

Using Custom Spark Runtime Docker Images via API/CLI (Preview)

Date published: 2022-09-06

Date modified: 2025-06-30

Legal Notice

© Cloudera Inc. 2025. 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.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

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
Prerequisites	4
Entitlement for Customer Docker images	4
Docker repository credentials	4
Steps	5
Accessing ECR repository from different AWS environments	11
Troubleshooting	13

This is a detailed user guide that demonstrates how to run Spark jobs using custom Spark runtime Docker images via API/CLI. Custom Spark runtime Docker images are used when custom packages and libraries need to be installed and used when executing Spark jobs. These custom packages and libraries can be proprietary software packages like RPMs that need to be compiled to generate the required binaries. Docker images allow you to pre-bake these dependencies into a self-contained Docker file that can be used across multiple Spark jobs.

Important: If you're using Cloudera Data Engineering (CDE) 1.18 and above, you can't use Custom Spark Runtime Docker Images with CDE 1.17 and below.

Prerequisites

Entitlement for Customer Docker images

In order to use custom Spark runtime Docker images, your CDP tenant must have the DE_CUSTOM_RUNTIME entitlement enabled. If not yet in place, this should be requested via your Cloudera Account Team (primarily your Solution Engineer). They will fulfill the request internally and confirm to you when the entitlement has been applied. Please allow 24 hours for fulfillment.

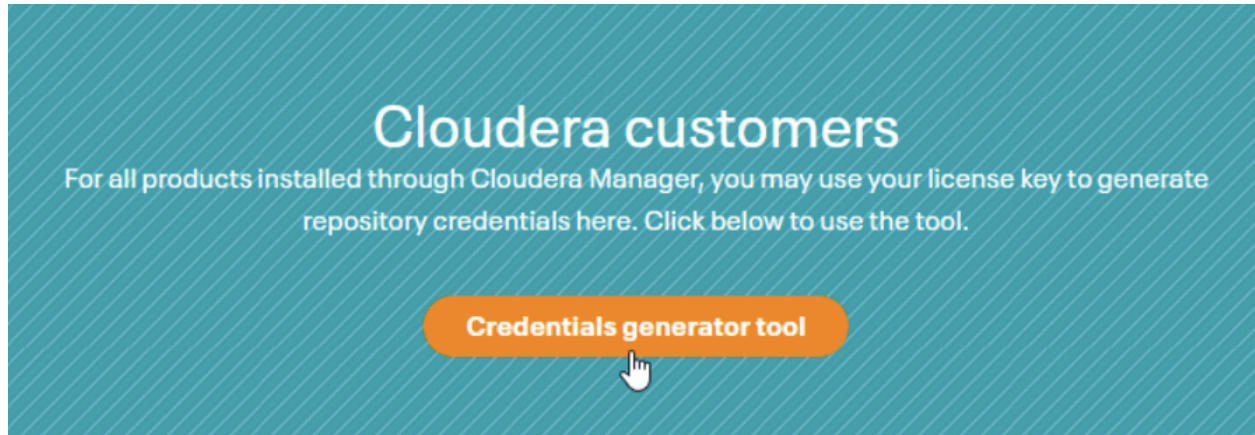
Note that the DE_CUSTOM_RUNTIME entitlement must be enabled before a Virtual Cluster is created in order to be able to create custom runtime resources within that cluster.

Docker repository credentials

In order to pull the base Docker image, you must have credentials and authenticate your docker client to `docker.repository.cloudera.com`. To get credentials:

1. Raise an "Admin" type case with Cloudera Support and request a **License Key**.
2. Once the License Key is received, navigate to <https://www.cloudera.com/downloads.html> and **log in** with your MyCloudera credentials.
3. Use the Credentials Generator tool - this is about half-way down the page (you will see an orange "Sign In" button if not already logged in).

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.



As directed on the page, copy and paste the entire contents of your license file into the text box and press the Get Credentials button to generate your username and password.

Steps

1. Create a custom Docker image

Build the “custom-spark-dex-runtime” image based on the “dex-spark-runtime” image of the Cloudera Data Engineering (CDE) version.

The image should be based on the dex-spark-runtime of the current dex version.

The relevant dex-spark-runtime images are as follows.

- Spark 3 Cloudera security hardened images

```
<registry-host>/cloudera/dex/dex-spark-runtime-<spark version>-<cdh
version>:<CDE version>
```

Example: DockerFile for DEX 1.24.0-b711, Spark 3.3.2 and CDP Runtime version 7.1.9.1015

Unset

```
FROM
docker.repository.cloudera.com/cloudera/dex/dex-spark-runtime-3.3.2-7.1.9
.1015:1.24.0-b711
USER root
RUN apk add --no-cache git
RUN pip3 install virtualenv-api
USER ${DEX_UID}
```

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.

- Spark 3 Redhat (insecure and deprecated) images

<registry-host>/cloudera/dex/dex-spark-runtime-<spark version>-<cdh version>-compat:<CDE version>

Example: DockerFile for DEX 1.24.0-b711, Spark 3.3.2 and CDP Runtime version 7.1.9.1015

Unset

```
FROM
docker.repository.cloudera.com/cloudera/dex/dex-spark-runtime-3.3.2-7.1.9
.1015-compat:1.24.0-b711
USER root
RUN yum install -y git && yum clean all && rm -rf /var/cache/yum
RUN pip2 install virtualenv-api
RUN pip3 install virtualenv-api
USER ${DEX_UID}
```

- Spark 2 Redhat (insecure and deprecated) images

<registry-host>/cloudera/dex/dex-spark-runtime-<spark version>-<cdh version>:<CDE version>

Example: DockerFile for DEX 1.24.0-b711, Spark 2.4.8 and CDP Runtime version 7.1.9.1015

Unset

```
FROM
docker.repository.cloudera.com/cloudera/dex/dex-spark-runtime-2.4.8-7.1.9
.1015:1.24.0-b711
USER root
RUN yum install -y git && yum clean all && rm -rf /var/cache/yum
RUN pip2 install virtualenv-api
RUN pip3 install virtualenv-api
USER ${DEX_UID}
```

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.

2. Build the docker image by tagging it with the custom registry and push it to the custom registry.

Example:

Unset

```
mac@local:$ docker build --network=host -t
docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.8
-7.2.14.0:1.15.0-b117-custom . -f Dockerfile
mac@local:$ docker push
docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.8
-7.2.14.0:1.15.0-b117-custom
```

Here, the custom registry is “docker.my-company.registry.com” and the registry namespace is “custom-dex”.

3. Create a custom runtime image resource.
 - Register “custom-spark-dex-runtime” docker image as a resource of type “Custom-runtime-image”.
 - a. Create a resource for the registries which do not require any authentication. If using a public docker registry or If the docker registry is in the same environment, for example, the same AWS account or Azure subscription where the CDE service is running, then you do not need to create credentials.

CLI

Unset

```
mac@local:$ cde resource create --name custom-image-resource
--image
docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.8
-7.2.14.0:1.15.0-b117-custom --image-engine spark2 --type
custom-runtime-image
```

Note: To obtain \$CDE_TOKEN to execute the REST API examples, follow the [Getting a Cloudera Data Engineering API access token](#) document.

REST API

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.

Unset

```
curl -X POST -k 'https://<dex-vc-host>/dex/api/v1/resources \
  -H "Authorization: Bearer ${CDE_TOKEN}" \
  -H 'accept: application/json' \
  -H 'Content-Type: application/json' \
  --data '{
    "customRuntimeImage": {
      "engine": "spark2",
      "image":
        "docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.
        8-7.2.14.0:1.15.0-b117
    }
  },
  "name": "custom-image-resource",
  "type": "custom-runtime-image"
}'
```

Once done, skip to [step 4](#) to submit the job.

- b. Create a resource which requires the credentials to access the registry. Use the command below or the API request to create the credentials. These credentials are stored as a secret.

CLI

Unset

```
mac@local:$ ./cde credential create --name docker-creds --type
docker-basic --docker-server docker-sandbox.infra.cloudera.com
--docker-username my-username
```

REST API

Unset

```
curl -X POST -k 'https://<dex-vc-host>/dex/api/v1/credentials' \
  -H "Authorization: Bearer ${CDE_TOKEN}" \
  -H 'accept: application/json' \
  -H 'Content-Type: application/json' \
  --data '{
    "dockerBasic": {
```

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.


```
"password": "password123",
"server": "docker-sandbox.infra.cloudera.com",
"username": "my-username"
},
"name": "docker-creds",
"type": "docker-basic"
}'
```

- c. Register “custom-spark-dex-runtime” docker image as a resource of type “custom-runtime-image” by specifying the name of the credential created above.

CLI

```
Unset
mac@local:$ ./cde resource create --name custom-image-resource
--image
docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.8
-7.2.14.0:1.15.0-b117
--custom --image-engine spark2 --type custom-runtime-image
--image-credential docker-creds
```

REST API

```
Unset
curl -X POST -k 'https://<dex-vc-host>/dex/api/v1/resources \
-H "Authorization: Bearer ${CDE_TOKEN}" \
-H 'accept: application/json' \
-H 'Content-Type: application/json' \
--data '{
  "customRuntimeImage": {
    "credential": "docker-creds",
    "engine": "spark2",
    "image":
      "docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.
      8-7.2.14.0:1.15.0-b117
    -custom"
```

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.

```

    },
    "name": "custom-image-resource",
    "type": "custom-runtime-image"
  }'

```

4. Submit a job by setting the “custom-spark-dex-runtime” image as a resource using the CDE CLI.

SPARK COMMAND

Unset

```

mac@local:$ ./cde --user cdpuser1 spark submit
/Users/my-username/spark-examples_2.11-2.4.4.jar
--class org.apache.spark.examples.SparkPi 1000
--runtime-image-resource-name=custom-image-resource

```

JOB COMMAND

Unset

```

mac@local:$ ./cde --user cdpuser1 resource create --name
spark-jar
mac@local:$ ./cde --user cdpuser1 resource upload --name
spark-jar --local-path spark-examples_2.11-2.4.4.jar
mac@local:$ ./cde --user cdpuser1 job create --name
spark-pi-job-cli --type spark --mount-1-resource spark-jar
--application-file spark-examples_2.11-2.4.4.jar --class
org.apache.spark.examples.SparkPi --user cdpuser1 --arg 22
--runtime-image-resource-name custom-image-resource

```

5. The spark driver/executor pods should use this image and you can confirm it by opening a shell into those pods and verifying if the external installed libraries or files exist.

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.

Accessing ECR repository from different AWS environments

Use Case:

There are separate AWS environments for different business units and there is an ECR repository which contains the common docker images that needs to be accessed from these environments (having different AWS accounts).

By default, The ECR will be accessible to the AWS services running in the same AWS environment. The same ECR can be accessed from different AWS environments by updating the ECR repository permissions.

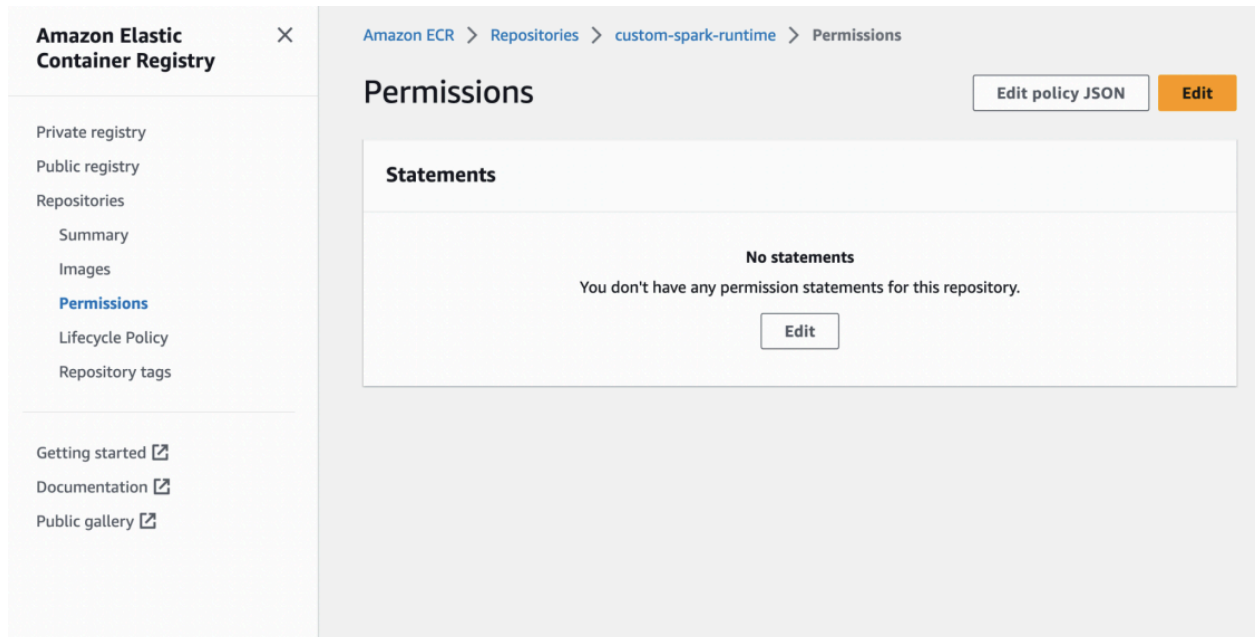
For Example:

- ECR repository is in env1 (AWS-ACCOUNT-1)
- Sales Team services are running in env2 (AWS-ACCOUNT-2)
- Finance Team services are running in env3 (AWS-ACCOUNT-3)

Below steps can be followed to update the ECR permissions to allow the access for both Sales and Finance teams' services.

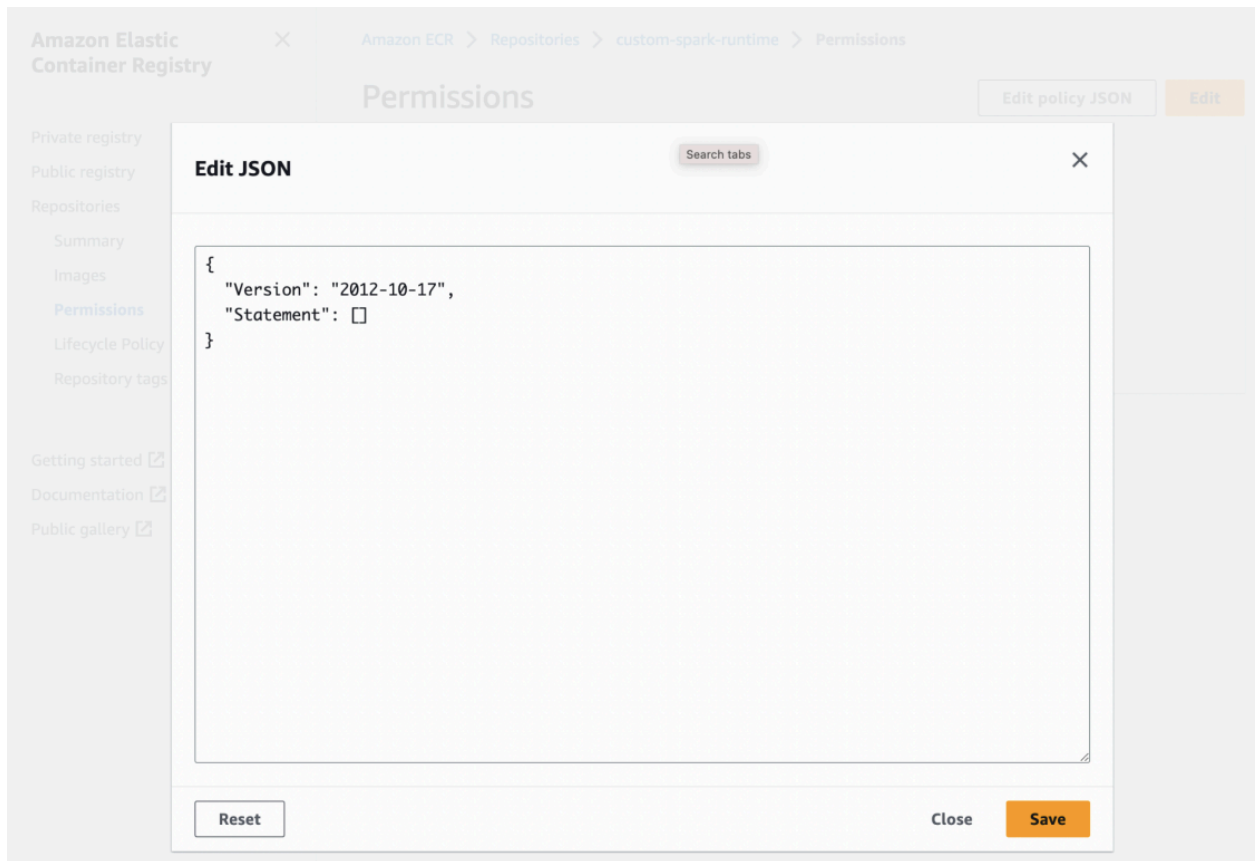
Steps:

1. Login into AWS console and open the desired ECR repository.
2. Click on the Permissions tab.



3. Click on Edit Policy Json. This will show you the current JSON policy document for the repository.

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. Modify the current JSON policy document by adding a new statement. If the repository has no permissions set yet then you can simply copy and paste the below JSON policy document (Update the required AWS account ids in this document)

Unset

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "crossAccountAllowRead",
      "Effect": "Allow",
      "Principal": {
        "AWS": [
          "arn:aws:iam::<AWS-ACCOUNT-2>:root",
          "arn:aws:iam::<AWS-ACCOUNT-3>:root"
        ]
      },
      "Action": [
```

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.

```

        "ecr:BatchCheckLayerAvailability",
        "ecr:BatchGetImage",
        "ecr:DescribeImages",
        "ecr:GetDownloadUrlForLayer",
        "ecr:ListImages"
    ]
}
]
}

```

5. Click on Save button.
6. Verify If the policy has been set successfully by clicking on the Permissions tab again.
The new permissions and the AWS accounts IDs should be there.

Note: Ensure the AWS account IDs are correct otherwise you may see errors while saving the policy.

Once the permissions are set successfully on ECR for different AWS accounts then the services running in these AWS environments can access the ECR without any changes in the services or environment itself. There is no need to provide any credentials to access the ECR from these AWS accounts.

Troubleshooting

Error: Custom image resource with missing or wrong credentials

Creating a custom image resource with missing or wrong credentials should result in the below error which can be seen in the logs or in Kubernetes pod events.

```

Unset
Failed to pull image
"docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.
8-7.2.14.0:1.15.0-b117-custom":
rpc error: code = Unknown desc = Error reading manifest
1.15.0-b117-custom in
docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.8
-7.2.14.0:errors: denied: requested access to the resource is

```

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.

```
denied unauthorized: authentication required
```