

# Statistics

Since v5.31/25

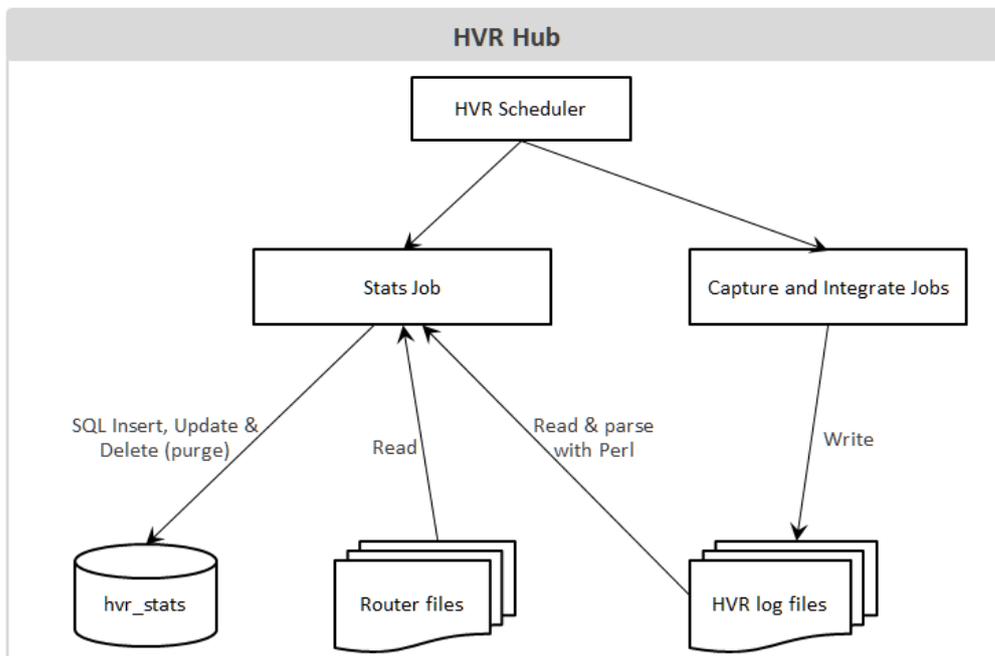
## Contents

- [Stats Job](#)
  - [Aggregation of Statistics Data](#)
  - [Statistics History Retention](#)
  - [Viewing Stats Job and Stat Job Log](#)
- [Launching Statistics from HVR GUI](#)
- [Statistics View](#)
  - [Graphs in Statistics](#)

Insight **Statistics** is a real-time, web-based dashboard which allows you to monitor replication performance. HVR's statistics dashboard helps you stay in control of your replication activity by giving you visibility into your most important data and metrics. It can display several types of statistical information such as number of rows replicated, time taken for capture and integrate, compression rates, and size of data processed on the replication system(s). The dashboard displays statistics for data updates very frequently, sometimes even on a minute-by-minute or second-by-second basis.

## Stats Job

The stats job (**hvrstats**) generates the information required for **Topology** and **Statistics** and saves it into the [catalog table \(HVR\\_STATS\)](#) which is responsible for maintaining the statistics data. So, the stats job must be running to display live details in **Topology** or **Statistics**. The stats job is created along with the catalog tables during HVR's installation and automatically started when starting the **HVR Scheduler**.



To generate the information required for **Topology** and **Statistics**,

1. The stats job reads data from the HVR log files and the router transaction files.
2. The stats job then modifies (using **insert**, **update**, and **delete** SQL statements) the [catalog table - hvr\\_stats](#) based on the data read from the HVR log files and the router transaction files. It also [aggregates the statistics data](#) that are written in the [hvr\\_stats](#) table.

The **hvr\_stats** table consists of a number of columns that store statistical information about data replication. In particular, the **hvr\_stats** table include the **metric\_name** and **metric\_value** columns storing data on a variety of metrics captured by HVR, such as capture/integrate latency, captured row counts, integrated change counts. For more information, see [Metrics for Statistics](#).

Since HVR 5.7.0/17 and 5.7.5/11, by default, only the basic/important metrics that are gathered by the **hvrstats** job from the HVR log files. The behavior of the **hvr\_stats** job gathering metrics can be configured by defining [Scheduling /StatsMetrics](#).

## Aggregation of Statistics Data

The stats job performs two types of aggregations (grouping) when writing statistics data to the **hvr\_stats** table:

### 1. Scope aggregation

The metrics information received from the HVR log files are written into the **hvr\_stats** table based on the scope defined by a channel name (column **chn\_name**), location name (column **loc\_name**), and table name (column **tbl\_name**), which can be either named explicitly or regarded as '\*' (which means applies to all channels, locations, tables).

For example, if there are 5 'captured inserts' with **chn\_name='chn1'**, **loc\_name='src'** and **tbl\_name='tbl1'** and 5 'captured inserts' with **chn\_name='chn1'**, **loc\_name='src'** and **tbl\_name='tbl2'**. The **hvr\_stats** table will store these values, but it will also store value 10 for **tbl\_name='\*'**, the sum of both values ('captured inserts').

For more information on various scopes that can be defined, see [hvrstats \(option -s\)](#).

### 2. Time granularity aggregation

Metrics are gathered/output with a per-minute granularity. For example, the value of 'captured inserts' for one-minute granularity means the number of rows inserted within that minute. These values can be aggregated up to 10 minutes, 1 hour and 1 day. For more information on the time granularity, see [hvrstats \(option -T\)](#).

## Statistics History Retention

By default, the old statistics data in the **hvr\_stats** table is purged automatically to save the disk space needed for the hub database. The default statistics history retention size is **MEDIUM** (described in [Scheduling /StatsHistory](#)).

To change default statistics history retention size, define action [Scheduling /StatsHistory](#).

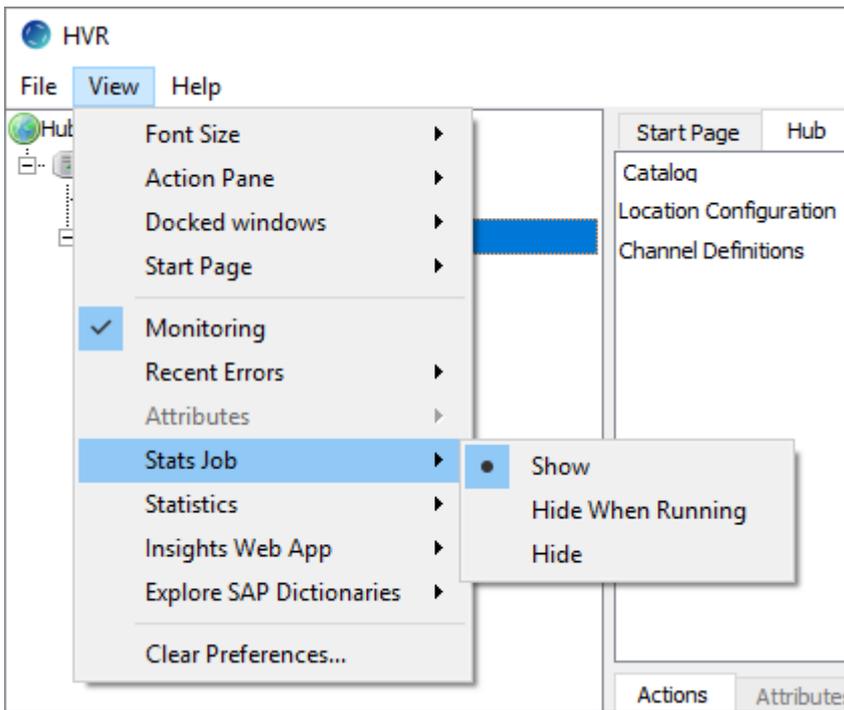
To purge the statistics data immediately (as a one-time purge) from the **hvr\_stats** table, use the command [hvrstats \(option -p\)](#).

## Viewing Stats Job and Stat Job Log

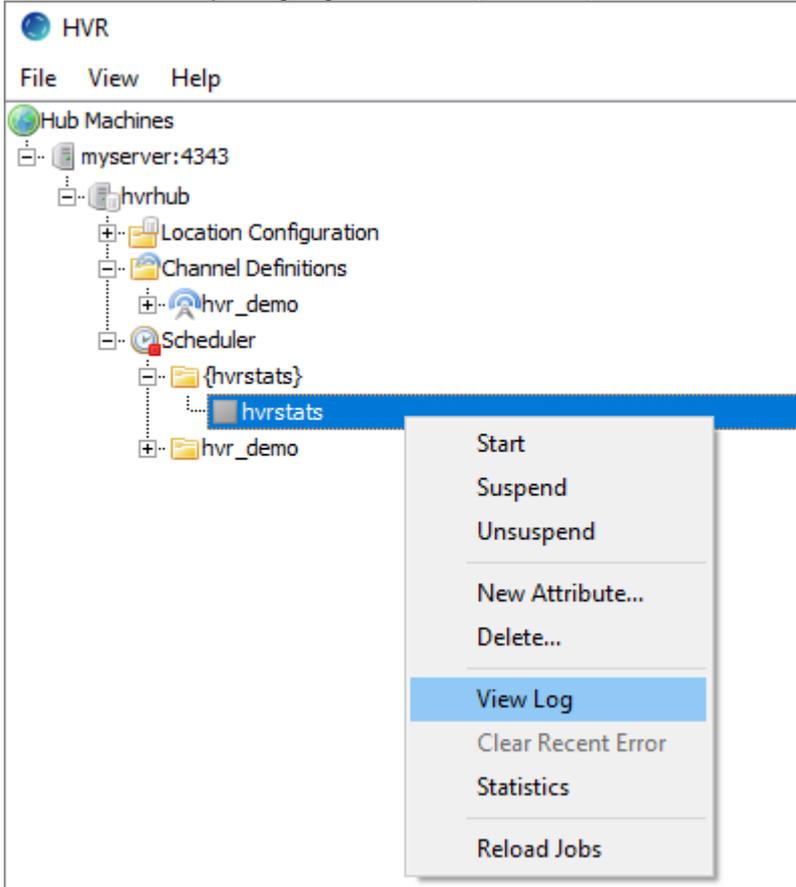
In HVR GUI, the stats job is displayed under the **Scheduler** node.

By default, if the stats job is running, it is not displayed under the **Scheduler** node. The display settings for the stats job can be configured from the menu bar.

- In the menu bar, click **View Stats JobShow**, to always display the stats job under the **Scheduler** node.



To view the stats job log, right-click the **{hvrstats}** and select **View Log**.



## Launching Statistics from HVR GUI

Statistics can be launched only from [HVR GUI](#), and it can only be viewed within the **Insights** interface in a web browser.

### Settings for Viewing Insights

The **Insights** interface for **Topology**, **Statistics**, and **Events** can be viewed only in the web browser.

Following are the two available options that can be configured (**View Insights Web App**) for viewing the **Insights** interface:

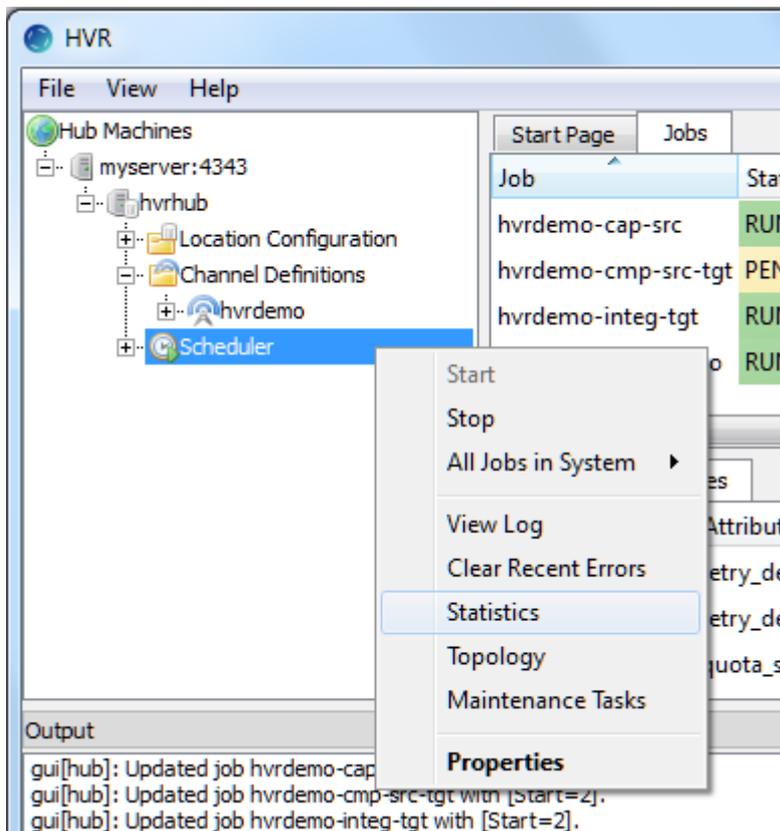
- **Open in local web browser** (default) - Automatically open the respective **Insights** interface in the default web browser.

When this option is used and if the URL is shared with other users, only the users who are connected to the machine on which the **HVR GUI** is executed can access the **Insights** interface.

- **Show URL only** - A prompt is displayed with an option to copy the **Insights** URL, which can be then pasted into any web browser's address bar to view the respective **Insights** interface. This option is used in any of the following situations:
  - If the machine on which the **HVR GUI** is executed does not have a web browser installed.
  - To share the URL with all other users. All other users can access the **Insights** interface even if they are not connected to the machine on which the **HVR GUI** is executed.

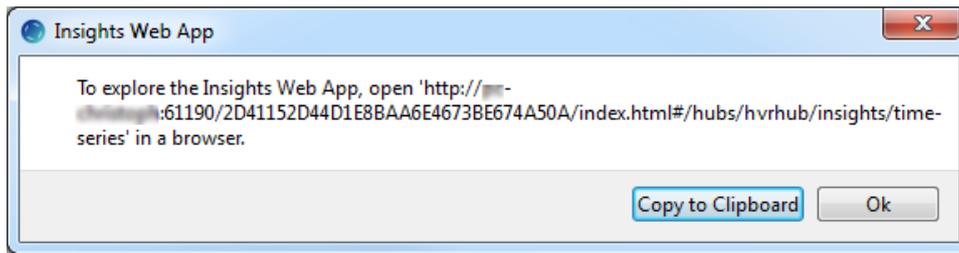
For **Insights** to work, **HVR GUI** should be running. If **HVR GUI** is closed with the **Insights** window open, then an error message is displayed in the **Insights** window.

- To launch **Statistics**, right-click **Scheduler** and select **Statistics**.



- To launch **Statistics**, when the viewing option for **Insights Web App** is set to **Show URL only**, the following needs to be performed:
  1. Right-click **Scheduler** and select **Statistics**

## 2. Click **Copy to Clipboard**.



## 3. Paste the URL in the web browser's address bar and press **Enter**.

## Statistics View



The **Statistics** view contains the following user interface components:

1. **Hub** - name of the hub for which the statistics is displayed.
2. **Scope** - filter the dashboard based on the selected location, channel, or table.
3. **Time Selector** - filter the dashboard based on the selected start date/time and end date/time. The other method to select time range is to zoom in on a custom time range on the graph. Click-and-drag the mouse across the chart to select a time range. The newly selected time range is applied to all graphs available in the dashboard. For example, to view statistics from 25th to 27th June, click-and-drag the mouse across the chart as shown below.



- a. **Granularity** - Sets the aggregation of data for the dashboard. The granularity of time refers to the size in which data fields are sub-divided. The default granularity is an **hour**. However, this aggregation can be changed to a 'fine granularity' level like **1 minute** or to a 'coarse granularity' level like **1 day**. Changing the granularity is immediately reflected in the x-axis of all charts displayed in the dashboard. Finer granularity has overheads like ex

cess computation time, memory, or other resources that are required to generate statistics dashboard. For this reason the data stored for each granular level is limited/purged at certain point of time - [Scheduling /StatsHistory](#).

The available options are -

- **minute**: Indicates 1 minute granularity. The period displayed in the **Time Selector** is limited to the last 14 days.
  - **10 min**: Indicates 10 minute granularity. The period displayed in the **Time Selector** is limited to the last 30 days.
  - **hour**: Indicates 1 hour granularity. The period displayed in the **Time Selector** is limited to the last 6 months.
  - **day**: Indicates 1 day granularity. The period displayed in the **Time Selector** is limited to the last 3 years.
4. **Add Graph**- Add graph to the dashboard. For information about the available graphs in **Statistics**, see section [Graphs](#).
  5. Options to customize each graph.
    - **Split**: create a graph using the selected graph's **Metric** and **Scope** (Location, Channel, or Table and Channel). In **Split** drop-down, you can select a **Metric** available in the graph and split it by **Scope**. The split graph will be inserted next to the graph containing a metric you want to split.
    - **Rate unit**: change the Y-axis unit (per second, per minute, per hour, or per day). By default, the rate unit matches the **Granularity**: per minute for **minute** and **10 minute** granularity, per hour for **hour** granularity and per day for **day** granularity.
    - **+**: Expand/collapse graph (when collapsed, only the graph header is displayed).
    - **X**: Remove graph from the dashboard.
  6. **Legend/Metrics** - Indicates the metrics for graph. Each graph displays different metrics. For more information about metrics displayed in Statistics, see [Metrics for Statistics](#).
    - **MORE**: displays more metrics in the graph.
    - **LESS**: displays less metrics.
    - **SELECT**: allows you to select particular metrics that needs to be displayed in a graph. This option is displayed only when a graph has more than 5 metrics associated with it.

## Graphs in Statistics

The dashboard contains a set of graphs that indicate the state of different aspects of the replication. Following are the default graphs displayed in **Statistics**:

Graph Name	Description
<b>Latency</b>	<p>Displays the latency information for capturing and integrating changes. Latency is the time (in seconds) taken for a transaction committed on the source system to be replicated (or committed) on the target system. This graph allows you to analyze the delay in data replication.</p> <p>The following <a href="#">Latency Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Capture Latency Min</b></li> <li>• <b>Capture Latency Max</b> (default)</li> <li>• <b>Integrate Latency Min</b></li> <li>• <b>Integrate Latency Max</b> (default)</li> <li>• <b>Capture Rewind Interval</b></li> </ul>
<b>Captured Changes Counts</b>	<p>Displays the total number of changes captured.</p> <p>The following <a href="#">Captured Row Counts Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Captured Inserts</b> (default)</li> <li>• <b>Captured Updates</b> (default)</li> <li>• <b>Captured Deletes</b> (default)</li> <li>• <b>Captured Changes</b></li> <li>• <b>Captured Changes Backdated</b></li> </ul>

	<ul style="list-style-type: none"> <li>• <b>Captured Skipped Rows</b></li> <li>• <b>Augmented Rows</b></li> <li>• <b>SAP Augment Selects</b></li> <li>• <b>Captured DDL Statements</b></li> </ul>
<b>Integrated Changes (split by Table /Channel)</b>	Displays the total number of changes integrated. This is basically a <a href="#">Integrated Change Counts</a> graph which is split using the metric <b>Integrated Changes</b> and scope <b>Table/Channel</b> .

The following graphs can be added by clicking on the **Add Graph** button.

<b>Graph Type</b>	<b>Description</b>
<b>Integrated Change Counts</b>	<p>Displays the total number of changes integrated.</p> <p>The following <a href="#">Integrated Change Counts Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Integrated Inserts</b> (default)</li> <li>• <b>Integrated Updates</b> (default)</li> <li>• <b>Integrated Deletes</b> (default)</li> <li>• <b>Integrated Changes</b></li> <li>• <b>Integrated Skipped Changes</b></li> <li>• <b>Changes Coalesced Away</b></li>   <li>• <b>Failed Inserts Saved</b></li> <li>• <b>Failed Updates Saved</b></li> <li>• <b>Failed Deletes Saved</b></li> <li>• <b>Failed Changes Saved</b></li> <li>• <b>Collision Changes Discarded</b></li> <li>• <b>Empty Updates Discarded</b></li> </ul>
<b>Transactions</b>	<p>Displays the total number of transactions.</p> <p>The following <a href="#">Transactions Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Captured Transactions</b></li> <li>• <b>Captured Transactions Backdated</b></li> <li>• <b>Integrated Transactions</b></li> </ul>
<b>Durations</b>	<p>Displays the time taken for replication in seconds.</p> <p>The following <a href="#">Duration Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Capture Duration Total</b></li> <li>• <b>Capture Duration Max</b></li> <li>• <b>Capture Duration Average</b></li> <li>• <b>Integrate Duration Total</b></li> <li>• <b>Integrate Duration Max</b></li> <li>• <b>Integrate Duration Average</b></li>   <li>• <b>Integrate Burst Load Duration Total</b></li> <li>• <b>Integrate Burst Load Duration Max</b></li> <li>• <b>Integrate Burst Load Duration Average</b></li> <li>• <b>Integrate Burst Setwise Duration Total</b></li> <li>• <b>Integrate Burst Setwise Duration Max</b></li> <li>• <b>Integrate Burst Setwise Duration Average</b></li> </ul>

<b>Speed</b>	<p>Displays the speed of replication. The unit of speed is indicated by the number of captured and integrated rows per selected <a href="#">granularity</a> of time.</p> <p>The following <a href="#">Speed Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Capture Speed Max</b> (default)</li> <li>• <b>Capture Speed Average</b></li> <li>• <b>Integrate Speed Max</b> (default)</li> <li>• <b>Integrate Speed Average</b></li>   <li>• <b>Integrate Burst Load Speed Max</b></li> <li>• <b>Integrate Burst Load Speed Average</b></li> <li>• <b>Integrate Burst Setwise Speed Max</b></li> <li>• <b>Integrate Burst Setwise Speed Average</b></li> </ul>
<b>Cycles</b>	<p>Displays the number of <a href="#">Capture</a> and <a href="#">Integrate</a> cycles.</p> <p>The following <a href="#">Cycles Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Capture Cycles</b></li> <li>• <b>Integrate Cycles</b></li> </ul>
<b>Byte I/O</b>	<p>Displays the size of captured row data and files in bytes before and after compression.</p> <p>The following <a href="#">Byte I/O Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Routed Bytes Written</b></li> <li>• <b>Routed Bytes Written Uncompressed</b></li> <li>• <b>Captured File Size</b></li> <li>• <b>Capture DbmsLog Bytes</b></li> <li>• <b>Capture DbmsLog Bytes Backdated</b></li> </ul>
<b>Compression</b>	<p>Displays the compression ratio for row data. Captured row data is compressed when sent from the capture location to the hub and from the hub to the integrate location.</p> <p>The following <a href="#">Compression Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Compression Ratio Max</b></li> <li>• <b>Compression Ratio Average</b></li> </ul>
<b>Replicated Files</b>	<p>Displays the number of rows captured and integrated into file locations.</p> <p>The following <a href="#">Replicated Files Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Captured Files</b></li> <li>• <b>Integrated Files</b></li> <li>• <b>Failed Files Saved</b></li> </ul>
<b>Errors /Warnings</b>	<p>Displays the total number of errors and warnings.</p> <p>The following <a href="#">Errors/Warnings Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b></li> <li>• <b>^Errors</b></li> <li>• <b>Errors F_J*</b></li> <li>• <b>^Errors F_J*</b></li> </ul>

	<ul style="list-style-type: none"> <li>• <b>Warnings</b></li> <li>• <b>^Warnings</b></li> <li>• <b>Warnings W_J*</b></li> <li>• <b>^Warnings W_J*</b></li> </ul>
<b>Router Rows</b>	<p>Displays the total number of rows in the transaction files.</p> <p>The following <a href="#">Router Rows Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Capture Router Rows</b></li> <li>• <b>Integrate Router Rows</b></li> </ul>
<b>Router Bytes</b>	<p>Displays the total size of the transaction files in bytes.</p> <p>The following <a href="#">Router Bytes Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Capture Router Bytes</b></li> <li>• <b>Integrate Router Bytes</b></li> </ul>
<b>Router Files</b>	<p>Displays the total number of transaction files.</p> <p>The following <a href="#">Router Files Stats Metrics</a> are displayed in this graph:</p> <ul style="list-style-type: none"> <li>• <b>Capture Router Files</b></li> <li>• <b>Integrate Router Files</b></li> </ul>