

Spectrum™ Routing for Big Data

Version 4.0

Spectrum™ Routing for Big Data Install Guide: Hortonworks



Table of Contents

1 - Welcome

Introducing Spectrum™ Routing for Big Data	4
System Requirements and Dependencies	5

2 - Data Installation

Before You Begin	7
Installing the SDK	8
Distributing Reference Data	9

3 - Routing Setup

Hive Setup	12
Spark Setup	16

4 - Advanced Setup Options

Download Permissions	21
----------------------	----

1 - Welcome

In this section

Introducing Spectrum™ Routing for Big Data	4
System Requirements and Dependencies	5

Introducing Spectrum™ Routing for Big Data

Spectrum™ Routing for Big Data is a toolkit that can be used for processing enterprise data for large-scale spatial analysis. Billions of records from a single file can be processed in parallel using Hive's cluster processing framework, yielding results faster than ever. Unlike the traditional processing techniques that took weeks to process data, using this product, data processing can now be done in a just a few hours.

System Requirements and Dependencies

Spectrum™ Routing for Big Data is a collection of jar files that can be deployed to your Hadoop system. This product is verified on the following Hadoop distribution:

- Hortonworks 3.x

To use these jar files, you must be familiar with configuring Hadoop in Hortonworks and developing applications for distributed processing. For more information, refer to [Hortonworks](#) documentation.

To use the product, the following must be installed on your system:

for Hive:

- Hive version 1.2.1 or above

for Hive Client

- Beeline, for example

for Spark and Zeppelin Notebook:

- Java JDK version 1.8 or above
- Hadoop version 2.6.0 or above
- Spark version 2.0 or above

2 - Data Installation

In this section

Before You Begin	7
Installing the SDK	8
Distributing Reference Data	9

Before You Begin

Before you can use the Routing function, you must install routing data following the instructions in the Release Notes that accompany the data. See the following sections for specific data installation requirements for Hive and Spark.

Data Installation for Hive

A data resource file (dbList.json) that contains the datasets and database configurations is required. A sample dbList.json file is provided in the distribution at `spectrum-bigdata-routing-version.zip/resources/config`.

The `datasets paths` element must point to each extracted dataset, or `.spd` files if the data is stored in a remote location such as HDFS or S3.

```
{
  "defaultDatabase": "US",
  "datasets": [{
    "id": "US",
    "paths": ["hdfs:///pb/routing/data/US_Driving.spd"]
  }
],
  "databases": [{
    "name": "US",
    "datasets": ["US"]
  }
]
}
```

Note: If you are using S3, change the paths in the dbList.json accordingly.

Data Installation for Spark

The dbList.json data resource file is not required for the Spark API. You only need to install the routing data, and note the path to it for use later when setting up the distribution of the data in your cluster.

Installing the SDK

For the purposes of this guide, we will:

- use a user called pbuser
- install everything into /pb
- use installed routing datasets

Perform the following steps from a node in your cluster, such as the master node.

1. Create the install directory and give ownership to pbuser.

```
sudo mkdir /pb
sudo chown pbuser:pbuser /pb
```

2. Add the routing distribution zip to the node at a temporary location, for example:

```
/pb/temp/spectrum-bigdata-routing-version.zip
```

3. Extract the routing distribution.

```
mkdir /pb/routing
mkdir /pb/routing/software
unzip /pb/temp/spectrum-bigdata-routing-version.zip -d /pb/routing/software
```

4. Create an install directory on HDFS and give ownership to pbuser.

```
sudo -u hdfs hadoop fs -mkdir -p hdfs:///pb/routing
sudo -u hdfs hadoop fs -chown -R pbuser:pbuser hdfs:///pb
```

5. Upload the distribution into HDFS.

```
hadoop fs -copyFromLocal /pb/routing/software hdfs:///pb/routing
```

6. You will now need to distribute the reference data. A remote method via HDFS or S3 is recommended (see [Distributing Reference Data Using HDFS](#) on page 9 or [Distributing Reference Data Using S3](#) on page 9). If you wish to manually distribute the reference, see [Installing Reference Data Manually on Local File System](#) on page 10.

Distributing Reference Data

Distributing Reference Data Using HDFS

Now that the SDK is installed and the routing reference data is configured the reference data must be distributed around the cluster.

For the purposes of this guide, we will:

- continue using pbuser
- install the reference data into `hdfs:///pb/routing/data`

1. Create an install directory on hdfs and give ownership to pbuser.

```
sudo -u hdfs hadoop fs -mkdir hdfs:///pb
sudo -u hdfs hadoop fs -chown pbuser:pbuser hdfs:///pb
```

2. Upload the reference data into HDFS.

```
hadoop fs -mkdir hdfs:///pb/routing
hadoop fs -copyFromLocal /pb/routing/data hdfs:///pb/routing/data
```

3. When the data node performs routing tasks, the node will download the reference data from HDFS and onto the local file system. This means a local directory needs to be set up on all data nodes. Perform the following commands on all data nodes and HiverServer nodes.

```
sudo mkdir /pb/downloads
sudo chown pbuser:hadoop /pb/downloads
sudo chmod 775 /pb/downloads
```

Distributing Reference Data Using S3

Now that the SDK is installed and the routing reference data is configured the reference data must be distributed around the cluster. Your cluster must be set up for S3 access.

Use the s3a URI scheme on Hortonworks distributions.

For the purposes of this guide, we will install the reference data into `s3a://<your-bucket>/pb/routing/data`.

1. Upload the reference data into S3.

```
hadoop fs -mkdir -p s3a://<your-bucket>/pb/routing
hadoop fs -copyFromLocal /pb/routing/data
s3a://<your-bucket>/pb/routing/data
```

2. When the data node performs routing tasks, the node will download the reference data from S3 and onto the local file system. This means a local directory needs to be set up on all data nodes. Perform the following commands on all data nodes and HiverServer nodes.

```
sudo mkdir /pb/downloads
sudo chown pbuser:hadoop /pb/downloads
sudo chmod 775 /pb/downloads
```

Installing Reference Data Manually on Local File System

After the SDK is installed and the routing reference data is configured, the reference data must be distributed around the cluster. If you do not wish to use HDFS to do this, you can manually distribute the data instead.

1. Create the install directories on every data node and HiverServer2 node.

```
sudo mkdir -p /pb/routing/data/
sudo mkdir -p /pb/routing/software/resources/
sudo chown -R pbuser:pbuser /pb
```

2. Copy the data directory to the same path on every data node and HiveServer2 node.

```
/pb/routing/data/
```

Note: While it is recommended you use HDFS for the configuration directories and preference files, it is not a requirement. You may use local file system paths but you will have to make sure the file or directory is present and current on every node in the cluster. If you do this then you do not need to set `pb.download.location`.

3 - Routing Setup

In this section

Hive Setup	12
Spark Setup	16

Hive Setup

To set up routing for Hive, perform the following steps from the same node used in [Installing the SDK](#) on page 8:

1. On the HiveServer2 node, create the Hive auxlib folder if one does not already exist.

```
sudo mkdir /usr/hdp/current/hive-server2/auxlib/
```

2. Copy the Hive routing jar to the folder on the HiveServer2 node created in step 1:

```
sudo cp
/pb/routing/software/hive/lib/spectrum-bigdata-routing-hive-version.jar
/usr/hdp/current/hive-server2/auxlib/
```

3. Restart all Hive services.
4. Launch Beeline, or some other Hive client, for the remaining steps:

```
beeline -u jdbc:hive2://localhost:10000/default -n pbuser
```

5. Create the Routing functions and set the Hive variables. Add the `temporary` keyword after `create` if you want a temporary function (this step would need to be redone for every new Hive session). For more information, see [Hive Variables](#) on page 13.

```
create function PointToPointRoute as 'com.pb.bigdata.spatial.routing.hive.PointToPointRoute';
create function IsoChrono as 'com.pb.bigdata.spatial.routing.hive.IsoChrono';
create function IsoDistance as 'com.pb.bigdata.spatial.routing.hive.IsoDistance';
set pb.routing.config.location=hdfs://pb/routing/software/resources/config/dbList.json;
set pb.routing.allowFallback=true;
set pb.download.location = /pb/downloads;
set pb.download.group= pbdownloads;
```

6. Test the Routing function.

```
SELECT IsoDistance(-77.088217, 38.937072, 3, map('distanceUnit',
'km'));
```

Note:

- The first time you run the job may take a while if the reference data has to be downloaded remotely from HDFS or S3. It may also time out when using a large number of datasets that are stored in remote locations such as HDFS or S3. If you are using Hive with the MapReduce engine, you can adjust the value of the `mapreduce.task.timeout` property.
- Some types of queries will cause Hive to evaluate UDFs in the HiveServer2 process space instead of on a data node. The Routing UDFs in particular use a significant amount of memory and can shut down the Hive server due to memory constraints. To

process these queries, we recommend increasing the amount of memory available to the HiveServer2 process (for example, by setting HADOOP_HEAPSIZE in hive-env.sh).

This query should return a polygon geometry comprising all the points that lie at the specified distance from the starting point.

Hive Variables

The Routing UDFs use Hive variables to set various properties, which take precedence over the system properties. They can also be set as part of an option in a UDF query, where they will take precedence over the Hive variables as well as the system properties.

Variable	Required	Example
<code>pb.routing.config.location</code> Location of dbList.json containing the routing configuration.	Yes	<code>hdfs://pb/routing/software/resources/config/dbList.json;</code>
<code>pb.routing.engine.timeout</code> Specifies the duration (in milliseconds) after which a GRA engine shuts down if it does not receive any request and is idle. The default value for this property is 900000 (15 minutes).	No	600000
<code>pb.routing.routeTimeout</code> Specifies the amount of time (in milliseconds) allowed for a point-to-point routing job to be processed and completed. If the job is not processed or completed within the specified duration, the SQL returns null to the console and exception to logs for each record processed. This property can be used for the PointToPointRoute UDF.	No	10000
<code>pb.routing.isoTimeout</code> Specifies the amount of time (in milliseconds) allowed for a boundary routing job to be processed and completed. If the job is not processed or completed within the specified duration, the SQL returns null to the console and exception to logs for each record processed. This property can be used for the IsoDistance and IsoChrono UDFs	No	17000

Variable	Required	Example
<code>pb.routing.allowFallback</code>	No	false
<p>This Boolean property controls whether the engine can use major roads in case other roads cannot be used.</p> <p>If set to true, the engine can use major roads.</p>		
<code>pb.download.location</code>	No	/pb/downloads
<p>Note: Use only if reference data was distributed remotely via HDFS or S3.</p> <p>Location of the directory where reference data will be downloaded to. This path must exist on every data node and the HiveServer2 node.</p>		
<code>pb.download.group</code>	No	pbdownloads
<p>Note: Use only if reference data was distributed remotely via HDFS or S3.</p> <p>This is an optional property and only specific to POSIX-compliant platforms like Linux. It specifies the operating system group which should be applied to the downloaded data on a local file system, so that each Hadoop service can update the data when required. This group should be present on all nodes in the cluster and the operating system user executing the Hadoop service should be a part of this group.</p> <p>For more information, see Download Permissions on page 21.</p>		
<code>pb.routing.error.limit</code>	No	1
<p>The number of routing errors to allow before failing a task for "too many errors". This prevents a task (and thus the job) from continuing in the case of a likely configuration error. The default value is 10.</p>		
<code>pb.routing.error.limit.disabled</code>	No	true
<p>Disables the error limit. All errors will be logged but will not cause the task or job to fail.</p>		

Setting Variables

Since these are Hive configuration variables, you can set them permanently by editing the `hiveserver2-site.xml` file in your cluster.

Example (using HDFS)

```
pb.routing.config.location=hdfs:///pb/routing/software/resources/config/dbList.json;
pb.download.location = /pb/downloads
pb.download.group= pbdownloads
```

Alternatively you can set them temporarily in each Hive session that you open.

For example, if you wish to perform routing in Beeline you can set the variables in the following way:

Example (using HDFS)

```
set pb.routing.config.location=hdfs:///pb/routing/software/resources/config/dbList.json;
set pb.routing.allowFallback=true;
set pb.download.location = /pb/downloads;
set pb.download.group= pbdownloads;
```

Variables set in the local session will override any hive-site variables. This means you can have default values set in the hiveserver2-site file and override them in Beeline when necessary.

System Variables (Deprecated)

Hive variables used to exist in the system namespace. This is deprecated as of version 3.1 and no longer the recommended method.

Spark Setup

Routing Jar for Routing Operations

If you want to use routing operations in Spark, a Routing jar file is available which supports compiling and running routing programs.

The Routing jar includes the Download Manager API, which provides the capability to manage downloading the routing datasets programmatically. If your routing datasets are stored remotely, then use the DownloadManager API to create a `DownloadManager` instance and pass it to the `GRAInstanceBuilder`, which will download and manage the datasets for you. In addition, the Download Manager can set permissions in the case where you have multiple users interacting with the data.

- For more information about the Download Manager API, see the Javadocs located in:
`/pb/routing/software/spark2/sdk/javadocs/downloadmanager`

The Routing jar also includes a Scala class called `GRAInstanceBuilder` that constructs an `IGRAInstance` to use for Routing requests.

- For more information about the builder, see the Scaladocs located in:
`/pb/routing/software/spark2/sdk/scaladocs/bigdata`
- For more information about the Global Routing API (GRA), see the Javadocs located in:
`/pb/routing/software/spark2/sdk/javadocs/gra`

There are several options for deploying the Routing jar; refer to the following table to proceed according to your use case:

Application	Do this
Spark job	Installing and setting up the Routing Jar for a Spark Job on page 17
Zeppelin notebook	Integrating the Routing Jar with Zeppelin on page 17
Hue notebook	Integrating the Routing Jar with Hue on page 18

Installing and setting up the Routing Jar for a Spark Job

Perform the following steps:

1. Deploy the `spectrum-bigdata-routing-sdk-spark2-version.jar` to the Hadoop cluster.
2. Start the Spark job using the following command:

Note: This example assumes the product is installed to `/pb/routing/software` as described in [Installing the SDK](#) on page 8.

```
spark-submit --class com.example.spark.app.MyDriver
--master yarn --deploy-mode cluster
--jars
/pb/routing/software/spark2/sdk/lib/spectrum-bigdata-routing-sdk-spark2-version.jar
/path/to/custom/driver.jar driverParameter1 driverParameter2 ...
```

where:

- `--jars` The path to the jar file.

Integrating the Routing Jar with Zeppelin

The Routing jar can be loaded and referenced in a Zeppelin notebook as an external library.

Perform the following steps to use the Routing jar in Zeppelin:

1. On the navigation bar, using the **Settings** drop-down menu and click **Interpreter**. The **Interpreters** page displays.
2. Go to the **Spark2** section, click **Edit**. The editing view displays.
3. Go to the **Dependencies** section. Enter the full local path to the jar in the **artifact** field:

Note: This example assumes the product is installed to `/pb/routing/software` as described in [Installing the SDK](#) on page 8.

```
/pb/routing/software/spark2/sdk/lib/spectrum-bigdata-routing-sdk-spark2-version.jar
```

4. Click **Save**. This will restart the interpreter with the loaded library.
5. To return to your notebook, click the **Notebook** drop-down menu and select your notebook. The Routing jar is now available to all notebooks.

Integrating the Routing Jar with Hue

The Routing jar can be loaded and referenced in a Hue notebook as an external library.

Prerequisite: To use the Routing jar in an interactive manner, you will need to use the Scala or Java interpreter which requires Livy.

Proceed according to your platform:

1. Cloudera - start here.

a) In Cloudera Manager, navigate to **Hue -> Configuration**.

b) Search for this property:

```
Hue Service Advanced Configuration Snippet (Safety Valve) for
hue_safety_valve.ini
```

c) Use the following 3 lines to set the property's value:

```
[desktop]
app_blacklist=
use_default_configuration=true
```

d) Click **Save**. Continue to step 3.

2. EMR - start here.

a) On the Master Node, open the file `/etc/hue/conf/hue.ini`.

b) Uncomment and change the following line from:

```
## use_default_configuration=false
to
use_default_configuration=true
```

c) Save the file. Continue to step 3.

3. Restart Hive.

4. Open the Hue editor.

5. Select one of the query editors (Scala or Java).

6. Establish a session. This can be accomplished by running any operation in the editor, executing `1+1` will suffice.

7. When the query completes, click the gears icon button (in some versions of Hue, this button becomes visible after you click the vertical ellipsis button).

8. In the session dialog, use the **Add Property** drop-down menu and select **Jars**.

9. In the text box, specify the location in HDFS where your Routing jar is located. For example:

```
hdfs:///pb/routing/software/spark2/sdk/lib/spectrum-bigdata-routing-sdk-spark2-version.jar
```

10. Click **Recreate**. This will recreate the session.

11. To return to your notebook, click any of the arrow icons or the gears button again. This will close any open dialogs and add the Routing jar to the session that was just created.

4 - Advanced Setup Options

In this section

Download Permissions

21

Download Permissions

Setting the download permissions allows multiple services to download and update the downloaded data when required. You should have a common operating system group of which all the service users who need to download the data are part of. For example, if Hive and YARN jobs are required to download data and use the same download location, then both the Hive and YARN operating system users should be part of a common operating system group. The group of the download directory should be the common operating system group, one that has Read, Write, and Execute (775) permissions for the owner and group.

Your group should contain services and users that will run jobs in your cluster. You may skip services you will not use or do not have installed. Services include YARN, Hive, Zeppelin, and Hue.

You also should include all operating system users who will run jobs such as `pbuser` and `<myOtherUser>`.

1. Add the group.

```
sudo groupadd pbdownloads
```

2. Add users to the group.

```
sudo usermod -a -G pbdownloads hive
sudo usermod -a -G pbdownloads yarn
sudo usermod -a -G pbdownloads zeppelin
sudo usermod -a -G pbdownloads hue
sudo usermod -a -G pbdownloads pbuser
sudo usermod -a -G pbdownloads <myOtherUser>
```

3. Using a window where no job is running, restart all the services whose operating system users were added to the new group.
4. Using a window where no job is running, restart the session of all the operating system users that were added to new group (for example, `pbuser`).
5. Update the group to the common operating system group and update permissions to 775 for the download directory specified in `pb.download.location` property.

```
sudo chgrp pbdownloads /pb/downloads
sudo chmod 775 /pb/downloads
```

Copyright

Information in this document is subject to change without notice and does not represent a commitment on the part of the vendor or its representatives. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, without the written permission of Pitney Bowes Software Inc., 350 Jordan Road, Troy, New York 12180.

© 2020 Pitney Bowes Software Inc. All rights reserved. Location Intelligence APIs are trademarks of Pitney Bowes Software Inc. All other marks and trademarks are property of their respective holders.



3001 Summer Street
Stamford CT 06926-0700
USA

www.pitneybowes.com