

Hortonworks

Setup Guide for HDP Developer: Storm

Revision 1

Hortonworks University

Overview

The Hortonworks Training Course that you are attending is taught using a virtual machine (VM) for the lab environment. Before attending the course, you must complete the steps in this document. Here is an overview of the steps you will perform:

- Download the virtual machine file. The VM file is in a folder online that we will share with you. The file is very large (~ 4GB) and may take several hours to download if you have a slower Internet connection.
- You must install **VMWare Player** to open the VM file on Windows. On Mac, you can download a free trial of **VMWare Fusion** at <http://www.vmware.com/products/fusion/>.
- You must start the virtual machine after it is imported into VMWare to verify that it can run on your machine's hardware and software.
- It is possible that you will need to configure the BIOS on your Windows machine to enable virtualization.
- **THE FOLLOWING SETUP INSTRUCTIONS MUST BE COMPLETED AT LEAST 24 HOURS PRIOR TO THE START OF THE CLASS!**

IMPORTANT: It is critical for the success of the training that prior to the start of the class each student's computer has the VM file downloaded, imported into VMWare, and verified that the virtual machine starts on the student's computer and a MapReduce can execute successfully.

IMPORTANT: If any of the following steps fail, or if you have any issues or questions, please send an email to training-support@hortonworks.com.

System Requirements

The following is the recommended minimal system requirements:

- Reasonably powerful **x86/amd64 hardware**. Any **recent Intel (3rd generation i5/i7 or above preferred)** or equivalent AMD processor should do.
- RAM: At least **8 GB**
- Available Hard-disk: at least **50 GB**
- **64-bit OS** (Windows 7, Windows 8 and Mac OSX)
- At least 800 x 600 display
- VMWare Player for Windows (or VMWare Fusion for Mac)
- **Virtualization must be enabled on BIOS** for Windows based machine
 - Procedures to turn on virtualization settings in your computer's BIOS vary depending on the BIOS manufacturer. Check the information that came with your computer or go to the computer manufacturer's website.
 - Alternatively, to find out whether your computer's CPU is capable of hardware-assisted virtualization, you can download and run the Hardware-Assisted Virtualization Detection Tool (<http://www.microsoft.com/en-us/download/details.aspx?id=592>) provided by Microsoft.
- Internet connection must be available on the system.

Complete the Following Steps

Objective:	Install the necessary software for this course, and also import the VMWare image into VMWare Player and start the virtual machine successfully.
Before You Begin:	Make sure your hardware meets the requirements listed above.

Perform the following steps at least one day prior to the start of the course:

Step 1: Install VMWare Player

1.1. On Windows, download the latest 64-bit version of VMWare Player:

<http://www.vmware.com/support/download-player.html>

On a Mac, download a free trial of VMWare Fusion at

<http://www.vmware.com/products/fusion/>.

1.2. Run the downloaded installation file.

Step 2: Download the VM

2.1. Download the VM here. This file is about 5GB and may take several hours to download, depending on your internet connection:

<http://tinyurl.com/stormrev1>

Step 3: Unzip the Virtual Machine

3.1. Unzip the file you just downloaded, which may take a couple of minutes.

NOTE: Some versions of the built-in Windows extraction tool have been known to cause issues. ***You may have to unzip the file using a third-party tool like 7-Zip.***

Step 4: Open the VM with VMWare Player

4.1. Start the VMWare Player application and select the option to **Open a Virtual Machine**:




4.2. Browse to the folder you just unzipped, which should be named **HWX_classroom_Storm_RevN**.


4.3. Within the **HWX_classroom_Storm_RevN.vmware** folder, select the file **HWX_classroom_2.1_Rev2.vmx**.

4.4. Click the **Open** button.

4.5. If a dialog appears stating that the virtual machine may already be in use, click the **Take Ownership** button:



4.6. You should now see the VM in the VMWare Player library:



Step 5: Change the VM Name


- 5.1. Right-click on “**HWX_classroom_Storm_RevN**” in the list of VMs and select **Settings....**
- 5.2. Select the **Options** tab.
- 5.3. Change the **Virtual machine name** to **HWX_Storm**.
- 5.4. Select the **Hardware** tab. Make sure the VM is configured to use at least 4076MB of memory. If your hardware has 16GB or more of RAM, then feel free to increase the memory for this VM to 6 or 8GB. (This is not a requirement – just an option if your hardware has a large amount of physical memory.)
- 5.5. Click the **OK** button to close the **Virtual Machine Settings** dialog.

Step 6: Start the VM

6.1. Click the **Play virtual machine** link to start your **HWX_Java** VM. If a dialog appears asking if you moved or copied the VM, click the **“I copied it”** button:



6.2. Wait for the VM to start. When the VM has started successfully, you should see a login screen:



6.3. Login as the **root** user. The password is **hadoop**. When you are logged in, desktop for the VM should display:



Step 7: Run a Smoke Test

7.1. To verify that your machine has enough memory, try starting a Hadoop cluster on your VM. Enter the following command, which starts up a 5 node Hadoop cluster in your VM:

```
root@ubuntu:~# storm_cluster.sh
```

7.2. Wait for few minute for your cluster to start.

7.3. To verify all the components of the cluster are running, run the following command:

```
root@ubuntu:~# hdfs dfsadmin -report
```

You should see a list of two DataNodes in your cluster.

7.4. Enter the following command:

```
root@ubuntu:~# yarn node -list
```

You should see 2 NodeManager nodes in your cluster.

Step 8: Cleanup the Smoke Test

8.1. Enter the following command to stop and delete the cluster:

```
root@ubuntu:~# destroy_cluster.sh
```

Step 9: Suspend the VM

9.1. Suspend the virtual machine by clicking the “X” to close the window and a dialog will appear. Select “**Suspend**”.

RESULT: You should now have the classroom VM downloaded and imported into VMWare Player (or Fusion). The VM should be able to start successfully and run a Hadoop cluster.

IMPORTANT: If any of these steps failed or if you have any issues or questions, please send an email to training-support@hortonworks.com.