

## Specify a datastore

You must specify a datastore where you want to store the data collected by InsightIQ. The specific steps depend on whether you want to store data on the local datastore or on an NFS-mounted datastore.

If you install InsightIQ as a virtual appliance, the local datastore is a virtual hard drive that is included with InsightIQ. If you install InsightIQ through RPM, the local datastore is the hard drive of the Linux machine. If you want to store InsightIQ data on the local datastore, you must ensure that the virtualization host or Linux machine contains at least 70 GB of free disk space available.

The NFS datastore can be either an Isilon cluster or another NFS-mounted server. If you want to store InsightIQ data on an NFS datastore, you must ensure that the datastore meets certain requirements; for more information, see [NFS datastore requirements on page 30](#).

- ◆ If you want to store InsightIQ data on the local datastore that is included with InsightIQ, see [Specify the local datastore on page 28](#).
- ◆ If you want to store InsightIQ data on an NFS-mounted datastore, see [Specify an NFS datastore on page 29](#).

## Specify the local datastore

You can configure InsightIQ to store monitored-cluster data on the local datastore.

**⚠ CAUTION**

**If you modify the datastore location after InsightIQ has been installed and configured, InsightIQ does not migrate any previously collected data from the old datastore to the new datastore. Any data collected on the old datastore will not be available while InsightIQ is connected to the new datastore.**

---

**Procedure**

1. Click **Settings > Datastore**.

The **Configure Datastore Path** page appears and displays the current location of the datastore.

**Note**

If you are installing InsightIQ for the first time, the **Configure Datastore Path** page might automatically appear.

2. Click **Local Datastore**.
3. Specify the path of the local datastore.

Options	Description
<b>For virtual appliance</b>	In the <b>Datastore Path</b> field, type <code>/datastore</code> .
<b>For RPM</b>	In the <b>Datastore Path</b> field, type the path of an empty directory on the Linux machine.

4. Click **Submit Changes**.

All data subsequently collected by InsightIQ is stored in the local datastore.

## Specify an NFS datastore

You can configure InsightIQ to store monitored-cluster data on an NFS datastore.

**CAUTION**

If you modify the datastore location after InsightIQ has been installed and configured, InsightIQ does not migrate any previously collected data from the old datastore to the new datastore. Any data collected on the old datastore will not be available while InsightIQ is connected to the new datastore.

**Procedure**

1. Click **Settings > Datastore**.

The **Configure Datastore Path** page appears and displays the current location of the datastore.

**Note**

If you are installing InsightIQ for the first time, the **Configure Datastore Path** page might automatically appear.

2. Click **NFS Mounted Datastore**.
3. In the **Datastore NFS Server** field, type the host name or IP address of the server or Isilon cluster on which collected performance data will be stored.
4. In the **Datastore NFS Server Path** field, type the absolute path, beginning with a slash mark (/), to the directory on the server or cluster where you want the collected data to be stored. This field must only contain ASCII characters.
5. Click **Submit Changes**.

All data subsequently collected by InsightIQ is stored in the specified NFS datastore.

## InsightIQ datastore size

You must provide an adequate amount of space in the InsightIQ datastore to contain the data that InsightIQ creates. The amount of data that InsightIQ creates is dependent on the number of nodes being monitored and the amount of time you want InsightIQ to retain data for.

You must reserve at least 1.5 GB of space per node that you monitor. If you do not reserve the minimum amount of space, the InsightIQ datastore might become full before InsightIQ begins deleting older data to make room for newer data. Reserving 1.5 GB of space per node will allow InsightIQ to retain approximately two weeks worth of data.

If you want to retain more than two weeks worth of data, it is recommended that you increase the size of the InsightIQ datastore and reserve 3 GB per node per month that you want to retain data for. Use the following equations to determine the size of your datastore:

$$\begin{aligned} \text{<minimum\_datastore\_size>} &= \text{<number\_of\_nodes>} * 1.5 \\ \text{<larger\_datastore\_size>} &= \text{<number\_of\_nodes>} * \text{<number\_of\_months>} * 3 \end{aligned}$$

For example, if you monitor 8 nodes, you must reserve a minimum of 12 GB of space in your datastore. However, if you want to retain three months worth of data, you must reserve 72 GB.

## NFS datastore requirements

As an alternative to storing InsightIQ data on the local datastore, you can store InsightIQ data on the monitored cluster itself, on a different Isilon cluster, or on any NFS-mounted server. If you want to store InsightIQ data on an Isilon cluster or other NFS-based server, you must provide adequate space in which to store the data that InsightIQ collects.

If you want to store InsightIQ data on an Isilon cluster or other NFS-based server, you must ensure that the cluster or server includes a properly configured NFS export rule that exports the datastore path and grants write access to the "root" user for the specified InsightIQ IP address; this enables InsightIQ to mount the cluster or server and create the necessary directories and files on the cluster or server.

InsightIQ might cache permissions for failed attempts to mount the NFS export. If InsightIQ continues to report that it does not have sufficient rights to create the path on the NFS host, reboot InsightIQ to clear the previous mount attempts and then try again.

## Datastore requirements for an Isilon cluster

If you want to use an Isilon cluster to store data collected by InsightIQ, the Isilon cluster must meet specific requirements.

---

### Note

This information is applicable only if you are storing data on an Isilon cluster.

If you store InsightIQ data on an Isilon cluster other than the monitored cluster itself, you do not need to license InsightIQ on that cluster. In addition, if you store the InsightIQ data on a different Isilon cluster, that cluster can be running any version of the Isilon OneFS operating system that is supported by InsightIQ. If you store the InsightIQ data on the cluster that InsightIQ is monitoring, the cluster will appear as a client of itself in InsightIQ.

You must ensure that the cluster includes a properly configured NFS export rule that exports the datastore path and grants write access to the "root" user. 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, you must create a new NFS export rule that allows write access for InsightIQ.

## View datastore settings

You can view the configured settings for the server or Isilon cluster on which InsightIQ stores collected performance data for all monitored clusters.

### Procedure

1. Click **Settings** > **Datastore**.

The **Configure Datastore Path** page appears.

2. Review the datastore settings:

#### Datastore path

The directory on the virtual or Linux machine that InsightIQ data is stored in.

#### Datastore NFS server

The host name or IP address of the server or cluster that InsightIQ data is stored in.

#### Datastore NFS server path

The path of the directory on the NFS server that InsightIQ data is stored in.

## View data set sizes

You can view the size of each data set collected for a monitored cluster.

### Procedure

1. Click **Settings** > **Monitored clusters**.
2. In the **Actions** column for the cluster that the data was collected for, click **Configure**.
3. Click **Data Set Sizes**.
4. In the **Data Set Sizes** table, view data set sizes.

## Managing File System Analytics result sets

The InsightIQ File System Analytics feature enables you to view and analyze file-system reports. If you want to enable the File System Analytics functionality, the monitored cluster must be running OneFS 6.0 or later.

When File System Analytics 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 File System Analytics job through OneFS; you can also configure the level of detail displayed in file-system reports through InsightIQ.

---

### Note

When enabled, the File System Analytics job consumes computing resources on the monitored cluster and can affect cluster performance. If cluster performance is negatively affected, you can disable the File System Analytics feature; this prevents the job from running.

---

## Pin or unpin a File System Analytics result set

You can pin a File System Analytics result set so that 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.

### Procedure

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** page 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.

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

## Delete a File System Analytics result set

You can manually delete a File System Analytics result set. Any result set can be deleted manually, even if the result set is pinned.

### Procedure

1. Click **Settings** > **Monitored Clusters**.

The **Monitored Clusters** page appears.

2. In the **Actions** column for the cluster whose File System Analytics result set you want to delete, click **Configure**.

The **Configuration** page 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 **Action** column for the result set that you want to delete, click **delete**.

## View File System Analytics result sets

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

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 [View a live file system report on page 58](#).



### Procedure

1. Click **Settings > Monitored Clusters**.

The **Monitored Clusters** page appears.

2. In the **Actions** column for the cluster whose File System Analytics result sets you want to view, click **Configure**.

The **Configuration** page 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.

## File System Analytics result set information

The **Manage FSA Result Sets** section on the **Manage FSA Result Sets** tab of the **Configuration** page 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.

## Configuring File System Analytics settings

You can configure File System Analytics settings that enable and disable File System Analytics, and how File System Analytics collects data on the cluster. File System

Analytics settings also control how long result sets are retained on the cluster before being automatically deleted.

## Configure File System Analytics settings

You can configure settings for the InsightIQ File System Analytics feature, including whether InsightIQ generates File System Analytics result sets from snapshots taken on the monitored cluster.

### Procedure

1. Click **Settings** > **Monitored Clusters**.

The **Monitored Clusters** page appears.

2. In the **Actions** column for the cluster whose File System Analytics settings you want to configure, click **Configure**.

The **Configuration** page appears.

3. Click the **FSA Configuration** tab.

The **FSA Configuration** tab appears.

4. Configure the **Result Set Options** settings:

- 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 reports.

- c. In the **Path squash depth** field, type an integer that represents the maximum number of directory levels to include in a path, including `/ifs`.

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. Enable or disable snapshots for the File System Analytics job.

- If you want InsightIQ to generate File System Analytics result from snapshots on the monitored cluster, select the **Take snapshot** check box.

- File System Analytics result sets generated directly from the cluster, clear the **Take snapshot** check box.

5. Configure the **Result Set Retention** settings:

- 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. Any pinned result sets are not deleted, even if they are older than the specified setting.

If you want 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 an integer that represents, in days, how long to retain each result set before automatically deleting it.

If you do not want to limit the number of result sets that are retained, type `0`. You can manually delete a result set at any time.

6. Click **Submit**.

## Enable or disable File System Analytics

Before you can view and analyze data-usage and data-properties information through InsightIQ, you must enable the File System Analytics feature.

### Procedure

1. Click **Settings** > **Monitored Clusters**.

The **Monitored Clusters** page appears.

2. In the **Actions** column for the cluster for which you want to enable or disable File System Analytics, click **Configure**.

The **Configuration** page appears.

3. Click the **Enable FSA** tab.

The **Enable FSA** tab appears.

4. Enable or disable File System Analytics.

Options	Description
<b>Enable File System Analytics.</b>	Select the <b>Enable</b> check box.
<b>Disable File System Analytics.</b>	Clear the <b>Enable</b> check box.

5. Click **Submit**.

## View File System Analytics settings

You can view the configured settings for the InsightIQ File System Analytics feature.

### Procedure

1. Click **Settings** > **Monitored Clusters**.

The **Monitored Clusters** page appears and displays a list of all clusters that InsightIQ is configured to monitor.

2. In the **Actions** column for the cluster whose File System Analytics settings you want view, click **Configure**.

The **Configuration** page appears.

3. Click the **Enable FSA** tab.

The **Enable FSA** tab appears.

4. Review the **Job Control** setting, which indicates whether the File System Analytics feature is enabled or disabled.

5. Click the **FSA Configuration** tab.

The **FSA Configuration** tab appears, and displays the current File System Analytics settings.

## Managing InsightIQ virtual appliance settings

If you are running InsightIQ as a virtual appliance, you can modify network and time zone settings through the VMware Studio interface. You cannot view or modify these settings through the InsightIQ web application.

If you installed InsightIQ through Red Hat Package Manager (RPM), you can manage equivalent settings through the Linux interface. For more information, see the CentOS or Red Hat Enterprise Linux documentation.

### Modify the InsightIQ virtual appliance time zone

You can modify the time zone of the InsightIQ virtual appliance. You might want to do this if report schedules are not generating reports at the scheduled times.

#### Procedure

1. In a supported web browser, connect to the VMware Studio interface for InsightIQ at `https://<IPAddressOrHostName>:5480`, where `<IPAddressOrHostName>` is the IP address or the host name of the InsightIQ virtual appliance.
2. Click **System** > **Time Zone**.
3. Select a time zone from the **System Time Zone** list.
4. Click **Save Settings**.

### Modify the InsightIQ virtual appliance administrator password

You can modify the InsightIQ administrator password through the InsightIQ virtual machine (VM) console.

The InsightIQ administrator user name is administrator. You cannot modify the InsightIQ administrator user name.

#### Procedure

1. Open the virtualization-environment console for the InsightIQ virtual appliance and login as administrator.
2. Run the following command:
 

```
passwd
```
3. Type the currently configured administrator password, and then press ENTER.
 

The `Enter new UNIX password` prompt appears.
4. Type a new password for the administrator and then press ENTER.

#### Note

The password must contain at least one character.

The `Retype new UNIX password` prompt appears.

5. Retype the new password for the administrator and then press ENTER.

## Modify InsightIQ virtual appliance network settings

You can modify the network settings for the InsightIQ virtual appliance.

### Procedure

1. In a supported web browser, connect to the VMware Studio interface for InsightIQ at `https://<IPAddressOrHostName>:5480`, where `<IPAddressOrHostName>` is the IP address or the host name of the InsightIQ virtual appliance.
2. Click **Network > Address**.
3. Specify which networking mode you want to run InsightIQ in.

Options	Description
To configure DHCP networking mode	From the <b>IPv4 Address Type</b> list, select <b>DHCP</b> .
To configure a static IP address	<ul style="list-style-type: none"> <li>• From the <b>IPv4 Address Type</b> list, select <b>Static</b>.</li> <li>• In the <b>Network Address Settings</b> area, specify static network address settings.</li> </ul>

4. Click **Save Settings**.

## View InsightIQ virtual appliance network settings

You can view the network settings for the InsightIQ virtual appliance.

### Procedure

1. In a supported web browser, connect to the VMware Studio interface for InsightIQ at `https://<IPAddressOrHostName>:5480`, where `<IPAddressOrHostName>` is the IP address or the host name of the InsightIQ virtual appliance.
2. Click **Network > Address**.
3. In the **Network Status** area, view network settings.

## Configure the InsightIQ virtual appliance to monitor clusters through a proxy server

You can allow the InsightIQ virtual appliance to monitor clusters through a proxy server. This can be useful if you want to monitor a cluster that is on a network other than the network that the InsightIQ virtual appliance is connected to.

### Procedure

1. In a supported web browser, connect to the VMware Studio interface for InsightIQ at `https://<IPAddressOrHostName>:5480`, where `<IPAddressOrHostName>` is the IP address or the host name of the InsightIQ virtual appliance.
2. Click **Network > Proxy**
3. Select the **Use a proxy server** check box.
4. In the **HTTP Proxy Server** field, enter the IP address of a proxy server.
5. In the **Proxy Port** field, enter the number of the port you want clients to access the proxy server through.
6. Optional: If the proxy server requires a username and password, in the **Proxy Username (Optional)** and **Proxy Password (Optional)** fields, enter a valid username and password.

## Monitoring the status of InsightIQ

You can view the current status of InsightIQ, including which clusters InsightIQ is currently able to monitor. You can configure InsightIQ to send alert emails about the status of InsightIQ.

### View InsightIQ status

You can view the current status of InsightIQ, including which clusters InsightIQ is currently able to monitor.

#### Procedure

1. Click **Settings** > **Status**.
2. On the **InsightIQ Status** page, view the status of the InsightIQ application.

### Enable and configure InsightIQ status alert emails

You can configure InsightIQ to send alert emails about the status of InsightIQ.

#### Procedure

1. Click **Settings** > **Status**.
2. Select **Enable InsightIQ status email alerts**.
3. In the **Send InsightIQ status email alerts to** field, type the name of an email address you want to send alerts to.
4. Click **Submit**.

### Disable InsightIQ status alert emails

You can disable InsightIQ status alert emails.

#### Procedure

1. Click **Settings** > **Status**.
2. Clear **Enable InsightIQ status email alerts**.

### InsightIQ cluster monitoring errors

InsightIQ can encounter the following errors:

#### Data Retrieval Delayed

The most recent data that InsightIQ has retrieved from the cluster is more than 15 minutes old. If InsightIQ recently began or resumed monitoring the cluster, this issue will most likely resolve itself as InsightIQ retrieves more recent data. If this issue is accompanied by one or more InsightIQ errors, resolve the errors. If the error persists after all other InsightIQ errors have been resolved, contact Isilon Technical Support.

#### License Error

InsightIQ is not licensed on the monitored cluster. To resolve this issue, configure a valid InsightIQ license on the cluster.

#### Authorization Error

InsightIQ is not authorized to communicate with the monitored cluster. To resolve this issue, ensure that InsightIQ is authenticating with the correct username and password.

**Connection Error**

InsightIQ cannot connect to the monitored cluster. To resolve this issue, ensure network connectivity between InsightIQ and the monitored cluster. InsightIQ will attempt to reconnect with the cluster until InsightIQ can resume monitoring.

**Cluster Add Error**

InsightIQ is unable to obtain sufficient information from the cluster to start monitoring. To resolve this issue, stop monitoring the cluster and then add the cluster again. This will destroy all InsightIQ data about the cluster. However, this will not affect any data outside of the InsightIQ datastore.

**GUID Error**

The globally unique identifier (GUID) associated with the monitored cluster has changed. The hostname or IP address associated with the cluster might have been reassigned to another cluster. To resolve this issue, ensure that the hostname or IP address of the monitored cluster matches the hostname or IP address that InsightIQ is accessing.

**Datastore Full**

InsightIQ datastore does not have sufficient space to continue monitoring. To resolve this issue, free space on the InsightIQ datastore and then restart InsightIQ.

## Managing read-only users

You can create and delete read-only user accounts as well as modify the password of existing accounts.

### Add a read-only user account

During the read-only user account setup process, you must specify a user name and password for the user; InsightIQ stores these credentials. The read-only users must submit these credentials in order to log in to the InsightIQ application.

**Before you begin**

- ◆ You must be logged in as the administrator user.

**Procedure**

1. Click **Settings > Users**.

The **Read-Only Users** page appears, and displays a list of all configured read-only users.

2. Click **Add User**.

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

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

Make a note of the password you configured. InsightIQ does not enable the administrator user or read-only users to retrieve a lost password.

5. Click **Submit**.

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

## Modify the password of a read-only user account

The administrator user can modify the password of a read-only user account through the InsightIQ web application. Read-only users cannot modify passwords for their accounts.

### Procedure

1. Click **Settings** > **Users**.

The **Read-Only Users** page appears, and displays a list of all configured read-only users.

2. In the **Actions** column for the user whose password you want to modify, click **Change Password**.
3. In both the **New password** and **Confirm password** field, type a new password for the user.

Make a note of the password you configured. InsightIQ does not enable the administrator user or read-only users to retrieve a lost password.

4. Click **Submit**.

## Delete a read-only user account

You can delete a read-only user account. This can be useful if, for example, you want to prevent certain previously authorized users from logging in to the InsightIQ application.

### Procedure

1. Click **Settings** > **Users**.

The **Read-Only Users** page appears, and displays a list of all configured read-only users.

2. In the **Actions** column for the read-only user account that you want to delete, click **Delete**.

A confirmation dialog box appears.

3. Click **Yes**.

The user account is removed, and users can no longer log in to the InsightIQ application with the associated credentials.

## Disable gathering event data

If you do not regularly view information about cluster events, you can configure InsightIQ to stop gathering data for the Events data module. Gathering event data consumes InsightIQ resources, and disabling this operations can improve InsightIQ performance.

### Procedure

1. Open an SSH connection to the InsightIQ appliance and log in.
2. Open the `/etc/isilon/isightiq.ini` file in a text editor and set the value of `enable_event_monitoring` to `false`.
3. Restart InsightIQ by running the following command:

```
iiq_restart
```



## Configure LDAP authentication

You can configure InsightIQ to accept authentication credentials from an LDAP server. If a user logs in through LDAP credentials, the user will have read-only privileges. The user can access InsightIQ through a web browser or an SSH connection.

### Procedure

1. Open an SSH connection to the InsightIQ appliance and log in.

2. Enable SSSD by running the following commands:

```
chkconfig sssd --levels 3456 on
chkconfig sssd --levels 012 off
```

3. Create a new file called `/etc/sss/sss.conf`.

4. Copy the contents of `/etc/sss/sss.conf.template` to `/etc/sss/sss.conf`

5. Open the `/etc/sss/sss.conf` file in a text editor and specify values for the following settings:

Setting	Value
<code>ldap_uri</code>	The URIs of the LDAP server that you want to authenticate through. Specify URIs in the order of preference.
<code>ldap_search_base</code>	The default base distinguished name (DN) to search on the LDAP server.
<code>ldap_default_bind_dn</code>	The default bind DN user that has permission to search the LDAP directory.
<code>ldap_default_authtok</code>	The password of the default bind DN user.
<code>ldap_schema</code>	The schema type of the LDAP server.

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

6. Save the `/etc/sss/sss.conf` file.

7. Ensure that the `/etc/sss` directory has the proper permission settings by running the following command:

```
sudo chmod 600 -R /etc/sss
```

8. Commit your changes by running the following command:

```
sudo authconfig --enablesssd --enablesssdauth --enablelocauthorize --update
```

9. Restart the SSSD service by running the following command:

```
sudo service sssd restart
```

10. Restart InsightIQ by running the following command:

```
iiq_restart
```

## Specify a custom SSL certificate

Although InsightIQ includes a default SSL certificate, you can specify a custom SSL certificate to ensure that clients connect to the correct InsightIQ web address.

### Procedure

1. Open an SSH connection to the InsightIQ appliance and log in.
2. On the InsightIQ appliance, save a copy of the SSL certificate you want to specify.

The certificate must be of the file type `.pem`

3. Open the `/etc/isilon/insightiq.ini` file in a text editor and specify the value of the `ssl_pem` setting as the path of the SSL certificate.

For example, the following text specifies an SSL certificate located at `/home/administrator/certificates/`

```
ssl_pem = /home/administrator/certificates/ssl.pem
```

4. Restart InsightIQ by running the following command:

```
iiq_restart
```

## Specify the InsightIQ port

You can specify the port through which you connect to InsightIQ. InsightIQ communicates through the port with HTTPS. The default port is 443.

Before you import data from another instance of InsightIQ, you must ensure that the other InsightIQ instance is configured to use the same port number as the instance you are importing to.

### Procedure

1. Open an SSH connection to the InsightIQ appliance and log in.
2. Open the `/etc/isilon/insightiq.ini` file in a text editor and specify values for the following settings (the value must be the same for both settings):

Setting	Value
<code>port</code>	The port through which you want to connect to InsightIQ.
<code>redirect_to_port</code>	

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

3. Save and close the `/etc/isilon/insightiq.ini` file.
4. Restart InsightIQ by running the following command:

```
iiq_restart
```

## Configure email settings

If you want to send PDF reports generated from a report schedule by email, you must enable and configure InsightIQ to send outbound email through a specified email server.

### Procedure

1. Click **Settings > Email**.

The **Configure Email Settings (SMTP)** page appears.

2. In the **SMTP server** field, type the host name or IP address of an SMTP server that handles email for your organization.
3. In the **SMTP port** field, type the number of the port 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 will appear in the **From** field of email messages sent by InsightIQ.

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** check box.
7. Click **Submit**.

## Test email configuration

You can test your InsightIQ email configuration by sending yourself a test email. If the email is delivered successfully, InsightIQ has been configured to interact with your email server correctly.

### Procedure

1. Click **Settings > Email**.
2. In the **Send a test email** field, type the name of an email address.
3. Click **Send**.
4. Check your email inbox.

If you have configured InsightIQ correctly, you will receive a test email.

## Create a permalink

You can create a permalink in InsightIQ that can later be accessed by you or another InsightIQ user.

### Procedure

1. View and configure the report that you want to create a permalink for, applying breakouts and filters as needed.
2. In the upper-right corner of the InsightIQ application, click **Permalink**.

InsightIQ generates and displays a URL for the currently displayed report configuration for the specified time period.

3. Copy the URL and then paste it in a location that is convenient for you, such as a text file, and then click **OK**.

**After you finish**

To view the permalink at a later time, you must log in to the InsightIQ application with the configured user credentials, and then go to the generated URL.

# CHAPTER 3

## Monitoring cluster status summaries

You can view cluster data that InsightIQ collects through reports.

You can create, view, and customize performance reports to display specific information about cluster performance. If the File System Analytics feature is enabled, you can also view reports that display information about data usage and data properties. For all reports, you can apply specialized breakouts and filters to view even more detailed data; for example, you can break out node activity by a specific protocol or client.

---

### Note

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

---

- ◆ [Monitor aggregated-cluster status](#)..... 46
- ◆ [Monitor the status of an individual cluster](#)..... 46

## Monitor aggregated-cluster status

You can the status of all monitored clusters as a whole.

### Procedure

1. Click **Dashboard**.

The **InsightIQ Dashboard** page appears.

2. In the **Aggregated Cluster Overview** section, view the status of all monitored clusters as a whole.

## Monitor the status of an individual cluster

You can view the status of an individual cluster.

### Procedure

1. Click **Dashboard**.

The **InsightIQ Dashboard** page appears.

2. In the **Cluster-by-cluster Overview** section, view the status of a cluster.

# CHAPTER 4

## Monitoring cluster status through reports

You can view cluster data that InsightIQ collects through the cluster-monitoring summary page. The cluster-monitoring summary page enables you to review, on a single page, the overall status of all clusters that InsightIQ is monitoring.

- ◆ [Managing live performance reports](#)..... 48
- ◆ [Managing performance report schedules](#).....51
- ◆ [Generating performance reports](#)..... 55
- ◆ [View a live file-system report](#)..... 58
- ◆ [Managing rules and filters](#)..... 59
- ◆ [Exporting InsightIQ data](#)..... 61

## Managing live performance reports

You can create, view, schedule, modify and delete live performance reports.

### Create a live performance report

You can create custom live performance reports.

#### Procedure

1. Click **Performance Reporting > Create a New Performance Report**.

The **Create a New Performance Report** report template selection page appears.

2. Specify a template to create the live performance report from.

Options	Description
<b>Create a live performance report from a template based on the default settings</b>	Click <b>Create from Blank Report</b> .
<b>Create a live performance report from a template based on a user-created performance report</b>	Click the name of a report in the <b>Saved Reports</b> area.
<b>Create a live performance report from a template based on one of the standard reports included with InsightIQ</b>	Click the name of a report in the <b>Standard Reports</b> area.

The **Create a New Performance Report** configuration page appears.

3. In the **Build Your New Performance Report** area, in the **Performance Report Name** field, type a name for the live performance report.
4. Select the **Live Performance Reporting** check box.
  - If you also want to create a performance report schedule with the same properties as this live performance report, select both the **Live Performance Reporting** and **Scheduled Performance Report** check boxes.
5. In the **Insert the Data You Want to See** area, specify the performance modules that you want to view in the report.

Options	Description
<b>Add a new performance module</b>	<ol style="list-style-type: none"> <li>a. If all existing performance module areas have been assigned, click <b>Add another performance module</b>.</li> </ol> <hr/> <p><b>Note</b></p> <p>Unassigned performance module areas contain a <b>Select a Module for this Position</b> list.</p> <hr/> <ol style="list-style-type: none"> <li>b. In an unassigned performance module area, from the <b>Select a Module for this Position</b> list, select a performance module.</li> </ol>



Options	Description
<b>Modify an existing performance module</b>	<ul style="list-style-type: none"> <li>In an assigned performance module area, from the performance module list, select a different performance module.</li> </ul>

Repeat this step for each performance module that you want to include.

- Save the report.
  - If you want to save the report but continue editing it, click **Save**.
  - If you want to save the report and then immediately apply it to a cluster and view it, click **Finish**.

## View a live performance report

You can view and configure a live performance report.

### Procedure

- Click **Performance Reporting > Live Performance Reporting**.

The **Live Performance Reporting** page appears.

- In the **Select Your Performance Report** area, from the **Select a report** list, select the cluster that you want to view a report for.
- In the **Select Your Performance Report** area, from the **Apply to Cluster** list, select the live performance report that you want to view.
- In the **Date Range** section, in the upper time and day fields, specify the first point in time that you want to view data about.
- In the lower time and day fields, specify the last point in time that you want to view data about.
- Click **View Report**.

The report appears below the **Select Your Performance Report** section.

- If you want to apply a breakout, in a performance module area, in the **Breakout by area**, click the name of a breakout.
 

The breakouts appear below the chart in the performance module section.
- If you want to apply a filter, in a performance module area, click the green button for the filter you want to apply.

For example, if you want to filter by write events, click **write**.

If you want to remove a filter rule from a report, click the red button for the filter you want to remove in either the **Data Filters** section or a performance module section where the filter is applied.

For example, if you want to stop filtering by write events, click **FS Event: write**.

- View the live performance report.

---

### Note

The information displayed in the **Cluster Status** area reflects the status of the cluster at the time displayed in the **Current Status** area. This information does not reflect the cluster status during the time period that the report is displaying data about.

---

## Schedule a live performance report

You can configure InsightIQ to periodically generate and distribute static PDF versions of a live performance report.

### Procedure

1. Click **Performance Reporting > Manage Performance Reporting**.

The **Manage Performance Reporting** page appears with the **Live Performance Reporting** tab displayed. A list of all saved live performance reports is displayed in the **Live Reports** table.

2. In the row for the live performance report that you want to schedule, click **Schedule this report**.

The **Schedule this Performance Report** configuration page appears.

3. Configure schedule and delivery settings for the performance report.
4. Click **Save**.

## Modify a live performance report

You can add and remove performance modules from a live performance report.

### Procedure

1. Click **Performance Reporting > Manage Performance Reporting**.

The **Manage Performance Reporting** page appears with the **Live Performance Reporting** tab displayed. A list of all saved live performance reports is displayed in the **Live Reports** area.

2. In the row for the live performance report you want to modify, click **Edit**.

The **Edit this Performance Report** configuration page appears.

3. Modify the report settings as necessary, and then click save.

## Delete a live performance report

If you delete a live performance report, that report can no longer be viewed.

### Procedure

1. Click **Performance Reporting > Manage Performance Reporting**.

The **Manage Performance Reporting** page appears with the **Live Performance Reporting** tab displayed. A list of all saved live performance reports is displayed in the **Live Reports** area.

2. Specify the report or reports that you want to delete.

Options	Description
<b>Delete one report.</b>	In the row for the report that you want to delete, click <b>Delete</b> .
<b>Delete multiple reports.</b>	In the rows for reports that you want to delete, select the check boxes, and then, in the <b>Select an action</b> list, click <b>Delete selected reports</b> .

Options	Description
<b>Delete all live performance reports</b>	Select the highest check box in the <b>Live Reports</b> area, and then, in the <b>Select an action</b> list, click <b>Delete selected reports</b> .

A dialog box appears and prompts you to confirm that you want to delete the performance report or reports.

3. Click **Delete**.

## Managing performance report schedules

You can create, disable, enable, modify, and delete performance report schedules.

### Create a performance report schedule

You can create a report schedule that specifies when InsightIQ generates static PDF reports. You can configure a report schedule to specify one or more email addresses that InsightIQ will send the generated reports to.

#### Procedure

1. Click **Performance Reporting > Create a New Performance Report**.

The **Create a New Performance Report** page appears.

2. Specify a template to create the report schedule from.

Options	Description
<b>Create a report schedule from a template based on the default settings</b>	Click <b>Create from Blank Report</b> .
<b>Create a report schedule from a template based on a user-created performance report</b>	In the <b>Saved Reports</b> area, click the name of a report.
<b>Create a report schedule from a template based on one of the standard reports included with InsightIQ</b>	In the <b>Standard Reports</b> area, click the name of a report.

The **Create a New Performance Report** page reloads and displays report configuration options.

3. In the **Build Your New Performance Report** area, in the **Performance Report Name** field, type a name for the report schedule.
4. Select the **Scheduled Performance Report** check box.

If you want to also create a live performance report that has the same properties as this static performance report schedule, select both the **Scheduled Performance Report** and **Live Performance Reporting** check boxes.

Additional scheduling options appear.

5. Under **Configure Your Scheduled Report**, from the **Subject of Report** list, select the cluster that you want to apply this report to.
6. From the **Reporting Period** list, select the length of time that you want each report to cover.

For example, if you select **24 Hours**, and a report is generated at 10:00 AM on 08/02/2011, the report would include data collected from 10:00 AM on 08/01/2011 to 10:00 AM on 08/02/2011.

7. Under **Schedule and Delivery**, from the **Report Frequency** list, specify how often you want this report to be generated.

Options	Description
<b>Generate one or more reports every day</b>	Select <b>Hourly</b> , specify the number of hours that will elapse before the next report is generated, and specify the time when each report will be generated for the first time each day. For example, if you configure InsightIQ to generate a report every seven hours starting at 1:30 PM, the report will be generated daily at 1:30 PM and 8:30 PM.
<b>Generate no more than one report per day. Optionally, suspend report generation for a number of days.</b>	Select <b>Daily</b> , specify the number of days that will elapse before the next report is generated, and specify the time of day when the reports will be generated.
<b>Generate no more than one report per day. Optionally, suspend report generation for a number of weeks.</b>	Select <b>Weekly</b> , specify the number of weeks that will elapse before the next report is generated, and specify one or more days of the week and the time of day when the reports will be generated.
<b>Generate no more than one report per month</b>	Select <b>Monthly</b> , specify the number of months that will elapse before the next report is generated, and specify the day of the month on which the reports will be generated. Reports are always generated at 11:59 PM on the specified day.

8. If you want to send reports generated from this schedule to one or more email addresses, specify the addresses.
  - a. Select the **Email this report as a PDF attachment each time it is generated** check box.
  - b. Type one or more email addresses in the **Report Recipients** box.  
 Separate each address with a comma, a space, or a semi-colon. You can specify up to 10 email addresses.
9. In the **Insert the Data You Want to See** area, specify the performance modules that you want to view in the report.

Options	Description
<b>Add a new performance module</b>	a. If all existing performance module areas have been assigned, click <b>Add another performance module</b> . <hr/> <b>Note</b> Unassigned performance module areas contain a <b>Select a Module for this Position</b> list.

Options	Description
	b. In an unassigned performance module area, from the <b>Select a Module for this Position</b> list, select a performance module.
<b>Modify an existing performance module</b>	<ul style="list-style-type: none"> <li>In an assigned performance module area, from the performance module list, select a different performance module.</li> </ul>

Repeat this step for each performance module that you want to include.

10. If you want to apply a breakout to a performance module, in a performance module area, in the **Display Options** area, click the name of the breakout that you want to apply.

Repeat this step for each performance module that you want to apply a breakout to.

11. Save the report schedule.

- If you want to save the report but want to continue modifying it, click **Save**.
- If you want to save the report and complete the configuration process, click **Finish**.

## Disable a performance report schedule

You can disable a performance report schedule if you want to temporarily prevent InsightIQ from generating reports from, but do not want to permanently delete the schedule. A disabled report can be enabled, after which InsightIQ will begin generating reports from it again.

### Procedure

1. Click **Performance Reporting > Manage Performance Reporting**.

The **Manage Performance Reporting** page appears with the **Live Performance Reporting** tab displayed.

2. Click the **Scheduled Performance Reports** tab.

The **Scheduled Performance Reports** tab appears. A list of all saved report schedules is displayed in the **Scheduled Reports** table.

3. In the **Scheduled Reports** table, specify which report schedule you want to disable.

Options	Description
<b>Disable a single report schedule.</b>	In the row of the report schedule you want to disable, click <b>Pause</b> .
<b>Disable multiple report schedules.</b>	In the rows of the report schedules you want to disable, select the check boxes, and then, in the <b>Select an action</b> list, select <b>Pause selected reports</b> .
<b>Disable all report schedules.</b>	In the header row, select the check box next to <b>Report</b> , and then, in the <b>Select an action</b> list, select <b>Pause selected reports</b> .

In the **Next Run** columns of the report schedules, the **Paused** icon appears; InsightIQ will not generate reports from the specified schedules.

## Enable a performance report schedule

If a performance report schedule has been disabled, you can enable it. After a performance report has been enabled, InsightIQ will resume generating reports from it.

### Procedure

1. Click **Performance Reporting > Manage Performance Reporting**.

The **Manage Performance Reporting** page appears with the **Live Performance Reporting** tab displayed.

2. Click the **Scheduled Performance Reports** tab.

The **Scheduled Performance Reports** tab appears. A list of all saved report schedules is displayed in the **Scheduled Reports** table.

3. In the **Scheduled Reports** table, specify which report schedule you want to enable.

Options	Description
<b>Enable a single report schedule.</b>	In the row of the report schedule you want to enable, click <b>Resume</b> .
<b>Disable multiple report schedules.</b>	In the rows of the report schedules you want to enable, select the check boxes; then, in the <b>Select an action</b> list, select <b>Resume selected reports</b> .
<b>Disable all report schedules.</b>	In the header row, select the check box next to <b>Report</b> ; then, in the <b>Select an action</b> list, select <b>Resume selected reports</b> .

The next date that InsightIQ will generate reports for the specified report schedules appears in the **Next Run** column of the schedules.

## Modify a performance report schedule

You can add and remove data elements from a performance report schedule. You can also modify the schedule on which the static PDF reports are generated, and the list of recipient email addresses.

### Procedure

1. Click **Performance Reporting > Manage Performance Reporting**.

The **Manage Performance Reporting** page appears with the **Live Performance Reporting** tab displayed.

2. Click the **Scheduled Performance Reports** tab.

The **Scheduled Performance Reports** tab appears, and displays a list of all saved report schedules.

3. In the **Scheduled Reports** area, in the row for the report schedule you that want to modify, click **Edit**.

The **Edit this Performance Report** page appears.

4. Modify the schedule settings as necessary, and then click save.

## Delete a performance report schedule

If you no longer want InsightIQ to generate PDF reports from a specific report schedule, you can permanently delete the report schedule. Deleted report schedules cannot be

recovered; however, any reports that InsightIQ already generated based on the schedule will not be deleted as a result of deleting the schedule.

#### Procedure

1. Click **Performance Reporting > Manage Performance Reporting**.

The **Manage Performance Reporting** page appears with the **Live Performance Reporting** tab displayed.

2. Click the **Scheduled Performance Reports** tab.

The **Scheduled Performance Reports** tab appears, and displays a list of all saved report schedules.

3. In the **Scheduled Reports** area, specify which report schedule or schedules you want to delete.

Options	Description
<b>Delete one report schedule</b>	In the row for the report schedule that you want to delete, click <b>Delete</b> .
<b>Delete multiple report schedules</b>	In the rows for the report schedules that you want to delete, select the check boxes, and then, in the <b>Select an action</b> list, select <b>Delete selected reports</b> .
<b>Delete all report schedules</b>	Select the highest check box in the <b>Scheduled Reports</b> area, and then, in the <b>Select an action</b> list, click <b>Delete selected reports</b> .

A dialog box appears and prompts you to confirm that you want to delete the report schedule or schedules.

4. Click **Delete**.

## Generating performance reports

You can generate scheduled performance reports, after which you can send, delete and view the generated performance report. You can also modify the schedules of the generated reports.

### Generate a performance report on demand

You can manually generate a report from a report schedule at any time.

#### Procedure

1. Click **Performance Reporting > Manage Performance Reporting**.

The **Manage Performance Reporting** page appears with the **Live Performance Reporting** tab displayed.

2. Click the **Scheduled Performance Reports** tab.

The **Scheduled Performance Reports** tab appears, and displays a list of all saved report schedules.

3. In the **Scheduled Reports** area, in the row for the schedule that you want to generate a report from, click **Generate now**.

The **Generated Reports Archive** tab appears, and displays the newly generated report in the **Generated Reports** area.

## Send a generated performance report

You can manually send a generated PDF report to up to 10 email addresses at any time.

### Procedure

1. Click **Performance Reporting > Manage Performance Reporting**.

The **Manage Performance Reporting** page appears with the **Live Performance Reporting** tab displayed.

2. Click the **Generated Reports Archive** tab.

The **Generated Reports Archive** tab appears, and displays a list of all generated reports.

3. In the **Generated Reports** area, specify which generated reports you want to send.

Options	Description
<b>Send one generated report</b>	In the row for the generated report that you want to send, click <b>Send</b> .
<b>Send multiple reports in a single email</b>	In the rows for the reports you that want to send, select the check boxes, and then, in the <b>Select an action</b> list, select <b>Send Selected reports</b> .
<b>Send all reports in a single email</b>	Select the highest check box in the <b>Generate Reports</b> area, and then, in the <b>Select an action</b> list, click <b>Delete selected reports</b> .

The **Send Report** dialog box appears.

4. In the **Report Recipients** field, type up to 10 email addresses that InsightIQ will send this report to. Separate each address with a comma, a space, or a semi-colon.
  - If you want the specified email addresses to receive the latest version of this report each time it is generated in the future, select the **Add each of these recipients...** check box.
5. If you want to specify a custom subject line for the email that will contain the PDF report, type the subject line in the **Message Subject** field.
6. If you would like to include a message with the email being sent, type the message in the **Message** field.
7. Click **Send**.

InsightIQ immediately sends the report, and the **Message Sent** dialog box appears.

8. Click **Close**.

## Delete a generated performance report

You can permanently delete a generated PDF report.

### Procedure

1. Click **Performance Reporting > Manage Performance Reporting**.

The **Manage Performance Reporting** page appears with the **Live Performance Reporting** tab displayed.



2. Click the **Generated Reports Archive** tab.

The **Generated Reports Archive** tab appears, and displays a list of all generated reports.

3. In the **Generated Reports** area, specify which generated reports you want to delete.

Options	Description
<b>Delete one generated report</b>	In the row of the generated report you want to delete, click <b>Delete</b> .
<b>Delete multiple generated reports</b>	In the rows of the reports you want to delete, select the check boxes, and then, in the <b>Select an action</b> list, select <b>Delete Selected reports</b> .
<b>Send all reports in a single email.</b>	Select the highest check box in the <b>Generated Reports</b> area, and then, in the <b>Select an action</b> list, click <b>Delete selected reports</b> .

A confirmation dialog box appears and prompts you to confirm that you want to delete the report.

4. Click **OK**.

## Modify the schedule of a generated performance report

You can modify the report schedule of a static PDF report. You can add and remove data elements from a performance report schedule. You can also modify the schedule on which the static PDF reports are generated, and the list of recipient email addresses.

### Procedure

1. Click **Performance Reporting > Manage Performance Reporting**.

The **Manage Performance Reporting** page appears with the **Live Performance Reporting** tab displayed.

2. Click the **Generated Reports Archive** tab.

The **Generated Reports Archive** tab appears, and displays a list of all generated reports.

3. In the **Generated Reports** area, in the row for the generated report whose schedule you want to modify, click **Edit scheduled report**.

The **Edit this Performance Report** page appears.

4. Modify the schedule settings as needed, and then click **save**.

## View a generated performance report

You can view a generated PDF report at any time.

### Procedure

1. Click **Performance Reporting > Manage Performance Reporting**.

The **Manage Performance Reporting** page appears with the **Live Performance Reporting** tab displayed.

2. Click the **Generated Reports Archive** tab.

The **Generated Reports Archive** tab appears, and displays a list of all generated reports.

3. In the **Generated Reports** area, in the row for the generated report that you want to view, click **View**.

The report PDF appears in appears.

## View a live file-system report

You can view and configure a live file-system report.

### Before you begin

The File System Analytics feature must be enabled.

### Procedure

1. Click **File System Reporting**.

The **File System Reporting** page appears.

2. From the **Apply to cluster** list, select the name of the cluster that you want to view the report for.

3. From the **Apply this report** list, select the name of the report that you want to view.

If you selected **Data Properties Report**, the **Compare to** field appears.

4. In the **Select a reporting date** list, select the time and day that you want to view the report for.

The time and day references a File System Analytics result set.

5. If you are viewing a data properties report, and want to compare the data for the selected reporting date with data for another reporting date, from the **Compare to** list, select the other reporting date.

6. Click **View Report**.

The file-system modules associated with the specified report appear.

7. If you want to apply a breakout, in a file-system module area, in the **Breakout by** area, click the name of a breakout.

The breakouts appear below the chart in the file-system module section.

8. If you want to apply a filter, in a file-system module area, click the green button of the filter you want to apply.

For example, if you want to filter by write events, click **write**.

If you want to remove a filter rule from a report, click the red button of the filter you want to remove in either the **Data Filters** section or a file-system module section where the filter is applied.

For example, if you want to stop filtering by write events, click **Event:write**.

9. View the live file-system report.

## Managing rules and filters

You can create, modify, and delete filters applied to performance reports.

### Create a filter

You can create and save a custom filter that consists of one or more filter rules. After you specify the filter rules to include, you can save the filter and then apply it to reports as necessary.

#### Procedure

1. On the **Live Performance Reporting** page, click **Create/manage data filters**.

The **Data Filters** area appears.

2. Click **Add Rule**.
3. From the **Type** list, select a breakout that you want to filter by.
4. In the **Match** field, select or type a filter value.
5. Repeat steps 2 through 4 as needed until you have built and applied all of the filter rules that you want to include in the filter.
6. Click **Apply**.

The filters are applied to all data modules in the report.

7. From the **Manage** menu, click **Save**.

The **Save Filter As** dialog box appears.

8. In the **Please enter a filter name** field, type a name for the filter, and then click OK.

The filter is saved, and the name of the filter appears in the **Current filter** field, indicating that the filter is applied globally across all data modules in the report.

### Modify a filter

You can add filter rules to a filter; you can also modify and remove filter rules contained in a filter. You must apply a filter before you can modify it.

#### Procedure

1. On the **Live Performance Reporting** page, click **Create/manage data filters**.

The **Data Filters** area appears.

2. On the **Manage** menu, point to **Load filter**, and then click the name of the filter that you want to modify.

The individual filter rules appear in the **Filter** section, and InsightIQ applies the filter globally across all data modules in the report.

3. Modify filter settings as needed.

Options	Description
<b>Add a filter rule.</b>	<ol style="list-style-type: none"> <li>a. Click <b>Add Rule</b>.</li> <li>b. From the <b>Type</b> list, select the breakout that you want to filter by.</li> </ol>

Options	Description
	c. In the <b>Match</b> field for the filter rule, select or type a filter value.
<b>Modify a filter rule.</b>	a. In the <b>Match</b> field for the filter rule that you want to modify, select or type a filter value.
<b>Remove a filter rule.</b>	a. Click either the <b>Type</b> list or <b>Match</b> field for the filter rule you want to remove. b. Click <b>Delete Rule</b> .

## 4. Save the filter.

Options	Description
<b>Save the filter under its existing name.</b>	a. Click <b>Apply</b> . b. On the <b>Manage</b> menu, click <b>Save</b> . The <b>Save Filter</b> dialog box appears. c. Click <b>Yes</b> .
<b>Save a new instance of the filter with a different name.</b>	a. On the <b>Manage</b> menu, click <b>Save as</b> . The <b>Save Filter</b> dialog box appears. b. In the <b>Please enter a filter name</b> field, type a name for the filter, and then click <b>OK</b> .

## Delete a filter

You can permanently delete a filter. You must apply a filter to a report before the filter can be deleted.

### Procedure

1. On the **Live Performance Reporting** page, click **Create/manage data filters**.

The **Data Filters** area appears.

2. On the **Manage** menu, point to **Load Filter**, and then click the name of the filter that you want to delete.

The individual filter rules appear in the **Filter** section, and InsightIQ applies the filter globally across all data modules in the report.

3. On the **Manage** menu, click **Delete**.

---

### Note

If you want to remove an individual filter rule from a filter while keeping the rest of the filter intact, with the filter loaded, click the rule that you want to delete, click the **Delete Rule** button at the top of the rule list, and then save the filter.

The **Delete Filter** dialog box appears and prompts you to confirm that you want to delete the filter.

4. Click **Yes**.

## Filter values

To create a filter rule, you must specify a valid match value. For some filters, InsightIQ automatically generates a list of values. For other filters, you must type a valid value.

You must type values for the following filters:

Filter	Match value
Client	The IP address of a client. To filter by all clients exceeding the top 1000 for a given data module, specify <b>Other clients</b>
Disk Pool	The name of a disk pool on the cluster.
Directory	The path of a directory within <code>/ifs</code> . For example, <code>/ifs/data/media</code> The backslash character ( <code>\</code> ) is not valid.
File Extension	The name of a file extension. For example <code>.xml</code> To filter files without file extensions, specify the following: (none) The backslash character ( <code>\</code> ) is not valid.
Job ID	The numeric ID of a job. For example, 5
Path	The path of a directory within <code>/ifs</code> . For example, <code>/ifs/data/media</code> The backslash character ( <code>\</code> ) is not valid.
User Attribute	The name of a user attribute defined on the cluster. All characters are valid.

You can select values for the following filters from an automatically generated list:

- ◆ Accessed Time
- ◆ Disk
- ◆ Direction
- ◆ FS Event
- ◆ Interface
- ◆ Job Name
- ◆ Logical Size
- ◆ Modified Time
- ◆ Node
- ◆ Node (Device ID)
- ◆ Op Class
- ◆ Physical Size
- ◆ Protocol

## Exporting InsightIQ data

You can export InsightIQ data if you want to access the data through a third-party application, such as Microsoft Excel. After you import the data to a third-party application, you can analyze and display the data.

InsightIQ outputs data to a comma-separated value (CSV) file. InsightIQ averages data before displaying it through the InsightIQ web administration interface. However, when

you export InsightIQ data, InsightIQ does not average the data before creating the CSV file. Because of this, exported data points might not exactly match the data as it appears in the InsightIQ web administration interface.

## Export InsightIQ data to a CSV file

You can export InsightIQ data to a comma-separated value (CSV) file.

### Procedure

1. View a live performance report that displays the data that you want to export.
2. In the area of a data module you want to export data for, click **Download as CSV**.  
InsightIQ exports the data that currently appears in the data-module area. Any applied filters and breakouts are reflected in the CSV file.
3. Save the CSV file.

# CHAPTER 5

## InsightIQ commands

You can configure some InsightIQ settings and export InsightIQ data through InsightIQ commands.

- ◆ [iiq\\_data\\_export export](#)..... 64
- ◆ [iiq\\_data\\_export list](#)..... 69
- ◆ [iiq\\_data\\_export describe](#).....71
- ◆ [iiq\\_restart](#)..... 73
- ◆ [iiq\\_stop](#).....73
- ◆ [iiq\\_start](#)..... 73
- ◆ [iiq\\_network](#)..... 73

## iiq\_data\_export export

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

### Syntax

```
iiq_data_export export --data-modules <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-modules | -d} <data-module>...
```

Exports the specified data module.

**Table 1** Data module values

Module description	Module value
Active Jobs	active_job
Inactive Jobs	inactive_job
Active Job Workers	active_worker
Assigned Job Workers	assigned_worker
Active Clients	active
Connected Clients	connected
CPU % Use	cpu
Average Disk Hardware Latency (in seconds)	disk_adv_access_latency
Slow Disk Access Rate (in number per second)	disk_adv_access_slow
Disk Activity (percent)	disk_adv_busy
Disk Throughput Rate (in bytes per second)	disk_adv_bytes
Pending Disk Operations Latency (in seconds)	disk_adv_io_latency
Average Pending Disk Operations Count	disk_adv_io_queue
Disk Operations Rate (in number per second)	disk_adv_op_rate
Average Disk Operation Size (in bytes)	disk_adv_op_size
External Network Errors (in number per second)	ext_errors
Protocol Operations Average Latency (in seconds)	ext_latency
External Network Throughput Rate (in bytes per second)	ext_net
External Network Packets Rate (in number per second)	ext_packets



**Table 1** Data module values (continued)

Module description	Module value
Blocking File System Events Rate (in number per second)	ifs_blocked_heat
Contended File System Events Rate (in number per second)	ifs_contended_heat
Deadlocked File System Events Rate (in number per second)	ifs_deadlocked_heat
Total HDD Storage Capacity (in bytes)	ifs_hdd_total
Used HDD Storage Capacity (in bytes)	ifs_hdd_used
File System Events Rate (in number per second)	ifs_heat
Locked File System Events Rate (in number per second)	ifs_lock_heat
Total SSD Storage Capacity (in bytes)	ifs_ssd_total
Used SSD Storage Capacity (in bytes)	ifs_ssd_used
File System Output Rate (in bytes per second)	ifs_storage
Total Storage Capacity (in bytes)	ifs_total
Used Storage Capacity (in bytes)	ifs_used
L2 Cache Throughput Rate Hits (in bytes per second)	l2_data_read_hits
L2 Cache Throughput Rate Starts (in bytes per second)	l2_data_read_starts
Cached Data Age (in seconds)	oldest_page_cache
Protocol Operations Rate (in number per second)	op_rate
All data modules	all

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

Exports data about the specified cluster. If InsightIQ is monitoring only one cluster, that cluster's name is the default value.

`{--path | -p} <path>`

Specifies where to create the directory that will contain the .csv files. Specify a full directory path. The default value is the full path of the working directory.

`{--name | -n} <name>`

Specifies a name for the directory that will contain the .csv files. A time stamp is appended to the specified name. The default name is `export`.

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

Applies the specified breakouts to the exported data.

**Table 2** Breakout values

Breakout description	Breakout value
Direction	direction
Disk	disk
Event	event_name

**Table 2** Breakout values (continued)

Breakout description	Breakout value
Interface	interface
Job ID	job_id
Job Name	job_name
Node (logical node number)	node
Operation Class	op_class
Path	path
Protocol	proto_name
Client	remote_addr

```
{--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 3** Filter values

Breakout description	Breakout value	Filter values
Direction	direction	<ul style="list-style-type: none"> <li>out</li> <li>in</li> </ul> <p>For example: --filter-rule direction:out</p>
Disk	disk	<p>The disk number in the following format: &lt;devid&gt;/bay &lt;number&gt;</p> <p>For this filter, you must enclose both the breakout value and the filter value in quotation marks. For example: --filter-rule "disk:1/bay 1"</p>
Event	event_name	<ul style="list-style-type: none"> <li>lookup</li> <li>getattr</li> <li>read</li> <li>write</li> <li>setattr</li> <li>rename</li> <li>link</li> <li>unlink</li> </ul> <p>For example: --filter-rule event_name:write</p>

**Table 3** Filter values (continued)

Breakout description	Breakout value	Filter values
Interface	interface	<p>The interface number and type in the following format:</p> <pre>&lt;integer&gt;/&lt;type&gt;</pre> <p>where <i>&lt;type&gt;</i> is one of the following values:</p> <ul style="list-style-type: none"> <li>• ext-1</li> <li>• ext-2</li> <li>• 10gige-1</li> <li>• ext-agg</li> </ul> <p>For example:</p> <pre>--filter-rule interface:1/ext-1</pre>
Job ID	job_id	<p>The numerical ID of a job. For example:</p> <pre>--filter-rule job_id:4</pre>
Job Name	job_name	<ul style="list-style-type: none"> <li>• AutoBalance</li> <li>• AutoBalanceLin</li> <li>• AVScan</li> <li>• Collect</li> <li>• Dedupe</li> <li>• DedupeAssessment</li> <li>• DomainMark</li> <li>• FlexProtect</li> <li>• FlexProtectLin</li> <li>• FSAnalyze</li> <li>• IntegrityScan</li> <li>• MediaScan</li> <li>• MultiScan</li> <li>• PermissionRepair</li> <li>• QuotaScan</li> <li>• SetProtectPlus</li> <li>• ShadowStoreDelete</li> <li>• SmartPools</li> <li>• SnapRevert</li> <li>• SnapshotDelete</li> <li>• TreeDelete</li> </ul> <p>For example:</p> <pre>--filter-rule job_name:AVScan</pre>

**Table 3** Filter values (continued)

Breakout description	Breakout value	Filter values
Node	node	The logical node number (LNN) of a node. For example: --filter-rule node:1
Operation Class	op_class	<ul style="list-style-type: none"> <li>• read</li> <li>• write</li> <li>• other</li> <li>• namespace_read</li> <li>• file_state</li> <li>• create</li> <li>• namespace_write</li> <li>• delete</li> <li>• session_state</li> </ul> For example: --filter-rule op_class:delete
Path	path	The path of a directory starting with /ifs. For example: --filter-rule path:/ifs/data/media
Protocol	proto_name	<ul style="list-style-type: none"> <li>• nfs3</li> <li>• nfs4</li> <li>• smb2</li> <li>• http</li> <li>• iscsi</li> <li>• papi</li> <li>• siq</li> </ul> For example: --filter-rule proto_name:siq
Client	remote_addr	The host name or IP address of a client. For example: --filter-rule remote_addr:client.ip.com

`{--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 client, specifying `--number-breakouts 5` would export data about the top 5 clients. The default number is 3.

`{--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 following units are valid:

**Y**  
Specifies years

**M**  
Specifies months

**D**  
Specifies weeks

**H**  
Specifies hours

**m**  
Specifies minutes

**s**  
Specifies seconds

The default interval is 1H.

```
{--fmt-time | -f}
```

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

```
{--min-max | -x}
```

Includes the maximum and minimum values for each downsampled data point. Downsampled data points represent 10 minute intervals.

### Examples

The following command exports data about the disk activity over the past hour broken out by node:

```
iiq_data_export export --cluster cluster1 --data-modules
disk_adv_busy --path /home/exports --name disk-act --breakout-by lnn
```

The following command exports the total amount of disk space used during January of 2013:

```
iiq_data_export export --cluster cluster1 --data-modules ifs-used --
path /home/exports --name used --end 2013-02-01T24:00 --interval 1M
```

## iiq\_data\_export list

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

### Syntax

```
iiq_data_export 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. Each module supports a subset of breakouts.

```
{--list-breakouts | -l} <data-module>
```

Displays the names of all breakouts that are available for the specified data module. The following values are valid:

**Table 4** Data module values

Value	Data module
active_job	Active Jobs
inactive_job	Inactive Jobs
active_worker	Active Job Workers
assigned_worker	Assigned Job Workers
active	Active Clients
connected	Connected Clients
cpu	CPU % Use
disk_adv_access_latency	Average Disk Hardware Latency
disk_adv_access_slow	Slow Disk Access Rate
disk_adv_busy	Disk Activity
disk_adv_bytes	Disk Throughput Rate
disk_adv_io_latency	Pending Disk Operations Latency
disk_adv_io_queue	Average Pending Disk Operations Count
disk_adv_op_rate	Disk Operations Rate
disk_adv_op_size	Average Disk Operation Size
ext_errors	External Network Errors
ext_latency	Protocol Operations Average Latency
ext_net	External Network Throughput Rate
ext_packets	External Network Packets Rate
ifs_blocked_heat	Blocking File System Events Rate
ifs_contended_heat	Contended File System Events Rate
ifs_deadlocked_heat	Deadlocked File System Events Rate
ifs_hdd_total	Total HDD Storage Capacity (in bytes)
ifs_hdd_used	Used HDD Storage Capacity (in bytes)
ifs_heat	File System Events Rate
ifs_lock_heat	Locked File System Events Rate
ifs_ssd_total	Total SSD Storage Capacity (in bytes)
ifs_ssd_used	Used SSD Storage Capacity (in bytes)
ifs_storage	File System Output Rate
ifs_total	Total Storage Capacity (in bytes)

**Table 4** Data module values (continued)

Value	Data module
ifs_used	Used Storage Capacity (in bytes)
l2_data_read_hits	L2 Cache Throughput Rate Hits
l2_data_read_starts	L2 Cache Throughput Rate Starts
oldest_page_cache	Cached Data Age
op_rate	Protocol Operations Rate
all	All data modules

`--all-modules`

Displays the names of all available data module.

**Examples**

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

```
iiq_data_export list --clusters
```

The system displays output similar to the following example:

```
The list of clusters monitored on this server:
cluster_1
cluster_2
```

# iiq\_data\_export describe

Displays descriptions of data modules.

**Syntax**

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

**Options**

If no options are specified, displays descriptions of all data modules.

```
--data-module <data-module>
```

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

**Table 5** Data module values

Value	Data module
active_job	Active Jobs
inacitve_job	Inactive Jobs
active	Active Clients
connected	Connected Clients
cpu	CPU % Use
disk_adv_access_latency	Average Disk Hardware Latency
disk_adv_access_slow	Slow Disk Access Rate

**Table 5** Data module values (continued)

Value	Data module
disk_adv_busy	Disk Activity
disk_adv_bytes	Disk Throughput Rate
disk_adv_io_latency	Pending Disk Operations Latency
disk_adv_io_queue	Average Pending Disk Operations Count
disk_adv_op_rate	Disk Operations Rate
disk_adv_op_size	Average Disk Operation Size
ext_errors	External Network Errors
ext_latency	Protocol Operations Average Latency
ext_net	External Network Throughput Rate
ext_packets	External Network Packets Rate
ifs_blocked_heat	Blocking File System Events Rate
ifs_contended_heat	Contended File System Events Rate
ifs_deadlocked_heat	Deadlocked File System Events Rate
ifs_hdd_total	Total HDD Storage Capacity (in bytes)
ifs_hdd_used	Used HDD Storage Capacity (in bytes)
ifs_heat	File System Events Rate
ifs_lock_heat	Locked File System Events Rate
ifs_ssd_total	Total SSD Storage Capacity (in bytes)
ifs_ssd_used	Used SSD Storage Capacity (in bytes)
ifs_storage	File System Output Rate
ifs_total	Total Storage Capacity (in bytes)
ifs_used	Used Storage Capacity (in bytes)
l2_data_read_hits	L2 Cache Throughput Rate Hits
l2_data_read_starts	L2 Cache Throughput Rate Starts
oldest_page_cache	Cached Data Age
op_rate	Protocol Operations Rate
all	All data modules

**Examples**

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

```
iiq_data_export describe --data-module disk_adv_busy
```

The system displays output similar to the following example:

```
Entity Descriptions:
  Label: Disk Activity
  Key: disk_adv_busy
```



```

        Format: fmtPercent
    Description: Displays the average percentage of time that disks in
the
                cluster spend performing operations instead of
                sitting idle.
    Breakouts: ['lnn', 'disk']

```

## iiq\_restart

Stops and then restarts InsightIQ. This can be useful if the InsightIQ web application is currently unable to communicate with InsightIQ.

### Syntax

```
iiq_restart
```

### 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.

## iiq\_network

Opens the prompt through which you can configure network settings for an InsightIQ virtual appliance. This command is available for only the InsightIQ virtual appliance. You cannot run this command if you installed InsightIQ through Red Hat Package Manager (RPM).

### Syntax

```
iiq_network
```

### Options

There are no options for this command.



# CHAPTER 6

## InsightIQ troubleshooting

InsightIQ troubleshooting issues typically fall into one of the following categories: configuration issues, datastore issues, OneFS Platform API issues, InsightIQ web application issues, cluster import issues, or report schedule issues.

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

The current 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 datastore (for example, because the datastore is full), or InsightIQ could not contact the InsightIQ virtual machine. InsightIQ will not collect additional data until the issue is resolved.

◆ <a href="#">Configuration issues</a> .....	76
◆ <a href="#">Datastore issues</a> .....	76
◆ <a href="#">OneFS Platform API issues</a> .....	77
◆ <a href="#">Web application issues</a> .....	77
◆ <a href="#">Report schedule issues</a> .....	78

## Configuration issues

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

### VM configuration issues

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

If you cannot log in to the InsightIQ web application, open the InsightIQ virtual appliance console, configure InsightIQ networking to run in DHCP networking mode, and then try accessing the InsightIQ web application through the DHCP-generated IP address. This can indicate whether the problem 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 the ping attempt fails, verify that the node at the specified IP address is operating properly 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 host name or IP address. If the problem persists, configure a SmartConnect zone on the monitored cluster, and then try to connect to that SmartConnect zone instead of a host name 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

- ◆ If the monitored cluster is running OneFS 6.0 or later, the corresponding local user on the monitored cluster is the InsightIQ user named **insightiq**.
  - ◆ If the monitored cluster is running a version of OneFS earlier than 6.0, the corresponding local user on the monitored cluster is the OneFS Platform API user named **api**.
- 

## 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 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; therefore, the configured NFS export must grant the "root" user write access for the specified InsightIQ virtual-appliance IP address.

If InsightIQ is configured to use a local datastore and a permissions error message appears, connect to the virtual machine (VM) console and verify that the parent directory of the datastore is configured with a permissions setting of 755 or higher.

If InsightIQ cannot write to the datastore, review the permissions settings for the datastore directory and for all of the files contained in the directory. All of the files in the datastore directory must be configured with a permissions setting of 744 or higher. If the problem 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**.

InsightIQ begins deleting older data from the datastore when the datastore becomes 92% full. If InsightIQ is unable to delete data from the datastore, and the datastore continues to grow, InsightIQ will stop monitoring clusters when the InsightIQ datastore becomes 98% full. If this happens, you must manually delete data from the datastore before InsightIQ can continue monitoring clusters. InsightIQ is unable to automatically delete the following types of data from the datastore, and you must delete these types of data manually:

**Most recent data**

InsightIQ will not delete data less than 12 days old.

**Non-InsightIQ data**

InsightIQ will not delete data from the datastore that was not generated by InsightIQ.

**Generated performance reports**

InsightIQ will not delete generated performance reports.

## OneFS Platform API issues

InsightIQ connects to the monitored cluster through the Isilon OneFS Platform API.

If InsightIQ displays a `Fetching` error message, verify that the monitored cluster is online. If InsightIQ is connecting to the monitored cluster by a specific node's IP address, verify that the specified node is operating properly and that the IP address is valid.

You may also need to 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.

If the monitored cluster is running OneFS 6.0 or later, the corresponding local user on the monitored cluster is the InsightIQ user named **insightiq**. If the monitored cluster is running a version of OneFS earlier than 6.0, the corresponding local user on the monitored cluster is the OneFS Platform API user named **api**.

## Web application issues

If you are unable to access the InsightIQ web application through a web browser, your browser might be attempting to access InsightIQ through an incorrect port or the web application might not be able to communicate with InsightIQ.

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

If the problem 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 console and restart InsightIQ by running the following command:

```
iiq_restart
```

## Report schedule issues

If InsightIQ is not generating reports from a report schedule or the generated reports are not reaching the specified email addresses, verify the following settings.

### Report generation issues

If InsightIQ is not generating reports according to the times specified in a report schedule, InsightIQ might be generating reports based on an incorrect time zone. To correct this, modify the InsightIQ time zone as needed.

### Report delivery issues

If a generated report is not being received by the specified email address or addresses, perform the following tasks.

- ◆ Configure InsightIQ email settings. InsightIQ might not be able to send outbound email through your organization's email server.
- ◆ Reduce the size of the generated reports by removing breakouts from the report schedule. Many email servers reject emails that are larger than a certain limit. Your email server might be rejecting the reports because they are too large.
- ◆ Configure your email program to allow emails sent by InsightIQ. Your email program might be filtering emails sent by InsightIQ.