Installing the HBase shell for Cloud Bigtable

The HBase shell is a command-line tool that performs administrative tasks, such as creating and deleting tables. The Cloud Bigtable HBase client for Java makes it possible to use the HBase shell to connect to Cloud Bigtable.

This page describes how to install the HBase shell, along with the Cloud Bigtable HBase client for Java, on a Compute Engine instance or on your own machine. You do not need to install any other HBase dependencies, such as Apache ZooKeeper. The Cloud Bigtable HBase client for Java does not require these additional dependencies.

You can also run the HBase shell on Google Cloud Shell (/shell). See Quickstart Using HBase Shell (/bigtable/docs/quickstart-hbase) for instructions.

If you want to set up a cluster of Compute Engine instances that can run Hadoop MapReduce jobs, see Creati
p Cluster for Cloud Bigtable (/bigtable/docs/creating-hadoop-cluster) for instructions.

Before you begin

Before you begin, you'll need to complete the following tasks:

- Create a Cloud Bigtable instance (/bigtable/docs/creating-instance). Be sure to note the project ID and instance ID.
- Install the `gcloud` command-line tool. See the Cloud SDK setup instructions (/bigtable/docs/installing-cloud-sdk) for details.
- If you're installing the HBase shell on a Compute Engine instance, create an instance (/bigtable/docs/creating-compute-instance) that has the correct scopes for Cloud Bigtable.

Installing Java

To use the HBase shell with the Cloud Bigtable HBase client for Java, you must install a Java 8 runtime environment. Other versions of Java are not supported.

On Debian GNU/Linux or Ubuntu, you can install the Java 8 JDK by running the following commands:

```
apt-get update
apt-get install openjdk-8-jdk-headless
```

For other operating systems, you can download and install the Java JDK (http://www.oracle.com/technetwork/java/javase/downloads/).

Installing Maven

To run the HBase shell, you must install Apache Maven (https://maven.apache.org/). The HBase shell uses Maven to download the required dependencies.

On Debian GNU/Linux or Ubuntu, you can install Maven by running the following commands:

```
apt-get update
apt-get install maven
```

On macOS, you can install Maven using Homebrew (https://brew.sh/):

```
install maven
```

For other operating systems, see the instructions on the Maven website (https://maven.apache.org/).

Obtaining credentials

When you install the HBase shell on your own machine, you need to obtain user access credentials for your Google Cloud Platform resources.

If you're installing the shell on a Compute Engine instance, you don't need to obtain credentials.

To obtain credentials, run the following command:

```
d auth application-default login
```

The command opens your web browser. Choose your Google account if prompted, then click **Allow** to authorize access to your account.

Setting up the HBase shell

Downloading required files

From the command line, clone the GitHub repository that contains the HBase shell quickstart:

```
lone https://github.com/GoogleCloudPlatform/cloud-bigtable-examples.git
```

The repository is cloned to the `cloud-bigtable-examples` directory.

Setting the JAVA_HOME environment variable

Your shell's `JAVA_HOME` environment variable should be set to the location where your Java runtime environment is located.

To find the current value of `JAVA_HOME`:

```
${JAVA_HOME}
```

If this command does not print any output, the variable is not set.

To set the value on Debian GNU/Linux or Ubuntu:

```
t JAVA_HOME=$(update-alternatives --list java | tail -1 | sed -E 's/\/bin\/java//')
```

A typical value on Debian GNU/Linux or Ubuntu is `/usr/lib/jvm/java-8-openjdk-amd64/jre`.

To set the value on macOS (requires Xcode's command-line tools):

```
t JAVA_HOME=$(/usr/libexec/java_home)
```

A typical value on macOS is `/Library/Java/JavaVirtualMachines/jdk1.8.0_131.jdk/Contents/Home`.

After you set the value of `JAVA_HOME`, run the following command to ensure that the value refers to the appropriate Java version:

```
A_HOME}/bin/java -version
```

Connecting to a Cloud Bigtable instance

You can use the HBase shell to connect to your Cloud Bigtable instance, get information about your existing tables, and create new tables.

To start the HBase shell, run the following commands in the directory where you cloned the `cloud-bigtable-examples` GitHub repository:

```
ickstart
ckstart.sh
```

If your project has only one Cloud Bigtable instance, the HBase shell will automatically connect to that instance. If your project has multiple Cloud Bigtable instances, the HBase shell will display a list of instance IDs; type the number next to your instance ID, then press Enter.

The first time you start the HBase shell, there is a delay while Maven downloads the required JAR files. This de
s only the first time you start the HBase shell.

After the HBase shell connects to your instance, it displays a series of log messages, which is normal. The shell prompt then appears:

```
(main):001:0>
```

Now that you've started the HBase shell, try running the following commands:

- To list your tables, type `list` and press Enter. If you haven't created any tables yet, the HBase shell displays a message similar to the following:

  ```
  TABLE
  0 row(s) in 1.3580 seconds

  => []
  ```

- To create a table, use the `create` command. For example, to create a table named `my-table`, with one column family named `cf1`:

  ```
  create 'my-table', 'cf1'
  ```

  The shell displays a message similar to the following:

  ```
  0 row(s) in 1.5210 seconds

  => Hbase::Table - my-table
  ```

- To add data to a table, use the `put` command. For example, to put the value `test-value` in the row `r1`, using the column family `cf1` and the column qualifier `c1`:

  ```
  put 'my-table', 'r1', 'cf1:c1', 'test-value'
  ```

  You can then use the `scan` command to scan the table and read the data you added:

```
scan 'my-table'
```

The shell displays output similar to the following:

```
ROW      COLUMN+CELL
 r1      column=cf1:c1, timestamp=1430687836046, value=test-value
1 row(s) in 0.6260 seconds
```

- To drop a table, use the `disable` command to disable the table, followed by the `drop` command. For example, to drop the table `my-table`:

```
disable 'my-table'
drop 'my-table'
```

- To exit the HBase shell, type `exit` and press Enter.

While the HBase shell is running, you can type `help` to get usage instructions. See the HBase documentation (https://hbase.apache.org/book.html#shell) for additional information.

Last updated 2020-06-25 UTC.