



Quick Deployment

Step-by-step instructions to deploy Oracle Big Data Lite Virtual Machine

Version 4.1.0

Please note: This appliance is for testing and educational purposes only; it is unsupported and not to be used in production.

CONTENTS

Introduction	1
Installation Steps	2
Troubleshooting Tips	6
Useful Tips	8

INTRODUCTION

Oracle Big Data Lite Virtual Machine provides an integrated environment to help you get started with the Oracle Big Data platform. Many Oracle Big Data platform components have been installed and configured - allowing you to begin using the system right away. See the [Big Data Lite landing page](http://www.oracle.com/technetwork/database/bigdata-appliance/oracle-bigdatalite-2104726.html) (<http://www.oracle.com/technetwork/database/bigdata-appliance/oracle-bigdatalite-2104726.html>) on OTN for tutorials, videos, white papers and more.

The following components are included on Oracle Big Data Lite:

- Oracle Enterprise Linux 6.5
- Oracle Database 12c Release 1 Enterprise Edition (12.1.0.2) - including Oracle Big Data SQL-enabled external tables, Oracle Multitenant, Oracle Advanced Analytics, Oracle OLAP, Oracle Partitioning, Oracle Spatial and Graph, and more.
- Cloudera Distribution including Apache Hadoop (CDH5.3.0)
- Cloudera Manager (5.3.0)

- Oracle Big Data Connectors 4.1
 - Oracle SQL Connector for HDFS 3.1.0
 - Oracle Loader for Hadoop 3.3.0
 - Oracle Data Integrator 12c
 - Oracle R Advanced Analytics for Hadoop 2.4.1
 - Oracle XQuery for Hadoop 4.1.0
- Oracle NoSQL Database Enterprise Edition 12cR1 (3.2.5)
- Oracle JDeveloper 12c (12.1.3)
- Oracle SQL Developer and Data Modeler 4.0.3
- Oracle Data Integrator 12cR1 (12.1.3)
- Oracle GoldenGate 12c
- Oracle R Distribution 3.1.1
- Oracle Perfect Balance 2.3.0
- Oracle CopyToBDA 1.1
-

Big Data Lite includes software products that are optional on the Oracle Big Data Appliance (BDA), including Oracle NoSQL Database Enterprise Edition and Oracle Big Data Connectors.

To get started, you will perform the following high level tasks:

- Ensure your host computer meets the requirements for running the virtual machine
- Download and install VirtualBox – which is a cross-platform virtualization application
- Download the zip files from OTN and use 7-zip to extract the appliance file
- Create the VM by import the BigDataLite-xxx.ova file into VirtualBox
- Start the machine!
- Log into Big Data Lite using:

user id: **oracle**

password: **welcome1**

- Open the Start Here document on the desktop for details about the environment.

INSTALLATION STEPS

1. Prepare your host system.

- Minimum 8GB of real memory; more is better. 5G memory will be dedicated to the VM. Note: if you want to use Cloudera Manager (which is not required), you will need to dedicate 10G memory to the VM.
- Turn on Virtual Assist features in the BIOS (usually done by default). Refer to [Troubleshooting tips](#) in Step 12 for more details.

- ~51GB disk space needed to download and install. (This includes the 15.6 GB zipped .ova file and ~35GB of the imported image). Once the image is imported, the 15.6GB .ova file can be removed and the imported appliance of 35GB is all that you need to run the image.
- Download and install [7Zip](#).
- Download and install md5sum (optional)

2. Download and install [Oracle Virtual Box](#) (4.3 and above is supported).

It is highly recommended that you also install **the Virtual Box Guest Additions** – which provides enhanced mouse support, clipboard support and sharing of files between the virtual machine and its host.

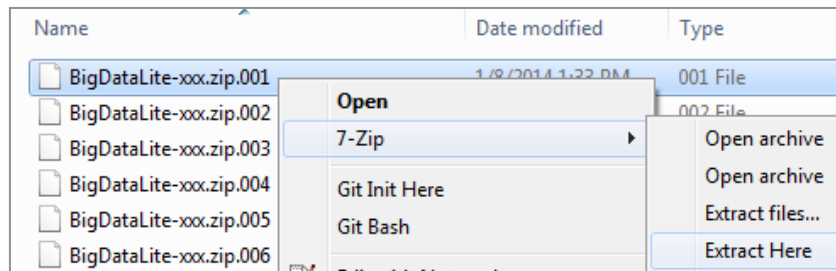
3. Download Big Data Lite files from OTN

- Go to the Big Data Lite landing page at:
<http://www.oracle.com/technetwork/database/bigdata-appliance/oracle-bigdatalite-2104726.html>
- Click to accept the OTN license.
- Download all of the files to a single directory.

4. After all zip files are downloaded, extract them with 7zip.

In Windows:

- Right-click on file **bigdatalite41.7z.001** file and select **7-zip->Extract Here**



In Linux:

- Run the following command to extract the file:

```
7za e bigdatalite41.7z.001
```

Depending on your computer, the extraction should take under 15 minutes.

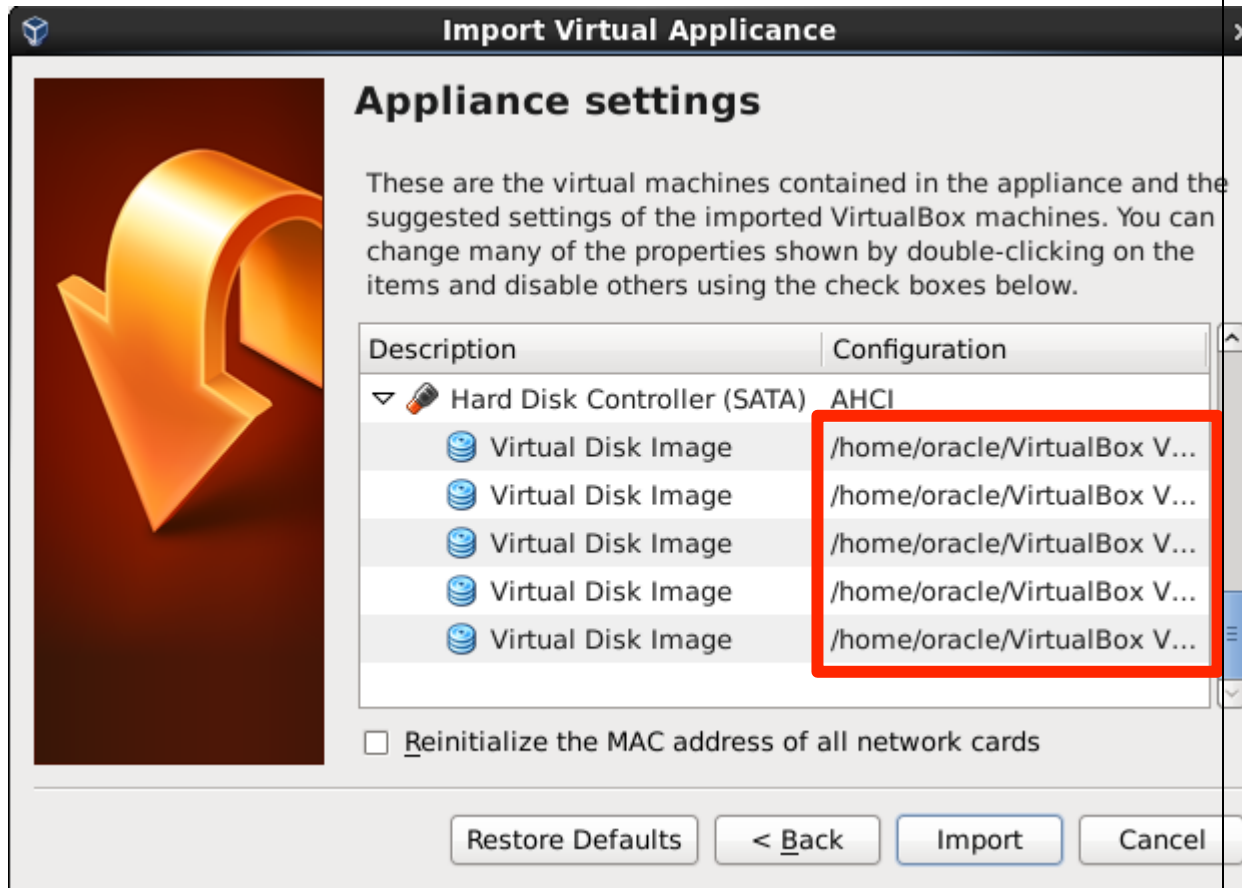
Result:

The extraction will create the **BigDataLite-xxx.ova** appliance file. This file will be used to create a new machine in Virtual Box. This single file contains the entire machine definition, including the physical disks as well as defaults for the machine configuration (e.g. network, CPUs, memory, etc.).

5. Start Oracle Virtual Box Manager and Import the Appliance

After starting Virtual Box, create the new virtual machine using the import wizard:

- Click **File** -> **Import Appliance** to launch the import wizard
- In the Import Virtual Appliance page, click **Open appliance...**
- Locate the **BigDataLite-xxx.ova** file and click **Open**. Click **Next**.
- **Appliance settings** give you an overview of the configuration. In this step, you may need to update the location of the **Virtual Disk Images**. These files will initially use approximately 30GB of disk – and can expand as you use the machine. Change the default location of the Virtual Disk Images if the current location does not have enough capacity. Note, the VM uses 4 separate disks. These disks should not be moved after importing the VM:




Click **Import**

- Review the Software License Agreement and click **Agree** to continue.

Depending on your computer, the extraction should take under 25 minutes.

Result:

A new BigDataLite-xxx Virtual Machine is available in the Virtual Box Manager:

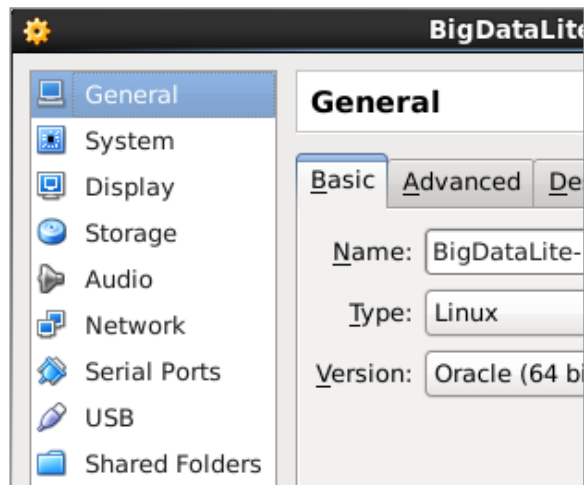


6. Virtual Machine Settings

Big Data Lite is configured with the following parameters, which is ideal for an 8GB host system:

- 5GB RAM
- 2 processors
- NAT Network Adapter

You can update these default settings by selecting the Big Data Lite VM in Virtual Box Manager and clicking **Settings**:



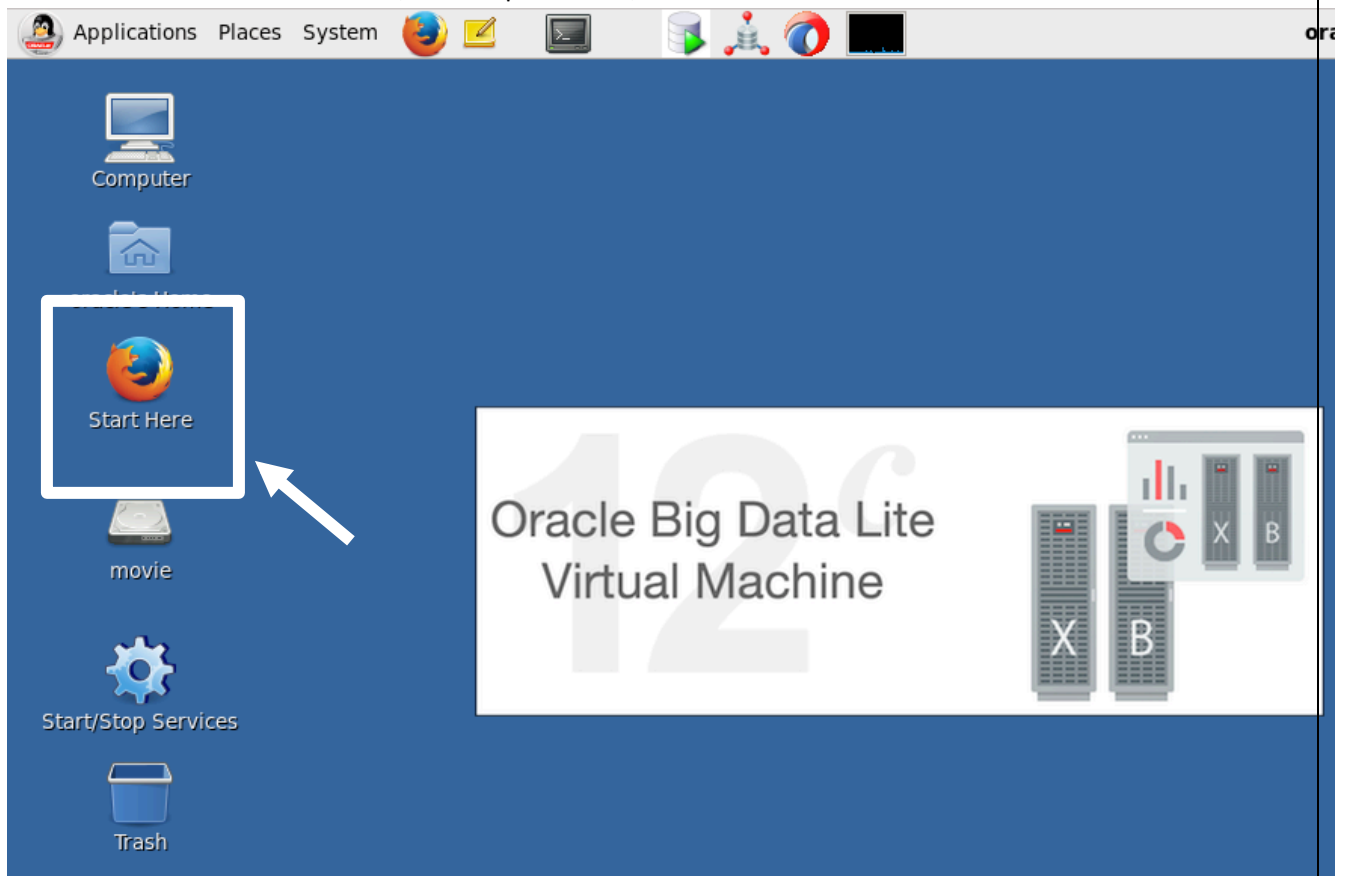
Review the Virtual Box documentation to learn details about updating the configuration.

7. Start your Big Data Lite!

Double-click on BigDataLite-xxx to start the VM. Log on as the oracle user to get started:

```
id:      oracle
password: welcome1
```

This will bring you to a Linux desktop. The **Start Here** document on the desktop will provide you details about the installed software, ids and passwords, and more.



Big Data Lite Desktop

TROUBLESHOOTING TIPS

1. ERROR: *"Failed to open a session for the virtual machine"* when attempting to start the machine

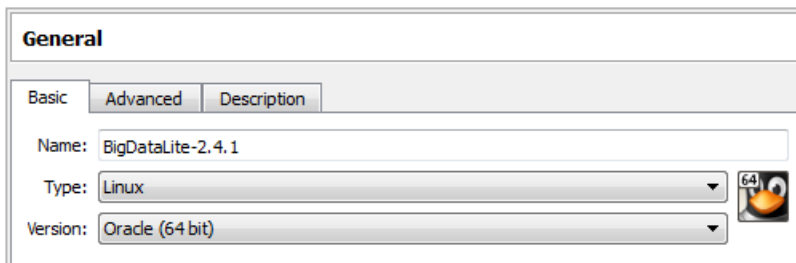
The VM requires that hardware virtualization is enabled for the host computer. If it is not enabled, you will receive the following error when attempting to start the machine:

*Failed to open a session for the virtual machine BigDataLite-xxx.
VT-x is disabled in the BIOS
(VERR_VMX_MSR_VMXON_DISABLED)*

You can update the BIOS at boot time for the host. Ensure that both "Intel(R) Virtualization Technology" and "Intel(R) VT-d Feature" are enabled. Please review the instructions provided by your hardware manufacturer that describe updating BIOS settings.

2. ERROR: "This kernel requires an x86-64 CPU, but only detected an i686 CPU. Unable to boot - please use a kernel appropriate for your CPU"

Ensure that your machine is running Oracle (64-bit) Linux in the VM settings select "General" and set the Version to Oracle (64-bit).




3. While the image is starting up, you initially see a few messages and then a blank screen

If this happens, then stop the image. Go back to Virtual Box Manager and click on the image's **Settings**. Go to **Display** and increase the **Video memory** to at least 6MB. Start the image and you should no longer see the blank screen.

4. Unable to copy/paste between the host and Big Data Lite

First, ensure that you have installed the Virtual Box Guest Additions. If you have installed Guest Additions and copy/paste still doesn't work – go to the **Devices** menu and ensure that **Shared Clipboard** is enabled.

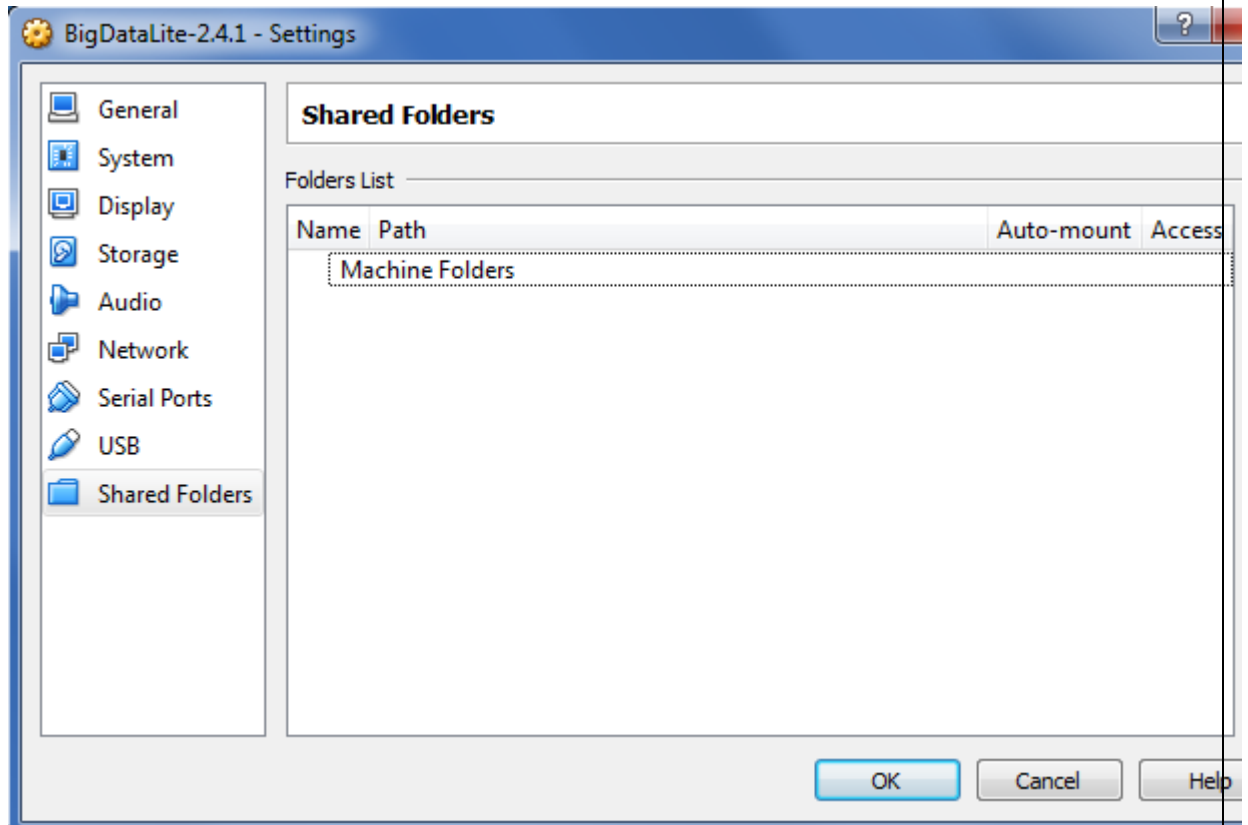


USEFUL TIPS


1. Sharing folders between Big Data Lite and its host

Virtual Box Guest Additions provides a useful feature that enables you to share a host's folder within the virtual machine. To configure a shared folder:

- Select BigDataLite-xxx from the list of machines and click **Settings**
- Go to **Shared Folders** and click **Add Folder**



- Specify the folder you would like to share and select **Auto-mount**. Then click **OK**. Below, the H:\ drive is being shared.



You can repeat this process for multiple shared folders.

- To access the shared folder from Big Data Lite, navigate to `/media/sf_<folder_name>`


```

oracle@bigdatalite:~
File Edit View Search Terminal Help
[oracle@bigdatalite ~]$ ls -l /media
total 8
drwxrwx--- 1 root vboxsf 8192 Jan 17 08:46 sf_H_DRIVE
[oracle@bigdatalite ~]$

```






2. Using the VM as a Hadoop Client

You can use Big Data Lite as a Hadoop client – allowing you to use it as a development environment for the Oracle Big Data Appliance. To connect the Big Data Lite to BDA, log into the Cloudera Manager instance that is managing your Hadoop cluster and select **Actions -> View Client Configuration URLs**:



This will take you to the Client Configuration URLs page. Select the appropriate Client Configuration (e.g. MapReduce):

Client Configuration URLs x

Name	Type	URL
 hdfs ▼	HDFS	/cmf/services/1/client-config
 mapreduce ▼	MapReduce	/cmf/services/2/client-config
 hive ▼	Hive	/cmf/services/5/client-config
 hbase1 ▼	HBase	/cmf/services/9/client-config
 solr1 ▼	Solr	/cmf/services/14/client-config

Close

This will download a zip file containing all the files required to connect to the cluster. After downloading the file, execute the following two steps:

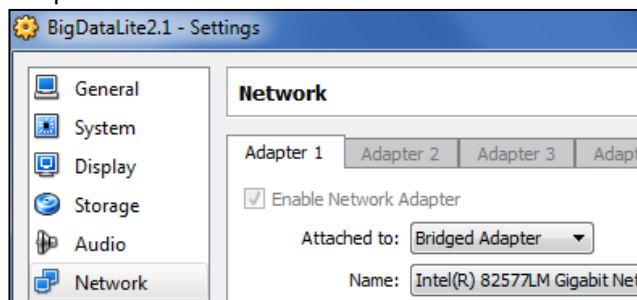
- Unzip the file into /home/oracle.
- Set the appropriate environment variable that will point Big Data Lite to the remote cluster. For example, if you downloaded both hive and mapreduce configurations:

```
export HADOOP_CONF_DIR=/home/oracle/hadoop-conf
export HIVE_CONF_DIR=/home/oracle/hive-conf
```

3. Connecting external clients to Big Data Lite

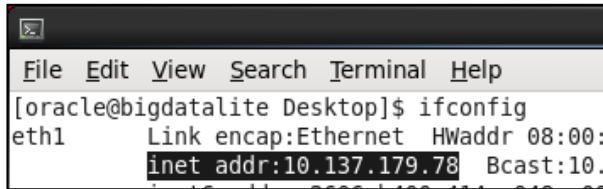
You may want to connect external clients to Big Data Lite – effectively using the VM as a server. This requires minor updates to your Networking setup.

- Stop the Virtual Machine and update its Network settings. Change the adapter to Bridged Adapter.



This will change the IP address of Big Data Lite – requiring updates to its configuration:

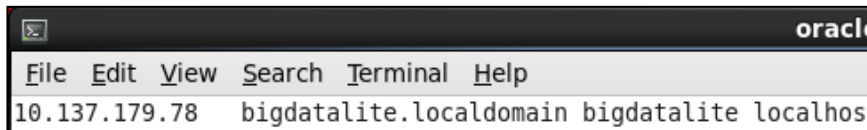
- Start the VM, log in and open a Terminal Window. Find the IP Address by typing `ifconfig` at the command line:



```

[oracle@bigdatalite Desktop]$ ifconfig
eth1      Link encap:Ethernet  HWaddr 08:00:
          inet addr:10.137.179.78  Bcast:10.
  
```

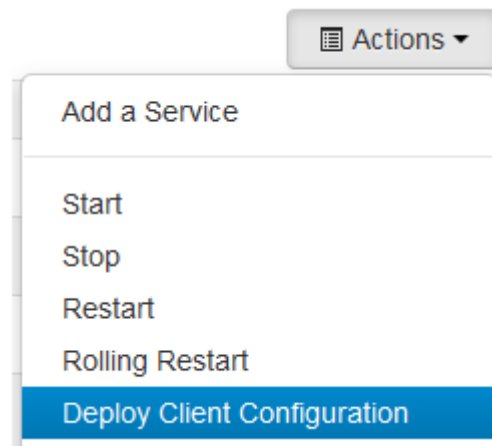
- Edit `/etc/hosts` (e.g. `sudo vi /etc/hosts`) and map this IP address to `bigdatalite.localdomain`. This should replace the `127.0.0.1` IP address



```

10.137.179.78  bigdatalite.localdomain bigdatalite localhost
  
```

- Reboot Big Data Lite
- Log into Cloudera Manager and **Deploy Client Configuration** from **the All Services** page:



- Reboot Big Data Lite

External clients can now access the VM using its IP address.