Configuring and Deploying Hadoop Cluster Deployment Templates

This chapter contains the following sections:

Creating an Instant Hadoop Cluster

Before You Begin

- Create a server pool in the Cisco UCS Manager account you plan to use for this cluster
- Create a MAC address pool

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

Step 2   Click the **Hadoop Cluster Deployment Templates** tab.

Step 3   Click **Instant Hadoop Cluster**.

Step 4   In the **Instant Hadoop Cluster Creation** 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
SSH (root) Password field	The SSH root password. **Note** The SSH username pertains to the root user.
Confirm SSH Password field	Enter the SSH root password.
Hadoop Manager Password field	The management console password.
Confirm Hadoop Manager Password field	Enter the management console password.
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.
Big Data Cluster Type drop-down list	Choose the type of Hadoop distribution to be used for 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.
Multi UCSM check box	Check the Multi UCSM check box if you use multiple UCSM accounts. The following workflows are created during an Instant/Customized Hadoop Cluster creation: • UCS CPA Multi-UCSM Hadoop Cluster WF • Single UCSM Server Configuration WF (This WF will be triggered per UCSM Account. For example, UCSM 120, UCSM121) • UCS CPA Node BareMetal (This WF will be 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.

Name	Description
Server UUID Pool drop-down list	Choose the UUID pool to be used for servers in this cluster.
	Note 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.
	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.
UUID Prefix drop-down list	Choose the UUID prefix.
UUID Pool-From field	To specify manually, enter the UUID pool to start from.
UUID Pool-Size field	To specify manually, enter the UUID pool size.
PXE VLAN ID field	Enter the PXE VLAN ID.
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 5 In the **vNIC Template** table, review and, if required, edit the vNIC templates available for the cluster.

Step 6 If you want to edit a vNIC template, select the row for that template and click **Edit**.

Step 7 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.

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 8   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 Hadoop Cluster Deployment Template

Before You Begin

• Create a Cisco UCS Service Profile Template for Big Data

• Create a Hadoop Cluster Profile Template

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

Step 2   Click the **Hadoop Cluster Deploy Templates** tab.

Step 3   Click **Add (+)**.

Step 4   In the **Add Hadoop Cluster Deploy Template** dialog box, complete the following fields:

Name	Description
Hadoop Cluster Deploy Template Name field	A unique name for the Hadoop cluster deployment template.
Description field	A short description of the template.
UCS SP Template for Big Data drop-down list	Choose the UCS service profile template for Big Data that you want to use to create service profiles for the servers in the Hadoop cluster.
Hadoop Cluster Profile Template drop-down list	Choose the Hadoop cluster profile template that you want to use to configure the cluster services.

Step 5   Click **Add**.

Creating a Customized Hadoop Cluster

Before You Begin

- Create a Hadoop cluster deployment template that Cisco UCS Director Express for Big Data will use to create the Hadoop cluster

- Create a server pool in the Cisco UCS Manager account you plan to use for this cluster

- Create a MAC address pool

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

Step 2   Click the **Hadoop Cluster Deployment Templates** tab.

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

Step 4   In the **Customized Hadoop Cluster Creation** 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.
SSH (root) Password field	The SSH root password. **Note** The SSH username pertains to the root user.
Confirm SSH Password field	Enter the SSH root password.
Hadoop Manager Password field	The management console password.
Confirm Hadoop Manager Password field	Enter the management console password.
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
Multi UCSM check box	Check the Multi UCSM check box if you use multiple UCSM accounts.
	The following workflows are created during an Instant/Customized Hadoop cluster creation:
	• UCS CPA Multi-UCSM Hadoop Cluster WF
	• Single UCSM Server Configuration WF (This WF will be triggered per UCSM Account. For example, UCSM 120, UCSM121)
	• UCS CPA Node BareMetal (This WF will be 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.
Server UUID Pool drop-down list	Choose the UUID pool to be used for servers in this cluster.
	Note 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.
	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.
UUID Prefix drop-down list	Choose the UUID prefix.
UUID Pool-From field	To specify manually, enter the UUID pool to start from.
UUID Pool-Size field	To specify manually, enter the UUID pool size.
PXE VLAN ID field	Enter the PXE VLAN ID.
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.

Name	Description
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 5   In the **vNIC Template** table, review and, if desired, edit the vNIC templates available for the cluster.

Step 6   If you want to edit a vNIC template, select the row for that template and click **Edit**.

Step 7   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 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.
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 8   Click **Submit**.

What to Do Next

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

Provisioning an Instant and Customized Hadoop Cluster

You must create and customize a Hadoop 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 On the menu bar, go to **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 that gets 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 BareMetal 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 BareMetal WF Monitor	Checks the status of baremetal OS install workflow.
Wait for Admin Host	Waits for the Admin host to restart.
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 will be informed that the provisioning has failed. To know more 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 gets processed sequentially.

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

Task Name	Description
Create Hadoop Cluster Profile	Configures Hadoop cluster services.
Setup Hadoop Cluster Env	Sets up the environment for cluster specific scripts and software files.
Muti BareMetal 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 BareMetal WF Monitor	Checks the status of baremetal OS install workflow.
Wait for Admin Host	Waits for the Admin host to restart.
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 will be informed that the provisioning has failed. To know more on how to monitor the workflow, see Monitoring Service Requests for Big Data.