



# Proactive Status Monitoring and Diagnostics

---

This chapter contains the following topics:

- [Aggregate CPU, Disk, and Network Bandwidth Utilization, page 1](#)
- [Monitoring Aggregate CPU, Disk, and Network Bandwidth Utilization, page 2](#)
- [Monitoring Top Jobs Based on CPU Utilization and Time, page 2](#)
- [Performance Metrics for CPU, Disk, and Network, page 3](#)
- [Viewing CPU, Disk, and Network Statistics for a Hadoop Cluster, page 3](#)
- [Analyzing Performance Bottlenecks through Historical Metrics, page 4](#)
- [Setting Alerts for Hadoop Cluster Service Failures, page 5](#)
- [Types of Disk and Network Failure Alerts, page 6](#)
- [Setting Alerts for Disk and Network Failures, page 7](#)
- [Setting Disk Utilization Threshold Alerts, page 8](#)

## Aggregate CPU, Disk, and Network Bandwidth Utilization

You can monitor the aggregate CPU, disk, and network bandwidth utilization across all the hosts in a cluster. The metrics are collected in the following way:

- **Aggregate CPU and Disk metrics:** For every host that is running the job, the PID collects the percentage of CPU and memory used by the job. The sum of all these percentages gives the aggregate CPU and disk metrics.
- **Aggregate network bandwidth metrics:** The network bandwidth on each network interface in a host is obtained, which is then added up to obtain aggregate for one node. Similarly network bandwidths are measured for all the nodes in the cluster. The sum of all these bandwidths provides the aggregate network bandwidth metrics for the cluster.
- **Duration of long-running jobs:** A Rest API collects the start time, lapsed time, and end time for each job identified on the cluster. The difference between start time and end time provides the duration of completed jobs. The lapsed time provides the duration of the jobs running currently.

# Monitoring Aggregate CPU, Disk, and Network Bandwidth Utilization

- 
- Step 1** On the menu bar, choose **Solutions > Big Data > Accounts**.
- Step 2** Click the **Big Data Accounts** tab.
- Step 3** Choose the Big Data Account and click **View Details**.
- Step 4** Click the **Hadoop Clusters** tab.
- Step 5** Choose the Hadoop cluster and click **View Reports**.
- Step 6** Click the **Monitoring** tab.  
Every time an inventory collection cycle is triggered, an entry containing the aggregate CPU, network bandwidth, and disk utilization metrics appears on the Monitoring Page.
- Step 7** Select the entry you want to analyze and click **View Details**.
- Click the **Aggregate CPU** tab to view the aggregate CPU utilization of all nodes for a particular time.
  - Click the **Aggregate Disks** tab to view the aggregate disk utilization and available memory across the cluster.
  - Click the **Aggregate Network Bandwidth Utilization** to view the aggregated network bandwidth across the cluster.
- Step 8** Click **Back** to return back to the **Monitoring** page.
- 

# Monitoring Top Jobs Based on CPU Utilization and Time

To monitor top jobs based on CPU utilization or time (both active and completed long-running jobs), do the following:

- 
- Step 1** On the menu bar, choose **Solutions > Big Data > Accounts**.
- Step 2** Click the **Big Data Accounts** tab.
- Step 3** Choose the Big Data Account and click **View Details**.
- Step 4** Click the **Hadoop Clusters** tab.
- Step 5** Choose the Hadoop cluster and click **View Reports**.
- Click the **Top 10 High CPU Jobs** tab to view top ten jobs based on CPU utilization.
  - Click the **Top 10 Long Running Active Jobs** tab to view the current top ten long-running jobs.
  - Click the **Top 10 Long Duration Jobs** tab to view the completed top ten long-running jobs.
- Step 6** Click **Back** to return back to the **Hadoop Clusters** page.
-

# Performance Metrics for CPU, Disk, and Network

You can collect CPU, disk, and network metrics and analyze these metrics to find and fix performance bottlenecks that occur in the compute, network, or Hadoop setup across the cluster.

The metrics reports are of the following types:

- **Pre-Cluster:** This metrics report is generated automatically for a server that has just been installed with Red Hat Linux. This report is created before the server becomes part of a cluster.
- **Post-Cluster:** This metrics report is generated on demand when you run the performance test for a Hadoop cluster.

When you run the performance test for a Hadoop cluster, the following metrics are shown in detail:

- **Memory metrics:** Memory metrics measure the memory utilization of each host on the Hadoop cluster. The report includes the triad rate, which is the average rate at which read, write, and copy operations take place. The triad rate is a standard measure of memory bandwidth.
- **Network metrics:** Network metrics measure the network bandwidth of the Hadoop cluster. The report displays the rates at which network packets are transferred between the client and the server in the Hadoop cluster.
- **Disk metrics:** Disk metrics identify how fast a disk can perform. The report lists the time taken to read and write a file, to rewrite to an existing file, and to randomly (non-sequentially) read and write files. The disk metrics are included only in the pre-cluster report.
- **DFSIO metrics:** The DFSIO test is a Hadoop benchmark that stress-tests the storage I/O (read and write) capabilities of the cluster. The report measures the bytes processed, execution time, the average I/O rate and throughput to read and write multiple files. The DFSIO metrics are included only in the post-cluster report.
- **TeraSort metrics:** The TeraSort test is a Hadoop benchmark that tests the memory of the cluster. The report lists the counters for generating input, sorting the generated input, and validating the sorted output. The TeraSort Metrics is included only in the post-cluster report.

## Viewing CPU, Disk, and Network Statistics for a Hadoop Cluster

You can collect and compare CPU, disk, and network metrics with the pre-cluster creation and post-cluster creation reports for a Hadoop cluster.

- 
- Step 1** On the menu bar, choose **Solutions > Big Data >> Accounts**.
  - Step 2** Click the **Big Data Accounts** tab.
  - Step 3** Choose the Big Data Account and click **View Details**.
  - Step 4** Click the **Hadoop Clusters** tab.
  - Step 5** Choose the Hadoop cluster and click **View Reports**.
  - Step 6** Click the **Performance** tab.
  - Step 7** Click **Run Test**.

The Performance tab displays a default Big Data Metrics Report that shows the statistics collected for each host before the Hadoop cluster creation and the reports post Hadoop cluster creation.

**Step 8** Click **Submit**, and then click **OK**.

For the following actions, choose the performance report:

Name	Description
<b>View</b>	Displays the metrics in the Big Data Metrics Report.
<b>Compare</b>	Compares and displays the metrics in the Big Data Metrics Report.
<b>View Graph Report</b>	Displays graphically the following reports from the <b>Summary</b> tab: <ul style="list-style-type: none"> <li>• Average TRIAD Rate (MB/Sec)</li> <li>• Average Network Bandwidth (MB/Sec)</li> <li>• Average DFSIO Write (MB/Sec)</li> <li>• Average DFSIO Read (MB/Sec)</li> </ul>
<b>Delete</b>	Deletes the Big Data Metrics Report.
<b>More Reports</b>	Displays the metrics on an hourly, daily, weekly, or monthly basis.

## Analyzing Performance Bottlenecks through Historical Metrics

Comparing a historical metrics report (especially a metrics report that was collected when the cluster was performing well at peak load) with a metrics report captured during a poor performance helps identify metrics that could be the cause of the performance bottleneck in the Hadoop cluster.

To compare and analyze two metrics reports, do the following:

- 
- Step 1** On the menu bar, choose **Solutions > Big Data > Accounts**.
  - Step 2** Click the **Big Data Accounts** tab.
  - Step 3** Choose the Big Data Account and click **View Details**.
  - Step 4** Click the **Hadoop Clusters** tab.
  - Step 5** Choose the Hadoop cluster and click **View Reports**.
  - Step 6** Click the **Performance** tab.
  - Step 7** Click **Run Test**.

The Performance tab displays a default Big Data Metrics Report that shows the statistics collected for each host before the Hadoop cluster creation and the reports post Hadoop cluster creation.

- Step 8** Choose two reports that you want to compare (for example, a report generated while the cluster was performing well and a report generated during poor performance) and click **Compare**.
- Step 9** Click **Submit**.

## Setting Alerts for Hadoop Cluster Service Failures

You can create an alert to monitor the health of the Hadoop cluster whenever Hadoop services go down. Based on the trigger conditions, you can also activate customized workflows that automatically take corrective action.

**Step 1** On the menu bar, choose **Policies > Orchestration**.

**Step 2** Click the **Triggers** tab.

**Step 3** Click **Add**.

In the **Trigger Information** page of the **Add Trigger** wizard, complete the following fields:

Name	Description
<b>Trigger Name</b> field	Name of the trigger.
<b>Is Enabled</b> check box	Check this box to enable the trigger.
<b>Description</b>	Description of the trigger.
<b>Frequency</b>	Choose the trigger rule validation frequency.
<b>Trigger Type</b>	Choose the type of the trigger. <ul style="list-style-type: none"> <li>• Stateful</li> <li>• Stateless</li> </ul>

**Step 4** Click **Next**.

**Step 5** In the **Specify Conditions** page of the **Add Trigger** wizard, click **Add a new entry to the table below (+)**, and complete the following fields in the **Add Entry to Conditions** dialog box:

- a) From the **Type of Object to Monitor** drop-down list, choose **BigData Cluster**.
- b) From the **Object** drop-down list, choose the Hadoop cluster to be monitored.
- c) From the **Parameter** drop-down list, choose the parameter against to validate.
- d) From the **Operation** drop-down list, choose **Equals** or **Not Equals**.
- e) From the **Value** drop-down list, choose **All Services Up** or **Any Service Down**.
- f) Click **Submit**.

g) From the **Trigger When** drop-down list, choose to satisfy all the conditions or any condition.

**Step 6** Click **Next**.

**Step 7** In the **Specify Workflow** page of the **Add Trigger** wizard, do the following when the Hadoop cluster service is down and when the trigger is reset:

- a) From the **Maximum Number of Invocations** drop-down list, choose the number of maximum invocations.
- b) Select a workflow for execution when the trigger state becomes active and check the **Pass Monitored Object** check box, if required.
- c) Select the workflow input.
- d) Select a workflow for execution when the trigger state becomes clear and check the **Pass Monitored Object** check box, if required.
- e) Select the workflow input.

**Step 8** Click **Next**.

**Step 9** In the **Specify Workflow Inputs** page of the **Add Trigger** wizard, enter the inputs for the selected workflows, and then click **Submit**.

---

## Types of Disk and Network Failure Alerts

You can create alerts to detect faults related to disks and networks in a cluster.

The alerts that you can create for memory faults are as follows:

- **fltMemoryUnitInoperable:** This alert is triggered when the number of correctable or uncorrectable errors that have occurred on a DIMM have reached a threshold. The DIMM might be inoperable.
- **fltMemoryUnitThermalThresholdNonRecoverable:** This alert is triggered when the temperature of a memory unit on a server has been out of the operating range, and the issue is not recoverable.
- **fltMemoryArrayVoltageThresholdCritical:** This alert is triggered when the memory array voltage exceeds the specified hardware voltage rating.
- **fltMemoryArrayVoltageThresholdNonRecoverable:** This alert is triggered when the memory array voltage exceeds the specified hardware voltage rating, with potential memory hardware damage.
- **fltMemoryBufferUnitThermalThresholdCritical:** This alert is triggered when the temperature of a memory buffer unit on a blade or rack server exceeds a critical threshold value.
- **fltMemoryBufferUnitThermalThresholdNonRecoverable:** This alert is triggered when the temperature of a memory buffer unit on a blade or rack server has been out of the operating range, and the issue is not recoverable.
- **fltMemoryUnitDisabled:** This alert is triggered when the server BIOS disables a DIMM. The BIOS could disable a DIMM for several reasons, including incorrect location of the DIMM or incompatible speed.

The alerts that you can create for disk faults are as follows:

- **fltStorageItemCapacityExceeded:** This alert is triggered when the partition disk usage exceeds 70% but is less than 90%.

- **fltStorageItemCapacityWarning:** This alert is triggered when the partition disk usage exceeds 90%.
- **fltStorageLocalDiskInoperable:** This alert is triggered when the local disk has become inoperable.
- **fltStorageLocalDiskSlotEpUnusable:** This alert is triggered when the server disk drive is in a slot that is not supported by the storage controller.
- **fltStorageLocalDiskMissing:** This alert is triggered when a disk is missing.
- **fltStorageLocalDiskDegraded:** This alert is triggered when the local disk has degraded. The fault description contains the physical drive state, which indicates the reason for the degradation.

The alerts that you can create for network faults are as follows:

- **fltAdaptorUnitMissing:** This alert is triggered when the network adapter is missing or when the server cannot detect or communicate with the adapter.
- **fltAdaptorHostIfLink-down:** This alert is triggered when the fabric interconnect is in End-Host mode and all uplink ports had failed, when the server port to which the adapter is pinned had failed, or when a transient error caused the link to fail.
- **fltAdaptorExtIfLink-down:** This alert is triggered when the adapter's connectivity to any of the fabric interconnects cannot be validated, or when a node reports a vNIC down or link down event on the adapter link.

## Setting Alerts for Disk and Network Failures

You can create alerts for disk or network failures in the Hadoop cluster. Alerts help you in pro-active cluster maintenance. Based on the trigger conditions, you can activate customized workflows that automatically take corrective action.

**Step 1** On the menu bar, choose **Policies > Orchestration**.

**Step 2** Click the **Triggers** tab.

**Step 3** Click **Add**.

In the **Trigger Information** page of the **Add Trigger** wizard, complete the following fields:

Name	Description
Trigger Name field	Name of the trigger.
Is Enabled check box	Check this box to enable the trigger.
Description	Description of the trigger.
Frequency	Choose the trigger rule validation frequency.
Trigger Type	Choose the type of the trigger. <ul style="list-style-type: none"> <li>• Stateful</li> <li>• Stateless</li> </ul>

- Step 4** Click **Next**.
- Step 5** In the **Specify Conditions** page of the **Add Trigger** wizard, click **Add a new entry to the table below (+)**, and complete the following fields in the **Add Entry to Conditions** dialog box:
- From the **Type of Object to Monitor** drop-down list, choose **BigData Nodes**.
  - From the **Object** drop-down list, choose the disk to be monitored.
  - From the **Parameter** drop-down list, choose the parameter against which to validate.
  - From the **Operation** drop-down list, choose the type of operation.
  - From the **Value** drop-down list, choose the value against which to validate.
  - Click **Submit**.
  - From the **Trigger When** drop-down list, choose to satisfy all the conditions or any condition.
- Step 6** Click **Next**.
- Step 7** in the **Specify Workflow** page of the **Add Trigger** wizard, do the following when there is a network or disk failure and when the trigger is reset:
- From the **Maximum Number of Invocations** drop-down list, choose the number of maximum invocations.
  - Select a workflow for execution when the trigger state becomes active and check the **Pass Monitored Object** check box, if required.
  - Select the workflow input.
  - Select a workflow for execution when the trigger state becomes clear and check the **Pass Monitored Object** check box, if required.
  - Select the workflow input.
- Step 8** Click **Next**.
- Step 9** In the **Specify Workflow Inputs** page of the **Add Trigger** wizard, enter the inputs for the selected workflows, and then click **Submit**.

## Setting Disk Utilization Threshold Alerts

You can set an alert when the disk capacity reaches a threshold. This helps you to pro-actively plan for capacity expansions.

- Step 1** On the menu bar, choose **Policies > Orchestration**.
- Step 2** Click the **Triggers** tab.
- Step 3** Click **Add**.  
In the **Trigger Information** page of the **Add Trigger** wizard, complete the following fields:

Name	Description
<b>Trigger Name</b> field	Name of the trigger.
<b>Is Enabled</b> check box	Check this box to enable the trigger.

Name	Description
Description	Description of the trigger.
Frequency	Choose the trigger rule validation frequency.
Trigger Type	Choose the type of the trigger. <ul style="list-style-type: none"> <li>• Stateful</li> <li>• Stateless</li> </ul>

**Step 4** Click **Next**.

**Step 5** In the **Specify Conditions** page of the **Add Trigger** wizard, click **Add a new entry to the table below (+)**, and complete the following fields in the **Add Entry to Conditions** dialog box:

- a) From the **Type of Object to Monitor** drop-down list, choose **BigData Cluster**.
- b) From the **Object** drop-down list, choose the disk to be monitored.
- c) From the **Parameter** drop-down list, choose the **Disk Utilization (%)**.
- d) From the **Operation** drop-down list, choose the type of operation.
- e) From the **Value** drop-down list, choose the threshold value against which to validate.
- f) Click **Submit**.
- g) From the **Trigger When** drop-down list, choose to satisfy all the conditions or any condition.

**Step 6** Click **Next**.

**Step 7** In the **Specify Workflow** page of the **Add Trigger** wizard, do the following when the disk utilization reaches the threshold value and when the trigger is reset:

- a) From the **Maximum Number of Invocations** drop-down list, choose the number of maximum invocations.
- b) Select a workflow for execution when the trigger state becomes active and check the **Pass Monitored Object** check box, if required.
- c) Select the workflow input.
- d) Select a workflow for execution when the trigger state becomes clear and check the **Pass Monitored Object** check box, if required.
- e) Select the workflow input.

**Step 8** Click **Next**.

**Step 9** In the **Specify Workflow Inputs** page of the **Add Trigger** wizard, enter the inputs for the selected workflows, and then click **Submit**.

