



## Overview

---

This chapter contains the following sections:

- [Cisco UCS Director Express for Big Data, page 1](#)
- [Cisco UCS Integrated Infrastructure for Big Data, page 2](#)
- [Managing Cisco UCS Director and Cisco UCS Director Express for Big Data Personalities, page 2](#)
- [Creating User Roles, page 3](#)
- [Supported Hadoop Distributions, page 9](#)
- [Supported Splunk Distribution, page 9](#)
- [High-level Workflow to Create an Instant Hadoop Cluster, page 10](#)
- [High-level Workflow to Create a Customized Hadoop Cluster, page 10](#)

## Cisco UCS Director Express for Big Data

Cisco UCS Director Express for Big Data is a single-touch solution within Cisco UCS Director that automates deployment of Big Data infrastructure. Cisco UCS Director Express for Big Data provides a single management pane across physical infrastructure and across Hadoop and Splunk Enterprise software. It supports key Hadoop distributions, including Cloudera, MapR, and Hortonworks.

Cisco UCS Director Express for Big Data delivers end-to-end automation of Hadoop cluster deployment, allowing you to spin up and expand clusters on-demand. The physical infrastructure configuration is handled automatically, with minimal user input. The configuration includes compute, internal storage, network, and installation of operating system, Java packages, and Hadoop, along with the provisioning of Hadoop services. This is achieved through Cisco UCS service profiles wherein both the physical infrastructure and Hadoop configuration are incorporated into a Hadoop cluster deployment profile.

Cisco UCS Director Express for Big Data also delivers end-to-end automation of Splunk cluster deployment, with minimal user input. This is achieved through Cisco UCS service profiles wherein both the physical infrastructure and Splunk configuration are incorporated into a Splunk cluster deployment profile.

# Cisco UCS Integrated Infrastructure for Big Data

Cisco UCS Integrated Infrastructure for Big Data is an industry leading architecture designed to meet various Big Data workloads. It scales as processing and storage needs grow without increasing management challenges and delivers predictable performance along with reduced total cost of ownership (TCO).

Cisco UCS Integrated Infrastructure consists of the following components:

- Cisco UCS Fabric Interconnects
- Cisco UCS 2200 Series Fabric Extenders
- Cisco UCS C-Series Rack-Mount Servers
- Cisco UCS Virtual Interface Cards (VICs)
- Cisco UCS Manager

You can read more about the Cisco UCS Integrated Infrastructure for Big Data in the [Data Center Designs Cloud Computing - Design Zone for Big Data](#) .

## Managing Cisco UCS Director and Cisco UCS Director Express for Big Data Personalities

Cisco UCS Director is the default personality made available after deployment, but you can choose to use only Cisco UCS Director Express for Big Data, or use both Cisco UCS Director and Cisco UCS Director Express for Big Data.

You can manage personalities here: **Administration > License > License Keys > Manage Personalities**.

**Table 1: Personality Switch Behavior**

Personality Selection	Cisco UCS Director Features	Cisco UCS Director Express for Big Data Features
Cisco UCS Director, Release 6.0 (Default)	Yes	No
Cisco UCS Director Express for Big Data, Release 3.0	No	Yes
Cisco UCS Director, Release 6.0, and Cisco UCS Director Express for Big Data, Release 3.0	Yes	Yes

**Note**

Depending on the personality, you start with and the personality selection, Cisco UCS Director and Cisco UCS Director Express for Big Data features are enabled or disabled with the restart of services on the appliance.

## Creating User Roles

You can create user roles that are specific to Cisco UCS Director Express for Big Data, and define menu settings and permissions for those users. Ensure that you create a group before you add users to any role.

**Note**

You can determine the default roles only if the **Default Role** column on the **User Roles** page is marked with **Yes** in the system. Navigate to **Administration > System > User Roles**.

For example, you can create the following user roles, and then create users with those role:

- HadoopUser—A Hadoop user
- HadoopAdmin—A Hadoop administrator

For more information on Managing Users and Groups, see the latest *Cisco UCS Director Administration Guide*.

## Hadoop Administrator Permissions

A Hadoop administrator can:

- Read—Permission to only read a file.
- Write—Permission to read, write, and modify a file.
- Read/Write—Permission to read and/or write to a file

The following table shows a list of operations that a Hadoop administrator can do:

Operations	Permissions	
	Read	Write
Virtual Computing	Yes	Yes (Only VM Management Actions)
VM Label	Yes	—
Assign VM to VDC	Yes	—
Virtual Storage	Yes	Yes
Virtual Network	Yes	Yes

Operations	Permissions	
	Read	Write
Physical Computing	Yes	Yes
Physical Storage	Yes	Yes
Physical Network	Yes	Yes
Group Service Request	Yes	Yes
Approver Service Request	Yes	Yes
Budgeting	Yes	Yes
Resource Accounting	Yes	—
Chargeback	Yes	—
System Admin	Yes	Yes
Users and Groups	Yes	Yes
Virtual Accounts	Yes	Yes
Catalogs	Yes	Yes
VDC	Yes	Yes
Computing Policy	Yes	Yes
Storage Policy	Yes	Yes
Network Policy	Yes	Yes
Service Delivery	Yes	Yes
Resource Limit Report	Yes	Yes
Group Users	Yes	Yes
Cloudsense Reports	Yes	Yes
Cloudsense Assessment Reports	Yes	Yes
Orchestration	Yes	Yes
Open Automation Modules	Yes	Yes
CS Shared Reports	Yes	Yes

Operations	Permissions	
	Read	Write
Remote VM Access	—	Yes
Mobile Access Settings	Yes	Yes
End User Chargeback	Yes	—
Resource Groups	Yes	Yes
Tag Library	Yes	Yes
Big Data Infra	Yes	—
Big Data Accounts	—	Yes
Big Data Cluster Management	—	Yes
Big Data Node Management	—	Yes
Big Data Performance Test	—	Yes
Big Data Service Management	—	Yes
Big Data Role Management	—	Yes
Big Data UCS SP Template	—	Yes
Big Data Hadoop Profile Template	—	Yes
Big Data Hadoop Deploy Template	—	Yes
Big Data Cluster Deployment	—	Yes
Big Data License Upload	—	Yes
Big Data Configuration Parameters Template	—	Yes
Big Data Faults	—	Yes
Big Data Settings - QoS	—	Yes
Big Data Settings - IP Pool	—	Yes
Big Data Settings - Pre_Cluster Sanity	—	Yes

Operations	Permissions	
	Read	Write
Big Data Settings - Hadoop Software Upload	—	Yes
Big Data Settings - Configuration Check Rules	—	Yes
REST API access	Yes	Yes
Allow Change Password - Users	Yes	Yes

## Hadoop User Permissions

A Hadoop user can:

- Read—Permission to only read a file.
- Write—Permission to read, write, and modify a file.
- Read/Write—Permission to read and/or write to a file.

The following table shows a list of operations that a Hadoop user can do:

Operations	Permissions	
	Read	Write
Virtual Computing	Yes	—
VM Label	Yes	—
Assign VM to VDC	Yes	—
Virtual Storage	Yes	—
Virtual Network	Yes	—
Physical Computing	Yes	—
Physical Storage	Yes	—
Physical Network	Yes	—
Group Service Request	Yes	Yes
Approver Service Request	Yes	Yes

Operations	Permissions	
	Read	Write
Budgeting	Yes	—
Resource Accounting	Yes	—
Chargeback	Yes	—
System Admin	Yes	—
Users and Groups	Yes	—
Virtual Accounts	Yes	—
Catalogs	Yes	—
VDC	Yes	—
Computing Policy	Yes	—
Storage Policy	Yes	—
Network Policy	Yes	—
Service Delivery	Yes	—
Resource Limit Report	Yes	—
Group Users	Yes	—
Cloudsense Reports	Yes	—
Cloudsense Assessment Reports	Yes	—
Orchestration	—	—
Open Automation Modules	—	—
CS Shared Reports	—	—
Remote VM Access	—	—
Mobile Access Settings	—	—
End User Chargeback	—	—
Resource Groups	—	—

Operations	Permissions	
	Read	Write
Tag Library	—	—
Big Data Infra	Yes	—
Big Data Accounts	—	—
Big Data Cluster Management	—	—
Big Data Node Management	—	—
Big Data Performance Test	—	—
Big Data Service Management	—	—
Big Data Role Management	—	—
Big Data UCS SP Template	—	—
Big Data Hadoop Profile Template	—	—
Big Data Hadoop Deploy Template	—	—
Big Data Cluster Deployment	—	—
Big Data License Upload	—	—
Big Data Configuration Parameters Template	—	—
Big Data Faults	—	—
Big Data Settings - QoS	—	—
Big Data Settings - IP Pool	—	—
Big Data Settings - Pre_Cluster Sanity	—	—
Big Data Settings - Hadoop Software Upload	—	—
Big Data Settings - Configuration Check Rules	—	—
REST API access	Yes	—



Operations	Permissions	
	Read	Write
Allow Change Password - Users	—	—

## Supported Hadoop Distributions

Cisco UCS Director Express for Big Data supports the following Hadoop distributions:

- Cloudera 5.4, 5.5, 5.6, and 5.8
- MapR 3.1, 4.0, 4.1, 5.0, 5.1, and 5.2
- Hortonworks 2.1, 2.2, 2.3, 2.4, and 2.5

**Table 2: Java and JDK Software for Hadoop Distributions**

Hadoop Distribution	Supported Hadoop Distribution Versions	Installed Java and JDK
Cloudera	5.0.1, 5.0.6, 5.2.0, 5.2.1, 5.3.0, 5.4.1, 5.5.0, 5.6.0, and 5.8.0	oracle-j2sdk1.7
MapR	3.1.1, 4.0.1, 4.0.2, 4.1.0, 5.0.0, 5.1.0, and 5.2.0	java-1.7.0-openjdk
Hortonworks	2.1, 2.2, 2.3, 2.4, and 2.5	java-1.7.0-openjdk



**Note**

You can choose and install the required Oracle JDK version.

## Supported Splunk Distribution

Cisco UCS Director Express for Big Data supports the following Splunk distribution:

Splunk Distribution	Supported Splunk Distribution Version
Splunk	6.3.x and 6.4.x

## High-level Workflow to Create an Instant Hadoop Cluster

- 
- Step 1** Create a Cisco UCS Manager account. For more information, see [Adding a Cisco UCS Manager Account](#).
  - Step 2** Configure Big Data IP pools. For more information, see [Adding a Big Data IP Pool](#).
  - Step 3** Create an Instant Hadoop Cluster. For more information, see [Creating an Instant Hadoop Cluster](#).
- 

## High-level Workflow to Create a Customized Hadoop Cluster

- 
- Step 1** Configure a Cisco UCS Service Profile template for Big Data. For more information, see [Creating a Cisco UCS Service Profile Template for Big Data](#).
  - Step 2** Configure a Hadoop cluster profile template. For more information see, [Creating a Hadoop Cluster Template](#).
  - Step 3** Configure a Hadoop cluster deployment template. For more information, see [Creating a Hadoop Cluster Deploy Template](#).
  - Step 4** Create a customized Hadoop cluster. For more information, see [Creating a Customized Hadoop Cluster](#).
-