Applying Data Virtualization service patch v1.4.1.0-35

Contents

Introduction....1

Applying patches for Data Virtualization....2

Applying patch back to previous service version....9

`post-patching-jars-update.sh` script file....16

Introduction

This document describes how to install the Data Virtualization service patch `v1.4.1.0-35`.

The `v1.4.1.0-35` service patch includes fixes for the following issues:

- 6240650: On certain data sources, unable to preview or query virtual table if table name has a period in its prefix.

Optionally, this document provides instructions on how to patch down to service Version `v1.4.1.0-`*`33`*.

Applying patch for Data Virtualization

A Red Hat® OpenShift® project administrator can apply patches for Data Virtualization.

Before you begin

Required role: To install a patch, you must be an administrator of the project (namespace) where the software is deployed.

Ensure that the Data Virtualization service instance is running correctly:

```
oc get pods | grep dv
```

All service pods must be running and ready.

Ensure that the Mac OS or Linux machine where you will run the commands meets the appropriate requirements for your environment:

Requirements for the machine	Cluster is connected to the internet	Cluster is air-gapped
Can connect to the cluster.	✓	✓
Is connected to the internet.	✓	
Has the `oc` command-line interface. You can download the appropriate client tools for your operating system from Red Hat OpenShift: • Version 3.11: Get Started with the CLI 3.11 • Version 4.3: Getting started with the CLI 4.3 Ensure that the version is compatible with the version of Red Hat OpenShift on your cluster.	✓	✓
Has the Cloud Pak for Data command-line interface. See Obtaining the installation files. Use the same version of the command-line interface each time you run the commands.	✓	✓
Has the updated `repo.yaml` file in the same directory as the Cloud Pak for Data command-line interface. See Obtaining the installation files.	✓	
Has the `cpd-`*`Operating_System`*`-workspace` directory, which contains the required files. See Preparing for air-gapped installations.		✓

Ensure that you have the following information from your Red Hat OpenShift cluster administrator:

Required information	Description
OpenShift_URL:port	The URL and port number to use when logging in to your Red Hat OpenShift cluster. Ensure that you have the appropriate credentials to log into the cluster using `oc login`.

Required information	Description
	Value: Your cluster administrator should tell you whether your cluster is connected to the internet or is air-gapped.
Project	The project where the software is currently installed. **Value:**
Assembly_version	The version of the software that is currently installed. **Value:** 1.4.1

Registry_location	The location to store the updated images on the registry server.
	If you are patching the software when you are connected to the internet, ensure that you have the appropriate credentials to *push* images to the registry server.
	Value:
	Guidance for Red Hat OpenShift registry users:
	• To determine the external route to the registry, run the appropriate command for your environment:
	– OpenShift 3.11:
	`oc get route/docker-registry -n default --template {{.spec.host}}`
	The command returns a route similar to `docker-registry-default.apps.`*`my_cluster_address`*
	Append the *project* name to the route. For example:
	`docker-registry-default.apps.`*`my_cluster_address`*`/`*`project`*
	– OpenShift 4.3:
	`oc get route/default-route -n openshift-image-registry --template='{{ .spec.host }}'`
	The command returns a route similar to `default-route-openshift-image-registry.apps.`*`my_cluster_address`*.
	Append the *project* name to the route. For example:
	`default-route-openshift-image-registry.apps.`*`my_cluster_address`*`/`*`project`*

Required information	Description
	• When you specify a value for the *Registry_location* variable, ensure that you include the *project* name.

Registry_from_cluster	The location from which pods on the cluster can *pull* images.
	Value:
	Guidance for Red Hat OpenShift registry users:
	• This is the internal name of the registry service. The default service name is:
	– OpenShift 3.11:
	`docker-registry.default.svc:5000/project`
	– OpenShift 4.3:
	`image-registry.openshift-image-registry.svc:5000/project`
	• When you specify a value for the *Registry_from_cluster* variable, ensure that you include the *project* name.

Procedure

To apply a Data Virtualization patch:

1. If you have an existing `/cpd-Operating_System-workspace` directory, for example, `cpd-linux-workspace` directory, rename this directory.

   The service patch must create a new `/cpd-Operating_System-workspace` directory to store the required files.
2. If you have more than one `dv-worker` pod running, scale the service down to a single `dv-worker` pod:

   ```
   oc scale statefulset dv-worker --replicas 1
   ```
3. Run the appropriate command for **`cpd patch`** your environment.

   Tip: For a list of all available options, enter the command: `./cpd-Operating_System --help`.

   - **To apply patches on a cluster that can connect to the internet:**

     a. Change to the directory where you placed the Cloud Pak for Data command-line interface and the `repo.yaml` file.

     b. Log in to your Red Hat OpenShift cluster as a project administrator:

     ```
     oc login OpenShift_URL:port
     ```

     c. Run the following command to patch the service:

     ```
     ./cpd-Operating_System patch --repo ./repo.yaml \
     --assembly dv \
     --namespace Project \
     --patch-name Patch_name \
     --transfer-image-to Registry_location \
     --cluster-pull-prefix Registry_from_cluster \
     --ask-push-registry-credentials
     ```

     Replace the following values:

Variable	Replace with
Operating_System	For Linux, specify `linux`. For Mac OS, specify `darwin`.
Project	Specify the project (namespace) where the software that you want to patch is deployed.
Patch_name	Specify the name of the patch that you want to install. This information is included in the patch description. **Value:** v1.4.1.0-35
Registry_location	Use the value specified by your cluster administrator or the value that you used when you installed the software.
Registry_from_cluster	Use the value specified by your cluster administrator or the value that you used when you installed the software.

d. Run the following command to patch service instances:

```
./cpd-Operating_System patch --repo ./repo.yaml \
--assembly dv \
--namespace Project \
--patch-name Patch_name \
--transfer-image-to Registry_location \
--cluster-pull-prefix Registry_from_cluster \
--insecure-skip-tls-verify \
--ask-push-registry-credentials \
--all-instances
```

Replace the following values:

Variable	Replace with
Operating_System	For Linux, specify `linux`. For Mac OS, specify `darwin`.
Project	Specify the project (namespace) where the software that you want to patch is deployed.
Patch_name	Specify the name of the patch that you want to install. This information is included in the patch description. **Value:** v1.4.1.0-35
Registry_location	Use the value specified by your cluster administrator or the value that you used when you installed the software.
Registry_from_cluster	Use the value specified by your cluster administrator or the value that you used when you installed the software.

- **To apply patches on an air-gapped cluster:**

a. On a machine that can connect to the internet, change to the directory where you extracted the Cloud Pak for Data installation command-line interface.

b. Run the following command to download the service patch to your local machine:

```
./cpd-Operating_System patch --repo ./repo.yaml \
--assembly dv \
--version Assembly_version \
--patch-name Patch_name \
--action download \
--insecure-skip-tls-verify
```

Replace the following values:

Variable	Replace with
Operating_System	For Linux, specify `linux`. For Mac OS, specify `darwin`.

Variable	Replace with
Assembly_version	Specify the version of the software that is currently installed.
Patch_name	Specify the name of the patch that you want to install. This information is included in the patch description. **Value:** v1.4.1.0-35

c. Transfer the following items to a machine that can connect to the cluster and to the registry server:

   – The cpd-*Operating_System*-workspace directory. Ensure that the directory structure remains unchanged.

   – A copy of the Cloud Pak for Data installation command-line interface. Ensure that the command-line interface is compatible with the machine that you are transferring the files to and that it is the same version as the command-line interface that you ran in the preceding steps.

d. Run the following command to push the images to the registry server:

```
./cpd-Operating_System patch \
--namespace Project \
--load-from Image_directory_location \
--assembly dv \
--patch-name Patch_name \
--transfer-image-to Registry_location \
--ask-push-registry-credentials \
--insecure-skip-tls-verify \
--cluster-pull-prefix Registry_from_cluster \
--action push
```

Replace the following values:

Variable	Replace with
Operating_System	For Linux, specify `linux`. For Mac OS, specify `darwin`.
Project	Specify the project (namespace) where the software that you want to patch is deployed.
Image_directory_location	The location of the cpd-*Operating_System*-workspace directory.
Patch_name	Specify the name of the patch that you want to install. This information is included in the patch description. **Value:** v1.4.1.0-35
Registry_location	Use the value specified by your cluster administrator or the value that you used when you installed the software.
Registry_from_cluster	Use the value specified by your cluster administrator or the value that you used when you installed the software.

e. Run the following command to download the instance patch to your local machine:

```
./cpd-Operating_System patch --repo ./repo.yaml \
--assembly dv \
--version Assembly_version \
--patch-name Patch_name \
--action download
--insecure-skip-tls-verify
--all-instances
```

Replace the following values:

Variable	Replace with
Operating_System	For Linux, specify `linux`. For Mac OS, specify `darwin`.
Assembly_version	Specify the version of the software that is currently installed.
Patch_name	Specify the name of the patch that you want to install. This information is included in the patch description. **Value:** v1.4.1.0-35

f. Run the following command to patch service instances:

```
./cpd-Operating_System patch \
--namespace Project \
--load-from Image_directory_location \
--assembly dv \
--patch-name Patch_name \
--transfer-image-to Registry_location \
--ask-push-registry-credentials \
--insecure-skip-tls-verify \
--cluster-pull-prefix Registry_from_cluster \
--all-instances \
--action push
```

Replace the following values:

Variable	Replace with
Operating_System	For Linux, specify `linux`. For Mac OS, specify `darwin`.
Project	Specify the project (namespace) where the software that you want to patch is deployed.
Image_directory_location	The location of the `cpd-Operating_System-workspace` directory.
Patch_name	Specify the name of the patch that you want to install. This information is included in the patch description. **Value:** v1.4.1.0-35
Registry_location	Use the value specified by your cluster administrator or the value that you used when you installed the software.
Registry_from_cluster	Use the value specified by your cluster administrator or the value that you used when you installed the software.

What to do next

After you patch the service, you must update JAR files and remote connectors:

1. Log in to the `dv-engine-0` service pod:

```
oc exec -it dv-engine-0 -- bash
```

2. Run the `post-patching-jars-update.sh` script:

```
sh  /opt/dv/current/post-patching-jars-update.sh
```

3. Ensure the following files have a more recent timestamp than `June 2020` if they exist in the persistent volume:

```
/mnt/PV/versioned/ibm/home-bigsql/datavirtualization/DATAVIRTUALIZATION.jar
/mnt/PV/versioned/opt/ibm/qpendpoint/sysroot/lib/DATAVIRTUALIZATION.jar
/mnt/PV/versioned/uc_dsserver_shared/config/lib/GAIANDB.jar
/mnt/PV/versioned/ibm/home-bigsql/datavirtualization/GAIANDB.jar
```

```
/mnt/PV/versioned/opt/ibm/qpendpoint/sysroot/lib/GAIANDB.jar\
/mnt/uc_dsserver_shared/config/jars/DATAVIRTUALIZATION_config.jar
/mnt/PV/versioned/ibm/home-bigsql/datavirtualization/icu4j-66_1.jar
/mnt/PV/versioned/opt/ibm/qpendpoint/sysroot/lib/icu4j-66_1.jar
/home/bigsql/sqllib/function/jar/QPLEXSYS/GAIANDB.jar
/home/bigsql/sqllib/function/jar/QPLEXSYS/QUERYPLEX_ADMIN.jar
/mnt/PV/versioned/uc_dsserver_shared/config/DATAVIRTUALIZATION_ENDPOINT_*.tar.gz
/mnt/PV/versioned/uc_dsserver_shared/config/DATAVIRTUALIZATION_ENDPOINT_*.zip
/mnt/PV/versioned/dv_data/DATAVIRTUALIZATION_ENDPOINT_*.zip
/mnt/PV/versioned/dv_data/DATAVIRTUALIZATION_ENDPOINT_*.tar.gz
/mnt/PV/versioned/QUERYPLEX_INSTALL/DATAVIRTUALIZATION_ENDPOINT_*.zip
/mnt/PV/versioned/QUERYPLEX_INSTALL/DATAVIRTUALIZATION_ENDPOINT_*.tar.gz
```

4. To ensure that the changes work correctly, delete the `dv-engine-0` service pod:

```
oc delete pod dv-engine-0
```

5. If you had multiple `dv-worker` pods before the patch, you can scale the number of `dv-worker` pods back up by running the following command:

```
oc scale statefulset dv-worker --replicas=n
```

   Replace *n* with the number of `dv-worker` pods you had before the patch.

6. To patch remote connectors:

   a. Locate script to patch remote connectors:

      - In Linux® and Mac, find the **dv_remoteupgrade_linux.sh** in the `/opt/ibm/qpserver/install/DATAVIRTUALIZATION_INSTALL` directory of the `dv-engine` service pod.
      - In Microsoft Windows, find the `dv_remoteupgrade_windows.bat` batch file in the `/opt/ibm/qpserver/install/DATAVIRTUALIZATION_INSTALL` directory of the `dv-engine` service pod.

   b. Copy the corresponding script to the `dvendpoint` directory on the remote connector. Ensure that the `datavirtualization.env` file is located in that directory.

   c. If you have an `upgrade` directory under the `dvendpoint` directory, rename it to `upgrade_old`.

   d. Specify the value of the *_ADMIN_PORT* variable. The value of the *_ADMIN_PORT* variable is displayed in the script you generated to install the remote connectors. For more information, see Installing remote connectors.

   e. In Windows, close the command prompt of the remote connector that you want to patch.

   f. Run the script:

      - In Linux and Mac, run the following command:

```
chmod +x dv_remoteupgrade_linux.sh
```

      - In Windows, ensure you have permission to run batch file. Close the command prompt on which remote connector is running.

      Note: In Linux or Mac operating systems, if you have more than one remote connector endpoints on the same machine, the patching script stops the other endpoints. Once the patching is complete, you must start these endpoints manually:

```
sh DVENDPOINT_DIRECTORY/sysroot/launchGaianServer.sh
```

Applying patch back to previous service version

A Red Hat OpenShift project administrator can apply patches back to a previous version of Data Virtualization.

Before you begin

Required role: To install a patch, you must be an administrator of the project (namespace) where the software is deployed.

Ensure that the Data Virtualization service instance is running correctly:

```
oc get pods | grep dv
```

All service pods must be running and ready.

Ensure that the Mac OS or Linux machine where you will run the commands meets the appropriate requirements for your environment:

Requirements for the machine	Cluster is connected to the internet	Cluster is air-gapped
Can connect to the cluster.	✓	✓
Is connected to the internet.	✓	
Has the `oc` command-line interface. You can download the appropriate client tools for your operating system from Red Hat OpenShift: • Version 3.11: Get Started with the CLI 3.11 • Version 4.3: Getting started with the CLI 4.3 Ensure that the version is compatible with the version of Red Hat OpenShift on your cluster.	✓	✓
Has the Cloud Pak for Data command-line interface. See Obtaining the installation files. Use the same version of the command-line interface each time you run the commands.	✓	✓
Has the updated `repo.yaml` file in the same directory as the Cloud Pak for Data command-line interface. See Obtaining the installation files.	✓	
Has the `cpd-`*`Operating_System`*`-workspace` directory, which contains the required files. See Preparing for air-gapped installations.		✓

Ensure that you have the following information from your Red Hat OpenShift cluster administrator:

Required information	Description
OpenShift_URL:port	The URL and port number to use when logging in to your Red Hat OpenShift cluster.

Required information	Description
	Ensure that you have the appropriate credentials to log into the cluster using `oc login`. **Value:** Your cluster administrator should tell you whether your cluster is connected to the internet or is air-gapped.
Project	The project where the software is currently installed. **Value:**
Assembly_version	The version of the software that is currently installed. **Value:** 1.4.1

Registry_location	The location to store the updated images on the registry server.
	If you are patching the software when you are connected to the internet, ensure that you have the appropriate credentials to *push* images to the registry server.
	Value:
	Guidance for Red Hat OpenShift registry users:
	• To determine the external route to the registry, run the appropriate command for your environment:
	– OpenShift 3.11:
	`oc get route/docker-registry -n default --template {{.spec.host}}`
	The command returns a route similar to `docker-registry-default.apps.`*`my_cluster_address`*
	Append the *project* name to the route. For example:
	`docker-registry-default.apps.`*`my_cluster_address`*`/`*`project`*
	– OpenShift 4.3:
	`oc get route/default-route -n openshift-image-registry --template='{{ .spec.host }}'`
	The command returns a route similar to `default-route-openshift-image-registry.apps.`*`my_cluster_address`*.

Required information	Description
	Append the *project* name to the route. For example: `default-route-openshift-image-registry.apps.`*`my_cluster_address`*`/`*`project`* • When you specify a value for the *Registry_location* variable, ensure that you include the *project* name.
Registry_from_cluster	The location from which pods on the cluster can *pull* images. **Value:** **Guidance for Red Hat OpenShift registry users:** • This is the internal name of the registry service. The default service name is: – OpenShift 3.11: `docker-registry.default.svc:5000/`*`project`* – OpenShift 4.3: `image-registry.openshift-image-registry.svc:5000/`*`project`* • When you specify a value for the *Registry_from_cluster* variable, ensure that you include the *project* name.

Procedure

To apply a Data Virtualization patch:

1. If you have an existing `/cpd-`*`Operating_System`*`-workspace` directory, for example, `cpd-linux-workspace` directory, rename this directory.

   The service patch must create a new `/cpd-`*`Operating_System`*`-workspace` directory to store the required files.
2. If you have more than one `dv-worker` pod running, scale the service down to a single `dv-worker` pod:

   ```
   oc scale statefulset dv-worker --replicas 1
   ```
3. Run the appropriate command for **`cpd patch`** your environment.

   Tip: For a list of all available options, enter the command: `./cpd-`*`Operating_System`*` --help`.

   - **To apply patches on a cluster that can connect to the internet:**

     a. Change to the directory where you placed the Cloud Pak for Data command-line interface and the `repo.yaml` file.

     b. Log in to your Red Hat OpenShift cluster as a project administrator:

     ```
     oc login OpenShift_URL:port
     ```

     c. Run the following command to patch the service:

     ```
     ./cpd-Operating_System patch --repo ./repo.yaml \
     --assembly dv \
     ```

```
--namespace Project \
--patch-name Patch_name \
--transfer-image-to Registry_location \
--cluster-pull-prefix Registry_from_cluster \
--ask-push-registry-credentials
```

Replace the following values:

Variable	Replace with
Operating_System	For Linux, specify `linux`. For Mac OS, specify `darwin`.
Project	Specify the project (namespace) where the software that you want to patch is deployed.
Patch_name	Specify the name of the patch that you want to install. This information is included in the patch description. **Value:** v1.4.1.0-33
Registry_location	Use the value specified by your cluster administrator or the value that you used when you installed the software.
Registry_from_cluster	Use the value specified by your cluster administrator or the value that you used when you installed the software.

d. Run the following command to patch service instances:

```
./cpd-Operating_System patch --repo ./repo.yaml \
--assembly dv \
--namespace Project \
--patch-name Patch_name \
--transfer-image-to Registry_location \
--cluster-pull-prefix Registry_from_cluster \
--insecure-skip-tls-verify \
--ask-push-registry-credentials \
--all-instances
```

Replace the following values:

Variable	Replace with
Operating_System	For Linux, specify `linux`. For Mac OS, specify `darwin`.
Project	Specify the project (namespace) where the software that you want to patch is deployed.
Patch_name	Specify the name of the patch that you want to install. This information is included in the patch description. **Value:** v1.4.1.0-33
Registry_location	Use the value specified by your cluster administrator or the value that you used when you installed the software.
Registry_from_cluster	Use the value specified by your cluster administrator or the value that you used when you installed the software.

- **To apply patches on an air-gapped cluster:**

a. On a machine that can connect to the internet, change to the directory where you extracted the Cloud Pak for Data installation command-line interface.

b. Run the following command to download the service patch to your local machine:

```
./cpd-Operating_System patch --repo ./repo.yaml \
--assembly dv \
--version Assembly_version \
--patch-name Patch_name \
```

```
--action download \
--insecure-skip-tls-verify
```

Replace the following values:

Variable	Replace with
Operating_System	For Linux, specify `linux`. For Mac OS, specify `darwin`.
Assembly_version	Specify the version of the software that is currently installed.
Patch_name	Specify the name of the patch that you want to install. This information is included in the patch description. **Value:** v1.4.1.0-33

c. Transfer the following items to a machine that can connect to the cluster and to the registry server:

- The `cpd-`*`Operating_System`*`-workspace` directory. Ensure that the directory structure remains unchanged.
- A copy of the Cloud Pak for Data installation command-line interface. Ensure that the command-line interface is compatible with the machine that you are transferring the files to and that it is the same version as the command-line interface that you ran in the preceding steps.

d. Run the following command to push the images to the registry server:

```
./cpd-Operating_System patch \
--namespace Project \
--load-from Image_directory_location \
--assembly dv \
--patch-name Patch_name \
--transfer-image-to Registry_location \
--ask-push-registry-credentials \
--insecure-skip-tls-verify \
--cluster-pull-prefix Registry_from_cluster \
--action push
```

Replace the following values:

Variable	Replace with
Operating_System	For Linux, specify `linux`. For Mac OS, specify `darwin`.
Project	Specify the project (namespace) where the software that you want to patch is deployed.
Image_directory_location	The location of the `cpd-`*`Operating_System`*`-workspace` directory.
Patch_name	Specify the name of the patch that you want to install. This information is included in the patch description. **Value:** v1.4.1.0-33
Registry_location	Use the value specified by your cluster administrator or the value that you used when you installed the software.
Registry_from_cluster	Use the value specified by your cluster administrator or the value that you used when you installed the software.

e. Run the following command to download the instance patch to your local machine:

```
./cpd-Operating_System patch --repo ./repo.yaml \
--assembly dv \
--version Assembly_version \
--patch-name Patch_name \
--action download
```

```
--insecure-skip-tls-verify
--all-instances
```

Replace the following values:

Variable	Replace with
Operating_System	For Linux, specify `linux`. For Mac OS, specify `darwin`.
Assembly_version	Specify the version of the software that is currently installed.
Patch_name	Specify the name of the patch that you want to install. This information is included in the patch description. **Value:** v1.4.1.0-33

f. Run the following command to patch service instances and push images to the registry server:

```
./cpd-Operating_System patch \
--namespace Project \
--load-from Image_directory_location \
--assembly dv \
--patch-name Patch_name \
--transfer-image-to Registry_location \
--ask-push-registry-credentials \
--insecure-skip-tls-verify \
--cluster-pull-prefix Registry_from_cluster \
--all-instances \
--action push
```

Replace the following values:

Variable	Replace with
Operating_System	For Linux, specify `linux`. For Mac OS, specify `darwin`.
Project	Specify the project (namespace) where the software that you want to patch is deployed.
Image_directory_location	The location of the `cpd-Operating_System-workspace` directory.
Patch_name	Specify the name of the patch that you want to install. This information is included in the patch description. **Value:** v1.4.1.0-33
Registry_location	Use the value specified by your cluster administrator or the value that you used when you installed the software.
Registry_from_cluster	Use the value specified by your cluster administrator or the value that you used when you installed the software.

What to do next

After you patch the service, you must update JAR files and remote connectors:

1. Log in to the `dv-engine-0` service pod:

```
oc exec -it dv-engine-0 -- bash
```

2. Create a script file named `post-patching-jars-update.sh` by using content in post-patching-jars-update.sh script file.

3. Run the `post-patching-jars-update.sh` script:

```
sh /opt/dv/current/post-patching-jars-update.sh
```

4. Restart `dv-engine-0` service pod:

```
oc delete pod dv-engine-0
```

5. If you had multiple `dv-worker` pods before the patch, you can scale the number of `dv-worker` pods back up by running the following command:

```
oc scale statefulset dv-worker --replicas=n
```

   Replace *n* with the number of `dv-worker` pods you had before the patch.

6. To patch remote connectors:

   a. Locate script to patch remote connectors:

      - In Linux and Mac, find the **dv_remoteupgrade_linux.sh** in the `/opt/ibm/qpserver/install/DATAVIRTUALIZATION_INSTALL` directory of the `dv-engine` service pod.
      - In Microsoft Windows, find the `dv_remoteupgrade_windows.bat` batch file in the `/opt/ibm/qpserver/install/DATAVIRTUALIZATION_INSTALL` directory of the `dv-engine` service pod.

   b. Copy the corresponding script to the `dvendpoint` directory on the remote connector. Ensure that the `datavirtualization.env` file is located in that directory.

   c. If you have an `upgrade` directory under the `dvendpoint` directory, rename it to `upgrade_old`.

   d. Specify the value of the *_ADMIN_PORT* variable. The value of the *_ADMIN_PORT* variable is displayed in the script you generated to install the remote connectors. For more information, see Installing remote connectors.

   e. In Windows, close the command prompt of the remote connector that you want to patch.

   f. Run the script:

      - In Linux and Mac, run the following command:

```
chmod +x dv_remoteupgrade_linux.sh
```

      - In Windows, ensure you have permission to run batch file. Close the command prompt on which remote connector is running.

      Note: In Linux or Mac operating systems, if you have more than one remote connector endpoints on the same machine, the patching script stops the other endpoints. Once the patching is complete, you must start these endpoints manually:

```
sh DVENDPOINT_DIRECTORY/sysroot/launchGaianServer.sh
```

post-patching-jars-update.sh script file

Use the following content to create the `post-patching-jars-update.sh` script file. You can use this script to update Data Virtualization JAR files after you patch the service.

```
#!/bin/bash
################################################################################################
#
# Licensed Materials - Property of IBM
#
# "Restricted Materials of IBM"
#
# (C) COPYRIGHT IBM Corp. 2020 All Rights Reserved.
#
# US Government Users Restricted Rights - Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
################################################################################################
#  This script is used to update QP binaries post patching
################################################################################################

DV_VERSION="V1.4.1.0"
ORIGINAL_BUILD="*"

# pre check
sudo ls -al /opt/ibm/qp_artifacts/archives/DATAVIRTUALIZATION_ENDPOINT_${DV_VERSION}_*_*.tar.gz
> /dev/null
if [ $? -ne 0 ]; then
  echo "Patching failed, DATAVIRTUALIZATION_ENDPOINT_${DV_VERSION}_*_* tar.gz file not exist"
  exit 1
fi

update_jars_if_exist() {
  while true
  do
    test -z "${1}" && break
    file_dir="${1}"; shift
    base=$(basename ${file_dir})
    sudo find /mnt/PV/versioned|grep -i "${base}" | while read target_dir; do
      cp -f ${target_dir} /mnt/PV/versioned/backup
      cp -f ${file_dir} ${target_dir}
    done
  done
}

save_and_delete_jars() {
  while true
  do
    test -z "${1}" && break
    file_dir="${1}"; shift
    cp -f ${file_dir} /mnt/PV/versioned/backup
    rm -f ${file_dir}
  done
}

# patching starts here
# create backup dir
sudo mkdir -p /mnt/PV/versioned/backup
sudo chmod 777 /mnt/PV/versioned/backup
update_jars_if_exist "/opt/ibm/qpserver/sysroot/lib/DATAVIRTUALIZATION.jar" \
  "/opt/ibm/qpserver/sysroot/lib/GAIANDB.jar" \
  "/opt/ibm/qpendpoint/sysroot/jars/DATAVIRTUALIZATION_config.jar"

db2 connect to bigsql
db2 "CALL sqlj.replace_jar(\"file:/opt/ibm/qpserver/install/DATAVIRTUALIZATION_INSTALL/
GAIANDB.jar\", \"QPLEXSYS.GAIANDB\")" 2>&1
db2 "CALL sqlj.replace_jar(\"file:/opt/ibm/qpserver/install/DATAVIRTUALIZATION_INSTALL/
DATAVIRTUALIZATION_dashproc.jar\", \"QPLEXSYS.QUERYPLEX_ADMIN\")" 2>&1
db2 connect reset

mkdir /tmp/dv_qp_tmp
sudo tar -zxf /opt/ibm/qp_artifacts/archives/DATAVIRTUALIZATION_ENDPOINT_$
{DV_VERSION}_*_*.tar.gz -C /tmp/dv_qp_tmp
update_jars_if_exist "/tmp/dv_qp_tmp/sysroot/lib/icu4j-66_1.jar"

chmod +x /opt/ibm/qp_artifacts/archives
```

```
save_and_delete_jars "/mnt/PV/versioned/uc_dsserver_shared/config/DATAVIRTUALIZATION_ENDPOINT_$
{DV_VERSION}_${ORIGINAL_BUILD}.tar.gz" \
  "/mnt/PV/versioned/dv_data/DATAVIRTUALIZATION_ENDPOINT_${DV_VERSION}_$
{ORIGINAL_BUILD}.tar.gz" \
  "/mnt/PV/versioned/QUERYPLEX_INSTALL/DATAVIRTUALIZATION_ENDPOINT_${DV_VERSION}_$
{ORIGINAL_BUILD}.tar.gz"
cp -f /opt/ibm/qp_artifacts/archives/DATAVIRTUALIZATION_ENDPOINT_$
{DV_VERSION}_*_*.tar.gz /mnt/PV/versioned/uc_dsserver_shared/config/
cp -f /opt/ibm/qp_artifacts/archives/DATAVIRTUALIZATION_ENDPOINT_$
{DV_VERSION}_*_*.tar.gz /mnt/PV/versioned/dv_data/
cp -f /opt/ibm/qp_artifacts/archives/DATAVIRTUALIZATION_ENDPOINT_$
{DV_VERSION}_*_*.tar.gz /mnt/PV/versioned/QUERYPLEX_INSTALL/

save_and_delete_jars "/mnt/PV/versioned/uc_dsserver_shared/config/DATAVIRTUALIZATION_ENDPOINT_$
{DV_VERSION}_${ORIGINAL_BUILD}.zip" \
  "/mnt/PV/versioned/dv_data/DATAVIRTUALIZATION_ENDPOINT_${DV_VERSION}_${ORIGINAL_BUILD}.zip" \
  "/mnt/PV/versioned/QUERYPLEX_INSTALL/DATAVIRTUALIZATION_ENDPOINT_${DV_VERSION}_$
{ORIGINAL_BUILD}.zip"
cp -f /opt/ibm/qp_artifacts/archives/DATAVIRTUALIZATION_ENDPOINT_${DV_VERSION}_*_*.zip /mnt/PV/
versioned/uc_dsserver_shared/config/
cp -f /opt/ibm/qp_artifacts/archives/DATAVIRTUALIZATION_ENDPOINT_${DV_VERSION}_*_*.zip /mnt/PV/
versioned/dv_data/
cp -f /opt/ibm/qp_artifacts/archives/DATAVIRTUALIZATION_ENDPOINT_${DV_VERSION}_*_*.zip /mnt/PV/
versioned/QUERYPLEX_INSTALL/
chmod -x /opt/ibm/qp_artifacts/archives
```

Related information

Virtualizing data

Installing remote connectors

Troubleshooting Data Virtualization