Cisco UCS CPA Workflows

This chapter contains the following sections:

Workflows for Big Data

Cisco UCS Director Express for Big Data defines a set of workflows in the **UCS CPA** folder under **Orchestration** page.

- UCS CPA Multi-UCSM Hadoop Cluster WF—This workflow is triggered if you use multiple UCSM Hadoop accounts.

- UCS CPA Multi-UCSM Splunk Cluster WF—This workflow is triggered if you use multiple UCSM Splunk accounts.

- UCS CPA Single UCSM Server Configuration WF—This workflow is triggered for every UCSM account.

- UCS CPA Node Bare Metal—This workflow is triggered per node in the cluster.

- UCS CPA Node Bare Metal with Storage Profile WF—This workflow is triggered using Cisco UCS S3260 storage server having SSD Server Boot Drive under SAS RAID controller to deploy OS.

Note   Disk group policies are automatically created in the Cisco UCS Director Express for Big Data using workflows only for Cisco UCS Manager 3.2 or later version.

Table 1: Disk Group Policy

Servers Details	Disk Group Policy Required?	Disk Policy Allocation Details
Servers having PCH Controller and and Data drive slots with SSD disks	No	Not applicable

Servers Details	Disk Group Policy Required?	Disk Policy Allocation Details
SAS controller with dedicated boot drive slots (20x) and Data drive slots with SSD disks	Yes	For OS boot, if you want to manually allocate the SSD slot numbers to RAID 0 or RAID 1, BD_V1_SSM_RAID0 or BD_V1_SSM_RAID1 should be created. Currently dedicated boot drive slots 201 and 202 are supported only by Cisco UCS S3260 Storage Server.
SAS controller with dedicated boot drive slots (20x) with SSD disks and Data drive slots without SSD disks	Yes	For OS boot, if you want to automatically allocate the SSD slot numbers to RAID 0 or RAID 1, Boot_SSD_RAID0 or Boot_SSD_RAID1 should be created.
SAS controller and Data drive slots with HDD only	Yes	For OS boot, if you want to automatically allocate the HDD slot numbers to RAID 0 or RAID 1, Boot_HDD_RAID0 or Boot_HDD_RAID1 should be created. **Note** You can create Boot_HDD_RAID0 or Boot_HDD_RAID1 if SSD disks are available in Data drive slots in addition to HDD.
SAS controller and Data drive slots with SDD and HDD disks	Yes	During OS boot, if you want to manually allocate the slot numbers to RAID 1, BD_V1_ANYM_RAID1 should be created.

Note In some cases, `Wait for Complete Association` might fail for Cisco UCS 326x storage servers during UCS CPA Node Bare Metal with Storage Profile after a time out. In such cases, association process might have already failed after `Bind UCS Service Profile to Template` post 90% of FSM with Storage System initialisation error. You need to perform a re-acknowledgement on the Cisco UCS server and resubmit the Bare Metal workflow.

- UCS CPA Delete Node—This workflow is triggered if you delete a node from the Hadoop cluster.

- UCS CPA Cloudera Add Live Node—This workflow is triggered if you add a Live Node for Cloudera.

- UCS CPA MapR Add Live Node—This workflow is triggered if you add a Live Node for MapR.

- UCS CPA Hortonworks Add Live Node—This workflow is triggered if you add a Live Node for Hortonworks.

- UCS CPA Add New Disks—This workflow is triggered if you add new disks.

- UCS CPA Configure NFS Clients and Servers—This workflow is triggered if you configure NFS clients and servers.

- UCS CPA Enable HTTPS Modify Protocol and Port Number for Hadoop—This workflow is triggered if you modify protocol and port numbers for Hadoop.

- UCS CPA Enable HTTPS Modify Protocol and Port Number for Splunk—This workflow is triggered if you modify protocol and port numbers for Splunk.

- UCS CPA Migrate Splunk Cluster to Multi-Site—This workflow is triggered if you want to migrate an existing Splunk cluster to a multi-site Splunk cluster.

- UCS CPA Remove Disks—This workflow is triggered if you remove disks.

- UCS CPA Splunk Upgrade—This workflow is triggered if you upgrade the Splunk cluster.

- UCS CPA Instant Hadoop Cluster WF—This workflow is triggered if you create an instant Hadoop cluster based on the node count and other mandatory inputs such as the IP address range, memory, and number of interfaces. Cisco UCS Director Express for Big Data automatically creates one UCS service profile and a Hadoop cluster profile template at the back end that are required to create an instant Hadoop cluster. This saves you the effort of manually creating a service profile and a cluster profile.

- UCS CPA Customized Hadoop Cluster WF—This workflow is triggered if you choose to use a specific UCS service profile. A Hadoop cluster profile template with the specified number of nodes to be created in the Hadoop cluster.

- UCS CPA Disable High Availability WF—This workflow is triggered when you disable high availability.

- UCS CPA Enable High Availability WF—This workflow is triggered when you enable high availability.

- UCS CPA Shutdown Big Data Cluster WF—This workflow is triggered when you shut down the Hadoop cluster.

- UCS CPA Start Big Data Cluster WF—This workflow is triggered when you power up the Hadoop cluster.

- UCS CPA Splunk Add Live Archival Node—This workflow is triggered when if you add a Live Archival Node for Splunk.

- UCS CPA Upgrade Cluster WF—This workflow is triggered when you upgrade the Hadoop cluster.

- UCS CPA Cloudera Add New Service WF—This workflow is triggered when you add a new service for Cloudera.

- UCS CPA MapR Add New Service WF–This workflow is triggered when you add a new service for MapR.

- UCS CPA Hortonworks Add New Service WF—This workflow is triggered when you add a new service for Hortonworks.

- UCS CPA Multi-Node Bare Metal provisioning WF—This workflow is triggered when you deploy baremetal OS account using Cisco UCS server.

Table 2: Disk Group Policy

Servers Details	Disk Group Policy Required?	Disk Policy Allocation Details
Servers having PCH Controller	No	Disk group policy is not applicable for PCH controllers, for example, Lewisburg SATA Controller. The LUN gets created automatically by the PCH controller based on the controller definition.
Servers having SAS controller	Yes	The disks are configured for RAID 1. You can manually allocate two disks.
Servers having SATA controller	Yes	The disks are configured for RAID 1. You can manually allocate two disks.

• Bare-metal OS Deployment —This workflow is triggered per server to deploy baremetal OS. UCS CPA Multi-Node Bare Metal provisioning WF will automatically trigger the **Bare-metal OS Deployment** workflow when you deploy baremetal OS account per server and monitor it.

About Service Requests for Big Data

Cisco UCS Director Express for Big Data leverages Cisco UCS Director service requests and workflow orchestration for the overall deployment of Hadoop clusters. Each service request is a provisioning workflow that is created during a cluster creation.

For example, one UCS CPA Multi-UCSM Hadoop Cluster W/F, one Single UCSM Server Configuration W/F, and four UCS CPA Node Bare Metal W/Fs are created for a four-node Hadoop cluster. When the workflows are complete, the cluster is created under **Solutions** > **Big Data** > **Accounts** for that UCSM account.

A set of service requests is created under **Organizations** > **Service Requests** during a cluster creation.

• UCS CPA Multi-UCSM Hadoop Cluster WF—This workflow is triggered if you use multiple UCSM Hadoop accounts. It also applies when you create an instant or customized Hadoop cluster.

• UCS CPA Multi-UCSM Splunk Cluster WF—This workflow is triggered if you use multiple UCSM Splunk accounts. It also applies to when you create an instant or customized Splunk cluster.

• UCS CPA Single UCSM Server Configuration WF—This workflow is triggered for every UCSM account.

• UCS CPA Node Bare Metal—This workflow is triggered per node in the cluster.

The following service requests are created when you add a Bare Metal Node or a Live Node to the cluster.

• UCS CPA Cloudera Add Live Node—This workflow is triggered if you add a Live Node for Cloudera.

• UCS CPA MapR Add Live Node—This workflow is triggered if you add a Live Node for MapR.

• UCS CPA Hortonworks Add Live Node—This workflow is triggered if you add a Live Node for Hortonworks.

• UCS CPA BigInsights Add Live Node—This workflow is triggered if you add a Live Node for BigInsights.

The following service requests are created as explained, below:

• UCS CPA Disable High Availability WF—This workflow is triggered when you disable high availability.

• UCS CPA Enable High Availability WF—This workflow is triggered when you enable high availability.

• UCS CPA Shutdown Big Data Cluster WF—This workflow is triggered when you shut down the Hadoop cluster.

• UCS CPA Start Big Data Cluster WF—This workflow is triggered when you power up the Hadoop cluster.

• UCS CPA Upgrade Cluster WF—This workflow is triggered when you upgrade the Hadoop cluster.

• UCS CPA Cloudera Add New Service WF—This workflow is triggered when you add a new service for Cloudera.

• UCS CPA Hortonworks Add New Service WF—This workflow is triggered when you add a new service for Hortonworks.

• UCS CPA MapR Add New Service WF–This workflow is triggered when you add a new service for MapR.

• UCS CPA BigInsights Add New Service WF–This workflow is triggered when you add a new service for BigInsights.

For more information on the service requests and workflows, see the following guides:

• *Cisco UCS Director Self-Service Portal Guide*

• *Cisco UCS Director Orchestration Guide*

Monitoring Service Requests for Big Data

Before you begin

Create and customize a Cluster Deploy Template to monitor a service request.

Step 1 Choose **Organizations** > **Service Requests**.

Step 2 Click **Service Requests**.

Step 3 Select the service request that you want to monitor and click **View Details**.

One of the following Request Status is displayed:

• Complete
• In Progress
• Cancelled
• Failed

Viewing UCS CPA Workflow Tasks

From the **Service Request Status** screen, you can view the following:

- Workflow Status

- Log

- Objects Created and Modified

- Input/Output

Note   You can only modify inputs for failed service requests.

Step 1   Choose **Organizations** > **Service Requests**.

You can see the list of user-specific service requests added to a specific group. For example, All User Groups.

Step 2   Choose the **Service Request ID** that you want to view.

Step 3   Double-click the **Service Request ID** that opens the **Service Request Status** screen. (You can also choose the Service Request ID by the workflow name associated with it, and click **View Details**. For example, choose the **UCS CPA Node Bare Metal** workflow and click **View Details**).

On the **Service Request Status** screen, you can view the following tasks for the workflow:

UCS CPA Multi-UCSM Hadoop Cluster WF/UCS CPA Multi-UCSM Splunk Cluster WF	UCS CPA Single UCSM Server Configuration WF	UCS CPA Node Bare Metal
The following tasks are associated with the **UCS CPA Multi-UCSM Hadoop Cluster WF/UCS CPA Multi-UCSM Splunk Cluster WF**: a. Initiated by Admin b. Multi-UCSM Hadoop Cluster profile/Multi-UCSM Splunk Cluster profile c. Setup Big Data Cluster Env d. User (Compute) Approval Required e. User (OS) Approval Required f. Multi-UCSM Configuration WF g. Multi Bare Metal WF Monitor h. Synchronized Command Execution i. User (Hadoop) Approval Required j. Custom SSH Command k. Monitor Shell Script Status l. Provision Hadoop Cluster/Provision Splunk Cluster m. SSH Command n. Monitor Shell Script Status o. Edit Hadoop Account/Edit Splunk Account p. Assign Big Data Account to Group q. Custom SSH Command r. Complete	The following tasks are associated with the **Single UCSM Server Configuration WF**: a. Initiated by Admin b. Create UCS Service Profile Template c. Change Maintenance Policy UCS SP Template d. Muti-Bare Metal OS Install WF e. Create UCS Disk Group Policy f. Create UCS Disk Group Policy g. Create UCS Disk Group Policy h. Create UCS Disk Group Policy i. Multi-Bare Metal WF Monitor j. Change Maintenance Policy for UCS SP Template k. Complete	

UCS CPA Multi-UCSM Hadoop Cluster WF/UCS CPA Multi-UCSM Splunk Cluster WF	UCS CPA Single UCSM Server Configuration WF	UCS CPA Node Bare Metal
		The following tasks are associated with the **UCS CPA Node Bare Metal**:
		a. Initiated by Admin
		b. Modify Workflow Priority (High)
		c. Assign Bare Metal SR ID
		d. Create Service Profile from Template
		e. Unbind UCS Service Profile from Template
		f. Modify UCS Service Profile Boot Policy
		g. Add VLAN to Service Profile
		h. Associate UCS Service Profile
		i. Assign ServerIdentity
		j. PXE Boot With BMA Selection
		k. Setup RAID Commands
		l. UCS Blade Reset Action
		m. PXE Boot Wait
		n. Monitor RAID Configuration
		o. Custom Archive PXE Boot Request
		p. UCS Blade Power OFF Action
		q. PXE Boot With BMA Selection
		r. Setup RAID Commands
		s. Custom Wait for Complete Association of SP
		t. UCS Blade Reset Action
		u. PXE Boot Wait
		v. Modify UCS Service Profile Boot Policy
		w. Delete VLAN from Service Profile vNIC

UCS CPA Multi-UCSM Hadoop Cluster WF/UCS CPA Multi-UCSM Splunk Cluster WF	UCS CPA Single UCSM Server Configuration WF	UCS CPA Node Bare Metal	
		x.	Bind UCS Service Profile to Template
		y.	Wait for complete Association
		z.	UCS Blade Reset Action
		aa.	Assign IP Status
		ab.	Custom SSH Command
		ac.	Monitor Shell Script Status
		ad.	Linux Shell Script Execution
		ae.	Linux Shell Script Execution
		af.	UCS Blade Power OFF Action
		ag.	UCS Blade Power ON Action
		ah.	Synchronized Command Execution
		ai.	Assign UCS Server to Group
		aj.	Assign Service Profile to Group
		ak.	Complete

Step 4 Click **Close**.

Viewing UCS CPA Workflow Tasks for BareMetal OS

From the **Service Request Status** screen, you can view the following:

- Workflow Status

- Log

- Objects Created and Modified

- Input/Output

Note You can only modify inputs for failed service requests.

Step 1   Choose **Organizations** > **Service Requests**. You can see the list of user-specific service requests added to a specific group. For example, All User Groups.

Step 2   Choose the **Service Request ID** that you want to view.

Step 3   Double-click the **Service Request ID** that opens the **Service Request Status** screen. (You can also choose the Service Request ID by the workflow name associated with it, and click **View Details**. For example, choose the **UCS CPA Node Bare Metal** workflow and click **View Details**). On the **Service Request Status** screen, you can view the following tasks for the workflow:

UCS CPA Multi-Node Bare Metal provisioning WF	Bare-metal OS Deployment

UCS CPA Multi-Node Bare Metal provisioning WF	Bare-metal OS Deployment
The following tasks are associated with the **UCS CPA Multi-Node Bare Metal provisioning WF**:	The following tasks are associated with the **Bare-metal OS Deployment**:
a. Initiated by Admin	**a.** Initiated by Admin
b. Get UCSM account authentication	**b.** Create Service Profile from Template
c. Get UCS Organization	**c.** Unbind UCS Service Profile from Template
d. Create Multi UCS Policy for Linux Account	**d.** Modify UCS Service Profile Boot Policy
e. Create Maintenance Policy	**e.** Get UCS Organization
f. Create BIOS Policy	**f.** Add PXE VLAN to Service Profile (PXE VLAN is already created in the UCS CPA Multi-Node Bare Metal provisioning WF and passed here)
g. Create Local Boot Policy	
h. Create LAN Boot Policy	**g.** Associate Storage Profile to Service Profile (Storage Profile is already created in the UCS CPA Multi-Node Bare Metal provisioning WF and passed here)
i. Check Scrub Policy is required	
j. Create Scrub Policy (note only executed when check scrub policy is required is true)	
	h. Checking Set JBOD contains "true"
k. Check Selected slot is SATA	**i.** Set JBOD to Unconfigured Good (only executed when checking Set JBOD contains "true")
l. Create Local Disk Config Policy (note only executed when check Selected slot is SATA is true)	
	j. Select UCS Server
m. Create Service Profile Template	**k.** Check if Delete LUN is enabled
n. Get UCSM Account Name	**l.** Delete LUNs From Server (only execute when check if Delete LUN is enabled)
o. Create PXE VLAN	
p. Modify Org Permission for PXE VLAN	**m.** Reserve UCS IP address
q. Get PXE VLAN ID	**n.** Associate UCS Service Profile
r. Check if server have PCH	**o.** Check if server have SAS
s. Add UCS Storage Profile (only executed when Check if server have PCH is true)	**p.** Get UCS LUN ID (only execute when Check if server have SAS is true)
t. Create Storage Profile (only executed when Check if server have PCH is false)	**q.** Setup PXE Boot (OS Type: RHEL7.6) (only execute when Check if server have SAS is true)
u. Get UCS Storage Profile (only executed when Check if server have PCH is false)	**r.** Setup PXE Boot (OS Type: RHEL7.6) (only execute when Check if server have SAS is false)
v. Collect Inventory	
w. Request UCSM Inventory Collection By DN (for service profile template inventory)	**s.** Modify PXE Configuration for SATA Controller (only execute when Check if server have SAS is false)
x. Wait for PXE VLAN	
y. Check if server have PCH	

UCS CPA Multi-Node Bare Metal provisioning WF		Bare-metal OS Deployment	
z.	Request UCSM Inventory Collection By DN (for storage profile inventory & executed only when Check if server have PCH is false)	t.	Wait for Complete Association
		u.	UCS Blade Power ON Action
aa.	Wait for Service Profile Template Status	v.	Check if server have SAS
ab.	Install Multi Linux on UCS Server for Linux Account (it will automatically trigger the Bare-metal OS Deployment workflow per server)	w.	Monitor PXE Boot
		x.	Delete PXE VLAN from MGMT vNIC
		y.	Bind UCS Service Profile to Template
ac.	Multi Bare Metal Account WF Monitor (it will monitor the Bare-metal OS Deployment workflow SR's)	z.	UCS Blade Power OFF Action
ad.	Create Baremetal OS Account (once all the Bare-metal OS Deployment is successful, the baremetal OSs account will get created)	aa.	UCS Blade Power ON Action
		ab.	IPconfig
		ac.	vNIC Bonding
ae.	Complete	ad.	Complete

Step 4   Click **Close**.

Workflow Customization to Deploy a Hadoop or Splunk Cluster

You can customize the following UCS CPA workflows and use them to deploy a Hadoop or Splunk cluster. You can add installation packages (RPMs) required for your cluster environment in the Cisco UCS Director Express for Big Data Bare Metal Agent.

- Rename the UCS CPA Multi-UCSM Hadoop Cluster WF

- Rename the UCS CPA Multi-UCSM Splunk Cluster WF

- Rename the Single UCSM Server Configuration WF

- Rename the UCS CPA Node Bare Metal

Deploying a Hadoop or Splunk Cluster Through Workflow Customization

Before you begin

For more information on workflow orchestration, see the *Cisco UCS Director Orchestration Guide*.

- Customize UCS CPA Node Bare Metal workflows that you want to use in the cloned Single UCSM Server Configuration WF.

- Customize the Single UCSM Server Configuration WF that you want to use in the UCS CPA Multi-UCSM Hadoop Cluster WF or UCS CPA Multi-UCSM Splunk Cluster WF.

Step 1   Choose **Orchestration**.

Step 2	Click the **UCS CPA** folder from **Workflows**.
Step 3	Double-click the workflow that you want to customize in the **Workflow Designer**. For instance, double-click the **UCS CPA Multi-UCSM Hadoop Cluster WF**.
Step 4	Double-click the **Muti-UCSM Configuration WF** task in the **Workflow Designer**.
Step 5	Click **Next** on the **Workflow Task Basic Information** page.
Step 6	On the **User Input Mappings to Task Input Attributes** page, select the attributes that you want to map to the workflow input fields or provide values in the next step. If necessary, check the **Map to User Input** check box to provide user inputs.
Step 7	Click **Next** on the **User Input Mappings to Task Input Attributes** page.
Step 8	Enter the task values which are not mapped to workflow inputs. For example, enter the name of the cloned Single UCSM Server Configuration WF in the **Workflow Name** field.
Step 9	Click **Revalidate** to validate task input values.
Step 10	Click **Next**.
Step 11	On the **User Output mappings to Task Output Attributes** page, select the attributes that you want to map to the workflow output fields.
Step 12	Check the **Map to User Output** check box and choose the value from the **User Output** drop-down list.
Step 13	Click **Submit**.

Assigning Big Data Accounts to User Groups

Step 1	Choose **Orchestration** and click **Workflows**.
Step 2	Click **Add Workflow**.
Step 3	On the **Add Workflow Details** page, enter the workflow name and choose a folder. Click **Next**.
Step 4	On the **Workflow Task Basic Information** page, enter the required details and click **Next**.
Step 5	On the **User Input Mappings to Task Input Attributes** page, select the attributes that you want to map to the workflow input fields or provide values in the next step. If necessary, check the **Map to User Input** check box to provide user inputs..
Step 6	On the **User Input Mappings to Task Input Attributes** page, choose big data account type and big data account.
Step 7	Check the **Assign Users** checkbox and click **Submit**. The big data accounts are assigned to specified users if the **Allow resource assignment to users** option is enabled in the **Users Groups** page. When the **Assign Users** checkbox is not checked, the big data accounts are assigned to a specified group.

Unassigning Big Data Accounts

Before you begin

Step 1	Choose **Orchestration** and click **Workflows**.
Step 2	Click **Add Workflow**.
Step 3	On the **Add Workflow Details** page, enter the workflow name and choose a folder. Click **Next**.
Step 4	On the **Workflow Task Basic Information** page, enter the required details and click **Next**.

Step 5 On the **User Input Mappings to Task Input Attributes** page, select the attributes that you want to map to the workflow input fields or provide values in the next step. If necessary, check the **Map to User Input** check box to provide user inputs..

Step 6 On the **User Input Mappings to Task Input Attributes** page, choose big data account type and big data account.

Step 7 Check the **Unassign Users** checkbox and click **Submit**. The big data accounts are unassigned from the group.

Cloning UCS CPA Workflows

To customize cluster deployment through bare metal workflows, you can clone the following workflows in the UCS CPA folder:

- Clone the UCS CPA Node Bare Metal workflows.

- Rename the Single UCSM Server Configuration WF using the cloned UCS CPA Node Bare Metal workflows.

- Rename the UCS CPA Multi-UCSM Hadoop Cluster WF using the cloned Single UCSM Server Configuration WF.

Step 1 Choose **Orchestration**.

Step 2 Click the **UCS CPA** folder from **Workflows** and choose the workflow that you want to clone.

Step 3 Click **Clone Workflow**.

Step 4 On the **Workflow Details** page of the **Clone Workflow** wizard, complete the following fields:

Name	Description
Workflow Name field	A unique name for the workflow.
Version field	The current version of the workflow that you are cloning. This is a display-only field.
Description field	The description of the workflow.
Workflow Context drop-down list	The workflow context. Workflow Orchestration supports the following options: - **Any**—Enables you to use the workflow in any context. - **Selected VM**—Enables you to use the execute workflow. This option can be selected only when you choose a VM. - Check the **Save As Compound Task** check box to define the workflow as a compound task. - Check the **Place in New Folder** check box, and enter the folder name in the **Folder Name** field, to assign the workflow to a new folder other than the UCS CPA folder.

Name	Description
Select Folder drop-down list	Choose a folder. UCS CPA is the default folder for Big Data.
Notify status of execution to initiator User check box	Check the check box to notify the user through email, then enter appropriate email addresses in the **Additional User(s) to send Email Notification** field.

Step 5   Click **Next**.

Step 6   On the **Workflow User Inputs** page of the **Clone Workflow** wizard, complete the following fields:

Name	Description
Associate to Activity check box	If the check box is checked then any existing workflow's user input is overridden by any selected activity user input.
Activity drop-down list.	Choose an activity. The user-input table is updated based on the selected activity.
Workflow User Inputs table	On the **Workflow User Inputs** page: a. Click the + icon to add workflow input properties. b. On the **Add Entry to** screen, complete the following fields: 1. Enter the name for the activity workflow input in the **Input Label** field. 2. Enter the description for the activity workflow input in the **Input Description** field. 3. Check the **Optional** check box to set the input as optional during workflow execution. 4. Click **Select**. On the **Select** screen, click **Input Type**. 5. Click **Submit**.

Step 7   Click **Next**.

Step 8   On the **Workflow User Outputs** page of the **Clone Workflow** wizard, do the following:

Name	Description
Workflow User Outputs table	On the **Workflow User Outputs** page: a. Click the + icon to add workflow output properties. b. On the **Add Entry to** screen, complete the following fields: 1. Enter the name for the activity workflow output in the **Output Label** field. 2. Enter the description for the activity workflow output in the **Output Description** field. 3. Check the **Optional** check box to set the output as optional during workflow execution. 4. Click **Select**. On the **Select** screen, click **Output Type**. 5. Click **Submit**.

Step 9   Click **Submit**.