

Isilon InsightIQ

Version 4.1.3

User Guide

Copyright © 2009-2019 Dell Inc. or its subsidiaries. All rights reserved.

Published May 2019

Dell 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.” DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.
Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Chapter 1	Introduction to this guide	5
	About this guide.....	6
	Where to go for support.....	6
Chapter 2	InsightIQ overview	7
	Introduction to InsightIQ.....	8
	Monitoring clusters with InsightIQ Dashboard.....	8
	Performance reports and File System reports.....	9
	Report configuration.....	9
	InsightIQ feature compatibility with OneFS.....	9
Chapter 3	Cluster monitoring	11
	View cluster status.....	12
	View aggregate cluster status.....	12
	View the status of a cluster.....	12
	Monitoring errors.....	12
Chapter 4	Performance reports	15
	Performance reports overview.....	16
	InsightIQ downsampling.....	16
	Maximum and minimum data values.....	16
	Data Module overview.....	17
	Live Performance Reports overview.....	17
	Cluster Performance report overview.....	18
	Node Performance report overview.....	19
	Network Performance report overview.....	20
	Client Performance report overview.....	21
	Disk Performance report overview.....	22
	File System Performance report overview.....	23
	File System Cache Performance report overview.....	24
	Cluster Capacity Performance report overview.....	26
	Deduplication Performance report overview.....	26
	Cluster Events Performance report overview.....	27
	Jobs and Services Performance report overview.....	28
	Live Performance reports.....	29
	Creating a live Performance report.....	30
	View a Live Performance report.....	31
	Modify a live Performance report.....	32
	Delete a saved live Performance report.....	32
	Create a Performance report schedule from a live Performance report.....	33
	Publish reports.....	34
	View a published Performance report.....	34
	Send a published report.....	34
	Delete a published report.....	35
	Publish a report on demand.....	36
	Scheduling Performance reports.....	36
	Create a scheduled Performance report.....	37

	Disable a scheduled Performance report.....	39
	Enable a scheduled Performance report	39
	Modify a scheduled Performance report.....	40
	Delete a scheduled Performance report	40
	Troubleshooting scheduled Performance report issues.....	41
	Interpreting Performance reports.....	42
	Active Clients report.....	42
	Average Pending Disk Operations Count report.....	42
	External Network Throughput Rate report.....	43
	Overall Cache Hit Rate report.....	43
	Overall Cache Throughput Rate report.....	43
	Pending Disk Operations Latency report.....	44
	Protocol Operations Average Latency report.....	44
Chapter 5	File System reports	45
	File System reports.....	46
	File System Analytics result set snapshots.....	46
	File System report modules.....	46
	File System report breakouts.....	50
	Capacity reports overview.....	51
	Viewing a Capacity report.....	51
	Deduplication reports.....	52
	Viewing a Deduplication report.....	52
	File System reports overview.....	53
	Viewing a File System Analytics report.....	54
	Quota reports.....	54
	View a Quota report.....	55
Chapter 6	Creating custom reports	57
	Building reports overview.....	58
	Modules.....	58
	Breakouts.....	63
	Breakout types.....	64
	Applying a breakout.....	65
	Filters.....	65
	Filter rules.....	66
	Create a filter.....	72
	Modify a filter.....	73
	Delete a filter.....	74
Chapter 7	InsightIQ export data	75
	Export data overview.....	76
	Export InsightIQ data.....	76
	Permalinks.....	76
	Create a Permalink to share a report.....	77

CHAPTER 1

Introduction to this guide

This section contains the following topics:

- [About this guide](#).....6
- [Where to go for support](#).....6

About this guide

This guide describes how to monitor and analyze Isilon cluster data by using InsightIQ.

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

Where to go for support

If you have any questions about EMC Isilon products, contact EMC Isilon Technical Support.

Online Support	Live Chat Create a Service Request
Telephone Support	United States: 1-800-SVC-4EMC (800-782-4362) Canada: 800-543-4782 Worldwide: +1-508-497-7901 For local phone numbers for a specific country, see EMC Customer Support Centers .
Help with online support	For questions specific to EMC Online Support registration or access, email support@emc.com .

Support for InsightIQ

If you are running a free version of InsightIQ, community support is available through the [EMC Isilon Community Network](#). If you have one or more licenses of InsightIQ and have a valid support contract for the product, contact EMC Isilon Technical Support for assistance.

Support for IsilonSD Edge

If you are running a free version of IsilonSD Edge, community support through EMC community network is available. However, if you have purchased one or more licenses of IsilonSD Edge, you can contact EMC Isilon Technical Support for assistance, provided you have a valid support contract for the product.

CHAPTER 2

InsightIQ overview

This section contains the following topics:

- [Introduction to InsightIQ](#) 8
- [Monitoring clusters with InsightIQ Dashboard](#) 8
- [Performance reports and File System reports](#) 9
- [Report configuration](#) 9
- [InsightIQ feature compatibility with OneFS](#) 9

Introduction to InsightIQ

InsightIQ provides tools to monitor and analyze historical data from Isilon clusters.

With InsightIQ, in the InsightIQ web application, you can view standard or customized Performance reports and File System reports, to monitor and analyze Isilon cluster activity. You can create customized reports to view information about storage cluster hardware, software, and protocol operations. You can publish, schedule, and share reports, and you can export the data to a third-party application.

You can create and view specific InsightIQ reports to identify or confirm the cause of a performance issue. For example, if users report client connectivity issues during certain dates, you can create a report that indicates the cause of the issue. InsightIQ enables you to correlate data across present and historical conditions.

If you modify the Isilon cluster environment, you can measure the effects of those changes by creating an InsightIQ report that compares the past performance with the performance after the changes were made. For example, if you add 20 new clients, you can create a report to compare the performance before the change and after the change. This comparison allows you to determine the impact of the additional clients on system performance.

Create customized reports that help to identify bottlenecks or inefficiencies in Isilon cluster systems and workflows. For example, if you want to verify that all clients can access the cluster quickly and efficiently, you can create a report that measures which client connections are faster or slower. You can then modify the report to add breakouts and filters to identify the cause of the slower connections.

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

A review of past performance trends can help you predict future trends and needs. For example, if you deployed an Isilon cluster for a data-archival project 6 months ago, you might want to estimate when the cluster reaches its maximum capacity. You can customize a report to illustrate capacity usage by day, week, or month to estimate when the cluster reaches capacity.

You can install InsightIQ on a Linux computer or on a virtual machine. Configure settings such as the IP address of InsightIQ and the administrator password in the InsightIQ console interface.

Monitoring clusters with InsightIQ Dashboard

View the status of all the monitored clusters.

InsightIQ Dashboard, available through the InsightIQ web application, provides you with an overview of the status of all the monitored clusters. The Cluster Status summary includes information about cluster capacity, clients, throughput, and CPU usage. Current information about the monitored clusters appears alongside graphs that show the relative changes in the statistics over the past 12 hours.

The **Aggregated Cluster Overview** section displays the total or average values of the status information for the monitored clusters. This information can help you decide what to include in a Performance report. For example, if the total amount of network traffic for all the monitored clusters is higher than anticipated, a customized Performance report can show you the data about network traffic. The report can

show you the network throughput by direction, by using breakouts, to help you determine whether one direction of throughput is contributing to the total more than the other.

Performance reports and File System reports

Monitor EMC Isilon clusters with customizable reports.

Performance reports include information about the cluster activity and capacity. The information can help you confirm that storage clusters perform as expected, and help you identify the specific cause of a performance-related issue to investigate. File System reports include information about the cluster file system, including deduplication, quotas, and usable capacity. The information can help you identify the types of data that is stored and where the data is located.

With InsightIQ, you can create live reports and view them through the InsightIQ web application. You can create live reports for Performance reports and File System reports. You can modify the attributes of live reports as you view the reports, including the period, breakouts, and filters.

You can view Performance reports as a PDF file that is generated based on a report schedule that the administrator sets up. InsightIQ can be configured to send a PDF file as an email attachment. You can use a PDF file to verify cluster health or to distribute InsightIQ information to individuals who do not have access to the InsightIQ web application.

Report configuration

InsightIQ reports are configured by using modules, breakouts, and filter rules.

A module is a section of a report that displays information about a cluster. A breakout can be applied to a module so that users can view information about separate cluster components. A filter rule can be applied to a module so that users can view information about a specific component across the entire report. Filter rules can be combined into collections that are called filters. Filters can be saved and applied to multiple reports.

InsightIQ feature compatibility with OneFS

Certain InsightIQ features are supported on specific versions of OneFS.

OneFS feature compatibility

Certain InsightIQ features require specific versions of OneFS.

InsightIQ can be used with clusters that run versions of OneFS earlier than OneFS 8.0.1. However, InsightIQ includes features that OneFS versions earlier than OneFS 8.0.1 do not support.

Table 1 InsightIQ features supported by OneFS versions

InsightIQ feature	Requirement	Notes
Performance modules - Cache, Event Summary, Node pool, and Tier breakouts	OneFS 7.1.0 or later	To view data about L3 cache, the cluster must be running OneFS 7.1.1 or later.

Table 1 InsightIQ features supported by OneFS versions (continued)

InsightIQ feature	Requirement	Notes
Reports - Quota, Capacity, and Deduplication	OneFS 7.1.0 or later	To view more than 1000 quotas at a time, the cluster must be running OneFS 7.2.0 or later.
File System Analytics - Node pool and tier breakouts	OneFS 8.0.0 or later	None

CHAPTER 3

Cluster monitoring

This section contains the following topics:

- [View cluster status](#)..... 12
- [View aggregate cluster status](#).....12
- [View the status of a cluster](#)..... 12
- [Monitoring errors](#).....12

View cluster status

The InsightIQ Dashboard view shows information about all monitored clusters.

On the InsightIQ Dashboard, you can view information for all the monitored clusters or for individual monitored clusters. The InsightIQ Dashboard modules can help you identify issues and determine which cluster to investigate further. You can compare information between clusters to help identify whether the issue occurs on all clusters. You can apply a report to all the clusters to narrow down the source of a performance issue, and then create a Performance report to view more detailed information about the affected cluster. For example, if you identify an issue in the **Aggregated Cluster Overview**, you can review the information for individual clusters to identify which cluster is the source of the issue. You can then create a Performance report to view additional information about the affected cluster.

If Usable Capacity information is available for the cluster, the **Aggregated Cluster Overview** shows you information about the Usable Capacity available on the cluster rather than the remaining physical space that is available.

View aggregate cluster status

View the status of all monitored clusters.

Procedure

1. In the InsightIQ web application, click the **Dashboard** tab.
The **InsightIQ Dashboard** view appears.
2. In the **Aggregated Cluster Overview** section, view the status of all monitored clusters.

View the status of a cluster

View the status of an individual cluster.

Procedure

1. In the InsightIQ web application, click the **Dashboard** tab.
The **InsightIQ Dashboard** view appears.
2. Beneath the **Aggregated Cluster Overview** section is the status of all monitored clusters. To find the status for a specific cluster, scroll down the **InsightIQ Dashboard** view.
Expand the section to see more information about a cluster.

Monitoring errors

Identify and resolve problems that are encountered when monitoring a cluster.

To resolve the following errors, you must be logged in to InsightIQ with an administrator account.

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 is unable to connect to the monitored cluster. If the Monitoring Status of the cluster is yellow, InsightIQ is trying to resolve this issue. If the Monitoring Status of the cluster is red, InsightIQ is unable to repair the connection to the cluster and resolve this issue manually. To resolve this issue, reestablish network connectivity between InsightIQ and the monitored cluster. After the connection is reestablished, InsightIQ can connect to the monitored cluster automatically.

Database busy

InsightIQ is unable to add cluster data to the InsightIQ data store because the maximum number of data store connections are already in use. Cluster monitoring is temporarily paused until the necessary resources are freed. If this error persists for several minutes or more, open an SSH connection to InsightIQ and restart InsightIQ by running the following command:

```
iiq_restart
```

Data store connection error

InsightIQ is unable to connect to the InsightIQ data store. Cluster monitoring is suspended until this issue is resolved. To resolve this issue, ensure network connectivity between InsightIQ and the data store. InsightIQ tries to connect to the data store until the network connection is reestablished.

Data store full

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

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 commonly resolves itself as InsightIQ retrieves more recent data. If this issue also reports one or more InsightIQ errors, resolve the errors. If the error persists after all other InsightIQ errors are resolved, contact Isilon Technical Support.

FSA connection error

InsightIQ is unable to retrieve File System Analytics reports from the cluster. To resolve this issue, ensure that NFSv3 is enabled on the monitored cluster and that InsightIQ can connect to the cluster. After this issue is resolved, InsightIQ automatically resumes retrieving File System Analytics reports.

GUID error

The globally unique identifier (GUID) associated with the monitored cluster in InsightIQ has changed. The hostname or IP address that is associated with the cluster might be 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 accesses.

License error

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

Privilege error

The InsightIQ user does not have sufficient privileges to retrieve the specified data from the cluster. For information about adding privileges to InsightIQ, see [KB 194112](#).

CHAPTER 4

Performance reports

This section contains the following topics:

- [Performance reports overview](#) 16
- [Live Performance Reports overview](#) 17
- [Live Performance reports](#)..... 29
- [Publish reports](#).....34
- [Scheduling Performance reports](#).....36
- [Interpreting Performance reports](#)..... 42

Performance reports overview

View information about the activity of the nodes, networks, clients, disks, and IFS cache of a monitored cluster.

Performance reports provide you with information about the performance and capacity of a monitored 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 for a range of dates. 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 can create a Performance report that shows the changes in available capacity over the past year. You can create a Performance report schedule to send a weekly report that illustrates the cluster capacity over the past week.

Note

If InsightIQ 4.1.3 is monitoring a OneFS 8.2.0 cluster, filtering might not work properly on some reports.

InsightIQ downsampling

InsightIQ downsamples the data that it collects for Performance reports.

Downsampling means that InsightIQ converts higher-resolution data to lower-resolution data by adding or averaging several higher-resolution samples. InsightIQ collects data samples at different intervals, depending on the type of data collected. InsightIQ averages or adds samples together to create a single, low-resolution sample that represents a 10-minute period, and preserves the minimum and maximum values for that period. InsightIQ deletes the averaged samples after 24 hours, to limit the size of the data store.

Maximum and minimum data values

You can view the maximum and minimum values for the downsampled data in a report.

In reports, 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 monitored cluster activity varies over a date range. For example, if the maximum value exceeds the averaged value, and the minimum time value is close to the averaged value, it is likely that the value fluctuated significantly over a short time.

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

Table 2 Allowable data values by module

Module	Notes
Active Clients	You can view maximum and minimum values when you filter by protocol.
Connected Clients	

Table 2 Allowable data values by module (continued)

Module	Notes
Protocol Operations Average Latency	
Protocol Operations Rate	
File System Events Rate	You can view maximum and minimum values when you filter by file system event.
Cache Hits	You can view maximum and minimum values when you filter by service.
CPU Usage Rate	
Disk IOPS	
Deduplication Summary (Logical)	--
Deduplication Summary (Physical)	
HDD Storage Capacity	
Jobs	
Job Workers	
L1 and L2 Cache Prefetch Throughput Rate	
L1 Cache Throughput Rate	
L2 Cache Throughput Rate	
L3 Cache Throughput Rate	
Overall Cache Hit Rate	
Overall Cache Throughput Rate	
Storage Capacity	
Writable Capacity	

Data Module overview

Overview of the data modules used to create the performance reports.

Data modules collect cluster data and allow users to determine how that data is displayed in a performance report. Each standard performance report included with IIQ is a selection of related data modules. You can also use these same data modules to create your own custom performance reports.

Live Performance Reports overview

Live Performance reports provide information on the performance of the cluster.

Live Performance reports provide you with a quick and easy way to monitor a cluster's performance. You can focus the reports on different data types including cluster capacity, network throughput, clients, jobs and services, events, and cache. Each Live Performance report consists of a collection of modules that monitor a OneFS cluster and display live data on request.

Note

Modules, by default, automatically refresh every 30 seconds. If you expand the date range to greater than 6 hours, the automatic refresh is disabled.

Cluster Performance report overview

The Cluster Performance report displays the general health and status of the cluster. You can focus the modules on individual nodes or clients to view more detailed information about the affects on cluster performance.

Here is a list of modules that are included in this report:

- External Network Throughput Rate
- Protocol Operations Rate
- CPU % Use
- Active Clients
- Connected Clients
- Jobs
- Job Workers

Configure the Cluster Performance report

You can view cluster network activity throughput, protocol operations rate, CPU usage, client activity, and job activity within the Cluster Performance report.

For this example, you are researching into network latency on the cluster to determine the cause.

Procedure

1. Click the **Performance Reporting** tab.
2. On the **Live Performance Reporting** page, click the **Cluster** drop down and select the cluster that you want to monitor.
3. Click the **Report Type** drop down and select **Cluster Performance**.
4. In the **Date Range** fields, select a range to display, then click **View Report**.

The date range should include the entire period of time when the cluster was behaving unexpectedly. The default setting is to view the last hour of data.

The modules in the report are now populated with data for the cluster and date range you selected.

5. In the **External Network Throughput Rate** module, check the throughput rate during the time in question.

Note

High throughput values indicate times when the cluster's network was under heavy load.

6. In the **External Network Throughput Rate** module, click the **Client** breakout. To learn more about breakouts, refer to the [Breakouts](#) on page 63 topic. To remove focus on a breakout, click the **None** breakout.

On the left side of the module, you can now see a list of all clients that were using the cluster's network throughput during the date range you specified. The client at the top of the list was using the highest percentage of network throughput.

7. Click to select the client that was using the most throughput during the date range.

To remove focus on a client, click the client filter button at the top of the module.

All applicable modules in the report are now showing information that is focused on the client you selected, and data from all other clients are removed from the module.

Results

With the information you gathered, you can determine that a specific client might be causing performance issues on the cluster.

Node Performance report overview

The Node Performance report displays the health and status of each node in the cluster.

Users can focus the data modules on specific clients or protocols to view more detailed information about the affects on node performance.

Here is a list of data modules included in this report:

- Node Summary
- External Network Throughput Rate
- Protocol Operations Rate
- CPU % Use
- Disk Throughput Rate
- Disk Operations Rate
- Active Clients
- Connected Clients

Configure the Node Performance report

You can view a list of nodes in the cluster, network activity, disk operations, client activity, and protocols within the Node Performance report.

For this example, you are researching into a node with high I/O to determine the cause.

Procedure

1. Click the **Performance Reporting** tab.
2. On the **Live Performance Reporting** page, click the **Cluster** drop down and select the cluster that you want to monitor.
3. Click the **Report Type** drop down and select **Node Performance**.
4. In the **Date Range** fields, select a range to display, then click **View Report**.

The date range should include the entire period of time when the node was behaving unexpectedly. The default setting is to view the last hour of data.

The modules in the report are now populated with data for the node and date range you selected.

5. In the **Node Summary** module, check for nodes behaving differently than other nodes.

Note

Unusually high or low I/O on one node might indicate an issue.

6. In the **Node Summary** module, click the node with the highest Total I/O.

To remove focus on a node, click the node filter button at the top of the module.

All applicable modules in the report are now focused on the node you filtered. Data from all other nodes is removed from the modules.

7. In the **External Network Throughput Rate** module, click the **Client** breakout.

To learn more about breakouts, refer to the [Breakouts](#) on page 63 topic.

To remove focus on a breakout, click the **None** breakout.

All applicable modules in the report are now showing information that is focused on the node you filtered and the client breakout you selected.

Results

In this example, you can determine which node had high I/O, and focus the modules to determine which client has the most affect on the node performance.

Network Performance report overview

The Network Performance report displays the health and status of the cluster's network.

You can focus the data modules on specific protocols, errors, nodes, pools, or tiers to view more detailed information about network performance.

Here is a list of data modules included in this report:

- External Network Throughput Rate
- External Network Packets Rate
- Protocol Operations Average Latency
- External Network Errors

Configure the Network Performance report

You can view network throughput, packet rates, protocol latency, and external network errors within the Network Performance report.

For this example, you are researching into network latency to determine the cause.

Procedure

1. Click the **Performance Reporting** tab.
2. On the **Live Performance Reporting** page, click the **Cluster** drop down and select the cluster that you want to monitor.
3. Click the **Report Type** drop down and select **Network Performance**.

4. In the **Date Range** fields, select a range to display, then click **View Report**.

The date range should include the entire period of time when the cluster was behaving unexpectedly. The default setting is to view the last hour of data.

The modules in the report are now populated with data for the cluster and date range you selected.

5. In the **External Network Throughput Rate** module, check the throughput rate during the time in question.

Note

High throughput values indicate times when the cluster's network was under heavy load.

6. In the **External Network Throughput Rate** module, click the **Client** breakout.

To learn more about breakouts, refer to the [Breakouts](#) on page 63 topic.

To remove focus on a breakout, click the **None** breakout.

On the left side of the module, you can now see a list of all clients that were using the cluster's network throughput during the timeframe you specified. The client at the top of the list was using the highest percentage of network throughput.

7. Click to select the client that was using the most throughput during the date range.

To remove focus on a client, click the client filter button at the top of the module.

All applicable modules in the report are now showing information that is focused on the client you selected, and data from all other clients are removed from the module.

8. In the **Protocol Operations Average Latency** module, click the **Protocol** breakout.

On the left side of the module, you can now see a list of all protocols in use during the date range you specified, with the protocol at the top experiencing the most latency.

Results

With the information you gathered, you can determine which client and protocol might be causing network latency issues on the cluster.

Client Performance report overview

The Client Performance report is used to monitor clients connecting to the cluster.

You can focus the modules on specific clients, protocols, nodes, and tiers to see how they are affecting performance on the cluster.

Here is a list of data modules included in this report:

- Client Summary
- External Network Throughput Rate
- Protocol Operations Rate
- Protocol Operations Average Latency

- External Network Errors

Configure the Client Performance report

You can view a list of connected clients, network throughput, network errors, and protocol operations within the Client Performance report.

For this example, you are researching into clients experiencing performance issues through specific protocols.

Procedure

1. Click the **Performance Reporting** tab.
2. On the **Live Performance Reporting** page, click the **Cluster** drop down and select the cluster that you want to monitor.
3. Click the **Report Type** drop down and select **Client Performance**.
4. In the **Date Range** fields, select a range to display, then click **View Report**.

The date range should include the entire period of time when the cluster was behaving unexpectedly. The default setting is to view the last hour of data.

The modules in the report are now populated with data for the cluster and date range you selected.

5. In the **Client Summary** module, click the **Client** with the highest Total I/O.

The modules below the **Client Summary** are now populated with data for the client you selected.

6. In the **Protocol Operations Average Latency** module, click the **Protocol** breakout.

To learn more about breakouts, refer to the [Breakouts](#) on page 63 topic.

To remove focus on a breakout, click the **None** breakout.

On the left side of the module, you can now see a list of all protocols being used by the client you selected. The protocol at the top of the list experienced the highest average latency during the date range you specified.

Results

With the information you gathered, you can determine the client with the highest Total I/O, and the protocols the client used to access the cluster.

Disk Performance report overview

The Disk Performance report displays the general health of the disks on the cluster.

You can focus the modules on individual disks, nodes, and pools to determine the affect they have on the clusters performance.

Here is a list of the data modules included in this report:

- Protocol Operations Rate
- Disk Activity
- Disk Operations Rate
- Disk Throughput Rate
- Average Disk Operation Size
- Average Disk Hardware Latency

- Average Pending Disk Operations Count

Configure the Disk Performance report

You can view disk activity, disk throughput, and disk operations within the Disk Performance report.

For this example, you are researching into disk latency issues to determine what might have been the cause.

Procedure

1. Click the **Performance Reporting** tab.
2. On the **Live Performance Reporting** page, click the **Cluster** drop down and select the cluster that you want to monitor.
3. Click the **Report Type** drop down and select **Disk Performance**.
4. In the **Date Range** fields, select a range to display, then click **View Report**.

The date range should include the entire period of time when the cluster was behaving unexpectedly. The default setting is to view the last hour of data.

The modules in the report are now populated with data for the cluster and date range you selected.

5. In the **Disk Activity** module, click the **Disk** breakout.

To learn more about breakouts, refer to the [Breakouts](#) on page 63 topic.

To remove focus on a breakout, click the **None** breakout.

Note

High activity on a single disk or a few disks and low activity on the remaining disks might indicate performance issues on the OneFS cluster.

On the left side of the module, you can now see a list of all disks that were not idle during the date range you specified. The disk at the top of the list had the most activity during the date range you selected.

6. In the **Disk Throughput Rate** module, click the **Disk** breakout.

On the left side of the module, you can now see a list of all disks that were being written to or read from during the date range. The disk at the top of the list had the highest percentage throughput.

Results

With the information you gathered, you can determine which disk has the highest throughput and might be causing performance issues on the cluster.

File System Performance report overview

The File System performance report displays the general health of the file system and recent events.

You can focus the modules on specific nodes, pools, and events to view more detailed information about the affects on cluster performance.

Here is a list of the data modules included in this report:

- File System Events Rate

- Deadlocked File System Events Rate
- Contended File System Events Rate
- Locked File System Events Rate
- Blocking File System Events Rate

Configure the File System Performance report

You can view file system events and event service rates within the File System Performance report.

For this example, you are researching into a particular date range when the OneFS cluster behaved unexpectedly and determine what might have been the cause.

Procedure

1. Click the **Performance Reporting** tab.
2. On the **Live Performance Reporting** page, click the **Cluster** drop down and select the cluster that you want to monitor.
3. Click the **Report Type** drop down and select **File System Performance**.
4. In the **Date Range** fields, select a range to display, then click **View Report**.

The date range should include the entire period of time when the cluster was behaving unexpectedly. The default setting is to view the last hour of data.

The modules in the report are now populated with data for the cluster and date range you selected.

5. In the **File System Events Rate** module, click the **Node** breakout.

To learn more about breakouts, refer to the [Breakouts](#) on page 63 topic.

To remove focus on a breakout, click the **None** breakout.

On the left side of the module, you can now see a list of all nodes that were interacting with the file system during the date range you specified. The node at the top of the list had the highest percentage of file system events.

6. Select the Node with the highest percentage of file system events.
7. In the **File System Events Rate** module, click the **FS Event** breakout.

On the left side of the module, you can now see a list of all file system events that the selected node was sending and receiving during the date range you specified. The event at the top of the list was the most active.

Results

With the information you gathered, you can determine which node is servicing the most events and the rate of those events that might be causing performance issues on the cluster.

File System Cache Performance report overview

The File System Cache Performance report displays the general health of the file system cache.

You can focus on cache level, cache hits, and cache throughput to view more detailed information about the affects on cluster performance.

Here are the data modules included in this report:

- Overall Cache Hit Rate
- File System Throughput Rate
- Overall Cache Throughput Rate
- L1 Cache Throughput Rate
- L2 Cache Throughput Rate
- L3 Cache Throughput Rate
- L1 and L2 Cache Prefetch Throughput Rate
- Average Cached Data Age

Configure the File System Cache Performance report

You can view file system cache levels, throughput, and hit rate within the File System Cache Performance report.

For this example, you are researching into file system cache latency on the OneFS cluster to determine the cause.

Procedure

1. Click the **Performance Reporting** tab.
2. On the **Live Performance Reporting** page, click the **Cluster** drop down and select the cluster that you want to monitor.
3. Click the **Report Type** drop down and select **File System Cache Performance**.
4. In the **Date Range** fields, select a range to display, then click **View Report**.

The date range should include the entire period of time when the cluster was behaving unexpectedly. The default setting is to view the last hour of data.

The modules in the report are now populated with data for the cluster and date range you selected.

5. In the **Overall Cache Hit Rate** module, check the rate of L1, L2, and L3 cache during the time in question.

Note

If L3 cache is enabled and has a higher hit rate than L2 or L1 cache, cluster performance might be affected.

6. In the **Overall Cache Throughput Rate** module, select **L3 Starts** in the breakout dropdown and click the **Node** breakout.

To learn more about breakouts, refer to the [Breakouts](#) on page 63 topic.

To remove focus on a breakout, click the **None** breakout.

On the left side of the module, you can now see a list of all nodes that were requesting info from L3 during the date range you specified. The node at the top of the list made the most requests to L3.

Results

With the information you gathered, you can determine which node is making the most cache requests, what type of cache requests are being made, and how the cache requests affect performance on the OneFS cluster.

Cluster Capacity Performance report overview

The Cluster Capacity Performance report displays the capacity of the cluster.

You can focus the modules on capacity, node, and node pool to view more detailed information about the affect on cluster performance.

Here is a list of data modules included in this report:

- Cluster Capacity

Configure the Cluster Capacity Performance report

You can view total capacity, provisioned capacity, and writable capacity within the Cluster Capacity Performance report.

For this example, you are researching into the cluster's total data usage to determine if it is affecting performance on the cluster.

Procedure

1. Click the **Performance Reporting** tab.
2. On the **Live Performance Reporting** page, click the **Cluster** drop down and select the cluster that you want to monitor.
3. Click the **Report Type** drop down and select **Cluster Capacity**.
4. In the **Date Range** fields, select a range to display, then click **View Report**.

The date range should include the entire period of time when the cluster was behaving unexpectedly. The default setting is to view the last hour of data.

The modules in the report are now populated with data for the cluster and date range you selected.

5. In the **Cluster Capacity** module, check the total capacity and compare to the total usage during the time in question.

Note

If the total usage is close to the total capacity, you might be close to filling up the cluster. For more detailed information, view the **Capacity Reporting** page under the **File System Reporting** tab.

Results

With the information you gathered, you can determine how close you are to reaching the cluster's total capacity and estimate when additional capacity might be needed.

Deduplication Performance report overview

The Deduplication Performance report displays the amount of deduplicated data on the cluster.

You can focus the modules on disk capacity, deduplicated data, or saved storage space to view more detailed information on the affects on cluster performance.

Here is a list of modules that are included in this report:

- Cluster Capacity
- Deduplication Summary (Logical)

- Deduplication Summary (Physical)

Configure the Deduplication Performance report

You can view cluster capacity, deduplicated data, and saved storage space within the Deduplication Performance report.

For this example, you are researching into the amount of storage space deduplication has saved on the cluster.

Procedure

1. Click the **Performance Reporting** tab.
2. On the **Live Performance Reporting** page, click the **Cluster** drop down and select the cluster that you want to monitor.
3. Click the **Report Type** drop down and select **Deduplication**.
4. In the **Date Range** fields, select a range to display, then click **View Report**.
The date range should include the entire period of time when the cluster was behaving unexpectedly. The default setting is to view the last hour of data.

The modules in the report are now populated with data for the cluster and date range you selected.
5. In the **Deduplication Summary (Physical)** module, view the Space Saved.

Results

With the information you gathered, you can determine how much storage space (physical) you have saved with deduplication.

Cluster Events Performance report overview

The Cluster Events Performance report displays events and jobs on the cluster.

You can focus the modules on specific events, jobs, and nodes to view more detailed information about the affect on cluster performance.

Here is the list of modules that are included in this report:

- Event Summary
- Jobs
- Job Workers

Configure the Cluster Events Performance report

You can view cluster events and jobs within the Cluster Events Performance report.

For this example, you are researching into recent cluster events to determine what might have been the cause.

Procedure

1. Click the **Performance Reporting** tab.
2. On the **Live Performance Reporting** page, click the **Cluster** drop down and select the cluster that you want to monitor.
3. Click the **Report Type** drop down and select **Cluster Events**.
4. In the **Date Range** fields, select a range to display, then click **View Report**.
The date range should include the entire period of time when the cluster was behaving unexpectedly. The default setting is to view the last hour of data.

The modules in the report are now populated with data for the cluster and date range you selected.

5. In the **Event Summary** module, check the list of events, the severity, and the time noticed.

By default, the event summary list is sorted with the most recent event at the top.

Note

The bar at the top of the Event Summary module displays the date range that is specified, and highlights the times when events occurred. If you minimize the Event Summary module, you can better compare the highlighted events in the bar with other modules.

6. In the **Jobs** module, click the **Job Type** breakout.

To learn more about breakouts, refer to the [Breakouts](#) on page 63 topic.

To remove focus on a breakout, click the **None** breakout.

On the left side of the module, you can now see a list of all Job Types that were active during the date range you specified. The Job Type at the top of the list was the most active.

7. In the **Jobs** module, click the **Job ID** breakout.

To remove focus on a breakout, click the **None** breakout.

On the left side of the module, you can now see a list of all Job IDs that were active during the date range you specified. The Job ID at the top of the list was the most active.

Results

With the information you gathered, you can determine which job type and job id might have caused the event.

Jobs and Services Performance report overview.

The Jobs and Services Performance report displays jobs and services on the cluster.

Note

If InsightIQ 4.1.3 is monitoring a OneFS 8.2.0 cluster, the Job and Services related statistics, breakouts, and report data might not display properly.

You can focus the modules on individual jobs, services, and nodes to view more detailed information about the affects on cluster performance.

Here is a list of data modules included in this report:

- Jobs and Services Summary
- CPU Usage Rate
- Disk IOPS
- Cache Hits

Configure the Jobs and Services Performance report

You can view jobs, services, disk IOPS, cache hits, or CPU usage within the Jobs and Services Performance report.

For this example, you are researching into recent cluster latency to determine the cause.

Procedure

1. Click the **Performance Reporting** tab.
2. On the **Live Performance Reporting** page, click the **Cluster** drop down and select the cluster that you want to monitor.
3. Click the **Report Type** drop down and select **Jobs and Services**.
4. In the **Date Range** fields, select a range to display, then click **View Report**.

The date range should include the entire period of time when the cluster was behaving unexpectedly. The default setting is to view the last hour of data.

The modules in the report are now populated with data for the cluster and date range you selected.

5. In the **Jobs and Services Summary** module, check the service, the CPU usage rate, and the cache during the date range specified.

Note

High CPU usage rates indicate times when a job was using high levels of cluster resources.

6. Click to select the Service that was using the most CPU usage during the date range.

To remove focus, click the service filter button at the top of the module.

All applicable modules in the report are now showing information that is focused on the service you selected, and data from all other services are removed from the module.

7. In the **CPU Usage Rate** module, click the **Node** breakout.

To learn more about breakouts, refer to the [Breakouts](#) on page 63 topic.

To remove focus on a breakout, click the **None** breakout.

On the left side of the module, you can now see a list of all nodes that were using CPU time during the date specified. The node at the top of the list was using the highest percentage of CPU.

Results

With the information you gathered, you can determine which job and service had high CPU usage and what node originated the request.

Live Performance reports

Create, view, modify, delete, and schedule live Performance reports.

Live Performance reports allow you to modify a report as you view it. You can only view a live Performance report while using the InsightIQ web application.

When you view a live Performance report, you can modify the range of data that is displayed, as well as change breakouts and filters. Real-time Performance reports are useful when you want to investigate a performance issue that is affecting a cluster. You can view a live Performance report to see the individual performance of specific components.

In a live Performance report, you can view the status of the monitored cluster. This information includes the clusters current date and time, the number of nodes that are deployed, a client-activity overview, a network throughput overview, the capacity usage, and CPU utilization.

Creating a live Performance report

Creates custom live Performance reports.

Procedure

1. In the InsightIQ web application, click the **Performance Reporting** tab.
2. Click **Create a New Performance Report**.

The **Create a New Performance Report** report template view appears.

3. To create the Live Performance report, specify a template.

Option	Description
Create a live Performance report from a template that is based on the default settings.	In the Select a Starting Point for Your Performance Report section, click Create from Blank Report .
Create a live Performance report based on a Saved Performance report.	In the Saved Performance Reports section, click a saved report.
Create a live Performance report based on one of the standard reports.	<p>In the Standard Reports section, click the name of a Standard report.</p> <ul style="list-style-type: none"> • Node Performance report template • Cluster Performance report template • Client Performance report template • Network Performance report template • File System Performance report template • Disk Performance report template • Cluster Capacity report template • File System Cache Performance report template • Cluster Events report template • Deduplication report template • Jobs and Services report template

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

4. In the **Build Your New Performance Report** section, in the **Performance Report Name** field, type a name for the Performance report.
5. Select the **Live Performance Reporting** checkbox.

Note

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** checkboxes.

6. In the **Select the Data You Want to See** section, specify the Performance modules that you want to include in the report.

Option	Description
Add a Performance module	<ol style="list-style-type: none"> a. In any unassigned Performance module area, from the Select a Module for this Position drop-down list, select a performance module. b. If all existing Performance module areas are assigned, at the bottom of the section, click Add another performance module.
Modify a Performance module	<ul style="list-style-type: none"> • In an assigned Performance module area, from the Performance module drop-down list, select a different Performance module.
Delete a Performance module	<ul style="list-style-type: none"> • In an assigned Performance module area, click Delete this performance module.

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

7. Complete the Performance 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 the report to a monitored cluster and view the results, click **Finish**.

View a Live Performance report

View a Live Performance report.

Procedure

1. In the InsightIQ web application, click the **Dashboard** tab.
2. In the **InsightIQ Dashboard** view, locate the monitored cluster that you want to view information about.
3. Click the **Performance details** link that is located beside the cluster name.
The **Live Performance Reporting** view for the cluster appears.
4. (Optional) In the **Date Range** field, specify the start date and time and the end date and time for the period that you want to view.
5. Select the **Report Type** that you want to view from the list.
You can select a **Standard** report or a **Saved** report.
6. Click **View Report**.
7. Scroll down to view the report.

Results

Modify a live Performance report

Change the name and add or remove modules in a live Performance report.

Procedure

1. In the InsightIQ web application, click the **Performance Reporting** tab.
2. On the **Performance Reporting** ribbon, click **Manage Performance Reporting**.

The **Manage Performance Reporting** view appears.

3. Click the **Live Performance Reports** tab.

A list of the saved live performance reports is displayed in the **Live Performance Reporting** section.

4. Click **Edit** beside the name of the live performance report that you want to modify.

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

5. Specify the changes.

Option	Description
Change the name of the performance report	Type the new name in the Performance Report Name field.
Add a performance module	<ol style="list-style-type: none"> a. In the Select the Data You Want to See section, click Add another performance module. b. From the Select a Module for this Position list, select a performance module.
Modify a performance module	<ol style="list-style-type: none"> a. Locate the performance module that you want to modify. b. From the performance module list, select a new performance module.
Remove a performance module	<ol style="list-style-type: none"> a. Locate the performance module that you want to remove. b. In the performance module area, click Delete this performance module.

6. Complete changes to the performance 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 the report to a monitored cluster and view the results, click **Finish**.

Delete a saved live Performance report

Permanently delete a saved live performance report.

Procedure

1. In the InsightIQ web application, click the **Performance Reporting** tab.

- On the **Performance Reporting** ribbon, click **Manage Performance Reporting**.

The **Manage Performance Reporting** view appears and includes a list of all saved live performance reports.

- Click the **Live Performance Reports** tab.

A list of the saved live performance reports is displayed in the **Live Performance Reporting** section.

- Specify the report that you want to delete.

Option	Description
Delete one report	Beside the name of the report that you want to delete, click Delete .
Delete multiple reports at once	Beside the names of the reports that you want to delete, select the Report checkbox, and then, in the Select an action list, select Delete selected reports .
Delete all live Performance reports	Select the top-level checkbox that is next to the Report column title, and then, in the Select an action list, select Delete selected reports .

A **Confirm Delete** dialog box appears and prompts you to confirm that you want to delete the performance reports.

- Click **Delete**.

Create a Performance report schedule from a live Performance report

Create a performance report schedule based on a Live Performance report.

Procedure

- In the InsightIQ web application, click the **Performance Reporting** tab.
- On the **Performance Reporting** ribbon, click **Manage Performance Reporting**.

The **Manage Performance Reporting** view appears.

- Click the **Live Performance Reports** tab.

A list of the saved live performance reports is displayed in the **Live Performance Reporting** section.

- Beside the name of the live performance report that you want to schedule, click **Schedule this report**.

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

- To change the name of the live performance report schedule, update the text in the **Performance Report Name** field.
- To configure the scheduled performance report perform the following steps.
 - Select the name of the cluster in the **Subject of Report** list.
 - Select the time interval of data in the **Reporting Period** list.
 - Select the report generation interval in the **Report Frequency** list.
 - Configure the interval frequency by specifying a number in the **Generate Report Every** edit box.

- e. If the report frequency is **Weekly**, choose the day or days to generate the report each week.
- f. Configure the start time of the generated live performance reports by choosing a time in the **Generate Report at** list.
- g. Select **AM** or **PM**.
7. (Optional) To email each generated live performance report, select the **Email this report as a PDF attachment each time it is generated** checkbox.
Type up to 10 email addresses, which are separated by spaces, commas, or semicolons.
8. (Optional) Add, remove, or modify Performance modules to the generated live performance report schedule.
For more information, refer to the [Modify a live Performance report](#) task.
9. Click **Save**.

Publish reports

Publish a report as a PDF.

InsightIQ generates PDFs, sometimes referred to as a generated Performance report, at the times that are specified in a scheduled report. After the published report is complete, you can configure InsightIQ to send the PDF as an email attachment to up to 10 email recipients. You can download the PDFs or view the published reports in the InsightIQ web application.

View a published Performance report

You can view a PDF that is published for a performance 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 the 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.

Send a published report

You can manually send a PDF that is published for a Performance 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 the published reports.

3. In the **Generated Reports** area, specify the published reports that you want to send.

Option	Description
Send one published report	In the row for the published report that you want to send, click Send .
Send multiple reports in a single email	In the rows for the reports you that want to send, select the check boxes, and then, in the Select an action list, select Send Selected reports .
Send all reports in a single email	Select the top-level check box that is next to the Report column title, and then, in the Select an action drop-down list, select Send Selected reports .

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.
 - To specify multiple email addresses, separate each address with a comma, a space, or a semicolon.
 - To specify that these email addresses receive the latest version of the report each time it is published in the future, select the **Add each of these recipients** check box.
5. To specify the subject line for the email that will contain the PDF report, type the subject line in the **Message Subject** field.
6. To include a message with the email, type the message in the **Message** field.
7. Click **Send**.

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

8. Click **Close**.

Delete a published report

You can permanently delete a published 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 the published reports.

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

Option	Description
Delete one published report	Beside the name of the published report that you want to delete, click Delete .
Delete multiple published reports	Beside the names of the reports you want to delete, select the check boxes, and then, in the Select an action list, select Delete Selected reports .
Delete all published reports.	Select the top-level checkbox in the Generated Reports area, and then, in the Select an action list, click Delete selected reports .

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

- Click **OK**.

Publish a report on demand

Manually publish a scheduled performance report.

Procedure

- In the InsightIQ web application, click the **Performance Reporting** tab and then click **Manage Performance Reporting** on the **Performance Reporting** ribbon.

The **Manage Performance Reporting** view appears.

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

The **Scheduled Performance Reporting** view appears.

- In the **Report Schedules** table, in the **Actions** column, click **Generate** for the report that you want to publish.

The **Generate** link is unavailable while InsightIQ publishes the report.

- In the **Manage Performance Reporting** view, click the **Generated Reports Archive** tab.

Results

The published report is listed in the **Generated Reports** table.

Scheduling Performance reports

Create, disable, enable, modify, and delete Performance report schedules.

You can configure a Performance report schedule to generate PDF reports.

Performance report schedules specify what is in the reports, when the reports are generated, and, optionally, one or more email recipients. For each schedule, you specify when InsightIQ generates the report.

You can generate a Performance report that is based on a report schedule at any time. All reports are archived and can be viewed and sent to email recipients. InsightIQ does not automatically delete generated Performance reports. A user must manually delete unwanted Performance reports.

Use Performance report schedules to track cluster performance over a range of dates. For example, if you want to track the remaining unused capacity for a cluster, you can create a Performance report schedule that reports the capacity data for the cluster.

Note

Performance reports are based on the time settings of the computer or virtual machine that is running InsightIQ, not the time settings on the monitored cluster.

Create a scheduled Performance report

Create a performance report schedule and specify when InsightIQ publishes PDF reports.

You can create a performance report schedule to send the published performance report to one or more email recipients.

Procedure

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

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

2. Specify a template for the report schedule.

Option	Description
Create a report schedule from a template based on the default settings.	Click Create from Blank Report .
Create a report schedule from a template based on a user-created Performance report.	In the Saved Reports area, click the name of a report.
Create a report schedule from a template based on one of the standard reports included with InsightIQ.	In the Standard Reports 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** checkbox.

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** checkboxes.

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 January 10, 2017, the report includes data that is collected from 10:00 AM on January 9, 2017 through 10:00 AM on January 10, 2017.

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

Option	Description
Generate one or more reports every day	Select Hourly , specify the number of hours to elapse before the next report is generated, and specify the time when each report is generated for the first time each day. For example, if you configure InsightIQ to generate a report every 7 hours starting at 1:30 PM, the report is generated daily at 1:30 PM and 8:30 PM.
Generate no more than one report per day. Optionally, suspend report generation for a number of days.	Select Daily , specify the number of days that elapse before the next report is generated, and specify the time of day when the report is generated.
Generate no more than one report per day. Optionally, suspend report generation for a number of weeks.	Select Weekly , specify the number of weeks that elapse before the next report is generated, and specify one or more days of the week and the time of day when the report is generated.
Generate no more than one report per month	Select Monthly , specify the number of months that elapse before the next report is generated, and specify the day of the month on which the report is generated. Reports are always generated at 11:59 PM on the specified day.

8. If you want to send reports that are 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** checkbox.
 - 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 modules that you want to view in the report.

Option	Description
Add module	<ol style="list-style-type: none"> a. If all existing module areas are assigned, click Add another performance module. <hr/> <p>Note</p> <p>Unassigned modules contain a Select a Module for this Position list.</p> <hr/> <ol style="list-style-type: none"> b. In an unassigned module area, from the Select a Module for this Position list, select a module.
Modify a module	<ul style="list-style-type: none"> • In an assigned module area, from the module list, select a different module.

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

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

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

11. Save the report schedule.

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

Disable a scheduled Performance report

Disable a scheduled Performance report to temporarily prevent InsightIQ from publishing Performance reports.

You can disable a scheduled Performance report if you want to temporarily prevent InsightIQ from publishing Performance reports, but do not want to permanently delete the report. A disabled report can be enabled so that InsightIQ publishes Performance reports 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** view appears. A list of all saved report schedules is displayed in the **Report Schedules** table.

3. In the Report Schedules table, specify which report schedule you want to disable.

Option	Description
Disable a single report schedule.	In the row of the report schedule you want to disable, click the black double-bar icon in the Next Run (UTC) column.
Disable multiple report schedules.	In the rows of the report schedules you want to disable, select the check boxes, and then, in the Select an action list, select Pause selected report schedules .
Disable all report schedules.	In the header row, select the check box next to Report, and then, in the Select an action list, select Pause selected report schedules .

In the **Next Run (UTC)** column of the report schedules, the status changes to Paused. The icon changes to a right pointing triangle and InsightIQ does not generate reports for the specified schedules.

Enable a scheduled Performance report

Enable a scheduled Performance report.

Once a Performance report has been enabled, InsightIQ resumes Performance report publishing.

Procedure

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

The **Manage Performance Reporting** view 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 **Report Schedules** table, specify which report schedule you want to enable.

Option	Description
Enable a single report schedule.	In the row of the report schedule you want to enable, click the black right-pointing triangle icon in the Next Run (UTC) column.
Enable multiple report schedules.	In the rows of the report schedules you want to enable, select the checkboxes. In the Select an action list, select Resume selected report schedules .
Enable all report schedules.	In the header row, select the checkbox next to Report . In the Select an action list, select Resume selected report schedules .

The next date that InsightIQ will publishes a report appears in the **Next Run (UTC)** column of the schedules.

Modify a scheduled Performance report

Add or remove data elements from a Performance report schedule.

You can modify the scheduled Performance report to add or remove modules and change the list of email message recipients.

Procedure

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

The **Manage Performance Reporting** view appears with the **Live Performance Reporting** tab shown.

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 scheduled report that you that want to modify, click **Edit**.

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

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

Delete a scheduled Performance report

Permanently delete a scheduled Performance report.

If you no longer want InsightIQ to publish PDF reports on a configured report schedule, you can permanently delete the scheduled Performance report. Deleted

scheduled reports cannot be recovered. Any reports that InsightIQ publishes are not be deleted as a result of deleting the scheduled 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 **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.

Option	Description
Delete one report schedule	In the row for the report schedule that you want to delete, click Delete .
Delete multiple report schedules	In the rows for the report schedules that you want to delete, select the checkboxes, and then, in the Select an action list, select Delete selected report schedules .
Delete all report schedules	Select the highest checkbox in the Scheduled Reports area, and then, in the Select an action list, click Delete selected report schedules .

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

4. Click **Delete**.

Troubleshooting scheduled Performance report issues

Tips on troubleshooting scheduled InsightIQ performance reports.

Note

If InsightIQ 4.1.3 is monitoring a OneFS 8.2.0 cluster, some reports might not display or filter data properly.

Problems generating a scheduled Performance report

Review the **Report Schedules** table, found on the **Scheduled Performance Reports** tab on the **Manage Performance Reports** view. This table includes current information about saved report schedules, including a **Status** column that describes the current state of the report schedule.

If InsightIQ is not publishing reports from a report schedule or the published reports are not reaching the specified email recipients, verify the schedule settings.

If InsightIQ is not publishing reports at the times that are specified in a report schedule, InsightIQ might be publishing reports that are based on an incorrect time zone. To correct this problem, update the InsightIQ time zone to the correct time zone.

Problems delivering reports by email

If the specified email recipients do not receive a published report, perform the following actions.

- Verify the email address.
- Configure InsightIQ email settings. InsightIQ might not be authorized to send outbound email through the organization's email server.
- Reduce the size of the published reports by removing breakouts from the report schedule. Many email servers reject email messages that are larger than a certain limit. The email server might be rejecting the reports because they are too large.
- Configure the email program to allow InsightIQ to send email messages. The email program might be filtering email messages that InsightIQ sends.

Interpreting Performance reports

The following sections include recommendations for interpreting modules in Performance reports.

Active Clients report

The Active Clients report shows the number of unique client addresses that generate protocol traffic on the monitored cluster.

Clients that are connected, but not generating any traffic, are not counted in this report. You can optionally break out this data by node or protocol.

The recommended number of active clients that are supported for NFS is approximately 1,000 per node (OneFS 7.1.1 and earlier).

For SMB 2.0 and later, the number of active clients that are supported depends on the version of OneFS.

Table 3 Active clients by OneFS version

OneFS version	Number of active clients per node (SMB 2.0 or later)
7.0 Through 7.1	1,000
7.1.1 and later	1,500

Depending on usage details, exceeding the recommendations in OneFS 7.1.0 and earlier can lead to dropped connections or other performance problems. Generally speaking, if the number of active connections exceeds 90 percent of the limits, add nodes to the cluster when possible.

In OneFS 7.1.1 and later, over allocation is handled differently. Connection requests are denied when the total number of clients equals 90 percent of the total recommended limit for the cluster. If there many active clients, and if that quantity is near the supported limit, check each node to ensure that the load is evenly distributed. Verify that any one node is not oversubscribed.

Keep these limits in mind when you perform maintenance on a node so that you do not exceed the pool or cluster limit when you take nodes offline.

Average Pending Disk Operations Count report

The Average Pending Disk Operations Count report shows the average number of operations or transfers that are in the processing queue for each disk in the cluster.

You can break out this data by node or disk.

The graph focuses of on the queue depth of disk operations. The graph is most meaningful when viewing the disk detail breakout. Disk pending operation counts of 10

or less usually mean that the drives are operating as intended. Counts higher than 10 could indicate backlogs in disk operations, sometimes resulting in performance issues with the cluster. Viewing the disk breakout could identify a drive or node with excessively high counts that are compared to the others, which may indicate a hardware or configuration issue. If all the nodes/drives have high counts, the workflow is most likely spindle-bound.

External Network Throughput Rate report

The External Network Throughput Rate report shows the total amount of data that passed through the external network interfaces in the monitored cluster.

The information is straightforward to interpret, though it is more helpful to break out specifics by interface, direction, client, operation class, protocol, or node. For example, if you have an application that uses a particular SmartConnect zone, you can view the throughput on those nodes to observe that application's performance. This information can be useful, but it is not a good indicator of potential performance problems. There are many factors that affect the actual throughput rates over the network, many of which are unrelated to the cluster itself.

Overall Cache Hit Rate report

The Overall Cache Hit Rate report shows the percentage of data requests that returned hits.

A hit means that the requested data was available in cache. Higher hit rates indicate that more of the requested data is being retrieved from cache.

In the graph, you can plot data for the L1, L2, and L3 caches, or any combination of the three.

Table 4 Cache descriptions

Cache	Description
L1	Private to the node servicing input-output requests
L2	Global across all nodes
L3	Read only, data is stored on a solid state drive (SSD) instead of RAM

Break out the data for a selected cache by cache data type and node. If you want to determine if input-output is balanced across nodes or node pools, examining node details can be helpful.

Also, InsightIQ provides specific charts for throughput rates of the individual L1, L2, and L3 caches.

Overall Cache Throughput Rate report

The Overall Cache Throughput Rate report shows the amount of data that was requested from cache and how much of the requested data was available in cache.

In this chart, you can plot data for starts and hits on the L1, L2, and L3 caches (or any combination of the three). A start refers to the amount of data requested from a given cache, while a hit refers to the amount of data available in a given cache. You can also break out the data starts or hits for a particular cache, and, in addition, by cache data type or node for that cache.

Table 5 Cache by Type

Cache	Description
L1	Private to the node servicing input-output requests
L2	Global across all nodes
L3	Read only; stored on SSD instead of RAM

By comparing this graph to the disk, protocol, and network throughput rates, you can determine the actual impact of cache throughput on system performance.

InsightIQ also provides specific graphs for throughput rates of the individual L1, L2, and L3 caches.

Pending Disk Operations Latency report

The Pending Disk Operations Latency report shows the average amount of time that disk operations spend in the input-output scheduler.

You can also break out this information by node or disk, which yields more useful information than the aggregate values. The graph represents the amount of time that a disk operation remains in the input-output scheduler, not the actual latency of a disk operation (which is part of the total latency in every input-output request). Long times in the scheduler certainly indicate a busy system, and this may indicate the need to look more closely at cluster statistics to identify potential performance bottlenecks. In general, total latencies in the input-output system should be under 10 milliseconds.

Protocol Operations Average Latency report

The Protocol Operations Average Latency report displays the average amount of time that is required for protocols to process incoming operations.

You can also break out this information by node, protocol, client, or operation class. The Protocol Operations Average Latency report can be a helpful report, especially when you break out the details. For example, view the latency by protocol, then at the operation class for each protocol, to check for specific operations that might be causing excessive latency.

For the SMB protocol, the Change Notify operation can be the source of high latency. You might also see that namespace operations introduce significant latency, which can be alleviated with SSDs. HTTP transactions also exhibit high latency and skew the results. Filter out HTTP traffic, unless you are specifically interested in that data.

CHAPTER 5

File System reports

This section contains the following topics:

- [File System reports](#)..... 46
- [Capacity reports overview](#) 51
- [Deduplication reports](#)..... 52
- [File System reports overview](#) 53
- [Quota reports](#)..... 54

File System reports

View information about the cluster file system using File System reports.

There are four types of File System reports. These reports include Capacity reports, File System Analytics reports, Deduplication reports, and Quota reports.

File System Analytics result set snapshots

You can configure whether InsightIQ generates File System Analytics result sets from snapshots or the live file system. If you generate result sets from snapshots, the monitored cluster creates snapshots specifically for InsightIQ.

Generating result sets from snapshots enables you to view data that represents an exact point in time. However, it requires the monitored cluster to create snapshots, which consume cluster resources. Generating result sets from the live file system consumes less cluster resources; however, because the live file system is constantly changing, results sets generated from the live file system do not represent an exact point in time.

Note

Generating result sets from snapshots does not require you to activate a SnapshotIQ license on the monitored cluster.

File System report modules

Modules that are used in File System reports display information about the file system of a monitored cluster.

Note

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

Current Estimated Usable Capacity

Displays the amount data that can be added to the cluster before the cluster reaches capacity. The module displays the following values:

Allocated Capacity

The storage capacity of nodes that belong to node pools that include three or more nodes. In the OneFS command-line interface, Allocated Capacity is referred to as *size* in the output of the `isi_status` command. If all node pools include at least three nodes, the Allocated Capacity is the same as the Total Capacity.

Estimated Additional Protection Overhead

The estimated amount of data protection overhead that is required to protect data if the cluster reaches capacity.

Estimated Usable Capacity

The amount of data that can be added to the cluster before the cluster reaches capacity.

Overhead estimated based on FSA report

Specifies which FSA report to base the Estimated Additional Protection Overhead on. The estimated overhead is calculated based on the ratio of data protection that is required to protect user data at the time that the FSA report was created.

Remaining Capacity

The total amount of storage capacity that is unused on the cluster.

Reserved for Virtual Hot Spares

The amount of storage capacity that is reserved for virtual hot spares.

Snapshots Usage

The amount of data that is occupied by snapshots on the cluster.

Total Capacity

The total amount of storage capacity on the cluster.

Unallocated Capacity

The storage capacity of nodes that belong to node pools that include fewer than three nodes. These nodes are referred to as Unprovisioned Nodes in the OneFS web administration interface.

User Data Including Protection

The amount of storage capacity occupied by user data and protection for that user data.

Writable Capacity

The amount of storage capacity that user data can be written to on the cluster.

Capacity Forecast

Displays the amount data that can be added to the cluster before the cluster reaches capacity. This report module displays the following values.

Allocated Capacity

The storage capacity of nodes that belong to node pools that include three or more nodes. In the OneFS command-line interface, Allocated Capacity is referred to as *size* in the output of the `isi_status` command. If all node pools include at least three nodes, the Allocated Capacity is the same as the Total Capacity.

Estimated Additional Protection Overhead

The estimated amount of additional data protection overhead that is required to protect data if the cluster is filled to capacity.

Estimated Usable Capacity

The amount of data that can be added to the cluster before the cluster reaches capacity.

Forecast data

The breakout of information that is shown in the Forecast chart. The Forecast data includes the following information.

Calculation Range

The highlighted range of data that is used to calculate the forecast.

Forecast Usage

The projected Total Usage over time.

Standard Deviation

The standard deviation of the Forecast Usage calculation.

Outliers

Data points that are significantly outside of the range of the bulk of the Calculation Range data. Depending on the frequency and amount of variation, these points can have a major impact on the accuracy of the Forecast Usage data.

Overhead estimated based on FSA report

Specifies which FSA report to base the Estimated Additional Protection Overhead on. The estimated overhead is calculated based on the ratio of data protection that is required to protect user data at the time that the FSA report was created.

Provisioned Capacity

The storage capacity of nodes that belong to node pools.

Remaining Capacity

The total amount of storage capacity that is unused on the cluster.

Reserved for Virtual Hot Spares

The amount of storage capacity that is reserved for virtual hot spares.

Snapshots Usage

The amount of data occupied by snapshots on the cluster.

Total Capacity

The total amount of storage capacity on the cluster.

Total Usage

The total amount of user data and the associated Protection Overhead that is stored on the cluster.

User Data Including Protection

The amount of storage capacity that is occupied by user data and protection for that user data.

Writable Capacity

The amount of storage capacity that user data can be written to on the cluster.

Y-axis

The scale used on the Y-axis for the Forecast chart.

Percent

Forecast data as a percentage of total disk capacity.

Bytes

Forecast data expressed as actual bytes in use.

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. Resorting causes the chart on the left to reload and display a visual representation of the specified information. You can click pathnames 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.

Deduplication Current Summary (Physical)

Displays information about the current state of deduplication on the cluster. This module refers to the estimated physical space and data, meaning that file metadata and protection overhead are considered. The module displays the following values:

Deduplicated Data

The amount of data that has been deduplicated.

Space Saved

The amount of space that deduplication has saved on the cluster.

Space Used

The amount of space that deduplicated data currently occupies on the cluster.

Total Capacity

The total amount of storage capacity on the cluster.

Total Usage

The total amount of data that is stored on the cluster.

Note

The information that is displayed in this module always refers to the current state of the cluster and is not influenced by the selected date range.

Deduplicated Blocks per Job

Displays the number of blocks that were deduplicated compared to the number of blocks that were scanned by a given deduplication job.

File Count by Last Changed

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

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 changed time, node pool extension, or directory.

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 changed time, node pool, extension, or directory.

Files Scanned for Deduplication per Job

Displays the number of files that were scanned by a deduplication job.

Input/Output Bytes per Job

Displays the Input/Output byte throughput for a SmartDedupe job.

Input/Output Operations per Job

Displays the number of Input/Output operations that were performed while a SmartDedupe job was running.

Quota Browser

Deduplicated Data

The amount of data that has been deduplicated.

Paths

Displays whether a quota has been applied on or under a directory. The selected directory determines which quotas are displayed in the Quotas Applied On and Under directory module. This module reflects the current cluster configuration.

Quotas Applied On and Under Directory

Displays the quotas that have been applied on or under the selected directory. The **Select a report** list determines the point in time that this module represents, regardless of the Date Range selected.

Quota Limit Usage

Displays the usage limits for the selected quota along with the level of data usage for the directory that the quota is applied to. The **%Used** column shows the usage divided by the highest-level limit. If no hard limit is applied, the soft limit is used. If no soft limit is applied, the advisory limit is used. For example, assume that a quota has no hard limit, a soft limit of 2 TB, and an advisory limit of 1 TB. If the usage for the quota equaled 1 TB, the %Used would be 50%.

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.

File System report breakouts

Each breakout displays the contributions of individual component.

Accessed Time

Breaks out data by when a file was last accessed.

Directory

Breaks out data by the size of each directory.

File Extension

Breaks out data by file name extension.

Logical Size

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

Changed Time

Breaks out data by when a file was last changed.

Node Pool

Breaks out data by the size of each node pool.

Physical Size

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

Tier

Breaks out data by the size of each node tier.

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.

Capacity reports overview

Capacity reports show an estimate of available storage on a cluster.

Capacity reports show the estimated amount of data that can be added to a cluster before the cluster reaches capacity. This estimate is different than the physical free space that OneFS reports. The amount of physical free space does consider the amount of protection overhead that is required to protect the stored data. Because clusters have varied protection settings that can change over time, InsightIQ estimates the size of protection overhead that is required to protect data.

If the data protection changes on the cluster, the estimate of usable capacity changes, even if the data stored on the cluster remains constant. For example, assume that a cluster contains 3 TB of data. 1 TB is protected at +3, and the other 2 TB is protected at +2. If 2 TB of data is later protected at +3, and 1 TB is protected at +2, the estimated amount of protection overhead increases. This increase is because the ratio of user data to protection overhead changes.

You can choose the File System Analytics report that you want to create protection overhead estimates. Protection overhead is calculated based on the cluster configuration at the time that the job is run. This information can be useful if you expect the cluster configuration to resemble a previous configuration. If you do not create a File System Analytics report, protection overhead is not in Capacity reports.

If you view a Capacity report for clusters running OneFS 7.1.1 or earlier, the report does not include storage that is used by virtual hot spares or snapshots. Any remaining capacity appears greater than it actually is.

Capacity reports include the Cluster Capacity Performance module that can be viewed in a Live Performance report.

Note

Protection overhead estimates do not account for the `/ifs/.ifsvar` directory because the `/ifs/.ifsvar` directory is protected at a higher level than other directories. `/ifs/.ifsvar` does not reflect the level of the protection that is applied to the majority of the file system. Also, because the `/ifs/.ifsvar` directory stores only internal OneFS files, the directory is expected to scale at a much slower rate than other directories on the file system.

Viewing a Capacity report

View Capacity reports to see the estimate of when the cluster reaches capacity.

Procedure

1. Click **File System Reporting > Capacity Reporting**.
2. In the **Cluster** list, select the cluster that you want to view capacity information for.
3. Click **View Report**.

4. (Optional) From the **Overhead estimation based on FSA report** list, select the date of the File System Analytics report that you want to use as a basis for the Overhead Protection Estimation.

The framed control above the **Capacity Forecast** report module is the **Date Range** control.

5. (Optional) To configure the dates in the **Capacity Forecast** report module, in the **Date Range** control, specify a period and then click **View Report**.

The **Zoom Slider** is below the **Capacity Forecast** drop down control. The **Zoom Slider** shows a graph in light gray that represents the percentage of the total disk capacity in use over time for the date range that you typed in the **Date Range** control.

6. To calculate a **Capacity Forecast**, define a subset of data in the **Zoom Slider** by clicking and dragging the **Zoom Slider** handles.

The chart that is displayed in the **Capacity Forecast** report module is the **Capacity Forecast** chart. The **Capacity Forecast** chart shows a line graph of the data selected with the **Zoom Slider** handles.

7. To analyze a **Calculation Range**, drag on a region of the **Capacity Forecast** chart.

A light yellow highlight in the **Capacity Forecast** chart indicates the **Calculation Range**.

Results

Text below the **Capacity Forecast** chart provides a projected date when the cluster reaches critical capacity as well as details about the regression calculation.

Deduplication reports

Deduplication reports allow you to view historical and current information about data that has been deduplicated with the SmartDedupe software module.

You can view information about specific deduplication jobs or the current state of deduplication on the cluster. Deduplication job information is not cumulative. For example, if you run a deduplication job that deduplicates 10 GB of data, and then later run a deduplication job that deduplicates another 10 GB of data, both data points on the graph will report that only 10 GB of data was deduplicated.

To view deduplication reports, you must activate a SmartDedupe license on the monitored cluster. For more information about SmartDedupe, see the *OneFS Web Administration Guide* and *OneFS CLI Administration Guide*.

Note

Deduplication information is also included in Performance reports. In Performance reports, deduplication information is cumulative.

Viewing a Deduplication report

View information about the current state of deduplication on a monitored cluster.

Procedure

1. Click **File System Reporting > Deduplication Reporting**.
2. From the **Apply to cluster** list, select the name of the cluster that you want to view the report for.

3. In the **Date Range** section, select the days that you want to view deduplication information about.

Note

The Deduplication Summary module reflects the current state of the cluster. The specified date range does not affect the report.

4. Click **View Report**.

File System reports overview

File System reports include information about the files that are stored on a cluster.

The File System reports information includes the number of files that are stored on the cluster, the types of files that are stored on the cluster, and where files are located.

InsightIQ supports two kinds of File System Analytics reports: Data Usage reports and Data Properties reports. Data Usage reports include data about the individual files on a cluster, such as which files have remained unchanged for the longest amount of time. Data Properties reports include data about the overall cluster file system, such as how many files were last changed during a specific period. 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 used to manage storage space among multiple clusters. For example, if you have two clusters, you can view the Data Properties report to compare file utilization. You can compare that the files are changed regularly on one cluster and that the files on the second cluster are modified occasionally. If the Data Properties report reveals that a significant number of files on the cluster have not been accessed in a long time, you can view the Data Usage report to identify the inactive files. To archive storage, consider transferring the inactive files.

You can use a File System report to verify that the file system of a cluster is behaving as expected. For example, if you have a cluster that is designated for archiving data, you can use the data properties report to view the file count by last changed file system module. This report shows you the last change date of the files on the cluster. If the count of files that have been changed is particularly high, you can consider moving those files to a cluster intended for development.

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

Note

Do not apply a quota to the `/ifs/.ifsvvar/modules/fsa` directory through the SmartQuotas module. If you limit the size of the `/ifs/.ifsvvar/modules/fsa` directory through a quota, and the directory reaches its limit, the File System Analytics report process fails. A quota does not cause older File System reports to be deleted from the `/ifs/.ifsvvar/modules/fsa` directory to make room for newer reports.

Note

If your FSA report includes a directory name that has spaces, you must enclose the directory path in quotes.

Viewing a File System Analytics report

View and configure a File System Analytics report.

Before you begin

The File System Analytics (FSA) job must be enabled.

Procedure

1. Click **File System Reports > File System Analytics**.
The **File System Analytics** 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 that are 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**.
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 File System Analytics report.

Quota reports

Quota reports display information about quotas created through the SmartQuotas software module.

Quota reports can be useful if you want to compare the data usage of a directory to the quota limits for that directory over time. This information can help you predict when a directory is likely to reach its quota limit.

Historical quota data is generated according to the quota reports that are generated by OneFS. The granularity of historical quota data depends on how often quota reports are created. For information about configuring how frequent quota reports are generated, see the *OneFS Web Administration Guide* and *OneFS CLI Administration Guide*.

To view quota reports, activate a SmartQuotas license on the monitored cluster.

View a Quota report

View information about quotas on a monitored cluster.

Procedure

1. Click **File System Reporting > Quota Reporting**.
2. In the **Select Your Quota Report** area, from the **Apply to Cluster** list, select the cluster that you want to view quota information for.
 - The **Path** module displays directories on the cluster and indicates whether quotas have been applied on or under those directories. The **Path** module reflects the current cluster configuration.
 - The **Quotas Applied On and Under Path** module displays the quotas that have been applied on or under the path specified by the **Path** module. The **Quotas Applied On and Under Path** module reflects the time specified by the **Select a report list**.
 - The **Quota Limit Usage** module displays time-series data about the quota selected in the **Quotas Applied On and Under Path** module. The **Quota Limit Usage** module reflects the time specified by the **Date Range** section.
3. (Optional) To view quota limits and data usage for a specific directory over time, select a quota from the **Quotas Applied On and Under Path** module and then view the **Quota Limit Usage** module.
4. (Optional) To modify the time period shown in the **Quota Limit Usage** module, in the **Date Range** section, select the days that you want to view information for and then click **View Report**.

CHAPTER 6

Creating custom reports

This section contains the following topics:

- [Building reports overview](#).....58
- [Modules](#)..... 58
- [Breakouts](#).....63
- [Breakout types](#).....64
- [Applying a breakout](#)..... 65
- [Filters](#).....65

Building reports overview

Customize how information is shown in a report.

You can modify all reports based on a range of dates. You can create customized Performance reports by adding and removing Performance modules, breakouts, and filters.

Date ranges allow you to limit the data based on a start and end date. The date interval, also called a **Zoom Level**, can be automatically set to intervals of 30 minutes, 1 hour, 6 hours, 12 hours, 1 day, 3 days, 1 week, and 1 month. Custom date intervals can be created by specifying the beginning and ending day and time.

Performance modules are preconfigured collections of information about the monitored cluster. You can add and remove Performance modules to all Performance reports.

Breakouts divide a data set into component parts but do not limit the information presented.

Filters limit the information in a data set. With a filter, the data set is refined to simplify what is in a report. Filters can exclude data that can be repetitive or irrelevant. Different types of filters can be used to amend information results. A generated report often produces a large and complex data set. Filtering can make the report results more understandable. A filter can contain one or more filter rules. Filters can be saved and applied to more than one report.

Note

If InsightIQ 4.1.3 is monitoring a OneFS 8.2.0 cluster, the statistics, breakouts, filtering, and report data might not display properly.

Modules

Modules show collections of information about the cluster.

Modules are preconfigured collections of information about the monitored cluster. You can add and remove modules to all Performance reports.

Note

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

Table 6 Performance report module descriptions

Module	Description
Active Clients	Shows the number of unique client addresses that generate 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.

Table 6 Performance report module descriptions (continued)

Module	Description
	<p>Note</p> <p>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.</p>
Average Cached Data Age	Indicates the average amount of time that data has been in the L1 and L2 caches. Shorter times indicate that data is moving faster through cache and that older data is being replaced quickly.
Average Disk Hardware Latency	Shows 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	Shows 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	Shows 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	Shows the number of file blocking events occurring in the file system per second. You can optionally break out this data by path or node.
Cache Hits	Shows the number of hits to L2 and L3 cache per second. You can optionally break out this data by job, node, or service.
Cluster Capacity	<p>Shows the amount of storage on the cluster.</p> <ul style="list-style-type: none"> • Total Capacity - The total amount of storage capacity on the cluster. • Allocated Capacity - The storage capacity of nodes that belong to node pools that include three or more nodes. In the OneFS command-line interface, Allocated Capacity is referred to as "size" in the output of the <code>isi status</code> command. If all node pools include at least three nodes, the Allocated Capacity is the same as the Total Capacity. • Writable Capacity - The amount of storage capacity that user data can be written to on the cluster. • User Data Including Protection - The amount of storage capacity that is in use by user data and protection.
Connected Clients	<p>Shows 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.</p> <p>Note</p> <p>UDP connections do not appear as connected. Also, some short-lived TCP connections might not appear as connected, although they are active.</p>

Table 6 Performance report module descriptions (continued)

Module	Description
Contented File System Events Rate	Shows 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.
CPU % Use	<p>Shows 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.</p> <hr/> <p>Note</p> <p>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 January 10, 2017, the specified node was using 14.35% of the total available node CPU capacity.</p> <hr/>
CPU Usage Rate	Shows the rate of CPU time that is consumed on all cores per second. You can optionally break out this data by job, node, or service.
Deadlocked File System Events Rate	Shows the number of file system deadlock events that 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 detect and break them. You can optionally break out this data by path or node.
Deduplication Summary (Logical)	Shows the amount of space that deduplication has saved on the cluster and the amount of data that has been deduplicated. This module refers to the logical space and data. The file metadata and protection overhead are not considered. This module is available only for clusters running OneFS 7.1.1 or later.
Deduplication Summary (Physical)	Shows the amount of space that deduplication has saved on the cluster and the amount of data that has been deduplicated. This module refers to the estimated physical space and data. The file metadata and protection overhead are considered. This module is available only for clusters running OneFS 7.1.1 or later.
Disk Activity	Shows 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 IOPS	Shows the rate of disk read and write operations per second. You can optionally break out this data by job, node, or service.
Disk Operations Rate	Shows 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	Shows 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.

Table 6 Performance report module descriptions (continued)

Module	Description
External Network Errors	<p>Shows the number of errors that are generated for the external network interfaces. You can optionally break out this data by direction, interface, or node.</p> <hr/> <p>Note</p> <p>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 the network switch and other network infrastructure components.</p>
External Network Packets Rate	Shows 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	Shows 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	Shows 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 Throughput Rate	Shows the rate at which data is being read from and written to the file system.
Job Workers	Shows the number of active and assigned workers on the cluster. An active worker is a worker that is 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. You can optionally break out this data by job name or job ID.
Jobs	Shows the number of active and inactive jobs on the cluster. An active job is a system job that workers perform. An inactive job is a system job that has been assigned workers, but the workers are not currently performing the job. You can optionally break out this data by job name or job ID.
L1 and L2 Cache Prefetch Throughput Rate	<p>Shows the amount of data that was prefetched for L1 and L2 and how much of the prefetched data was requested.</p> <ul style="list-style-type: none"> Starts - Indicates the amount of data that was requested. Hits - Indicates the amount of requested data that was available.
L1 Cache Throughput Rate	<p>Shows the amount of data that was requested from the L1 cache and how much of the requested data was available in the L1 cache.</p> <ul style="list-style-type: none"> Starts - Indicates the amount of data that was requested. Hits - Indicates the amount of requested data that was available.

Table 6 Performance report module descriptions (continued)

Module	Description
	<ul style="list-style-type: none"> • Waits - Indicates the amount of requested data that existed in cache but was not available because the data was in use. • Misses - Indicates the amount of requested data that did not exist in cache. • Prefetch Hits - Indicates the amount of data that was requested from prefetch.
L2 Cache Throughput Rate	<p>Shows the amount of data that was requested from the L2 cache and how much of the requested data was available in the L2 cache.</p> <ul style="list-style-type: none"> • Starts - Indicates the amount of data that was requested. • Hits - Indicates the amount of requested data that was available. • Waits - Indicates the amount of requested data that existed in cache, but was not available because the data was in use. • Misses - Indicates the amount of requested data that did not exist in cache. • Prefetch Hits - Indicates the amount of data that was requested from prefetch.
L3 Cache Throughput Rate	<p>Shows the amount of data that was requested from the L3 cache and how much of the requested data was available in the L3 cache.</p> <ul style="list-style-type: none"> • Starts - Indicates the amount of data that was requested. • Hits - Indicates the amount of requested data that was available. • Misses - Indicates the amount of requested data that did not exist in cache.
Locked File System Events Rate	<p>Shows the number of file lock operations occurring in the file system per second. You can optionally break out this data by path or node.</p>
Overall Cache Hit Rate	<p>Shows the percentage of data requests that returned hits. A hit means that the requested data was available in cache. Higher hit rates indicate that more of the requested data is being retrieved from cache.</p> <ul style="list-style-type: none"> • L1 Hit Rate - Indicates the hit rate for L1 cache. • L2 Hit Rate - Indicates the hit rate for L2 cache. • L3 Hit Rate - Indicates the hit rate for L3 cache.
Overall Cache Throughput Rate	<p>Shows the amount of data that was requested from cache and how much of the requested data was available in cache.</p> <ul style="list-style-type: none"> • L1 Starts - Indicates the amount of data that was requested from the L1 cache. • L1 Hits - Indicates the amount of requested data that was available in the L1 cache. • L2 Starts - Indicates the amount of data that was requested from the L2 cache.

Table 6 Performance report module descriptions (continued)

Module	Description
	<ul style="list-style-type: none"> L2 Hits - Indicates the amount of requested data that was available in the L2 cache. L3 Starts - Indicates the amount of data that was requested from the L3 cache. L3 Hits - Indicates the amount of requested data that was available in the L3 cache.
Pending Disk Operations Latency	Shows 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 Operations Average Latency	Shows the average amount of time that is 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	Shows the total number of requests per client 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.
Slow Disk Access Rate	Shows the rate at which slow, or long-latency, disk operations occur. You can optionally break out this data by node or disk.

Breakouts

Refine information in reports with a breakout.

Modules in reports can provide a breakout of data by category to refine the scope of information.

You can apply breakouts to modules to view the individual contributions of various performance characteristics. You apply only one breakout to a module at a time.

Breakouts provide heat maps that display variations of color to represent each component's contribution to overall performance. The darker the color on a heat map, the greater the activity for that component. Heat maps help you to visualize performance trends and to identify periods of constrained performance. If you hover the mouse pointer over any location on a heat map, InsightIQ shows data for the specified component at that moment in time. Breakouts are sorted by component based on level of activity, with the most active elements at the top of the list.

Breakouts can be useful when trying to determine the source of a performance issue. For example, if you break out the CPU usage module by node, you can then view the individual CPU usage of each node.

Note

The sum of individual breakouts might not add up to the aggregated values of a specified category. For example, not all network traffic is associated with a protocol. Therefore, the sum of the various protocols might not match the total count of reported traffic.

Breakout types

Breakouts provide information on the individual components within modules.

Table 7 Breakout descriptions

Breakout	Description
Cached Data Type	Reports data by user data and metadata.
Client	Reports data by individual client. Only clients with the most activity are represented in the data that is displayed.
Direction	Reports data by the direction of cluster throughput. Only available in performance modules.
Disk	Reports data by individual disk.
Event	Reports data by individual network interface.
Interface	Reports data by individual network interface.
Job ID	The numeric ID of a job.
Job Type	The name of a job type.
Node	Reports data by individual node. Nodes are identified by their logical node number (LNN).
Node (devid)	Reports data by individual node. Nodes are identified by their device ID.
Node Pool	Breaks out data by node pool. <div> Note <ul style="list-style-type: none"> This breakout is available only for clusters running OneFS 7.1 or later. For clusters running OneFS 7.1.1 or earlier, accelerator nodes and unprovisioned nodes are grouped under the <code>Unclassified</code> node pool. </div>
Op Class (Operation Class)	Reports data by the type of operation performed. The following operation types are supported: <ul style="list-style-type: none"> 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

Table 7 Breakout descriptions (continued)

Breakout	Description
	<ul style="list-style-type: none"> • Namespace Read - attribute, statistic, ACL read, lookup, and directory read • Namespace Write - rename, set attribute, set permission, time, and write ACL • File State - open, close, lock, acquire, release, break, check, and notify. • Session State - negotiate, inquire, and change protocol connections and sessions • Other - file system information and other operations that are not categorized.
Path	Reports data by individual file name or directory name.
Protocol	Reports data by protocol (for example, NFS).
Service	Reports data by service.
Tier	Reports data by tier.

Applying a breakout

Apply a breakout to a report.

Procedure

1. In the InsightIQ web application, open any **Performance** report or a **File System Analytics** report.
2. To apply a breakout to a module, in the **Breakout by** list, click the name of a breakout.

The module shows the information broken out by component.

Filters

Refine a large data set with a filter.

You can create and manage filters. Apply filters to any Performance report and File System Analytics reports to provide information that is limited to specified parameters. You can create and apply filters, which contain one or more filter rules.

If you apply a filter rule to a report, all the Performance modules in the report display information that is constrained by the filter rule.

For example, you can apply a filter rule to show information about a specific node in a cluster. All the Performance modules in the report display data about that node.

While breakouts show all components in a category, filters can limit the number of contributors in a category. For example, you can breakout a Performance module by protocol. Breakouts appear beneath Performance modules but do not modify the Performance module. Filter rules limit what data is displayed in Performance modules. Filters are customized collections of filter rules that you can create, save, and apply to various reports.

A filter can 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 causes all Performance modules in a report to display information about the interactions for only that node and client. You can save filter rules and then apply them to specific reports.

Filter rules

Rules for InsightIQ filters.

To define a filter rule, specify the match parameter. For some rules, InsightIQ automatically generates a list of valid parameters. For other filters, type a valid parameter.

Note

If InsightIQ 4.1.3 is monitoring a OneFS 8.2.0 cluster, the statistics, breakouts, filters, and report data might not display properly.

Input values for the following Data Filter rules.

Table 8 Data Filter rules that require an input parameter

Filter	Match value
Client	The IP address of a client. To filter by all clients exceeding the top 1000 for a particular Performance module, specify Other clients
Node Pool	The name of a node 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	<p>The file name extension. For example <code>.xml</code> To filter files without file name extensions, specify the following:</p> <div>(none)</div> <p>The backslash character (<code>\</code>) is not valid.</p>
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 that is defined on the cluster. All characters are valid. Only available in the File System reports.

You can select conditional values for the following Filter rules.

Table 9 Filter rules that allow conditional values

Filter	Condition values	Description
Accessed Time	0 sec - 1 min	Includes only files that were accessed 0 seconds to 1 minute ago.
	1 min - 1 hour	Includes only files that were accessed 1 minute to 1 hour ago.
	1 hour - 1 day	Includes only files that were accessed 1 hour to 1 day ago.
	1 day - 7 days	Includes only files that were accessed 1 day to 7 days ago.
	7 days - 30 days	Includes only files that were accessed 7 days to 30 days ago.
	30 days - 60 days	Includes only files that were accessed 30 days to 60 days ago.
	60 days - 90 days	Includes only files that were accessed 60 days to 90 days ago.
	90 days - 180 days	Includes only files that were accessed 90 days to 180 days ago.
	180 days - 1 year	Includes only files that were accessed 180 days to 1 year ago.
	1 year - 2 years	Includes only files that were accessed 1 year to 2 years ago.
	2 years - Infinity	Includes only files that were accessed 2 years ago and older.
Direction	In	Includes only information that is related to inbound cluster traffic.
	Out	Includes only information that is related to outbound cluster traffic.
Disk	<devid>/Bay <number>	Includes only information that is related to the specified disk.
FS Event	blocked	Includes only information that is related to blocked file system events.
	contended	Includes only information that is related to contended file system events.
	deadlocked	Includes only information that is related to deadlocked file system events.
	getattr	Includes only information that is related to get attribute events.
	link	Includes only information that is related to linking events.
	lock	Includes only information that is related to locking events.

Table 9 Filter rules that allow conditional values (continued)

Filter	Condition values	Description
	lookup	Includes only information that is related to lookup events.
	read	Includes only information that is related to read events.
	rename	Includes only information that is related to rename events.
	setattr	Includes only information that is related to set attribute events.
	unlink	Includes only information that is related to unlinking events.
	write	Includes only information that is related to write events.
Interface	<i><devid>/ <type></i>	Includes only information that is related to the specified interface.
Job Type	AVScan	Includes only information that is related to the AVScan job engine job.
	AutoBalance	Includes only information that is related to the AutoBalance job engine job.
	AutoBalanceLin	Includes only information that is related to the AutoBalanceLin job engine job.
	ChangelistCreate	Includes only information that is related to the ChangelistCreate job engine job.
	Collect	Includes only information that is related to the Collect job engine job.
	Dedupe	Includes only information that is related to the Dedupe job engine job.
	DedupeAssessment	Includes only information that is related to the DedupeAssessment job engine job.
	DomainMark	Includes only information that is related to the DomainMark job engine job.
	FSAnalyze	Includes only information that is related to the FSAnalyze job engine job.
	FlexProtect	Includes only information that is related to the FlexProtect job engine job.
	FlexProtectLin	Includes only information that is related to the FlexProtectLin job engine job.
	IntegrityScan	Includes only information that is related to the IntegrityScan job engine job.
	MediaScan	Includes only information that is related to the MediaScan job engine job.

Table 9 Filter rules that allow conditional values (continued)

Filter	Condition values	Description
	MultiScan	Includes only information that is related to the MultiScan job engine job.
	PermissionRepair	Includes only information that is related to the PermissionRepair job engine job.
	QuotaScan	Includes only information that is related to the QuotaScan job engine job.
	SetProtectPlus	Includes only information that is related to the SetProtectPlus job engine job.
	ShadowStoreDelete	Includes only information that is related to the ShadowStoreDelete job engine job.
	ShadowStoreProtect	Includes only information that is related to the ShadowStoreProtect job engine job.
	SmartPools	Includes only information that is related to the SmartPools job engine job.
	SnapRevert	Includes only information that is related to the SnapRevert job engine job.
	SnapshotDelete	Includes only information that is related to the SnapshotDelete job engine job.
	TreeDelete	Includes only information that is related to the TreeDelete job engine job.
	WormQueue	Includes only information that is related to the WormQueue job engine job.
Logical Size	0 B - 8 KB	Includes only files that have a logical size of 0 B to 8 KB.
	8 KB - 128 KB	Includes only files that have a logical size of 8 KB to 128 KB.
	128 KB - 1 MB	Includes only files that have a logical size of 128 KB to 1 MB.
	1 MB - 10 MB	Includes only files that have a logical size of 1 MB to 10 MB.
	10 MB - 100 MB	Includes only files that have a logical size of 10 MB to 100 MB.
	100 MB - 1 GB	Includes only files that have a logical size of 100 MB to 1 GB.
	1 GB - 10 GB	Includes only files that have a logical size of 1 GB to 10 GB.
	10 GB - 100 GB	Includes only files that have a logical size of 10 GB to 100 GB.

Table 9 Filter rules that allow conditional values (continued)

Filter	Condition values	Description
	100 GB - 1 TB	Includes only files that have a logical size of 100 GB to 1 TB.
	1 TB - ∞	Includes only files that have a logical size of 1 TB to ∞ .
Changed Time	0 sec - 1 min	Includes only files that have been changed 0 seconds to 1 minute ago.
	1 min - 1 hour	Includes only files that have been changed 1 minute to 1 hour ago.
	1 hour - 1 day	Includes only files that have been changed 1 hour to 1 day ago.
	1 day - 7 day	Includes only files that have been changed 1 day to 7 day ago.
	7 days - 30 days	Includes only files that have been changed 7 days to 30 days ago.
	30 days - 60 days	Includes only files that have been changed 30 days to 60 days ago.
	60 days - 90 days	Includes only files that have been changed 60 days to 90 days ago.
	90 days - 180 days	Includes only files that have been changed 90 days to 180 days ago.
	180 days - 1 year	Includes only files that have been changed 180 days to 1 year ago.
	1 year - 2 years	Includes only files that have been changed 1 year to 2 years ago.
	2 years - Infinity	Includes only files that were accessed 2 years ago and older.
Node	<i><integer></i>	Includes only information that is related to the node with the specified LNN.
Node (Device ID)	<i><integer></i>	Includes only information that is related to the node with the specified device ID.
Op Class	create	Includes only information that is related to file create operations.
	delete	Includes only information that is related to file delete operations.
	file_state	Includes only information that is related to file access operations, such as open, close, lock, and unlock.
	namespace_read	Includes only information that is related to namespace read operations

Table 9 Filter rules that allow conditional values (continued)

Filter	Condition values	Description
	namespace_write	Includes only information that is related to write operations.
	other	Includes only information that is related to operations that do not fit in to any other operation class category.
	read	Includes only information that is related to file read operations.
	session_state	Includes only information that is related to operations that control session access, such as connect, disconnect, login, and logout.
	write	Includes only information that is related to file write operations.
Physical Size	0 B - 8 KB	Includes only files that have a logical size of 0 B to 8 KB.
	8 KB - 128 KB	Includes only files that have a logical size of 8 KB - 128 KB.
	128 KB - 1 MB	Includes only files that have a logical size of 128 KB - 1 MB.
	1 MB - 10 MB	Includes only files that have a logical size of 1 MB - 10 MB.
	10 MB - 100 MB	Includes only files that have a logical size of 10 MB - 100 MB.
	100 MB - 1 GB	Includes only files that have a logical size of 100 MB - 1 GB.
	1 GB - 10 GB	Includes only files that have a logical size of 1 GB - 10 GB.
	10 GB - 100 GB	Includes only files that have a logical size of 10 GB - 100 GB.
	100 GB - 1 TB	Includes only files that have a logical size of 100 GB - 1 TB.
	1 TB - ∞	Includes only files that have a logical size of 1 TB - ∞.
Protocol	ftp	Includes only information that is related to FTP.
	hdfs	Includes only information that is related to HDFS.
	http	Includes only information that is related to HTTP.
	nfs	Includes only information that is related to NFS.

Table 9 Filter rules that allow conditional values (continued)

Filter	Condition values	Description
	nfs3	Includes only information that is related to NFS3.
	nfs4	Includes only information that is related to NFS4.
	nlm	Includes only information that is related to NLM.
	papi	Includes only information that is related to PAPI.
	siq	Includes only information that is related to SynclQ.
	smb	Includes only information that is related to SMB.
	smb1	Includes only information that is related to SMBv1.
	smb2	Includes only information that is related to SMBv2.

Create a filter

Create a filter that consists of one or more rules.

Specify the rules to include in a filter, then save the filter and apply it to reports.

Procedure

1. On the **Live Performance Reporting** view, below the **Date Range** control, click **Create/manage data filters**.

The **Data Filters** control appears.

2. On the **Data Filters** control toolbar, click **Add Rule**.
3. From the **Type** drop-down list, select an attribute type that you want to filter on.
4. In the **Match** field, select or type a value for the attribute you selected.

Repeat steps 2 through 4 as needed until you have built and applied all the rules that you want to include in the filter.

5. Click **Apply**.

The filter is applied to all Performance modules in the report. Each rule is indicated

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

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

7. 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 Performance modules in the report.

Modify a filter

Modify and remove the filter rules that are contained in a filter.

Procedure

1. On the **Live Performance Reporting** view, click **Create/manage data filters**.
The **Data Filters** view 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 Performance modules in the report.

3. Modify filter settings as needed.

Option	Description
Add a filter rule.	<ol style="list-style-type: none"> a. Click Add Rule. b. From the Type list, select the breakout that you want to filter by. c. In the Match field for the filter rule, select or type a filter value.
Modify a filter rule.	<ol style="list-style-type: none"> a. In the Match field for the filter rule that you want to modify, select or type a filter value.
Remove a filter rule.	<ol style="list-style-type: none"> a. Click either the Type list or Match field for the filter rule you want to remove. b. Click Delete Rule.

4. Save the filter.

Option	Description
Save the filter under its existing name.	<ol style="list-style-type: none"> a. Click Apply. b. On the Manage menu, click Save. The Save Filter dialog box appears. c. Click Yes.
Save a new instance of the filter with a different name.	<ol style="list-style-type: none"> a. On the Manage menu, click Save as. The Save Filter dialog box appears. b. In the Please enter a filter name field, type a name for the filter, and then click OK.

Delete a filter

Permanently delete a filter.

Procedure

1. On the **Live Performance Reporting** view, click **Create/manage data filters**.
The **Data Filters** view 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 Performance modules in the report.
3. On the **Manage** menu, click **Delete**.

Note

To remove an individual 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**.

CHAPTER 7

InsightIQ export data

This section contains the following topics:

- [Export data overview](#)..... 76
- [Export InsightIQ data](#)..... 76
- [Permalinks](#)..... 76
- [Create a Permalink to share a report](#)..... 77

Export data overview

Use InsightIQ data with a third-party application.

You can export InsightIQ data if you want to access the data with 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.

For Performance reports, InsightIQ averages data before displaying it in the InsightIQ web application. However, when you export InsightIQ data, InsightIQ does not average the data before creating the CSV file. Because of this, exported datapoints for Performance reports might not exactly match the data as it appears in the InsightIQ web interface.

You can export data from Performance reports and File System Analytics reports. Exporting data from other reports is not supported.

CSV export format

The first two columns describe the data points displayed in the Performance module chart. The first column is the timestamp of when the sample period for that data point began. The second value describes the value for the sample period.

If you included minimum and maximum values, the next two columns describe the minimum and maximum values. The first column is the timestamp of when the ten-minute sampling period began. The second value describes the minimum value during the sample period. The third value describes the maximum value during the sample period.

If you included breakouts, the remaining columns describe the breakouts. Each breakout contains two columns. The first column describes the timestamp of when the sampling period of the data beginning point. The second value describes the value of the breakout during the sample period.

Export InsightIQ data

Export InsightIQ data to a CSV file.

Procedure

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

Permalinks

Share an InsightIQ report by using a permalink.

A permalink is a permanent link to an InsightIQ report. The permalink allows you to save or share an InsightIQ report with a specific configuration.

When you create a permalink, InsightIQ creates and displays a URL that represents a shortcut to the configured report. You can then use this URL at any time to view this report configuration.

For example, you want to share a report about the external network throughput of Node 5 from 11:20am to 12:20pm on February 15, 2001. You create the report and then create a permalink of the report to send to other InsightIQ users so they can view the information.

Note

When you use a permalink to access a report, any report that contains high-resolution data will degrade over time due to the InsightIQ data-retention and downsampling policies.

Create a Permalink to share a report

Create a Permalink.

Procedure

1. In the InsightIQ web application, view any report that you want to share.
2. In the upper-right corner of the report, click **Permalink**.
The **InsightIQ Permalink** dialog box displays and contains a URL for the currently configured report.
3. Put the cursor focus in the **Copy the URL below and use it for sharing this view** field.
4. Type **Ctrl + A** to select all the text.
5. Type **Ctrl + C** to copy the text.
6. Paste the text to a location that is convenient, such as a text file
7. Click **OK** to close the dialog box.

To view the report that is associated with the Permalink, you must be logged in to the InsightIQ web application.

Note

After an upgrade, IP addresses might have changed. This might require some permalinks to be updated with the new IP addresses.
