



## Managing a Hadoop Cluster

---

This chapter contains the following sections:

- [Managing a Hadoop Cluster](#), page 1
- [Managing Nodes in a Cluster](#), page 7
- [Delete Node and Delete Node to BareMetal Actions in Cloudera and Hortonworks](#), page 8
- [Adding Managed Nodes to the Hadoop Cluster](#), page 8
- [Adding Live Nodes to the Hadoop Cluster](#), page 9
- [Adding BareMetal Nodes to the Hadoop Cluster](#), page 10
- [Service Roles](#), page 12

## Managing a Hadoop Cluster

You can manage an existing cluster.

---

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

**Step 2** Click on the **Big Data Accounts** tab.

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

**Step 4** 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
Hadoop Clusters	

Name	Description
	<p>Displays the list of actions that you can perform on an existing Hadoop cluster.</p> <ul style="list-style-type: none"> <li>• (This tab is only for information) Click <b>Role Topology</b> to view the topology of the nodes.</li> <li>• (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.</li> <li>• Click <b>Shutdown Cluster</b> that stops all the services and power off all the nodes in the Hadoop cluster.</li> <li>• Click <b>Start Cluster</b> that powers on all the nodes in the Hadoop cluster and starts all the services in the Hadoop cluster.</li> <li>• Click <b>Restart Cluster</b> to power off and then power on all the nodes in the Hadoop cluster.</li> <li>• Click <b>Rebalance</b> to configure the threshold percentage to re-balance Hadoop clusters. For MapR cluster, configure the threshold percentage using CLDB Balancer <b>Disk Paused</b> and CLDB Balancer <b>Disk Max Switches in Nodes</b>.</li> <li>• Upgrade Hadoop distributions from the current version, if available. <ul style="list-style-type: none"> <li><b>Note</b> For a derived account, pssh and clush should be installed and configured for all the nodes in the Hadoop cluster and password less between nodes.</li> </ul> <ol style="list-style-type: none"> <li>1 Click <b>Upgrade Cluster</b>.</li> <li>2 Choose the Hadoop distribution that you want to upgrade from the current version from the <b>Available Version</b> drop-down list.</li> <li>3 Check the <b>Enable HA</b> check box to enable the Hadoop cluster for high availability.</li> <li>4 Click <b>Submit</b>.</li> </ol> </li> <li>• Enable Hadoop cluster high availability. <ol style="list-style-type: none"> <li>1 Click <b>Enable High Availability</b>.</li> <li>2 From the <b>Enable High Availability</b> dialog box, choose the Standby Name Node from the <b>Standby Name Node</b> drop-down list.</li> <li>3 Check a minimum of three nodes from the <b>Journal Nodes</b> table.</li> <li>4 Click <b>Submit</b>.</li> </ol> </li> <li>• Disable Hadoop cluster high availability. <ol style="list-style-type: none"> <li>1 Click <b>Disable High Availability</b>.</li> <li>2 From the <b>Disable High Availability</b> dialog box, choose the Name Node from the <b>Select Name Node</b> drop-down list.</li> <li>3 Check the <b>Node Name</b> check box from the <b>Select Secondary Name Node</b> table.</li> </ol> <p style="text-align: right;">Name Node and Secondary Name Node should not be mapped to a</p> </li> </ul>

Name	Description
	<p><b>Note</b> single node.</p> <p><b>4</b> Click <b>Submit</b>.</p> <ul style="list-style-type: none"> <li>• Click <b>Cluster Snapshot</b> to view the snapshot.</li> <li>• Click <b>View Reports</b> to view performance and monitoring reports.</li> </ul>
<b>Hadoop Services</b>	<p>Displays the list of Hadoop services and their status. You can do the following:</p> <ol style="list-style-type: none"> <li><b>1 Start All Services</b> - Start all Hadoop services depending on their current status.</li> <li><b>2 Stop All Services</b> - Stop all Hadoop services depending on their current status.</li> <li><b>3 Add New Service</b> - Add a new Hadoop service.</li> <li><b>4 Edit Service</b> - Start and stop a particular Hadoop service.</li> </ol>
<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><b>1</b> From the <b>Hadoop Service Name</b> drop-down list, choose the Hadoop service.</li> <li><b>2</b> From the <b>Role Type</b> drop-down list, choose the role type.</li> <li><b>3</b> From the <b>Node Name</b> drop-down list, choose the node name.</li> <li><b>4</b> From the <b>Role Name</b> drop-down list, choose the role name.</li> <li><b>5</b> 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><b>1</b> From the <b>Hadoop Service Roles</b> tab, choose the Hadoop service.</li> <li><b>2</b> Click <b>Start/Stop Role</b>.</li> <li><b>3</b> 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><b>1</b> From the <b>Hadoop Service Roles</b> tab, choose the Hadoop service.</li> <li><b>2</b> Click <b>Delete</b>.</li> <li><b>3</b> click <b>Submit</b>.</li> </ol>
<b>More Reports</b>	Additional reports that you can generate for data usage and CPU utilization.

## View Hadoop Cluster Details

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

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	Hadoop distribution that you used for this cluster.
Hadoop Distribution Version	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.
vNIC: eth0	Displays IPv4 network information for management interface and the management VLAN ID.
vNIC: eth1	Displays IPv4 network information for DATA1 interface and the VLAN ID.
vNIC: eth2	Displays IPv4 network information for DATA2 interface and the VLAN ID.

## Viewing a Cluster Snapshot

A cluster snapshots 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 **Big Data Accounts** tab.
  - Step 3** Choose the Big Data 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.  
Add new 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
<b>Workflow Inputs</b>	
<b>Big Data Account Name</b>	Choose the Big Data Account (Cloudera).
<b>Service Type</b>	Enter the Hadoop service for Cloudera.
<b>Role Assignment Pairs</b>	Enter Role Assignment Pairs separated by a comma. For example, RoleType1:hostname1, RoleType2:hostname2
<b>Dependant Services</b>	Enter the list of dependant services separated by a comma
<b>Pre Install Commands</b>	Enter the list of commands separated by a new line ("\n")
<b>Post Install Commands</b>	Enter the list of commands separated by a new line ("\n")

**Step 6** Click **Submit**.

## 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 reachable from the Hadoop cluster.
- **Baremetal Node**—Any node that is available and not associated with the Hadoop cluster.

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

**Step 2** Click on the **Big Data Accounts** tab.

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

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

Name	Description
<b>Refresh</b>	Refreshes the page.
<b>Favorite</b>	Adds the page to Favorites.
<b>Add Managed Node</b>	Add managed nodes to the Hadoop cluster.
<b>Add Live Node</b>	Add live nodes to the Hadoop cluster.
<b>Add BareMetal Nodes</b>	Add baremetal nodes to the Hadoop cluster.

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

Name	Description
<b>View Details</b>	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.
<b>Delete Node</b>	Deletes node from the cluster.
<b>Assign Rack</b>	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. <b>Note</b> 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.
<b>Host Mappings</b>	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 based on the Delete Node to BareMetal flag setting to 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.

This feature allows you to add nodes that are available only from the below 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 Sata > Accounts**.
- Step 2** Click on the **Big Data Accounts** tab.
- Step 3** Double-click on the Big Data 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 Sata > Accounts**.
- Step 2** Click on the **Big Data Accounts** tab.
- Step 3** Double-click on the Big Data 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** 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 Sata > Accounts**.
- Step 2** Click on the **Big Data Accounts** tab.
- Step 3** Double-click on 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
<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.
<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.

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

**Step 7** In the **vNIC Template** table, review and, if desired, edit 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
<b>vNIC Name</b> drop-down list	This field is for your information only.
<b>IP Pool</b> drop-down list	Choose the big data IP pool that you want to use for IP addresses assigned to this vNIC.
<b>MAC Address Pool</b> drop-down list	Choose the MAC address pool that you want to use for this cluster.

Name	Description
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 10** 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
DataNode	Yes	No	Yes
NodeManager	Yes	Yes	Yes
Ganglia Monitor	No	No	Yes