



# Managing a Hadoop Cluster

---

This chapter contains the following sections:

- [Managing a Hadoop Cluster](#), page 1
- [Managing Nodes in a Cluster](#), page 3
- [Delete Node and Delete Node to BareMetal Actions in Cloudera and Hortonworks](#), page 4
- [Adding Managed Nodes to the Hadoop Cluster](#), page 4
- [Adding Live Nodes to the Hadoop Cluster](#), page 5
- [Adding BareMetal Nodes to the Hadoop Cluster](#), page 5
- [Service Roles](#), page 7

## Managing a Hadoop Cluster

You can manage an existing cluster.

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

**Step 2** Select an account and click **View Details**.

**Step 3** Click on each tab to view the details.

Name	Description
Summary	Displays statistics data report for the selected Big Data Account and the high level report on the cluster and node account.
Hosts	See <a href="#">Managing Nodes in a Cluster</a> to know the functionality of this tab.

Name	Description
<b>Hadoop Clusters</b>	<p>This tab is only for information. Click <b>Role Topology</b> to view the topology of the nodes.</p> <p>This tab is only for information. Click <b>View Details</b> to view the inputs for the Hadoop cluster you have created and the virtual network interface configuration.</p>
<b>Hadoop Services</b>	<p>Displays the list of Hadoop services and their status. You can start or stop all Hadoop services depending on their current status using <b>Start All Services</b> and <b>Stop All Services</b>. To start and stop a particular service, use <b>Edit Service</b>.</p>
<b>Hadoop Service Roles</b>	<p>Displays the list of Hadoop services.</p> <p>To add a role to the cluster, do the following in the <b>Add Role</b> dialog box:</p> <ol style="list-style-type: none"> <li>1 From the <b>Hadoop Service Name</b> drop-down list, choose the Hadoop service.</li> <li>2 From the <b>Role Type</b> drop-down list, choose the role type.</li> <li>3 From the <b>Node Name</b> drop-down list, choose the node name.</li> <li>4 From the <b>Role Name</b> drop-down list, choose the role name.</li> <li>5 Click <b>Submit</b>.</li> </ol> <p>To start or stop any role that you have created, do the following in the <b>Start/Stop Role</b> dialog box:</p> <ol style="list-style-type: none"> <li>1 From the <b>Hadoop Service Roles</b> tab, choose the Hadoop service.</li> <li>2 Click <b>Start/Stop Role</b>.</li> <li>3 Click <b>Submit</b>.</li> </ol> <p>To delete a role in the cluster, do the following in the <b>Delete Role</b> dialog box:</p> <ol style="list-style-type: none"> <li>1 From the <b>Hadoop Service Roles</b> tab, choose the Hadoop service.</li> <li>2 Click <b>Delete</b>.</li> <li>3 click <b>Submit</b>.</li> </ol>
<b>More Reports</b>	<p>Additional reports that you can generate for data usage and CPU utilization.</p>

## Managing Nodes in a Cluster

You can add, delete, re-commission and decommission 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 a cluster to be installed.
- **Baremetal Node**—Any node that has no operating system and cluster installed.

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

**Step 2** Select an account and click **View Details**.

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

Name	Description
Refresh	Refreshes the page.
Favorite	Adds the page to Favorites.
Add Managed Node	Add managed nodes to the Hadoop cluster.
Add Live Node	Add live nodes to the Hadoop cluster.
Add BareMetal Nodes	Add baremetal nodes to the Hadoop 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. <b>Note</b> If you see a <b>License Status</b> tab, it indicates a licensing issue.
Delete Node	Deletes node from the cluster.
Assign Rack	Assigns the node to the rack server.

Name	Description
<b>Recommission Node/ Decommission Node</b>	Decommissioning or Re-commissioning a node depends on its current status.  When the node is in decommissioned status, it means that all the roles for that node have been withdrawn.
<b>Delete Node to BareMetal</b>	The node is removed from the cluster and disassociated from the service profile. The node becomes a BareMetal server.

## 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 except for MapR Hadoop distributions. 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 based on the Delete Node to BareMetal flag setting to true or false. In addition, the Rollback UCS CPA Node BareMetal is created.

## Adding Managed Nodes to the Hadoop Cluster

Add managed nodes to the Hadoop cluster.

- 
- Step 1** On the menu bar, choose **Solutions > Big Data Containers**.
  - Step 2** Double-click on the Big Data account.
  - Step 3** Click the **Hosts** tab.
  - Step 4** Click **Add Managed Node**.
  - Step 5** From the **Host Name** drop-down list, choose the host name.
  - Step 6** From the **Rack Name** drop-down list, choose the rack 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 Containers**.
  - Step 2** Double-click on the Big Data account.
  - Step 3** Click the **Hosts** tab.
  - Step 4** Click **Add Live Node**.
  - Step 5** Enter the IPv4 address in the **Host Management IPV4 Address** field.
  - Step 6** Enter the name of the rack server in the **Rack Name** field.
  - Step 7** Enter the password in the **(New Node) Password** field for that rack server.
  - Step 8** Choose the Cluster Manager Version for the Hadoop distribution from the **Cluster Management Version** drop-down list.
  - Step 9** Choose the operating system to be installed on the servers in this cluster from the **OS Version** drop-down list.
  - Step 10** 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 Containers**.
  - Step 2** Double-click on the Big Data account.
  - Step 3** Click the **Hosts** tab.
  - Step 4** Click **Add BareMetal Nodes**.
  - Step 5** In the **Add Baremetal Nodes** dialog box, complete the following fields:

Name	Description
<b>Big Data Account Name</b> field	The name of the Big Data account.
<b>UCSM Policy Name Prefix</b> field	The UCSM Policy Name prefix.
<b>Hadoop Cluster Name</b> field	A unique name for the Hadoop cluster.
<b>Hadoop Node Count</b> field	The number of nodes in the Hadoop cluster.
<b>Host Node Prefix</b> field	The Host Node prefix for the cluster.
<b>OS Version</b> drop-down list	Choose the operating system to be installed on the servers in this cluster.

Name	Description
<b>Hadoop Distribution</b> drop-down list	Choose the Hadoop distribution to be used for this cluster.
<b>Hadoop Distribution Version</b> drop-down list	Choose the Hadoop distribution version.
<b>Multi UCSM</b> 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/Customized Hadoop Cluster creation:</p> <ul style="list-style-type: none"> <li>• UCS CPA Multi-UCSM Hadoop Cluster WF.</li> <li>• Single UCSM Server Configuration WF (This WF will be triggered per UCSM Account. For example, UCSM 120, UCSM121).</li> <li>• UCS CPA Node BareMetal (This WF will be triggered per Node).</li> </ul>
<b>UCS Manager Account</b> drop-down list	Choose the Cisco UCS Manager account for this cluster.
<b>Organization</b> drop-down list	Choose the organization in which the servers for this cluster are located.
<b>Server UUID Pool</b> drop-down list	<p>Choose the UUID pool to be used for servers in this cluster.</p> <p><b>Note</b> A UUID pool is a collection of SMBIOS UUIDs that are available to be assigned to servers. The first number of digits that constitute the prefix of the UUID are fixed. The remaining digits, the UUID suffix, are variable. A UUID pool ensures that these variable values are unique for each server associated with a service profile which uses that particular pool to avoid conflicts.</p> <p>If you use UUID pools in service profiles, you do not have to manually configure the UUID of the server associated with the service profile.</p>
<b>UUID Prefix</b> drop-down list	Choose the UUID prefix.
<b>UUID Pool-From</b> field	To specify manually, enter the UUID pool to start from.
<b>UUID Pool-Size</b> field	To specify manually, enter the UUID pool size.
<b>PXE VLAN ID</b> field	Enter the PXE VLAN ID.
<b>UCS SP Template</b> table	Check the UCS Service Profile Template check box that you want to use and click <b>Submit</b> .

Name	Description
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.
Host Firmware Package table	The host firmware package that you want to use for this cluster. The Cisco UCS Manager account and the organization that you choose determine which host firmware package are displayed in this area.

**Step 6** In the **vNIC Template** table, review and, if desired, edit the vNIC templates available for the cluster.

**Step 7** If you want to edit a vNIC template, select the row for that template and click **Edit**.

**Step 8** In the **Edit vNIC Template Entry** dialog box, complete the following fields and click **Submit**.

Name	Description
vNIC Name drop-down list	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	To specify manually, enter the MAC address.
Size field	To specify manually, enter the size .
VLAN ID field	The VLAN ID for this cluster.

**Step 9** Click **Submit**.

## Service Roles

For each Hadoop distribution, specific roles are added for all nodes that are newly added by using Add Node BareMetal, Add Managed Node, and Add Live Node options to the Hadoop cluster.

Service Roles	Cloudera	MapR	Hortonworks
FileServer	No	Yes	No

<b>Service Roles</b>	<b>Cloudera</b>	<b>MapR</b>	<b>Hortonworks</b>
DataNode	Yes	No	Yes
NodeManager	Yes	Yes	Yes
Ganglia Monitor	No	No	Yes