



Managing Hadoop and Splunk Clusters

This chapter contains the following sections:

- [Managing a Hadoop Cluster](#), page 1
- [Managing Nodes in a Cluster](#), page 6
- [Delete Node and Delete Node to Baremetal Actions in Cloudera and Hortonworks](#), page 7
- [Adding Managed Nodes to the Hadoop Cluster](#), page 7
- [Adding Live Nodes to the Hadoop Cluster](#), page 9
- [Adding Baremetal Nodes to the Hadoop Cluster](#), page 9
- [Service Roles](#), page 11
- [Managing a Splunk Cluster](#), page 12

Managing a Hadoop Cluster

You can manage an existing cluster.

Step 1 On the menu bar, choose **Solutions > Big Data > Accounts**.

Step 2 Click the **Hadoop Accounts** tab.

Step 3 Select an account and click **View Details**.

Step 4 Click each tab to view the details.

| Name | Description |
|---------|--|
| Summary | Displays statistics data report for the selected Hadoop Account and the high-level report on the cluster and node account. |
| Hosts | See Managing Nodes in a Cluster to understand the functionality of this tab. |

| Name | Description |
|-----------------|---|
| Hadoop Clusters | <p>Displays the list of actions that you can perform on an existing Hadoop cluster.</p> <ol style="list-style-type: none"> 1 Click Role Topology to view the topology of the nodes. (This tab is not editable.) 2 Click View Details to view the inputs for the Hadoop cluster you have created, and to view the virtual network interface configuration. (This tab is not editable.) 3 Click Shutdown Cluster to stop all the services and power off all the nodes in the Hadoop cluster. 4 Click Start Cluster to power up all the nodes in the Hadoop cluster and start all services in the Hadoop cluster. 5 Click Restart Cluster to power off and then power up all nodes in the Hadoop cluster. 6 Click Rebalance to configure the threshold percentage to rebalance Hadoop clusters. For MapR cluster, configure the threshold percentage using CLDB Balancer Disk Paused and CLDB Balancer Disk Max Switches in Nodes. 7 Click Upgrade Cluster to upgrade Hadoop distributions from the current version, if available. <ul style="list-style-type: none"> Note For a derived account, install and configure pssh and clush for the nodes in the Hadoop cluster and password less between nodes. <ol style="list-style-type: none"> 1 Click Upgrade Cluster. 2 Choose the JDK version from the Oracle JDK drop-down list. 3 Choose the Hadoop distribution that you want to upgrade from the current version from the Available Version drop-down list. 4 Check the Enable HA check box to enable high availability for the Hadoop cluster, if Cloudera is the Hadoop distribution. 5 Click Submit. 8 Enable Hadoop cluster high availability. <ol style="list-style-type: none"> a Click Enable High Availability. b From the Enable High Availability dialog box, access the Standby Name Node drop-down list and choose the Standby Name Node. c Check a minimum of three nodes from the Journal Nodes table. d Click Submit. 9 Click Cluster Snapshot to view the snapshot. 10 Click View Reports to view performance and monitoring reports. |

| Name | Description |
|-----------------------------|--|
| Hadoop Services | <p>Displays the list of Hadoop services and their status. You can do the following:</p> <ol style="list-style-type: none"> 1 Start All Services - Start all Hadoop services, depending on their status. 2 Stop All Services - Stop all Hadoop services, depending on their status. 3 Add New Service - Add a new Hadoop service. 4 Edit Service - Start and stop a particular Hadoop service. |
| Hadoop Service Roles | <p>Displays the list of Hadoop services.</p> <p>To add a role to the cluster, do the following in the Add Role dialog box:</p> <ol style="list-style-type: none"> 1 From the Hadoop Service Name drop-down list, choose the Hadoop service. 2 From the Role Type drop-down list, choose the role type. 3 From the Node Name drop-down list, choose the node name. 4 From the Role Name drop-down list, choose the role name. 5 Click Submit. <p>To start or stop any role that you have created, do the following in the Start/Stop Role dialog box:</p> <ol style="list-style-type: none"> 1 From the Hadoop Service Roles tab, choose the Hadoop service. 2 Click Start/Stop Role. 3 Click Submit. <p>To delete a role in the cluster, do the following in the Delete Role dialog box:</p> <ol style="list-style-type: none"> 1 From the Hadoop Service Roles tab, choose the Hadoop service. 2 Click Delete. 3 Click Submit. |
| More Reports | Lists additional reports that you can generate about data usage and CPU utilization. |

View Hadoop Cluster Details

For each Big Data Account, use the **Hadoop Clusters** tab to view details of all Hadoop clusters associated with the account. See [Managing a Hadoop Cluster](#)

You can view the following details by clicking the **Hadoop Clusters** tab.

| Name | Description |
|---------------------------------|---|
| Big Data Account name | The name of the Big Data account. |
| UCS SP Template for Big Data | The UCS SP Template for Big Data that you used to create service profiles for the servers in the Hadoop cluster. |
| Hadoop Cluster Profile Template | The Hadoop cluster profile template that you used to configure the cluster services. |
| Hadoop Cluster Deploy Template | A unique name that you used for the Hadoop cluster deployment template. |
| UCSM Policy Name Prefix | The UCSM Policy Name prefix. |
| Hadoop Cluster Name | A unique name that you used for the Hadoop cluster. |
| Hadoop Node Count | The number of nodes in the Hadoop cluster. |
| Hadoop Node Prefix | The Host Node prefix for the cluster. |
| OS Version | The operating system that you installed on the servers for the Hadoop cluster. |
| Hadoop Distribution | The Hadoop distribution that you used for this cluster. |
| Hadoop Distribution Version | The Hadoop distribution version that used for this cluster. |
| PXE VLAN ID | The VLAN ID used for PXE boot of the servers. |
| UCS Service Profile Template | The UCS Service Profile Template that you specified parameters for installing the OS and Hadoop distribution software. |
| Host Maintenance Policy | – |
| Host Firmware Package | – |
| UCSM Version | Displays the Cisco UCS Manager version |
| vNIC: eth0 | Displays IPv4 network information for the management interface and the management VLAN ID. |
| vNIC: eth1 | Displays IPv4 network information for the DATA1 interface and the VLAN ID. This field is not displayed when the Use one vNIC check box is selected while creating a Cisco UCS Service Profile Template for Big Data. |

| Name | Description |
|------------|---|
| vNIC: eth2 | Displays IPv4 network information for the DATA2 interface and the VLAN ID. This field is not displayed when the Use one vNIC check box is selected while creating a Cisco UCS Service Profile Template for Big Data. |

Viewing a Cluster Snapshot

A cluster snapshot displays configuration details of a Hadoop cluster, such as hosts, roles, and services. To view the current snapshot of a cluster, do the following:

- Step 1** On the menu bar, choose **Solutions > Big Data > Accounts**.
- Step 2** Click the **Hadoop Accounts** tab.
- Step 3** Choose the Hadoop Account for which you want to view the snapshot and click **View Details**.
- Step 4** Click the **Hadoop Clusters** tab.
- Step 5** Choose the Hadoop cluster for which you want to view the snapshot and click **Cluster Snapshot**.
- Step 6** Click **Submit**.
You can view the snapshot for the selected Hadoop cluster.

Adding a New Hadoop Service

- Step 1** On the menu bar, choose **Solutions > Big Data > Accounts**.
- Step 2** Choose a Cloudera account for which you can add a new Hadoop service.
Adding a new Hadoop service is not supported for MapR and Hortonworks distributions.
- Step 3** Click **View Details**.
- Step 4** Click the **Hadoop Services** tab.
- Step 5** Click **Add New Service**.
In the **Add New Service** dialog box, complete the following fields:

| Name | Description |
|------------------------------|--|
| Workflow Inputs | |
| (Hadoop) Account Name | Choose the Cloudera Account. |
| Service Type | Enter the Hadoop service for Cloudera. |

| Name | Description |
|-----------------------|---|
| Role Assignment Pairs | Enter Role Assignment Pairs, separated by commas. For example, RoleType1:hostname1, RoleType2:hostname2 |
| Dependant Services | Enter the list of dependent services. Use commas to separate list entries. |
| Pre Install Commands | Enter the list of commands. Use the "\n" command to insert a new line after each list entry. |
| Post Install Commands | Enter the list of commands. Use the "\n" command to insert a new line after each list entry. |

Step 6 Click **Submit**.

Managing Nodes in a Cluster

You can add, delete, decommission, and recommission nodes in a cluster.

- **Managed Node**—Any node that was already a member (managed) in the cluster and deleted, which can be added again in the cluster.
- **Live Node**—Any node that has the operating system installed and is reachable from the Hadoop cluster.
- **Baremetal Node**—Any node that is available and is not associated with the Hadoop cluster.

Step 1 On the menu bar, choose **Solutions > Big Data > Accounts**.

Step 2 Click the **Hadoop Accounts** tab.

Step 3 Select an account and click **View Details**.

Step 4 Click the **Hosts** tab to perform the following actions:

| Name | Description |
|---------------------|---|
| Refresh | Refreshes the page. |
| Favorite | Adds the page to Favorites. |
| Add Managed Node | Adds managed nodes to the Hadoop cluster. |
| Add Live Node | Adds live nodes to the Hadoop cluster. |
| Add BareMetal Nodes | Adds baremetal nodes to the Hadoop cluster. |

Step 5 Select a host that allows you to perform the following actions:

| Name | Description |
|---|---|
| View Details | Displays the summary of the CPU usage, the I/O status of the hosts disks, and so on. Note If you see a License Status tab, it indicates a licensing issue. |
| Delete Node | Deletes node from the cluster. |
| Assign Rack | Assigns the node to the rack server. |
| Recommission Node/ Decommission Node | Decommissioning or Recommissioning a node depends on its status. Note When the node is in decommissioned status, it means that all the roles for that node have been withdrawn. |
| Delete Node to BareMetal | The node is removed from the cluster and disassociated from the service profile. The node becomes a baremetal server. |
| Host Mappings | Lists DNS entries of all the hosts in the Hadoop cluster. |

Delete Node and Delete Node to Baremetal Actions in Cloudera and Hortonworks

When you perform the Delete Node or Delete Node to BareMetal actions, the UCS CPA Delete Node (a new workflow) is created for Cloudera and Hortonworks. You can also execute this workflow for MapR to perform the delete node operation. If you execute the UCS CPA Delete Node workflow for Cloudera and Hortonworks, this workflow also provides the functionality of Delete Node to BareMetal action. This functionality is based on the Delete Node to BareMetal flag setting as true or false. In addition, the Rollback UCS CPA Node BareMetal workflow is created.

Adding Managed Nodes to the Hadoop Cluster

Add managed nodes to the Hadoop cluster.

**Note**

This **Add Managed Nodes** functionality is not supported for Hortonworks 2.3 and later versions, and MapR distribution, but you can use the Add Live Nodes functionality.

This feature allows you to add nodes that are available only from the following URLs, but not the members of the cluster.

- Cloudera—<http://serverip:7180/api/v6/hosts>, where the serverIP is the IPv4 address of the administration node.
- Hortonworks—<http://serverIP:8080/api/v1/hosts>, where the serverIP is the IPv4 address of the administration node

-
- Step 1** On the menu bar, choose **Solutions > Big Data > Accounts**.
- Step 2** Click the **Hadoop Accounts** tab.
- Step 3** Double-click the Hadoop account.
- Step 4** Click the **Hosts** tab.
- Step 5** Click **Add Managed Node**.
- Step 6** From the **Host Name** drop-down list, choose the host name.
- Step 7** Click **Submit**.
-

Adding Live Nodes to the Hadoop Cluster

Add live nodes to the Hadoop cluster.

-
- Step 1** On the menu bar, choose **Solutions > Big Data > Accounts**.
 - Step 2** Click the **Hadoop Accounts** tab.
 - Step 3** Double-click the Hadoop account.
 - Step 4** Click the **Hosts** tab.
 - Step 5** Click **Add Live Node**.
 - Step 6** Enter the IPv4 address in the **Host Management IPv4 Address** field.
 - Step 7** Enter the name of the rack server in the **Rack Name** field.
 - Step 8** Enter the password in the **(New Node) Password** field for that rack server.
 - Step 9** Choose the Cluster Manager Version for the Hadoop distribution from the **Cluster Management Version** drop-down list.
 - Step 10** Choose the operating system to be installed on the servers in this cluster from the **OS Version** drop-down list.
 - Step 11** Choose a MapR cluster template from the **Hadoop Template Name** drop-down list. This field is displayed only when you select the MapR cluster.
 - Step 12** Click **Submit**.
-

Adding Baremetal Nodes to the Hadoop Cluster

Add baremetal nodes to the Hadoop cluster.

-
- Step 1** On the menu bar, choose **Solutions > Big Data > Accounts**.
 - Step 2** Click the **Big Data Accounts** tab.
 - Step 3** Double-click the Big Data account.
 - Step 4** Click the **Hosts** tab.
 - Step 5** Click **Add BareMetal Nodes**.
 - Step 6** In the **Add Baremetal Nodes** dialog box, complete the following fields:

| Name | Description |
|--------------------------------------|--|
| Big Data Account Name field | The name of the Big Data account. |
| UCSM Policy Name Prefix field | The UCSM Policy Name prefix. |
| Hadoop Cluster Name field | A unique name for the Hadoop cluster. |
| Hadoop Node Count field | The number of nodes in the Hadoop cluster. |

| Name | Description |
|---|--|
| Host Node Prefix field | The Host Node prefix for the cluster. |
| OS Version drop-down list | Choose the operating system to be installed on the servers in this cluster. |
| Hadoop Distribution drop-down list | Choose the Hadoop distribution to be used for this cluster. |
| Hadoop Distribution Version drop-down list | Choose the Hadoop distribution version. |
| Oracle JDK Version drop-down list | Choose the Oracle JDK version. |
| External Database drop-down list | Choose an external database. You can also configure a new database from here. |
| Multi-UCSM check box | <p>Check the Multi-UCSM check box if you use multiple UCSM accounts.</p> <p>The following workflows are created during an Instant and Customized Hadoop Cluster creation:</p> <ul style="list-style-type: none"> • UCS CPA Multi-UCSM Hadoop Cluster WF. • Single UCSM Server Configuration WF. (This WF is triggered per UCSM Account. For example, UCSM 120, UCSM121.) • UCS CPA Node BareMetal. (This WF is triggered per Node.) |
| UCS Manager Account drop-down list | Choose the Cisco UCS Manager account for this cluster. |
| Organization drop-down list | Choose the organization in which the servers for this cluster are located. |
| SSD Boot Drives Available for OS check box | <p>Check this check box if you do not want to validate the server disk availability for RAID level OS disks. Ensure that the servers contain Solid-State Drive (SSD).</p> <p>If this check box is not selected, the server disk availability for RAID level OS disks and data.</p> |
| Hadoop Template Name drop-down list | Choose a template for a MapR cluster. This field is displayed only when you select the MapR cluster. |
| UCS SP Template table | Choose an existing UCS Service Profile Template for Hadoop cluster creation. |
| PXE VLAN ID field | Enter the PXE VLAN ID. |

| Name | Description |
|------------------------|---|
| UCSTemplate Name table | Check the UCS Service Profile Template check box that you want to use and click Submit to confirm the selection. |
| Server Pool table | The server pool that you want to use for this cluster. The Cisco UCS Manager account and the organization that you choose determine which server pools are displayed in this area. |

Step 7 In the **vNIC Template** table, verify the vNIC templates available for the cluster.

Step 8 If you want to edit a vNIC template, select the row for that template and click **Edit**.

Step 9 In the **Edit vNIC Template Entry** dialog box, complete the following fields and click **Submit**.

| Name | Description |
|---------------------------------|--|
| vNIC Name drop-down list | The vNIC name in the selected template. This field is for your information only. |
| IP Pool drop-down list | Choose the Big Data IP pool that you want to use for IP addresses assigned to this vNIC. |
| MAC Address Pool drop-down list | Choose the MAC address pool that you want to use for this cluster. |
| First MAC Address field | Enter the MAC address. |
| Size field | Enter the size. |
| VLAN ID field | The VLAN ID for this cluster. |

Step 10 Click **Submit**.

Service Roles

If you use Add Node BareMetal, Add Managed Node, and Add Live Node actions, the following nodes-specific roles are added for each Hadoop distribution.

| Service Roles | Cloudera | MapR | Hortonworks |
|---------------|----------|------|-------------|
| FileServer | No | Yes | No |
| DataNode | Yes | No | Yes |

| Service Roles | Cloudera | MapR | Hortonworks |
|-----------------|----------|------|-------------|
| NodeManager | Yes | Yes | Yes |
| Ganglia Monitor | No | No | Yes |

Managing a Splunk Cluster

You can manage the Splunk cluster from the **Hosts** tab.

Step 1 On the menu bar, choose **Solutions > Big Data > Accounts**.

Step 2 Click the **Splunk Accounts** tab.

Step 3 Click the **Hosts** tab to perform the following actions:

| Name | Description |
|----------------------------|---|
| Refresh | Refreshes the page. |
| Favorite | Adds the page to Favorites. |
| Add BareMetal Nodes | Add baremetal nodes to the Splunk cluster. You can add Indexer, Search Head, or Administrative node through Add Baremetal workflow. You need to provide the Replication Factor based on the Indexer count. |

Note You can also start, stop, or restart the Splunk cluster.

Step 4 Select a host that allows you to perform the following actions:

| Name | Description |
|---------------------|---|
| View Details | Displays the summary of the CPU usage, the I/O status of the hosts disks, and so on. Note If you see a License Status tab, it indicates a licensing issue. |
| Start | Starts the services on the node. |
| Stop | Stops the services on the node. |
| View Details | Restarts the services on the node. |
| Restart | Deletes node from the cluster. |

| Name | Description |
|---------------------------------|---|
| Delete Node to BareMetal | The node is removed from the cluster and disassociated from the service profile. The node becomes a BareMetal server. |

Step 5 Select an account and click **View Details**.
 You can see only the **Hosts** tab. You can start, stop, or restart the Splunk cluster.

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 Splunk cluster creation and the reports post Splunk cluster creation only when you check the **Memory Test**, **Network Test**, and **Disk Test** check boxes in the **Pre Cluster Performance Tests** section of the **Management** tab. If you enable the precluster disk test, it impacts Splunk 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) |
| Delete | Deletes the Big Data Metrics Report. |
| More Reports | Displays the metrics on an hourly, daily, weekly, or monthly basis. |

