



Big Data Cluster Configuration Settings

This chapter contains the following sections:

- [Creating an External Database Configuration, on page 1](#)
- [Creating a Hadoop Cluster Configuration Parameters Template, on page 3](#)
- [Updating Hadoop Cluster Configuration Parameters Template - Post Hadoop Cluster Creation, on page 4](#)
- [Quality of Service System Classes, on page 4](#)
- [Pre Cluster Performance Testing Settings, on page 7](#)
- [Approving Hadoop Cluster and Splunk Deployment Workflows, on page 7](#)
- [Adding NTP Server Details, on page 9](#)
- [Uploading Required OS and Big Data Software to Cisco UCS Director Bare Metal Agent , on page 9](#)
- [Cloudera, MapR, and Hortonworks RPMs on Cisco UCS Director Express for Big Data Bare Metal Agent, on page 13](#)
- [Cloudera and MapR RPMs for Upgrading Hadoop Cluster Distributions, on page 19](#)
- [Installation of User-Defined Software Post Hadoop Cluster Creation, on page 21](#)
- [Configuration Check Rules, on page 21](#)
- [Checking Hadoop Cluster Configuration, on page 22](#)
- [Fixing Configuration Violations, on page 22](#)

Creating an External Database Configuration

You can deploy each Hadoop cluster with its own external database for all Hadoop distributions (Cloudera, MapR, and Hortonworks) using instant Hadoop cluster and customized Hadoop cluster creation actions.

You can configure a new database or use an existing database in Cisco UCS Director Express for Big Data. The Oozie, Hive, and Hue services use configured database information that you have created using the **Create External Database Configurations** dialog.



Note MySQL is the only supported external database in Cisco UCS Director Express for Big Data.

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

Step 2 Click **External Database**.

Step 3 Click **Add**.

You can modify or delete any database you have previously created from the external database table.

Step 4 On the **Create External Database Configurations** screen, complete the following fields:

Name	Description
Database Name field	Enter a unique name for the database type you want to create.
Database Type field	Choose the database type from the list.
Server Name field	Enter an IPv4 address for the database server.
Port field	Enter a port number based on the database type.
User Name field	Enter a username to access the database server.
Password field	Enter the password to access the database server.
Confirm Password field	Confirm the password to access the database server.

Step 5 Click **Submit**.

What to do next

Deploy Hadoop clusters through instant Hadoop cluster and customized Hadoop cluster creation actions.

Default Databases Used in Hadoop Distribution Services

Default Databases for Cloudera (Service Names):

- Cloudera Manager—mysql
- Oozie—mysql
- Hive—mysql
- Hue—mysql

Default Databases for MapR (Service Names):

- Oozie—Derby
- Hive—mysql
- Hue—SQLite

Default Databases for Hortonworks (Service Names):

- Ambari—PostGres

- Oozie—Derby
- Hive—mysql

Creating a Hadoop Cluster Configuration Parameters Template

You can create the Hadoop Cluster Configuration Parameters Template only from the **Hadoop Config Parameters** tab on the menu bar here: **Solutions > Big Data > Settings** before triggering a Hadoop cluster. You can select the Hadoop cluster configuration parameters template to edit, clone, or delete.

-
- Step 1** Choose **Solutions > Big Data > Settings**.
- Step 2** Click **Hadoop Config Parameters**.
- Step 3** Click **Add**.
- Step 4** On the **Hadoop Config Parameters** page of the **Create Hadoop Cluster Configuration Parameters Template** wizard, complete the following fields:

Name	Description
Template Name field	A unique name for the Hadoop cluster configuration parameter template.
Template Description field	The description for the Hadoop cluster configuration parameter template.
Hadoop Distribution drop-down list	Choose the Hadoop distribution.
Hadoop Distribution Version drop-down list	Choose the Hadoop distribution version.

- Step 5** Click **Next**.
- Step 6** On the **Hadoop Config Parameters - HDFS Service** page of the **Create Hadoop Cluster Configuration Parameters Template** wizard, specify the Hadoop cluster HDFS service parameter name, value, and the minimum supported Hadoop distribution.
- Step 7** On the **Hadoop Config Parameters - YARN Service** page of the **Create Hadoop Cluster Configuration Parameters Template** wizard, configure the parameters.
- Step 8** On the **Hadoop Config Parameters - HBase Service** page of the **Create Hadoop Cluster Configuration Parameters Template** wizard, configure the parameters.
- Step 9** On the **Hadoop Config Parameters - MapReduce Service** page of the **Create Hadoop Cluster Configuration Parameters Template** wizard, configure the parameters.
- Step 10** On the **Hadoop Config Parameters - Zookeeper Service** page of the **Create Hadoop Cluster Configuration Parameters Template** wizard, configure the parameters.
- Step 11** On the **Hadoop Config Parameters - SmartSense Service** page of the **Create Hadoop Cluster Configuration Parameters Template** wizard, configure the parameters.
- Step 12** On the **Hadoop Config Parameters - Miscellaneous Parameters** page of the **Create Hadoop Cluster Configuration Parameters Template** wizard, configure the (ServiceLevel and RoleLevel) parameters.
- Step 13** Click **Submit**.
-

Updating Hadoop Cluster Configuration Parameters Template - Post Hadoop Cluster Creation

- Step 1** Choose **Solutions > Big Data > Accounts**.
- Step 2** Click the **Hadoop Accounts** tab and choose an existing Hadoop Account.
- Step 3** Click **Configure Cluster**.
- Step 4** On the **Hadoop Config Parameters** page of the **Update Hadoop Cluster Configuration Parameters Template** wizard, choose the Hadoop distribution.
- Step 5** Click **Next**.
- Step 6** On the **Hadoop Config Parameters - HDFS Service** page of the **Update Hadoop Cluster Configuration Parameters Template** wizard, specify the Hadoop cluster HDFS service parameter name, value, and the minimum supported Hadoop distribution version, if any.
- Step 7** On the **Hadoop Config Parameters - YARN Service** page of the **Update Hadoop Cluster Configuration Parameters Template** wizard, update the parameters as required.
- Step 8** On the **Hadoop Config Parameters - HBase Service** page of the **Update Hadoop Cluster Configuration Parameters Template** wizard, update the parameters as required.
- Step 9** On the **Hadoop Config Parameters - MapReduce Service** page of the **Update Hadoop Cluster Configuration Parameters Template** wizard, update the parameters as required.
- Step 10** On the **Hadoop Config Parameters - Miscellaneous Parameters** page of the **Update Hadoop Cluster Configuration Parameters Template** wizard, update the (ServiceLevel and RoleLevel) parameters as required.
- Step 11** Click **Submit**.
-

Quality of Service System Classes

For more information on Quality of Service and System Classes, see [QoS System Classes](#).

Quality of Service

Cisco Unified Computing System provides the following methods to implement quality of service (QoS):

- System classes that specify the global configuration for certain types of traffic across the entire system.
- QoS policies that assign system classes for individual vNICs.
- Flow control policies that determine how uplink Ethernet ports handle pause frames.

System Classes

Cisco UCS uses Data Center Ethernet (DCE) to handle all traffic inside a Cisco UCS domain. This industry-standard enhancement to Ethernet divides the bandwidth of the Ethernet pipe into eight virtual lanes. Two virtual lanes are reserved for internal system use and management traffic. You can configure quality of service (QoS) for the other six virtual lanes. System classes determine how the DCE bandwidth in these six virtual lanes is allocated across the entire Cisco UCS domain.

Each system class reserves a specific segment of the bandwidth for a specific type of traffic, which provides a level of traffic management, even in an oversubscribed system. For example, you can configure the Fibre Channel Priority system class to determine the percentage of DCE bandwidth allocated to FCoE traffic.

The following table describes the system classes that you can configure:

System Class	Description
Best Effort	<p>A system class that sets the quality of service for the lane reserved for basic Ethernet traffic. Some properties of this system class are preset and cannot be modified.</p> <p>For example, this class has a drop policy that allows it to drop data packets if necessary. You cannot disable this system class.</p>
<ul style="list-style-type: none"> • Platinum • Gold • Silver • Bronze 	<p>A configurable set of system classes that you can include in the QoS policy for a service profile. Each system class manages one lane of traffic. All properties of these system classes are available for you to assign custom settings and policies.</p>
Fibre Channel	<p>A system class that sets the quality of service for the lane reserved for Fibre Channel over Ethernet traffic. Some properties of this system class are preset and cannot be modified.</p> <p>For example, this class has a no-drop policy that ensures it never drops data packets. You cannot disable this system class.</p> <p>Note FCoE traffic has a reserved QoS system class that cannot be used by any other type of traffic. If any other type of traffic has a CoS value that is used by FCoE, the value is re-marked to 0</p>

Editing QoS System Classes

For more information on Quality of Service and System Classes, see [QoS System Classes](#).

-
- Step 1** Choose **Solutions > Big Data > Settings**.
- Step 2** Click **QoS System Class**.
- Step 3** Choose the QoS System Class (by Priority) that you want to edit and click **Edit**.
- Best Effort
 - Platinum
 - Gold

- Silver
- Bronze

Step 4 On the **Modify QoS System Class** screen, complete the following fields:

Name	Description
Enabled check box	<p>If checked, the associated QoS class is configured on the fabric interconnect and can be assigned to a QoS policy.</p> <p>If unchecked, the class is not configured on the fabric interconnect. Any QoS policies associated with this class default to Best Effort or, if a system class is configured with a CoS of 0, to the CoS 0 system class.</p> <p>This check box is checked for Best Effort and Fibre Channel.</p>
CoS drop-down list	<p>The class of service. You can enter an integer value between 0 and 6, with 0 being the lowest priority and 6 being the highest priority. We recommend that you do not set the value to 0, unless you want that system class to be the default system class for traffic if the QoS policy is deleted or the assigned system class is disabled.</p> <p>This field is set to 7 for internal traffic and to any for Best effort. Both of these values are reserved and cannot be assigned to any other priority.</p>
Packet Drop check box	<p>This field is unchecked for the Fibre Channel class, which never allows dropped packets, and is checked for Best Effort, which always allows dropped packets.</p>
Weight drop-down list	<p>A choice may be one of the following:</p> <ul style="list-style-type: none"> • An integer between 1 and 10. If you select an integer, Cisco UCS determines the percentage of network bandwidth assigned to the priority level as described in the Weight (%) field. • Best-effort. • None.
Multicast Optimized check box	<p>If checked, the class is optimized to send packets to multiple destinations simultaneously. This option is not applicable to the Fibre Channel.</p>

Name	Description
MTU drop-down list	<p>The maximum transmission unit for the channel. This can be one of the following:</p> <ul style="list-style-type: none"> • An integer between 1500 and 9216. This value corresponds to the maximum packet size. • fc—A predefined packet size of 2240. • Normal—A predefined packet size of 1500. • Specify Manually—A packet size between 1500 to 9216. <p>This field is always set to fc for Fibre Channel.</p>

Step 5 Click **Submit**.

Pre Cluster Performance Testing Settings

You can analyze memory, network, and disk metrics. A default Big Data Metrics Report provides the statistics collected for each host before creating any Hadoop cluster.

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

Step 2 Click **Management**.

Step 3 On the **Pre Cluster Performance Tests** page, check the check boxes for the following:

- Memory Test
- Network Test
- Disk Test

Note By default, the check boxes to run the memory, network, and the disk tests are unchecked. If you enable the pre cluster disk test, it impacts Hadoop cluster creation.

Step 4 Click **Submit**.

Approving Hadoop Cluster and Splunk Deployment Workflows

Before you begin

Choose **Administration > Users and Groups** and click **Users**, and add users with the following user roles:

- Network Admin (system default user role)

- Computing Admin (system default user role)
- Big Data User

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

Step 2 Click **Management**.

Step 3 Check the **Require OS User Approval** check box.

- a) From the **User ID** table, check the **Login Name** of the user against the Network Admin user role.
- b) Enter the **Number of Approval Request Reminders**.

Note Set the number of approval request reminders to zero if the reminder email has to be sent at a specified interval until the Network Admin approves or rejects the request.

- c) Enter the **Reminder Interval(s)** in hours.

Note Check the **Approval required from all the users** check box, if you want all users to approve or reject the request.

Step 4 Check the **Require Compute User Approval** check box.

- a) From the **User ID** table, select the **Login Name** of the user against the Computing Admin user role.
- b) Enter the **Number of Approval Request Reminders**.

Note Set the number of approval request reminders to zero if the reminder email has to be sent at a specified interval until the Computing Admin approves or rejects the request.

- c) Enter the **Reminder Interval(s)** in hours.

Note Check the **Approval required from all the users** check box, if you want the users to approve or reject the request.

Step 5 Check the **Require Accounts User Approval** check box.

- a) From the **User ID** table, select the **Login Name** of the user against the Hadoop User role.
- b) Enter the **Number of Approval Request Reminders**.

Note Set the number of approval request reminders to zero if the reminder email has to be sent at a specified interval until the Hadoop User approves or rejects the request.

- c) Enter the **Reminder Interval(s)** in hours.

Note Check the **Approval required from all the users** check box, if you want the users to approve or reject the request.

Step 6 Click **Submit**.

What to do next

Verify whether users of Network Admin, Computing Admin, and Big Data Accounts User roles have approved the request before deploying any Big Data software.

Adding NTP Server Details

- Step 1** Choose **Solutions > Big Data > Settings**.
- Step 2** Click **Management**.
- Step 3** Click **Add (+)**.
- Step 4** On the **Add Entry to Servers** screen, complete the following fields:

Name	Description
Server Name field	The IP address of NTP server.
Is Primary Server check box	Click the check box if you want the server to be a primary server.

- Step 5** Click **Submit**.

Uploading Required OS and Big Data Software to Cisco UCS Director Bare Metal Agent

You can upload (add) required RHEL or CentOS ISO files, Big Data software and common software, and Oracle JDKs to Cisco UCS Director Bare Metal Agent. You can upload the required files from your local system or any remote system, and the files are first uploaded to Cisco UCS Director. Click the **Submit** button in the **Create Software Catalogs** screen to move the required files to the target Cisco UCS Director Bare Metal Agent.

To upload the software packages using **Web Server**, create a directory on the web server path and place all the required software packages in the directory. For example, IPaddress/web/folder (containing the software).

Supported file formats for the **Upload Type** as **Desktop File**:

- Linux OS— [OS version].iso. For example, rhel-server-7.5-x86_64-dvd.iso, CentOS-7.5-x86_64-DVD-1708.iso
- Big Data software—[Big Data-Distro]-[Major version].[Minor version].[Patch version].zip (.gz or .tgz or .tar) For example, MapR-5.2.2.zip, cloudera-5.14.0.zip, and Splunk-7.3.3.zip.
- Common software—bd-sw-rep.zip (.gz or .tgz or .tar)
- JDK software—jdk-8u181-linux-x64 (.rpm or gz)



Tip If the required software column is empty for Big Data, then Cisco UCS Director Bare Metal Agent already contains all the files required.

You must add software catalogs to upload the required software packages. Perform the following to add software catalogs:

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

Step 2 Click **Software Catalogs**.

Step 3 Click **Add**.

Step 4 To upload the files from your local system, you can either drag and drop the required files or click **Select a File**.

Note Create a folder to include all the required files for the Big Data software, and compress the folders before uploading in the format specified.

Step 5 On the **Create Software Catalogs** page, complete the following fields:

Note Refresh the **Software Catalogs** page after 5 to 10 minutes to see new and modified catalogs.

Name	Description
Linux OS Upload	
OS Type drop-down list	Choose the OS image type that you want to install on the server. The drop-down list includes all OS types supported by the Cisco UCS Director Bare Metal Agent.
Catalog Name field	Operating System Name (for example, RHEL, CentOS)
Upload Type drop-down list	<p>Choose one of the following:</p> <ul style="list-style-type: none"> • Desktop file • The web server path that is reachable by the Cisco UCS Director Bare Metal Agent (For example, IPaddress/web/folder (containing the software)) • Mountpoint in Cisco UCS Director Bare Metal Agent (For example, /root/iso) • Path to ISO in Cisco UCS Director Bare Metal Agent (For example, /temp/rhel75/iso) <p>Note If you select the Desktop file or web server path option, the .iso file is uploaded to BMA and then it is mounted. If you select the Mountpoint or Path to ISO option, the .iso file is directly mounted since it is already available in the Bare Metal Agent.</p>
Location field	Location of the OS file.
Big Data Software Upload	
Distribution drop-down list	Choose the big data distribution. For example, Cloudera, MapR.

Name	Description
Distribution Version drop-down list	Choose the big data software version. For example, Hadoop Distribution (for example, Distribution_name-Major version.Minor version.Patch version) or splunk enterprise software (Splunk-.Major version.Minor version.Patch version)
OS Type drop-down list	Choose the required OS type. For MapR and Splunk, choose the type as Any option. The file path for the OS-specific software for Cloudera and Hontonwork are modified. For example, the file paths for the repository in Cisco UCS Director Bare Metal Agent is <code>/opt/cnsaroot/bd-sw-rep.</code>
Upload Type drop-down list	Choose one of the following: <ul style="list-style-type: none"> • Desktop file • The web server path that is reachable by the Cisco UCS Director Bare Metal Agent to upload remote software to Bare Metal Agent <p>Note If you select the Desktop file or web server path option, the big data software is copied to respective the folders in <code>/opt/cnsaroot/bd-sw-rep/</code> location.</p>
Location field	Location of the big data software.
Common Software Upload	
Upload Type drop-down list	Choose one of the following: <ul style="list-style-type: none"> • Desktop file • The web server path that is reachable by the Cisco UCS Director Bare Metal Agent to upload remote software to Bare metal Agent
Location field	Location of the common software.
JDK Upload	
JDK Version field	JDK version. For example, <code>jdk-8u60-linux-x64.rpm</code>
Upload Type drop-down list	Choose one of the following: <ul style="list-style-type: none"> • Desktop file • The web server path that is reachable by the Cisco UCS Director Bare Metal Agent to upload remote software to Bare metal Agent

Name	Description
Location field	Location of the JDK file.

Step 6 Click **Submit**.

What to do next

You can track software uploads here: **Administration** > **Integration**. Click **Change Record** to track the software upload in progress and verify its status.

Supported Oracle JDK Software Versions

This section lists the supported Oracle JDK software versions:

Supported Upgrade Scenarios for Cloudera

Hadoop Distribution Version to Upgrade	Supported Upgrade Version
Cloudera Enterprise 5.4.x, JDK 1.8	Cloudera Enterprise 5.5.0, JDK 1.8
Cloudera Enterprise 5.4.x, JDK 1.8	Cloudera Enterprise 5.6.x, JDK 1.8
Cloudera Enterprise 5.4.x, JDK 1.8	Cloudera Enterprise 5.8.x, JDK 1.8
Cloudera Enterprise 5.6.x, JDK 1.8	Cloudera Enterprise 5.8.x, JDK 1.8
Cloudera Enterprise 5.8.0, JDK 1.8	Cloudera Enterprise 5.10.0, JDK 1.8
Cloudera Enterprise 5.8.0, JDK 1.8	Cloudera Enterprise 5.11.1, JDK 1.8
Cloudera Enterprise 5.8.2, JDK 1.8	Cloudera Enterprise 5.13.1, JDK 1.8
Cloudera Enterprise 5.11.1, JDK 1.8	Cloudera Enterprise 5.13.1, JDK 1.8



Note For more information on the supported JDK versions, see Cloudera site.

Supported Upgrade Scenarios for MapR

Hadoop Distribution Version to Upgrade	Supported Upgrade Version
MapR 5.2.1, JDK 1.8	MapR 6.0.0, JDK 1.8
MapR 5.0.0, JDK 1.8	MapR 5.1.0, JDK 1.8
MapR 4.0.2, JDK 1.8	MapR 5.2.0, JDK 1.8



Note For more information on the supported JDK versions, see MapR site.

Supported Upgrade Scenarios for Hortonworks

Hadoop Distribution Version to Upgrade	Supported Upgrade Version
Hortonworks 2.2, JDK 1.7	Hortonworks 2.3, JDK 1.8
Hortonworks 2.2, JDK 1.7	Hortonworks 2.4, JDK 1.8



Note For more information on the supported JDK versions, see Hortonworks site.

Cloudera, MapR, and Hortonworks RPMs on Cisco UCS Director Express for Big Data Bare Metal Agent

Common Packages for Cloudera, MapR, and Hortonworks



Note For any Hadoop software that is not available, update the `/opt/cnsaroot/bigdata_templates/common_templates/HadoopDistributionRPM.txt` file with an appropriate file from the online repository of the vendor.



Note We recommend that you verify the supported versions from the Hadoop Vendor Support Documentation.

Download the following common packages to `/opt/cnsaroot/bd-sw-rep/`:

- `pssh-2.3.1.tar.gz` from <https://pypi.python.org/packages/source/p/pssh>
- `clustershell-1.7.1-1.el6.noarch.rpm`
- `clustershell-1.7.1-1.el7.noarch.rpm`

Common Packages for Cloudera

Download the following packages to `/opt/cnsaroot/bd-sw-rep/cloudera-X.X.X:`

- `ClouderaEnterpriseLicense.lic`—Get the license keys from Cloudera
- `userrpm1ist.txt`—For more packages lists

- `catalog.properties`—Provides the label name for the Cloudera version (x represents the Cloudera version on the Cisco UCS Director Express for Big Data Bare Metal Agent)
- `mysql-connector-java-5.1.39.tar.gz` from MySQL site
- `ext-2.2.zip` from <http://archive.cloudera.com/gplextras/misc/ext-2.2.zip>

Cloudera 5.14.0 Packages and Parcels

Download the following packages to `/opt/cnsaroot/bd-sw-rep/cloudera-5.14.0`:

- `CDH-5.14.0-1.cdh5.14.0.p0.24-el7.parcel` from <https://archive.cloudera.com/cdh5/parcels/5.14.0/>
- `CDH-5.14.0-1.cdh5.14.0.p0.24-el7.parcel.sha1` from <https://archive.cloudera.com/cdh5/parcels/5.14.0/>
- `cm5.14.0-centos7.tar.gz` from <https://archive.cloudera.com/cm5/repo-as-tarball/5.14.0/>
- `manifest.json` from <https://archive.cloudera.com/cdh5/parcels/5.14.0/>
- `mysql-connector-java-5.1.45.tar.gz` from MySQL site
- `ojdbc7.jar` from Oracle site
- `instantclient-basic-linux.x64-12.1.0.2.0.zip` from Oracle site
- `oracle-instantclient12.1-basic-12.1.0.2.0-1.x86_64.rpm` from Oracle site
- `oracle-instantclient12.1-sqlplus-12.1.0.2.0-1.x86_64.rpm` from Oracle site

Cloudera 6.0 Packages and Parcels

Download the following packages to `/opt/cnsaroot/bd-sw-rep/cloudera-6.0`:

- `CDH-6.0.0-1.cdh6.0.0.p0.537114-el7.parcel` from <https://archive.cloudera.com/cdh6/6.0.0/parcels/CDH-6.0.0-1.cdh6.0.0.p0.537114-el7.parcel>
- `CDH-6.0.0-1.cdh6.0.0.p0.537114-el7.parcel.sha256` from <https://archive.cloudera.com/cdh6/6.0.0/parcels/CDH-6.0.0-1.cdh6.0.0.p0.537114-el7.parcel.sha256>
- `manifest.json` from <https://archive.cloudera.com/cdh6/6.0.0/parcels/manifest.json>
- `cloudera-manager-daemons-6.0.0-530873.el7.x86_64.rpm` from https://archive.cloudera.com/cm6/6.0.0/redhat7/yum/RPMS/x86_64/cloudera-manager-daemons-6.0.0-530873.el7.x86_64.rpm
- `cloudera-manager-server-db-2-6.0.0-530873.el7.x86_64.rpm` from https://archive.cloudera.com/cm6/6.0.0/redhat7/yum/RPMS/x86_64/cloudera-manager-server-db-2-6.0.0-530873.el7.x86_64.rpm
- `cloudera-manager-server-6.0.0-530873.el7.x86_64.rpm` from https://archive.cloudera.com/cm6/6.0.0/redhat7/yum/RPMS/x86_64/cloudera-manager-server-6.0.0-530873.el7.x86_64.rpm
- `CDH-6.0.0-1.cdh6.0.0.p0.537114-sles12.parcel.sha256` from <https://archive.cloudera.com/cdh6/6.0.0/parcels/CDH-6.0.0-1.cdh6.0.0.p0.537114-sles12.parcel.sha256>

- `cloudera-manager-agent-6.0.0-530873.el7.x86_64.rpm` from https://archive.cloudera.com/cm6/6.0.0/redhat7/yum/RPMS/x86_64/cloudera-manager-agent-6.0.0-530873.el7.x86_64.rpm
- `mysql-connector-java-5.1.45.tar.gz` from MySQL site
- `ojdbc7.jar` from Oracle site
- `instantclient-basic-linux.x64-12.1.0.2.0.zip` from Oracle site
- `oracle-instantclient12.1-basic-12.1.0.2.0-1.x86_64.rpm` from Oracle site
- `oracle-instantclient12.1-sqlplus-12.1.0.2.0-1.x86_64.rpm` from Oracle site

Cloudera 6.1 Packages and Parcels

Download the following packages to `/opt/cnsaroot/bd-sw-rep/cloudera-6.1`:

- `CDH-6.1.0-1.cdh6.1.0.p0.770702-el7` from <https://archive.cloudera.com/cdh6/6.1.0/parcels/CDH-6.1.0-1.cdh6.1.0.p0.770702-el7.parcel>
- `CDH-6.1.0-1.cdh6.1.0.p0.770702-el7.parcel.sha256` from <https://archive.cloudera.com/cdh6/6.1.0/parcels/CDH-6.1.0-1.cdh6.1.0.p0.770702-el7.parcel.sha256>
- `manifest.json` from <https://archive.cloudera.com/cdh6/6.1.0/parcels/manifest.json>
- `cloudera-manager-daemons-6.1.0-769885.el7.x86_64.rpm` from https://archive.cloudera.com/cm6/6.1.0/redhat7/yum/RPMS/x86_64/cloudera-manager-daemons-6.1.0-769885.el7.x86_64.rpm
- `cloudera-manager-server-6.1.0-769885.el7.x86_64.rpm` from https://archive.cloudera.com/cm6/6.1.0/redhat7/yum/RPMS/x86_64/cloudera-manager-server-6.1.0-769885.el7.x86_64.rpm
- `cloudera-manager-server-db-2-6.1.0-769885.el7.x86_64.rpm` from https://archive.cloudera.com/cm6/6.1.0/redhat7/yum/RPMS/x86_64/cloudera-manager-server-db-2-6.1.0-769885.el7.x86_64.rpm
- `cloudera-manager-agent-6.1.0-769885.el7.x86_64.rpm` from https://archive.cloudera.com/cm6/6.1.0/redhat7/yum/RPMS/x86_64/cloudera-manager-agent-6.1.0-769885.el7.x86_64.rpm
- `oracle-j2sdk1.8-1.8.0+update141-1.x86_64.rpm` from https://archive.cloudera.com/cm6/6.1.0/redhat7/yum/RPMS/x86_64/oracle-j2sdk1.8-1.8.0+update141-1.x86_64.rpm
- `ClouderaEnterpriseLicense.lic` from Cloudera site
- `mysql-connector-java-5.1.45.tar.gz` from MySQL site
- `ojdbc7.jar` from Oracle site
- `instantclient-basic-linux.x64-12.1.0.2.0.zip` from Oracle site
- `oracle-instantclient12.1-basic-12.1.0.2.0-1.x86_64.rpm` from Oracle site
- `oracle-instantclient12.1-sqlplus-12.1.0.2.0-1.x86_64.rpm` from Oracle site

Common Packages for MapR

Download the following common packages to `/opt/cnsaroot/bd-sw-rep/MapR-X.X.X` directories:

- `libgenders-devel-1.14-2.el6.rf.x86_64.rpm` from <http://pkgs.repoforge.org/libgenders/>
- `libgenders-1.14-2.el6.rf.x86_64.rpm` from <http://pkgs.repoforge.org/libgenders/>
- `ext-2.2.zip` from Apache.Hadoop site.
- `sshpass-1.05-1.el6.x86_64.rpm` from http://ftp.pbone.net/mirror/download.fedora.redhat.com/pub/fedora/epel/6/x86_64
- `soci-mysql-3.2.1-1.el6.x86_64.rpm` from http://ftp.is.co.za/mirror/fedora.redhat.com/epel/6/x86_64
- `soci-3.2.1-1.el6.x86_64.rpm` from http://ftp.is.co.za/mirror/fedora.redhat.com/epel/6/x86_64
- `pdsh-2.27-1.el6.rf.x86_64.rpm` from <http://pkgs.repoforge.org/pdsh>
- `mapr-whirr-0.7.0.16780-1.noarch.rpm` from <http://archive.mapr.com/releases/ecosystem-all/redhat>
- `mapr-drill-0.7.0.29434-1.noarch.rpm` from <http://archive.mapr.com/releases/ecosystem/redhat>
- catalog.properties—Provides the label name for the MapR version (x represents the MapR version on the Cisco UCS Director Express for Big Data Bare Metal Agent)
- license.txt

MapR 5.2.2 Packages

Download the following packages to `/opt/cnsaroot/bd-sw-rep/MapR-5.2.2`

- `libgenders-1.22-2.el7.x86_64.rpm` from http://rpm.pbone.net/index.php3/stat/4/idpl/29487566/dir/redhat_el_7/com/libgenders-1.22-2.el7.x86_64.rpm.html
- `libgenders-devel-1.22-2.el7.x86_64.rpm` from https://centos.pkgs.org/7/epel-x86_64/libgenders-devel-1.22-2.el7.x86_64.rpm.html
- `mapr-ecosystem-5.x-20170802.rpm.tgz` from <http://archive.mapr.com/releases/ecosystem-all/redhat/mapr-ecosystem-5.x-20170802.rpm.tgz>
- `mapr-setup` from <http://package.mapr.com/releases/v5.2.2/redhat/mapr-setup>
- `mapr-v5.2.2GA.rpm.tgz` from <http://archive.mapr.com/releases/v5.2.2/redhat/mapr-v5.2.2GA.rpm.tgz>
- `mapr-whirr-0.7.0.16780-1.noarch.rpm` from <http://archive.mapr.com/releases/ecosystem-all/redhat/mapr-whirr-0.7.0.16780-1.noarch.rpm>
- `mysql-connector-java-5.1.44.tar.gz` from MySQL site.
- `pdsh-2.31-1.el7.x86_64.rpm` from http://mirrors.isu.net.sa/pub/fedora/fedora-epel/7/x86_64/p/

- **soci-3.2.3-1.el7.x86_64.rpm** from http://rpm.pbone.net/index.php3/stat/4/idpl/40962039/dir/redhat_el_7/com/soci-3.2.3-1.el7.x86_64.rpm.html
- **soci-mysql-3.2.3-1.el7.x86_64.rpm** from http://rpm.pbone.net/index.php3/stat/4/idpl/40172013/dir/redhat_el_7/com/soci-mysql-3.2.3-1.el7.x86_64.rpm.html
- **sshpass-1.06-1.el7.x86_64.rpm** from http://rpm.pbone.net/index.php3/stat/4/idpl/40962109/dir/redhat_el_7/com/sshpass-1.06-1.el7.x86_64.rpm.html

MapR 6.0.0 Packages

Download the following packages to `/opt/cnsaroot/bd-sw-rep/MapR-6.0.0`

- **mapr-v6.0.0GA.rpm.tgz** from <http://archive.mapr.com/releases/v6.0.0/redhat/>
- **mapr-mep-v4.0.0.201711161643.rpm.tgz** from <http://archive.mapr.com/releases/MEP/MEP-4.0.0/redhat/>
- **libgenders-1.22-2.el7.x86_64.rpm** from http://rpm.pbone.net/index.php3/stat/4/idpl/29487566/dir/redhat_el_7/com/libgenders-1.22-2.el7.x86_64.rpm.html
- **libgenders-devel-1.22-2.el7.x86_64.rpm** from https://centos.pkgs.org/7/epel-x86_64/libgenders-devel-1.22-2.el7.x86_64.rpm.html
- **mapr-whirr-0.7.0.16780-1.noarch.rpm** from <http://archive.mapr.com/releases/ecosystem-all/redhat/mapr-whirr-0.7.0.16780-1.noarch.rpm>
- **mysql-connector-java-5.1.44.tar.gz** from MySQL site.
- **pdsh-2.31-1.el7.x86_64.rpm** from http://mirrors.isu.net.sa/pub/fedora/fedora-epel/7/x86_64/p/
- **soci-3.2.3-1.el7.x86_64.rpm** from http://rpm.pbone.net/index.php3/stat/4/idpl/40962039/dir/redhat_el_7/com/soci-3.2.3-1.el7.x86_64.rpm.html
- **soci-mysql-3.2.3-1.el7.x86_64.rpm** from http://rpm.pbone.net/index.php3/stat/4/idpl/40172013/dir/redhat_el_7/com/soci-mysql-3.2.3-1.el7.x86_64.rpm.html
- **sshpass-1.06-1.el7.x86_64.rpm** from http://rpm.pbone.net/index.php3/stat/4/idpl/40962109/dir/redhat_el_7/com/sshpass-1.06-1.el7.x86_64.rpm.html
- **ext-2.2.zip** from Apache.Hadoop site.

MapR 6.1.0 Packages

Download the following packages to `/opt/cnsaroot/bd-sw-rep/MapR-6.1.0`

- **mapr-v6.1.0GA.rpm.tgz** from <http://archive.mapr.com/releases/v6.1.0/redhat/mapr-v6.1.0GA.rpm.tgz>
- **mapr-mep-v6.0.0.201810030946.rpm.tgz** from <http://archive.mapr.com/releases/MEP/MEP-6.0.0/redhat/mapr-mep-v6.0.0.201810030946.rpm.tgz>
- **libgenders-1.22-2.el7.x86_64.rpm** from http://rpm.pbone.net/index.php3/stat/4/idpl/29487566/dir/redhat_el_7/com/libgenders-1.22-2.el7.x86_64.rpm.html

- **libgenders-devel-1.22-2.el7.x86_64.rpm** from https://centos.pkgs.org/7/epel-x86_64/libgenders-devel-1.22-2.el7.x86_64.rpm.html
- **mapr-whirr-0.7.0.16780-1.noarch.rpm** from <http://archive.mapr.com/releases/ecosystem-all/redhat/mapr-whirr-0.7.0.16780-1.noarch.rpm>
- **mysql-connector-java-5.1.44.tar.gz** from MySQL site.
- **pdsh-2.31-1.el7.x86_64.rpm** from http://mirrors.isu.net.sa/pub/fedora/fedora-epel/7/x86_64/p/
- **soci-3.2.3-1.el7.x86_64.rpm** from http://rpm.pbone.net/index.php3/stat/4/idpl/40962039/dir/redhat_el_7/com/soci-3.2.3-1.el7.x86_64.rpm.html
- **soci-mysql-3.2.3-1.el7.x86_64.rpm** from http://rpm.pbone.net/index.php3/stat/4/idpl/40172013/dir/redhat_el_7/com/soci-mysql-3.2.3-1.el7.x86_64.rpm.html
- **sshpas-1.06-1.el7.x86_64.rpm** from http://rpm.pbone.net/index.php3/stat/4/idpl/40962109/dir/redhat_el_7/com/sshpas-1.06-1.el7.x86_64.rpm.html
- **ext-2.2.zip** from Apache.Hadoop site.

Common Package for Hortonworks

Download the following common package to `/opt/cnsaroot/bd-sw-rep/Hortonworks-X.X`:

- **openssl-1.0.1e-30.el6.x86_64.rpm**
- **ext-2.2.zip** from Apache.Hadoop site.
- **catalog.properties**—Provides the label name for the Hortonworks version (x represents the Hortonworks version on the Cisco UCS Director Express for Big Data Bare Metal Agent)

Hortonworks 2.6.4 Packages

Download the following packages to `/opt/cnsaroot/bd-sw-rep/Hortonworks-2.6.4`:

- **ambari-2.6.1.0-centos7.tar.gz** from <http://public-repo-1.hortonworks.com/ambari/centos7/2.x/updates/2.6.1.0/ambari-2.6.1.0-centos7.tar.gz>
- **HDP-2.6.4.0-centos7-rpm.tar.gz** from <http://public-repo-1.hortonworks.com/HDP/centos7/2.x/updates/2.6.4.0/HDP-2.6.4.0-centos7-rpm.tar.gz>
- **HDP-UTILS-1.1.0.22-centos7.tar.gz** from <http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz>
- **libtirpc-0.2.4-0.10.el7.x86_64.rpm** from http://rpm.pbone.net/index.php3/stat/4/idpl/38004637/dir/scientific_linux_7/com/libtirpc-0.2.4-0.10.el7.x86_64.rpm.html
- **libtirpc-devel-0.2.4-0.10.el7.x86_64.rpm** from http://rpm.pbone.net/index.php3/stat/4/idpl/37971478/dir/centos_7/com/libtirpc-devel-0.2.4-0.10.el7.x86_64.rpm.html
- **je-5.0.73.jar** from Oracle site.
- **ojdbc7.jar** from Oracle site.
- **oracle-instantclient12.1-basic-12.1.0.2.0-1.x86_64.rpm** from Oracle site.

- `oracle-instantclient12.1-sqlplus-12.1.0.2.0-1.x86_64.rpm` from Oracle site.

Hortonworks 3.0.0 Packages

Download the following packages to `/opt/cnsaroot/bd-sw-rep/Hortonworks-3.0.0`:

- `ambari-2.7.0.0-centos7.tar.gz` from <http://public-repo-1.hortonworks.com/ambari/centos7/2.x/updates/2.7.0.0/ambari-2.7.0.0-centos7.tar.gz>
- `HDP-3.0.0.0-centos7-rpm.tar.gz` from <http://public-repo-1.hortonworks.com/HDP/centos7/3.x/updates/3.0.0.0/HDP-3.0.0.0-centos7-rpm.tar.gz>
- `HDP-UTILS-1.1.0.22-centos7.tar.gz` from <http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz>
- `HDP-GPL-3.0.0.0-centos7-gpl.tar.gz` from <http://public-repo-1.hortonworks.com/HDP-GPL/centos7/3.x/updates/3.0.0.0/HDP-GPL-3.0.0.0-centos7-gpl.tar.gz>
- `libtirpc-0.2.4-0.10.el7.x86_64.rpm` from http://rpm.pbone.net/index.php3/stat/4/idpl/38004637/dir/scientific_linux_7/com/libtirpc-0.2.4-0.10.el7.x86_64.rpm.html
- `libtirpc-devel-0.2.4-0.10.el7.x86_64.rpm` from http://rpm.pbone.net/index.php3/stat/4/idpl/37971478/dir/centos_7/com/libtirpc-devel-0.2.4-0.10.el7.x86_64.rpm.html
- `je-5.0.73.jar` from Oracle site.
- `ojdbc7.jar` from Oracle site.
- `oracle-instantclient12.1-basic-12.1.0.2.0-1.x86_64.rpm` from Oracle site.
- `oracle-instantclient12.1-sqlplus-12.1.0.2.0-1.x86_64.rpm` from Oracle site.



Note Download the Splunk software from Splunk Suite.

Cloudera and MapR RPMs for Upgrading Hadoop Cluster Distributions

Cloudera 5.3.0 Packages and Parcels

- `cm5.3.0-centos6.tar.gz` from <http://archive.cloudera.com/cm5/repo-as-tarball/5.3.0>
- `CDH-5.3.0-1.cdh5.3.0.p0.30-el6.parcel` from <http://archive.cloudera.com/cdh5/parcels/5.3.0>
- `CDH-5.3.0-1.cdh5.3.0.p0.30-el6.parcel.sha1` from <http://archive.cloudera.com/cdh5/parcels/5.3.0>
- `manifest.json` from <http://archive.cloudera.com/cdh5/parcels/5.3.0>

Cloudera 5.4.1 Packages and Parcels

- `cm5.4.1-centos6.tar.gz` from <http://archive.cloudera.com/cm5/repo-as-tarball/5.4.1>
- `CDH-5.4.1-1.cdh5.4.1.p0.6-el6.parcel` from <http://archive.cloudera.com/cdh5/parcels/5.4.1>
- `CDH-5.4.1-1.cdh5.4.1.p0.6-el6.parcel.sha1` from <http://archive.cloudera.com/cdh5/parcels/5.4.1>
- `manifest.json` from <http://archive.cloudera.com/cdh5/parcels/5.4.1>

MapR 4.1.0 Packages

- `mapr-setup` from <http://package.mapr.com/releases/v4.1.0/redhat>
- `mapr-v4.1.0GA.rpm.tgz` from <http://package.mapr.com/releases/v4.1.0/redhat>
- `mysql-connector-java-5.1.26.tar.gz` from <http://cdn.mysql.com/archives/mysql-connector-java-5.1>

MapR 5.0.0 Packages

- `mapr-setup` from <http://package.mapr.com/releases/v5.0.0/redhat>
- `mapr-v5.0.0GA.rpm.tgz`: from <http://package.mapr.com/releases/v5.0.0/redhat>
- `mysql-connector-java-5.1.26.tar.gz` from <http://cdn.mysql.com/archives/mysql-connector-java-5.1>

MapR 5.2.0 Packages

- `mapr-setup` from <http://package.mapr.com/releases/v5.2.0/redhat/mapr-setup>
- `mapr-v5.2.0GA.rpm.tgz`: from <http://archive.mapr.com/releases/v5.2.0/redhat/mapr-v5.2.0GA.rpm.tgz>



Note `mapr-v5.2.0GA.rpm.tgz` contains the following `mapr-client-5.2.0.39122.GA-1.x86_64.rpm`, `mapr-posix-client-platinum-5.2.0.39122.GA-1.x86_64.rpm`, `mapr-posix-client-basic-5.2.0.39122.GA-1.x86_64.rpm`, `mapr-upgrade-5.2.0.39122.GA-1.x86_64.rpm`, `mapr-nfs-5.2.0.39122.GA-1.x86_64.rpm`, and `mapr-core-5.2.0.39122.GA-1.x86_64.rpm` files.

- `mysql-connector-java-5.1.26.tar.gz` from <https://downloads.mysql.com/archives/get/file/mysql-connector-java-5.1.26.tar.gz>
- `mapr-ecosystem-5.x-20160816.rpm.tgz` from <http://archive.mapr.com/releases/ecosystem-all/redhat/mapr-ecosystem-5.x-20160816.rpm.tgz>

Installation of User-Defined Software Post Hadoop Cluster Creation

Cisco UCS Director Express for Big Data provides an option to add user-defined installation packages (RPMs) post Hadoop cluster creation specific to a version. In Cisco UCS Director Express for Big Data, you cannot install more Hadoop related software other than what is required for the selected type of Hadoop distribution when creating an instant Hadoop cluster or customizing a Hadoop cluster.

To install user-defined Hadoop related software, you can specify a list of RPMs in the HadoopDistributionRPM.txt file. This modifiable list defines the required packages for each version of a Hadoop distribution. You can locate the HadoopDistributionRPM.txt here:

`/opt/cnsaroot/bigdata_templates/common_templates` in the

Cisco UCS Director Express for Big Data Bare Metal Agent

server.

For example, you can specify a list of RPMs in the HadoopDistributionRPM.txt file for MapR-5.0.0:

- `mapr-v5.0.0GA.rpm.tgz`
- `mapr-ecosystem-5.x-20150709.rpm.tgz`
- `mapr-whirr-0.8.1.18380-GA.noarch.rpm`

Configuration Check Rules

You can validate an existing cluster configuration by running a configuration check. The configuration check process involves comparing the current cluster configuration with reporting violations and configuration check rules.

Configuration check rules are predefined Cisco Validated Design (CVD) parameters for Hadoop clusters. Configuration check rules appear under **Solutions > Big Data > Settings**. After the configuration check is complete, violations appear in the **Faults** page under **Solutions > Big Data > Accounts**. You can enable or disable configuration check rules at any time, but you cannot add new rules.

Configuration Check Rule	Description
Parameter	The predefined CVD parameter of the configuration.
Enabled	The state of the configuration check rule, either enabled (true) or disabled (false).
Expected value	The value expected for a parameter as defined in the Cisco Validated Design (CVD).
Description	The description of the parameter of the configuration.
Distribution	The Hadoop distribution.
Minimum Supported Distribution	The minimum supported version of Hadoop distribution.

Configuration Check Rule	Description
Service	The Hadoop service.
Role	The Hadoop service role.
Type	The type of violation, either CVD or Inconsistent.
Fix Workflow	The reference to the workflow that can be triggered for fixing violations.

When the actual cluster configuration values differ from the expected values defined in the configuration check rules, then those configuration values are reported as violations. For example, CVD mandates that the NameNode heap size is 4 GB. But if the NameNode heap size in the cluster configuration is found to be 1 GB, then this is reported as a CVD violation. Also, inconsistent configuration parameters are reported. For example, NameNode heap size on both the primary and secondary nodes must be of the same size. If there is a mismatch in the size, then this parameter is reported as inconsistent.

Checking Hadoop Cluster Configuration

To validate the configuration of a cluster, do the following:

-
- Step 1** Choose **Solutions > Big Data > Accounts**.
 - Step 2** Click **Hadoop Accounts**.
 - Step 3** Choose the account for which you want to run the configuration check and click **Check Configuration**.
 - Step 4** Click **Submit**.
A page appears with the information that the configuration check is in progress.
 - Step 5** Click **OK**.
After the configuration check is complete, the violations appear under the **Faults** page for the selected Hadoop Account.
-

What to do next



Note You can track configuration checks here: **Administration > Integration**. Click **Change Record** to track the configuration checks in progress and verify if completed or failed.

Fixing Configuration Violations

After the configuration check is complete, the configuration violations appear in the **Faults** page for the selected big data account. You can either choose to fix these configuration violations manually on the **Big Data Cluster Configuration** page, or trigger a workflow. To trigger a workflow to fix the violation, create a workflow with the same name as the code specified in the violation.

To fix a configuration violation through a workflow, do the following:

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

Step 2 Click **Faults**.

Step 3 Choose the configuration violation you want to fix and click **Trigger Workflow**.

If a workflow exists with the same name as the code specified in the violation, then the workflow is triggered.

Step 4 Enter the required inputs for the workflow and click **Submit**.

A service request ID is generated after you submit the inputs. You can check the status of the service request on the **Service Requests** page.
