



Proactive Status Monitoring and Diagnostics

This chapter contains the following topics:

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

- **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:** For aggregate network bandwidth of one node, obtain the network bandwidth on each network interface, and then add them. 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, elapsed 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 elapsed time reports 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 big data cluster and click **View Reports**.

Step 6 Click the **Monitoring** tab.

Every time an inventory collection cycle is triggered, an entry listing the aggregate CPU, network bandwidth, and disk utilization metrics appears on the Monitoring Page.

Note For Splunk cluster, on clicking the **View Details** button, the **Monitoring** tab is displayed. Step 4 and Step 5 are specific to Hadoop cluster only.

Step 7 Select the entry you want to analyze and click **View Details**.

- a) Click the **Aggregate CPU** tab to view the aggregate CPU utilization of all nodes for a particular time period.
- b) Click the **Aggregate Disks** tab to view the aggregate disk utilization and available memory across the cluster.
- c) Click the **Aggregate Network Bandwidth Utilization** to view the aggregated network bandwidth across the cluster.

Step 8 Click **Back** to return 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**.

- a) Click the **Top 10 High CPU Jobs** tab to view the top ten jobs, based on CPU utilization.
- b) Click the **Top 10 Long Running Active Jobs** tab to view the current top ten long-running jobs.
- c) 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 find performance bottlenecks that occur in the compute, network, or Hadoop setup across the cluster. You can collect CPU, disk, and network metrics and analyze these metrics to fix bottlenecks.

The metrics reports are of the following types:

- **Pre-Cluster:** This metrics report is generated automatically for a server that has 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 disk metrics are included only in the pre-cluster report. The report lists the following:
 - The time taken to read and write a file.
 - The time taken to rewrite to an existing a file.
 - The time to randomly (nonsequentially) read and write files.
- **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 report is 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 report 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. This report shows the statistics collected for each host before the Hadoop cluster creation and the reports collected after Hadoop cluster creation.

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

For the following actions, choose the performance report:

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

Analyzing Performance Bottlenecks Through Historical Metrics

You can compare a metrics report generated while the cluster was performing well with a report generated during poor performance. It helps you identify a cause or causes of a 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. This report shows the statistics collected for each host before the Hadoop cluster creation and the reports collected after Hadoop cluster creation.

Step 8 Choose two reports that you want to compare, and click **Compare**.

You can compare a report generated while the cluster was performing well and a report generated during poor performance.

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**.

On 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 trigger. <ul style="list-style-type: none"> • Stateful • Stateless

Step 4 Click **Next**.

Step 5 On the **Specify Conditions** page of the **Add Trigger** wizard, click **Add a new entry to the table below (+)**.

Step 6 In the **Add Entry to Conditions** dialog box, complete the following fields:

- 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 to use in validation.
- 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, make a choice to satisfy all the conditions, or any individual condition.

Step 7 Click **Next**.

Step 8 On 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) Choose the maximum number of invocations from the **Maximum Number of Invocations** drop-down list.
- b) Select a workflow for execution when the trigger state becomes active, and check the **Pass Monitored Object** check box, if necessary.
- 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 necessary.
- e) Select the workflow input.

Step 9 Click **Next**.

Step 10 On 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:** Triggers when the number of correctable or uncorrectable errors have reached a threshold on a DIMM. The DIMM becomes inoperable.
- **fltMemoryUnitThermalThresholdNonRecoverable:** Triggers when the memory unit temperature on a server is out of the operating range. The issue is not recoverable.
- **fltMemoryArrayVoltageThresholdCritical:** Triggers when the memory array voltage exceeds the specified hardware voltage rating.
- **fltMemoryArrayVoltageThresholdNonRecoverable:** Triggers when the memory array voltage exceeds the specified hardware voltage rating, with potential memory hardware damage.
- **fltMemoryBufferUnitThermalThresholdCritical:** Triggers when the temperature of a memory buffer unit on a blade or rack server exceeds a critical threshold value.
- **fltMemoryBufferUnitThermalThresholdNonRecoverable:** Triggers when the temperature of a memory buffer unit on a blade or rack server is out of the operating range. The issue is not recoverable.
- **fltMemoryUnitDisabled:** Triggers 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:** Triggers when the partition disk usage exceeds 70% but is less than 90%.
- **fltStorageItemCapacityWarning:** Triggers when the partition disk usage exceeds 90%.
- **fltStorageLocalDiskInoperable:** Triggers when the local disk has become inoperable.

- **fltStorageLocalDiskSlotEpUnusable:** Triggers when the server disk drive is in a slot that the storage controller does not support.
- **fltStorageLocalDiskMissing:** Triggers when a disk is missing.
- **fltStorageLocalDiskDegraded:** Triggers 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:** Triggers when the network adapter is missing, or the server cannot detect or communicate with the adapter.
- **fltAdaptorHostIfLink-down:** Triggers—
 - When the fabric interconnect is in End-Host mode and all uplink ports have failed.
 - When the server port to which the adapter is pinned have failed.
 - When a transient error causes the link to fail.
- **fltAdaptorExtIfLink-down:** Triggers—
 - When the adapter's connectivity to any of the fabric interconnects cannot be validated.
 - When a node reports that a vNIC is down, or reports a 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 proactive 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**.

On 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 trigger. <ul style="list-style-type: none"> • Stateful • Stateless

- Step 4** Click **Next**.
- Step 5** On 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 to use in validation.
 - From the **Operation** drop-down list, choose the type of operation.
 - From the **Value** drop-down list, choose the value to use in validation.
 - Click **Submit**.
 - From the **Trigger When** drop-down list, make a choice to satisfy all conditions, or any individual condition.
- Step 6** Click **Next**.
- Step 7** On 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 maximum number of invocations.
 - Select a workflow for execution when the trigger state becomes active and check the **Pass Monitored Object** check box, if necessary.
 - Select the workflow input.
 - Select a workflow for execution when the trigger state becomes clear, and check the **Pass Monitored Object** check box, if necessary.
 - Select the workflow input.
- Step 8** Click **Next**.
- Step 9** On 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 to be delivered when the disk capacity reaches a threshold. This helps you to proactively plan for capacity expansions.

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

Step 2 Click the **Triggers** tab.

Step 3 Click **Add**.

On 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.

Name	Description
Trigger Type	Choose the type of trigger. <ul style="list-style-type: none"> • Stateful • Stateless

Step 4 Click **Next**.

Step 5 On the **Specify Conditions** page of the **Add Trigger** wizard, click +, 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 to use for validation.
- f) Click **Submit**.
- g) From the **Trigger When** drop-down list, make a choice to satisfy all conditions, or any individual condition.

Step 6 Click **Next**.

Step 7 On 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 maximum number for invocations.
- b) Select a workflow for execution when the trigger state becomes active, and check the **Pass Monitored Object** check box, if necessary.
- 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 necessary.
- e) Select the workflow input.

Step 8 Click **Next**.

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