



Managing Hadoop Clusters

This chapter contains the following sections:

- [Creating an Instant Hadoop Cluster, on page 1](#)
- [Creating a Customized Hadoop Cluster, on page 5](#)
- [Creating a Hadoop Cluster Using Workflow, on page 9](#)
- [Provisioning an Instant and Customized Hadoop Cluster, on page 9](#)
- [Managing a Hadoop Cluster, on page 11](#)
- [Managing Nodes in a Cluster, on page 15](#)
- [Delete Node and Delete Node to Bare Metal Actions in Cloudera and Hortonworks, on page 17](#)
- [Deleting an Unreachable Node from Hadoop Distribution, on page 17](#)
- [Adding Managed Nodes to the Hadoop Cluster, on page 19](#)
- [Adding Live Nodes to the Hadoop Cluster, on page 19](#)
- [Adding Bare Metal Nodes to the Hadoop Cluster, on page 20](#)
- [Adding Disks to the Hadoop Cluster, on page 22](#)
- [Service Roles, on page 23](#)

Creating an Instant Hadoop Cluster

Before you begin

- Create a service profile template .
- Create a server pool in the Cisco UCS Manager account that you plan to use for this cluster. See [Cisco UCS Director Management Guide for Cisco UCS Manager](#).
- Create a MAC address pool .

-
- Step 1** Choose **Solutions > Big Data > Containers**.
- Step 2** Click **Cluster Deploy Templates**.
- Step 3** Click **Instant Hadoop Cluster**.
- Step 4** On the **Instant Hadoop Cluster Creation** screen, complete the following fields.

Name	Description
Big Data Account Name field	Enter the name of the Big Data account.
UCSM Policy Name Prefix field	Enter the UCSM Policy Name prefix.
Hadoop Cluster Name field	Enter a unique name for the Hadoop cluster.
Hadoop Node Count field	Enter the number of nodes in the Hadoop cluster.
SSH (root) Password field	Enter the SSH root password. Special characters such as \$, %, and & are not supported. Note The SSH username pertains to the root user.
Confirm SSH Password field	Enter the SSH root password. Special characters such as \$, %, and & are not supported.
Hadoop Manager Password field	Enter the management console password. Special characters such as \$, %, and & are not supported.
Confirm Hadoop Manager Password field	Enter the management console password. Special characters such as \$, %, and & are not supported.
Host Node Prefix field	Enter 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.
Management Console Protocol drop-down list	Choose HTTP or HTTPS protocol. Note Only HTTPS protocol is supported in MapR.
Management Port Number field	Enter the port number. Enter an integer between 1024 and 65535. Usage of reserved ports by Hadoop services or Linux OS should be avoided so that the web server path is reachable. Note Management port number is supported only in MapR 6.0 version.
Oracle JDK 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.

Name	Description
<p>Multi-UCSM check box</p>	<p>Check the Multi-UCSM check box if you use multiple UCSM accounts.</p> <p>Note If you use the multiple UCSM accounts option, you can configure the Hadoop Server Roles as described in Step 5. You can add UCSM Specific Inputs in the Add Entry to UCSM Specific Inputs table.</p> <p>The following workflows are established 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 Bare Metal. (This WF is triggered per Node.) <p>The UCSM Specific Input area is displayed when the Multi-UCSM check box is checked.</p>
<p>UCS Manager Account drop-down list</p>	<p>Choose the Cisco UCS Manager account for this cluster.</p>
<p>Organization drop-down list</p>	<p>Choose the organization in which the servers for this cluster are located.</p>
<p>UCS SP Template table</p>	<p>Choose an existing UCS Service Profile Template for cluster creation.</p>
<p>SSD Boot Drives Available for OS check box</p>	<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 the check box is not selected, the disk availability for both the OS disk and data disk are validated based on their RAID level.</p>
<p>PXE VLAN ID field</p>	<p>Enter the PXE VLAN ID.</p>
<p>UCS SP Template table</p>	<p>Choose an existing UCS Service Profile Template for Hadoop cluster creation.</p>

Step 5 If you want to edit a Hadoop Server Role, select the row for that role, and click **Edit**.

Step 6 On the **Edit Hadoop Server Roles Entry** screen, complete the following fields and click **Submit**.

Name	Description
Node Type field	Displays the Hadoop node role. Note Kafka node is supported only for Cloudera and Hortonworks clusters. Note Cloudera clusters with Data node require a minimum of three Master nodes and three Data nodes. Also, the Data node should have minimum 14 disks since the tolerance value is set to 6.
Node Count field	The number of nodes in the Hadoop cluster for the selected node type.
SSD Boot Drives Available for OS check box	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 SSD. If the check box is not selected, the disk availability for both the OS disk and data disk are validated based on their RAID level. Note This check box is not displayed when the UCSM version is greater than or equal to 3.
Server Pool table	Enter the server pool that you want to use for the cluster for the selected node type. 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, review and, if necessary, 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 On the **Edit vNIC Template Entry** screen, 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. (This drop-down list is disabled if an existing UCS SP Template is selected.)
First MAC Address field	Enter the MAC address pool that you want to use for this cluster.

Name	Description
Size field	Enter the size. (This field is disabled if an existing UCS SP Template is selected.)
VLAN ID field	Enter the VLAN ID for this cluster. (This field is disabled if an existing UCS SP Template is selected.)

Step 10 Click **Submit**.

What to do next

You can view and monitor the workflow that is triggered after you create an instant Hadoop cluster.

Creating a Customized Hadoop Cluster

Before you begin

- Create a service profile template.
- Create a Hadoop cluster profile template.
- Setup the details for Hadoop Config Parameters.
- Create a Hadoop cluster deployment template that Cisco UCS Director Express for Big Data uses to create the Hadoop cluster.
- Create a server pool in the Cisco UCS Manager account you plan to use for this cluster. See [Cisco UCS Director Management Guide for Cisco UCS Manager](#).
- Create a MAC address pool.

Step 1 Choose **Solutions > Big Data > Containers**.

Step 2 Click **Cluster Deploy Templates**.

Step 3 Select the template that you want to use for the Hadoop cluster and click **Customized Hadoop Cluster**.

Step 4 On the **Customized Hadoop Cluster Creation** screen, complete the following fields.

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

Name	Description
SSH (root) Password field	Enter the SSH root password. Special characters such as \$, %, and & are not supported. Note The SSH username pertains to the root user.
Confirm SSH Password field	Enter the SSH root password. Special characters such as \$, %, and & are not supported.
Hadoop Manager Password field	Enter the management console password. Special characters such as \$, %, and & are not supported.
Confirm Hadoop Manager Password field	Enter the management console password. Special characters such as \$, %, and & are not supported.
Host Node Prefix field	Enter 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.
Management Console Protocol drop-down list	Choose HTTP or HTTPS protocol. Note Only HTTPS protocol is supported in MapR.
Management Port Number field	Enter the port number. Enter an integer between 1024 and 65535. Usage of reserved ports by Hadoop services or Linux OS should be avoided so that the web server path is reachable. Note Management port number is supported only in MapR 6.0 version.
Oracle JDK 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.

Name	Description
<p>Multi-UCSM check box</p>	<p>Click the Multi-UCSM check box if you use multiple UCSM accounts.</p> <p>Note If you use the multiple UCSM accounts option, you can configure the Hadoop Server Roles as described in Step 5. You can add UCSM Specific Inputs in the Add Entry to UCSM Specific Inputs table.</p> <p>The following workflows are created during an Instant Hadoop cluster creation 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 Bare Metal. (This WF is triggered per Node.) <p>The UCSM Specific Input area is displayed when the Multi-UCSM check box is checked.</p>
<p>UCS Manager Account drop-down list</p>	<p>Choose the Cisco UCS Manager account for this cluster.</p>
<p>Organization drop-down list</p>	<p>Choose the organization in which the servers for this cluster are located.</p>
<p>SSD Boot Drives Available for OS check box</p>	<p>Click this check box if you do not want to validate the server disk availability for RAID level OS disks. Ensure that the servers contain SSD.</p> <p>If the check box is not selected, the disk availability for both the OS disk and data disk are validated based on their RAID level.</p>
<p>PXE VLAN ID field</p>	<p>Enter the PXE VLAN ID.</p>
<p>UCS SP Template table</p>	<p>Choose an existing UCS Service Profile Template for Hadoop cluster creation.</p>

Step 5

If you want to edit a Hadoop Server Role, select the row for that role, and click **Edit**.

Step 6

On the **Edit Hadoop Server Roles Entry** screen, complete the following fields and click **Submit**.

Name	Description
<p>Node Type field</p>	<p>Displays the Hadoop node role.</p>
<p>Node Count field</p>	<p>The number of nodes in the Hadoop cluster for the selected node type.</p>

Name	Description
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 SSD.</p> <p>If the check box is not selected, the disk availability for both the OS disk and data disk are validated based on their RAID level.</p> <p>Note This check box is not displayed when the UCSM version is greater than or equal to 3.</p>
Server Pool table	<p>Enter the server pool that you want to use for the cluster for the selected node type.</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 On the **Edit vNIC Template Entry** screen, 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 field	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. (This drop-down list is disabled if an existing UCS SP Template is selected.)
First MAC Address field	Enter the MAC address.
Size field	Enter the size.
VLAN ID field	The VLAN ID for this cluster. (This field is disabled if an existing UCS SP Template is selected.)

Note When you use vNIC bonding, ensure that you assign IP Pool, MAC Address Pool, and VLAN ID to the first vNIC in the **vNIC Template** table.

Step 10 Click **Submit**.

What to do next

You can view and monitor the workflow that gets triggered after creating a customized Hadoop cluster.

Creating a Hadoop Cluster Using Workflow

In Cisco UCS Director Express for Big Data, administrator can map the advanced catalog option to Hadoop cluster creation workflow, with limited user inputs, so that the service end user can trigger cluster creation. See [Cisco UCS Director End User Portal Guide](#).

Before you begin

- Create a service profile template
- Create a server pool in the Cisco UCS Manager account that you plan to use for this cluster. See [Cisco UCS Director Management Guide for Cisco UCS Manager](#).
- Create a MAC address pool
- Create a user with user role as service end user.

-
- Step 1** Log into Cisco UCS Director Express for Big Data using admin credentials.
- Step 2** Choose **Orchestration** and click **Workflows**.
- Step 3** Click **Add Workflow**.
- Step 4** On the **Add Workflow Details** page, enter the workflow name and choose a folder. Click **Next**.
- Step 5** On the **Add User Inputs** page, enter the required details and click **Next**.
- Step 6** On the **Add User Outputs** page, enter the required details and click **Submit**.
- Step 7** Double-click the workflow in the **Workflow Designer**.
- Step 8** Add the Initiate Hadoop Cluster task.
- Step 9** Select the attributes that you want to map to the workflow input fields. Check the **Map to User Input** check box to provide user inputs, if required.
- Step 10** Enter required details in the **Hadoop Service Role** table and **vNIC Template** table, and click **Submit**.
- Step 11** Choose **Policies > Catalogs** and click **Add Catalog**.
- Step 12** On the **Add Catalog page**, choose the catalog type as Advanced and select a workflow. Click **Submit** to map the workflow to the catalog.
- Step 13** Log into Cisco UCS Director Express for Big Data using service end user credentials.
- Step 14** Choose **Catalogs**. The **Catalogs** page displays the list of catalogs available for the service end user.
- Step 15** Select a catalog and click **Create Request**. The **Create Server Request** page displays the mapped user inputs.
- Step 16** Specify the required details.
- Step 17** Click **Next** and enter the cluster details in the **Customize Workflow** page.
- Step 18** Click **Next** and view the cluster details in the **Summary** page.
- Step 19** Click **Submit** to trigger a workflow for creating a Hadoop cluster.
-

Provisioning an Instant and Customized Hadoop Cluster

Create and customize a Cluster Deploy Template to trigger the workflow.

Before you begin

- Create a UCS Service Profile template for a Customized Hadoop Cluster
- Create a Hadoop Cluster Profile template for a Customized Hadoop Cluster

Step 1 Choose **Policies > Orchestration**.

Step 2 Click the **UCS CPA** folder from the **Workflows** tab.

Step 3 Double-click the workflow to open the workflow designer and execute the workflow.

- a) When you open the workflow designer for an instant Hadoop Cluster, you get the following tasks, which are processed sequentially.

Task Name	Description
Instant Hadoop Cluster UCS SP	Cisco UCS Director Express for Big Data automatically specifies parameters for installing the OS and Hadoop distribution software at the back end.
Instant Hadoop Cluster Profile	Cisco UCS Director Express for Big Data automatically configures Hadoop cluster services at the back end.
Setup Hadoop Cluster Env	Sets up the environment for cluster-specific scripts and software files.
Muti Bare Metal OS Install WF	Attaches the UCS profile and sets up all boot files required to boot the operating system (Linux). When the Power ON task is executed, the boot files are picked up, and the operating system is installed successfully.
Multi Bare Metal WF Monitor	Checks the status of bare metal OS install workflow.
Synchronized Command Execution	—
Custom SSH Command	Installs and configures the Hadoop distribution software.
Provision Hadoop Cluster	Sends the Hadoop cluster properties to the Web Console.
Completed	The Hadoop cluster is provisioned successfully. Note If any of the tasks fail, you are informed that the provisioning has failed. For more information on how to monitor the workflow, see Monitoring Service Requests for Big Data .

- b) When you open the workflow designer for a customized Hadoop Cluster, you get the following tasks that get processed sequentially.

Task Name	Description
Create UCS Service Profile Template	Specifies parameters for installing the OS and Hadoop distribution software.
Create Hadoop Cluster Profile	Configures Hadoop cluster services.

Task Name	Description
Setup Hadoop Cluster Env	Sets up the environment for cluster-specific scripts and software files.
Muti Bare Metal OS Install WF	Attaches the UCS profile and sets up all boot files required to boot the operating system (Linux). When the Power ON task is executed, the boot files are picked up, and the operating system is installed successfully.
Multi Bare Metal WF Monitor	Checks the status of bare metal OS install workflow.
Synchronized Command Execution	—
Custom SSH Command	Installs and configures the Hadoop distribution software.
Provision Hadoop Cluster	Sends the Hadoop cluster properties to the Web Console.
Completed	The Hadoop cluster is provisioned successfully. Note If any of the tasks fail, you are informed that the provisioning has failed. For more information on how to monitor the workflow, see Monitoring Service Requests for Big Data .

Managing a Hadoop Cluster

You can manage an existing cluster.

- Step 1** Choose **Solutions > Big Data > Accounts**.
- Step 2** Click **Hadoop Accounts**.
- Step 3** Select an account and click **View Details**.
- Step 4** Click **Summary** to view the statistics data report for the selected Hadoop Account and the high-level report on the cluster and node account.
- Step 5** Click **Hosts** to view node details. For more information on **Hosts** page, see [Managing Nodes in a Cluster](#).
- Step 6** Click **Hadoop Clusters** to view the cluster details and list of actions that you can perform on an existing Hadoop cluster.
 - a) Click **Role Topology** to view the topology of the nodes. (This tab is not editable.)
 - b) Click **View Details** to view the inputs for the Hadoop cluster you have created, and to view the virtual network interface configuration. (This information is not editable.)
 - c) Click **Shutdown Cluster** to stop all the services and power off all the nodes in the Hadoop cluster.
 - d) Click **Start Cluster** to power up all the nodes in the Hadoop cluster and start all services in the Hadoop cluster.
 - e) Click **Restart Cluster** to power off and then power up all nodes in the Hadoop cluster.
 - f) 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**.
 - g) Click **Upgrade Cluster** to upgrade Hadoop distributions from the current version, if available.

Note For a derived account, install and configure pssh and clush for the nodes in the Hadoop cluster and password less between nodes.

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 or Hortonworks is the Hadoop distribution.

Note To enable high availability in a Hortonworks cluster, you require a minimum of four nodes and node 1, 2, and 3 should be selected as Journal nodes.

5. Click **Submit**.

h) Enable Hadoop cluster high availability.

1. Click **Enable High Availability**.
2. From the **Enable High Availability** screen, access the **Standby Name Node** drop-down list and choose the Standby Name Node.
3. Check a minimum of three nodes from the **Journal Nodes** table.

Note Journal nodes selection is recommended to be on first three Master nodes and should be increased by the node count in odd number only.

4. Click **Submit**.

i) Disable Hadoop cluster high availability. This action is supported only for Cloudera and Hortonworks.

1. Click **Disable High Availability**.
2. From the **Disable High Availability** screen, access the **Standby Name Node** drop-down list and choose the Standby Name Node.
3. Check a minimum of three nodes from the **Journal Nodes** table.

Note Journal nodes selection is recommended to be on first three Master nodes and should be increased by the node count in odd number only.

4. Click **Submit**.

j) Click **Cluster Snapshot** to view the snapshot.

k) Click **View Reports** to view performance and monitoring reports.

Step 7 Click **Hadoop Services** to view the list of Hadoop services and their status. You can do the following:

- a) Click **Start All Services** to start all Hadoop services, depending on their status.
- b) Click **Stop All Services** to stop all Hadoop services, depending on their status.
- c) Click **Add New Service** to add a new Hadoop service.
- d) Click **Edit Service** to start and stop a particular Hadoop service.

Step 8 Click **Hadoop Service Roles** to view the list of Hadoop services. You can do the following:

- a) To add a role to the cluster, do the following on the **Add Role** screen:
 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**.

- b) To start or stop any role that you have created, do the following on the **Start/Stop Role** screen:
 1. From the **Hadoop Service Roles** tab, choose the Hadoop service.
 2. Click **Start/Stop Role**.
 3. Click **Submit**.

- c) To delete a role in the cluster, do the following on the **Delete Role** screen:
 1. From the **Hadoop Service Roles** tab, choose the Hadoop service.
 2. Click **Delete**.
 3. Click **Submit**.

Step 9 Click **More Reports** to view the list of 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 **Hadoop Clusters**.

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.

Name	Description
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.
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** Choose **Solutions > Big Data > Accounts**.
 - Step 2** Click **Hadoop Accounts**.
 - Step 3** Choose the Hadoop Account for which you want to view the snapshot and click **View Details**.
 - Step 4** Click **Hadoop Clusters**.

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 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 **Hadoop Services**.

Step 5 Click **Add New Service**.

On the **Add New Service** screen, complete the following fields:

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

- Bare Metal Node—Any node that is available and is not associated with the Hadoop cluster.

- Step 1** Choose **Solutions > Big Data > Accounts**.
- Step 2** Click **Hadoop Accounts**.
- Step 3** Select an account and click **View Details**.
- Step 4** Click **Hosts** 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 Bare Metal Nodes	Adds bare metal nodes to the Hadoop cluster.
Add New Disks	Adds disks to the Hadoop cluster.
Remove Disks	Remove disks from the Hadoop cluster. Note This function is supported only in HDFS Data Node and Yarn Node Manager for Hortonworks and Cloudera clusters only. However, there is no restriction for MapR clusters and you can delete any disks from any nodes.
Disk Locator	Locates the disks in Cisco UCS Manager and turns on the LED on the selected disk of the server.
Report Metadata	

- 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.

Name	Description
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 Bare Metal	The node is removed from the cluster and disassociated from the service profile. The node becomes a bare metal server.
Host Mappings	Lists DNS entries of all the hosts in the Hadoop cluster.
Run Inventory	Collects the hardware inventory for the selected server. For example, disk details.

Delete Node and Delete Node to Bare Metal Actions in Cloudera and Hortonworks

When you perform the Delete Node or Delete Node to Bare Metal 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 Bare Metal action. This functionality is based on the Delete Node to Bare Metal flag setting as true or false. In addition, the Rollback UCS CPA Node Bare Metal workflow is created.

Deleting an Unreachable Node from Hadoop Distribution

This section explains about deleting the unreachable nodes.

Deleting an Unreachable Cluster Node from MapR Distribution

In a MapR distribution with four cluster node, when a cluster node is unreachable and the node status is displayed as **Critical** in the Cisco UCS Director Express for Big Data user interface, you can delete a cluster node by performing the following:

- Step 1** Click **Delete Node**. The status of the node becomes unknown and the node is deleted from the MapR user interface.
- Step 2** Click **Delete Node to Bare Metal**. The rollback of the node occurs.
- Step 3** Click **Delete Node** to delete the node from the Cisco UCS Director Express for Big Data user interface.

Note You can also delete a node by clicking **Delete Node to Bare Metal**. The node is not deleted but rollback of nodes occurs (refer CSCvg90939 bug). You need to manually delete the node from both the MapR user interface and Cisco UCS Director Express for Big Data user interface.

Deleting an Unreachable Cluster Node from Cloudera Distribution

In a Cloudera distribution with three fresh cluster node and one data node added through bare metal workflow, when a node is unreachable and the node status is displayed as **Bad**, you can delete the node by performing the following:

-
- Step 1** Click **Delete Node**. The status of the node becomes unknown and the commissioned state is displayed as deleted.
 - Step 2** Click **Delete Node to Bare Metal** to delete a Data or Edge node. The node is deleted from the Cloudera user interface and the status is not updated in the Cisco UCS Director Express for Big Data user interface.
 - Step 3** Click **Delete Node** to remove the node from the Cisco UCS Director Express for Big Data user interface (refer CSCvg90939 bug).

Note You can also delete a cluster node clicking **Delete Node to Bare Metal**. The node is deleted from both the Cloudera user interface and Cisco UCS Director Express for Big Data user interface.

Note In a four node Cloudera cluster, when one of the data node becomes unreachable you cannot delete the node as some cluster services require a minimum of three nodes.

Deleting an Unreachable Cluster Node from Hortonworks Distribution

In a Hortonworks distribution with three fresh cluster node and one data node added through bare metal workflow, when a node is unreachable and the node status is displayed as **Unknown**, you can delete the node by performing the following:

-
- Step 1** Click **Delete Node**. The status of the node becomes unknown and the commissioned state is displayed as deleted.
 - Step 2** Click **Delete Node to Bare Metal** to delete a Data node, Edge node, or a cluster node. The node is deleted from both the Ambari user interface and Cisco UCS Director Express for Big Data user interface.

Note You can also delete a Data node, Edge node, or a cluster node by clicking **Delete Node to Bare Metal**. The node is deleted from both the Ambari user interface and Cisco UCS Director Express for Big Data user interface.

Note In a four node Hortonworks cluster, when one of the data node becomes unreachable you cannot delete the node as some cluster services require a minimum of three nodes.

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



Note The **Add Managed Nodes** functionality is not supported in Cloudera when you select the node type as **Edge Node**.

-
- Step 1** Choose **Solutions > Big Data > Accounts**.
 - Step 2** Click **Hadoop Accounts**.
 - Step 3** Double-click the Hadoop account.
 - Step 4** Click **Hosts**.
 - 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** Choose **Solutions > Big Data > Accounts**.
 - Step 2** Click **Hadoop Accounts**.
 - Step 3** Double-click the Hadoop account.
 - Step 4** Click **Hosts**.
 - 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 the node type from the **Node Type** drop-down list.
- Step 12** 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 13** Click **Submit**.

Adding Bare Metal Nodes to the Hadoop Cluster

Add bare metal nodes to the Hadoop cluster.



Note To add bare metal nodes to the Hadoop clusters using RHEL 7.4 or CentOS7.4 (created prior to Release 6.6.0.1), create a service profile template in Cisco UCS Manager with UEFI boot option.

- Step 1** Choose **Solutions > Big Data > Accounts**.
- Step 2** Click **Hadoop Accounts**.
- Step 3** Double-click the Hadoop account.
- Step 4** Click **Hosts**.
- Step 5** Click **Add Bare Metal Nodes**.
- Step 6** Create a service profile template in Cisco UCS Manager with UEFI boot option, if you want to add bare metal nodes to the Hadoop clusters using RHEL 7.4 or CentOS7.4 (created prior to Release 6.6.0.1).
- Step 7** On the **Add Bare Metal Nodes** screen, 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.
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.

Name	Description
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.
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	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). If the check box is not selected, the disk availability for both the OS disk and data disk are validated based on their RAID level.
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.
UCSTemplate Name table	Check the UCS Service Profile Template check box that you want to use and click Submit to confirm the selection.

Step 8

If you want to edit a Hadoop Server Role, select the row for that role, and click **Edit**.

Step 9

On the **Edit Hadoop Server Roles Entry** screen, complete the following fields and click **Submit**.

Name	Description
Node Type field	Displays the Hadoop node role. Note Kafka node is supported only for Cloudera and Hortonworks clusters.
Node Count field	The number of nodes in the Hadoop cluster for the selected node type.
SSD Boot Drives Available for OS check box	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 SSD. If the check box is not selected, the disk availability for both the OS disk and data disk are validated based on their RAID level. Note This check box is not displayed when the UCSM version is greater than or equal to 3.

Name	Description
Server Pool table	Enter the server pool that you want to use for the cluster for the selected node type. The Cisco UCS Manager account and the organization that you choose determine which server pools are displayed in this area.

Step 10 In the **vNIC Template** table, verify the vNIC templates available for the cluster.

Step 11 If you want to edit a vNIC template, select the row for that template and click **Edit**.

Step 12 On the **Edit vNIC Template Entry** screen, 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 13 Click **Submit**.

Adding Disks to the Hadoop Cluster

Step 1 Choose **Solutions > Big Data > Accounts**.

Step 2 Click **Hadoop Accounts**.

Step 3 Double-click the Hadoop account.

Step 4 Click **Hosts**.

Step 5 From the **More Actions** drop-down list, choose **Add New Disks**.

Step 6 Expand **Select disks** and check each disk that you want to use.

Step 7 Choose the method to write the data in the disks from the **Write Mode** drop-down list.

Step 8 Choose the method to read data from the disks from the **Read Mode** drop-down list.

Step 9 Check **Use Cache** to use the RAID controller cache to read and write operations.

Step 10 Check **Use Cache if Bad BBU** to ensure that if the Battery Backup Unit (BBU) is not available for any reason, **Write back** will be disabled and **Write Through** will be enabled.

Step 11 Choose a strip size for each disk within a stripe from the **Stripe Size (MB)** drop-down list.

Step 12 Click **Submit**.

Service Roles

If you use Add Node Bare Metal, 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
NodeManager	Yes	Yes	Yes
Ganglia Monitor	No	No	Yes
NFS Gateway	No	Yes	No

