

Cluster Template Overrides (Preview)

Date published: 2022-03-08

Date modified: 2022-03-07

Legal Notice

© Cloudera Inc. 2022. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 ("ASLv2"), the Affero General Public License version 3 (AGPLv3), or other license terms.

Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER'S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Contents

Legal Notice	2
Contents	3
Overview	4
Limitations	7
Creating a Cluster Template Override	7
Using Cluster Template Override to create a new Data Hub Cluster	10
Beta CDP CLI Commands	11

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Overview

You can launch Data Hub clusters from a set of pre-defined cluster templates created for prescriptive use cases. These cluster templates are a “shared resource” that define the list of services that will be installed on the Data Hub, including their configurations.

For example, this is a portion of a default template:

```
{
  "services": [
    {
      "refName": "zookeeper",
      "serviceType": "ZOOKEEPER",
      "serviceConfigs": [
        {
          "name": "service_config_suppression_server_count_validator",
          "value": "true"
        }
      ],
      "roleConfigGroups": [
        {
          "refName": "zookeeper-SERVER-BASE",
          "roleType": "SERVER",
          "base": true
        }
      ]
    },
    {
      "refName": "hdfs",
      "serviceType": "HDFS",
      "serviceConfigs": [
        {
          "name": "hdfs_verify_ec_with_topology_enabled",
          "value": false
        },
        {
          "name": "core_site_safety_valve",
          "value":
"<property><name>fs.s3a.buffer.dir</name><value>${env.LOCAL_DIRS:-${hadoop.tmp.dir}}
/s3a</value></property><property><name>fs.s3a.committer.name</name><value>directory<
/value></property>"
        }
      ],
      "roleConfigGroups": [
        {
          "refName": "hdfs-NAMENODE-BASE",
          "roleType": "NAMENODE",
          "base": true,
          "configs": [
            {
```

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

CLOUDERA TECHNICAL PREVIEW DOCUMENTATION

```
    "name": "role_config_suppression_namenode_java_heapsize_minimum_validator",
    "value": "true"
  },
  {
    "name": "role_config_suppression_fs_trash_interval_minimum_validator",
    "value": "true"
  },
  {
    "name": "fs_trash_interval",
    "value": "0"
  },
  {
    "name": "fs_trash_checkpoint_interval",
    "value": "0"
  },
  {
    "name": "erasure_coding_default_policy",
    "value": " "
  }
]
},
{
  "refName": "hdfs-SECONDARYNAMENODE-BASE",
  "roleType": "SECONDARYNAMENODE",
  "base": true
},
{
  "refName": "hdfs-DATANODE-BASE",
  "roleType": "DATANODE",
  "base": true
},
{
  "refName": "hdfs-BALANCER-BASE",
  "roleType": "BALANCER",
  "base": true
},
{
  "refName": "hdfs-GATEWAY-BASE",
  "roleType": "GATEWAY",
  "base": true,
  "configs": [
    {
      "name": "dfs_client_use_trash",
      "value": false
    },
    {
      "name": "role_config_suppression_hdfs_trash_disabled_validator",
      "value": "true"
    },
    {
      "name": "hdfs_client_env_safety_valve",
      "value": "HADOOP_OPTS=\"-Dorg.wildfly.openssl.path=/usr/lib64
OOP_OPTS}\""
```

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

```

    }
  ]
}
]

```

In this section of an example cluster template, two primary types of service configurations are visible: `serviceConfigs` and configurations for various `roleConfigGroups`, in addition to a special type of service configuration called `safety_valves`. You can read about the details of these configs in the [Cloudera Manager Configuration Properties](#) reference.

Oftentimes you may want to modify the `serviceConfigs`, `roleConfigGroups`, and `safety_valves` present in a default Data Hub cluster template. Although you can [create a custom cluster template](#) by modifying the JSON of a default template, this process can be cumbersome and error-prone.

As an alternative to creating a custom template, you can specify custom configurations that override or append the properties in a default template. These custom configurations are saved as a shared resource called “cluster template overrides,” and can be used and re-used across Data Hub clusters in different environments. As a shared resource, they do not need to be attached to a specific Runtime version.

By using a default cluster template along with a cluster template override, you can create a customized Data Hub cluster, along with receiving improvements to the default templates that are present in newer Cloudera Runtime releases. Cluster template overrides can be used to override or append properties present in both the identified types of service configurations, as well as safety valves.

Note: If you customize safety valves, the provided value will be **appended** to the already existing value for that safety valve. Custom configs will not override the existing value.

You are not limited to customizing the configs present in a default cluster template; you can add any valid configuration recognized by Cloudera Manager for an included service. New configurations are appended to the cluster template when the cluster is launched.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided ‘as is’ without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Limitations

At present, there is no way to validate any individual property to ensure that it is valid and recognised by Cloudera Manager, so that it can be overridden to the desired value. Adding an incorrect config name can lead to errors while installing the CM template on the Data Hub cluster, or the property could be ignored by Cloudera Manager entirely.

An invalid value for a particular property can also cause errors during the cluster creation process. Carefully review the [Cloudera Manager Configuration Properties](#) reference in regards to configurations that you want to customize.

Cluster template overrides are only for overriding/appending the `serviceConfigs`, `roleConfigGroups`, and `safety_valves` in a Data Hub template. Unlike creating a custom template, you can not add a service to the template.

Note: Cluster template overrides can be applied to custom templates as well. If you want to apply a cluster template override to a custom template that contains [properties that are dynamically replaced during cluster creation](#), the cluster template override will override any dynamically-replaced properties when the two conflict.

Creating a Cluster Template Override

1. In the Management Console UI, click **Shared Resources > Cluster Template Overrides**.
2. Click **Create Custom Override**.

Cluster Template Overrides / Create

Name *

 ?

Cloudera Runtime Version

 ▼

Service Config Role Config

Config name/path *

 ?

Config value *

 ?

Service type *

 ?

Service Type	Role Type
No cluster template value override added.	

3. Provide the following information:

Property	Description
Name	Provide a unique name for the override.
Cloudera Runtime Version	Select a specific Runtime version to which this override applies, or select "Not Runtime specific."
Service Config / Role Config	Use the radio buttons to select whether you are

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

	customizing a service config or a role config.
Config name/path	Exact name/path of the config are you customizing.
Config value	New value for the config. Ensure the value is valid.
Service type	From the drop-down menu, select the service type for the service or role config that you are customizing.
Role type	For role configurations: the role type as it appears in the cluster template that you will apply the override to (or any valid role configuration that is recognized by Cloudera Manager).

For example:

Service Config
 Role Config

Config name/path *

oom_heap_dump_enabled ?

Config value *

false ?

Service type *

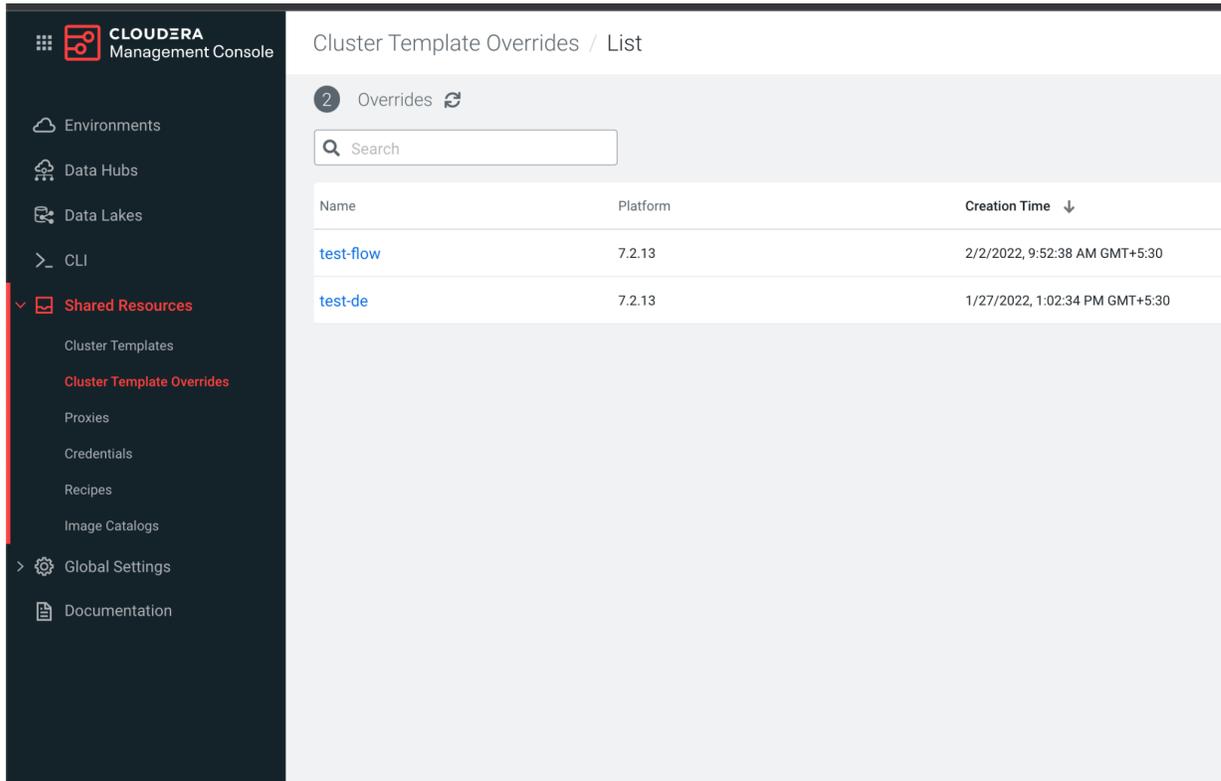
HIVE ▼

Role type *

HIVEMETASTORE ▼

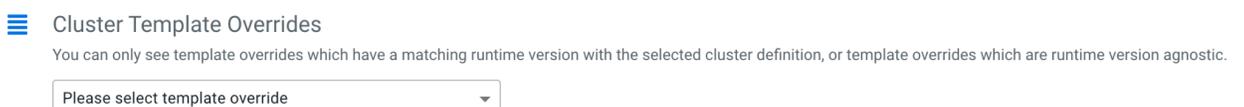
This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

4. Click **Add** and then repeat the previous step (excluding the Name and Runtime Version) for each service or role config that you want to customize in this template override.
5. When you are finished adding service and role configs to the template override, click **Save**. The override appears in the list of saved overrides, accessible from the *Shared Resources* menu:



Using Cluster Template Override to create a new Data Hub Cluster

1. Click **Data Hub Clusters > Create Data Hub Cluster**.
2. Enter the information required to create your cluster, then at the bottom of the cluster creation wizard, select **Advanced Options > Cluster Extensions**.
3. From the **Cluster Templates Overrides** drop-down box, select the template override that you previously created and want to apply when the cluster is provisioned. The only visible overrides are those with a Runtime version matching the selected cluster definition version, or those that are Runtime-agnostic.



4. Click **Provision Cluster**.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Beta CDP CLI Commands

If you do not have the beta CDP CLI installed, see the topic [Installing Beta CDP CLI](#).

Command	Description	Example
<pre>cdp datahub create-custom-configurations</pre>	<p>Creates a new cluster template override shared resource with an indefinite amount of properties.</p>	<pre>cdp datahub create-custom-configurations --custom-configurations-name "test-custom-configs" --configurations configName="name",configValue=" value",serviceType="hdfs" configName="name2",configValue= "value2",roleType="role2",serviceTy pe="hive_on_tez" --runtime-version "7.2.8"</pre>
<pre>cdp datahub list-custom-configurations</pre>	<p>Lists all the cluster template overrides for the account.</p>	
<pre>cdp datahub describe-custom-configuratio ns</pre>	<p>Returns the name, CRN, the properties, and the platform version associated with the cluster template override resource.</p>	<pre>cdp datahub describe-custom-configurations --custom-configurations "test-overrides"</pre>
<pre>cdp datahub delete-custom-configurations</pre>	<p>Deletes multiple custom configs given their names or CRNs.</p>	<pre>cdp datahub delete-custom-configurations --custom-configurations "test-override" "test-override-2" ...</pre>

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

CLUDERA TECHNICAL PREVIEW DOCUMENTATION

<pre> cdp datahub create-aws-cluster cdp datahub create-azure-cluster cdp datahub create-gcp-cluster </pre>	<p>Launches a Data Hub cluster with custom configs from a cluster template override. The parameter for specifying a cluster template override is:</p> <pre> --custom-configurations-name </pre>	<pre> cdp datahub create-aws-cluster --cluster-name "de-dh1" --environment-name "local-aws-env" --cluster-template-name "7.2.8 - Data Engineering: Apache Spark, Apache Hive, Apache Oozie" --cluster-definition-name "7.2.8 - Data Engineering for AWS" --custom-configurations-name "test-override" </pre>
---	---	--

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.