



Hortonworks Sandbox with VMware Fusion

April 2015

Overview

The Hortonworks Sandbox is delivered as a virtual appliance. The virtual appliance (indicated by an .ovf or .ova extension in the filename) runs in the context of a virtual machine (VM), a piece of software that appears to be an application to the underlying (host) operating system (OS), but that looks like a bare machine, including CPU, storage, network adapters, and so forth, to the operating system and applications that run on it.

To use the Hortonworks Sandbox, one of the supported virtual machine applications needs to be installed on your host machine:

- VirtualBox
- VMware Fusion & Player
- Hyper-V

This document describes importing the Hortonworks Sandbox virtual appliance into VMware.

Prerequisites

To use the Hortonworks Sandbox with VMware Fusion the following requirements need to be met:

- ✓ **VMware Fusion installed**
Version 5 or later (Version 7 recommended)
You can download VMware Fusion here:
https://my.vmware.com/web/vmware/info/slug/desktop_end_user_computing/vmware_fusion/7_0#product_downloads
- ✓ **Host Operating Systems**
Host operating system refers to the operating system of your computer. The following link gives list of operating systems supported to run VMware Fusion
<https://www.vmware.com/support/fusion/faq/requirements>
- ✓ **Hardware** (The newer the hardware the better):
 - ✓ A 64-bit machine with a multi-core CPU that supports virtualization.
Please look into your operating system's documentation to verify if you are running a 64 bit OS.
Mac OS X:
<https://support.apple.com/en-us/HT3696>
 - ✓ **BIOS** - Has been enabled for virtualization support. Please contact your specific computer vendor to determine how to enable/verify this feature in your machine's BIOS.

- ✓ **RAM** - At least 4 GB of RAM (The more, the better)
If you wish to enable services such as Ambari, HBase, Storm, Kafka, or Spark please ensure you have at least 10 Gb of physical RAM in order to run the VM using 8 GB.

More information for hardware requirements for VMware can be found here <https://www.vmware.com/support/fusion/faq/requirements>

- ✓ **Browsers**
 - Chrome 25+,
 - IE 9+ (Sandbox will not run on IE 10)
 - Safari 6+

- ✓ **Hortonworks Sandbox virtual appliance for VMware Fusion**
Download the correct virtual appliance file for your environment from <http://hortonworks.com/products/hortonworks-sandbox/#install>
The file extension for a virtual appliance for VMware Fusion should be .ova

Procedure

The steps provided describe how to import the Hortonworks Sandbox virtual appliance into VMware Fusion. The screenshots displayed are taken from Mac OS X machine running the VMware Fusion 7.1.1 software.

NOTE: The instructions provided is for a VMware Fusion application that currently does not have any existing virtual machines imported.

1. Open the **VMware Fusion** application.
You can do so by double clicking the icon:



On a Mac OS X machine, by default the application is installed within the Applications folder, hence the icon to launch the application can be found there.

2. The **VMware Fusion** window opens and prompts

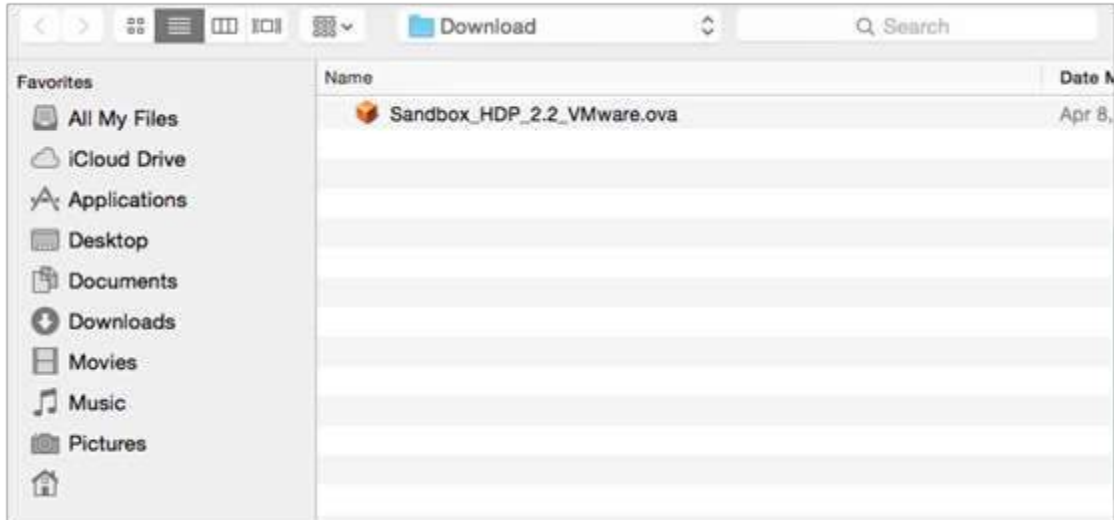


NOTE: If you currently have existing VM's installed you should get a different prompt from the above screenshot.

3. Select the "Import an existing virtual machine" and click the **Continue** button.



4. The file browser dialog opens. Select the virtual appliance file and click **Open**.




NOTE: The name of the file you have downloaded depends on the version of the Hortonworks Sandbox you have chosen to download. The above picture is referencing Sandbox HDP version 2.2

5. After clicking open, you will be brought back to this screen showing the file you have chosen highlighted. Click **Continue**.



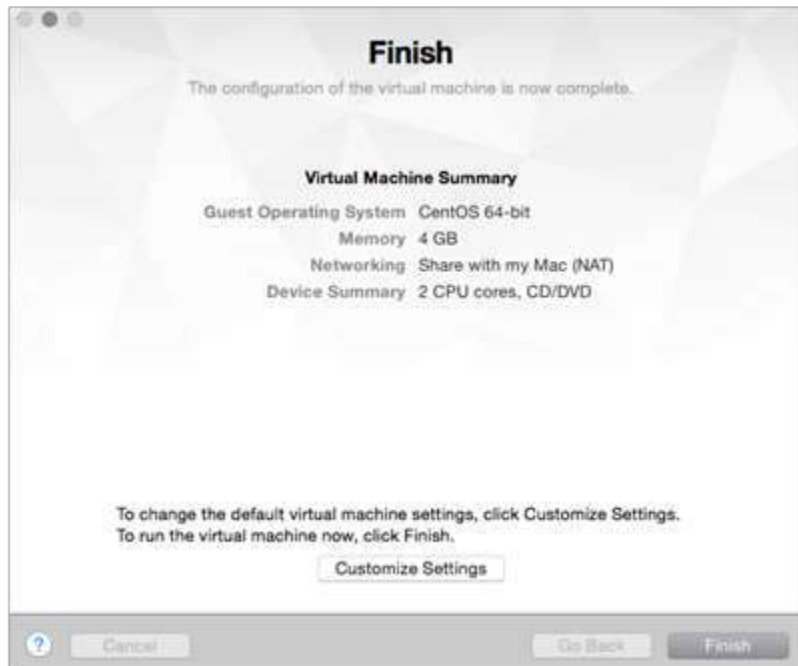
6. After clicking continue, a drop down window sheet displays.



Here you can choose to save the virtual machine with a different name. If you have no preference in renaming, you can just leave the default name and click save. You will get the importing progress dialog:



7. Once finished, the following screen is displayed:



Click the **Finish** button.

8. The next window is the console of the virtual machine, click on the **Play** icon within the console.



9. After clicking on the play button, you should see the console display boot up information.

```
Sandbox_HDP_2.2_VMware-
Updating IP...
Starting HDP ...
Starting mysql [ OK ]
Starting Knox ldap [ OK ]
Starting Postgre SQL [ OK ]
Starting secondary name node [ OK ]
Starting name node [ OK ]
Starting data node [ OK ]
Starting Oozie [ OK ]
Safe mode is OFF
Starting zookeeper nodes [ WARNINGS ]
JMX enabled by default
Using config: /etc/zookeeper/conf/zoo.cfg
Starting Hdfs nfs [ OK ]
Starting NFS portmap [ OK ]
Starting hive server [ OK ]
Starting Hiveserver2 [ OK ]
Starting yarn history server [ OK ]
Starting node manager [ OK ]
Starting mapred history server [ OK ]
Starting Falcon [ OK ]
Starting webhcat server [ OK ]
Starting Knox gateway [ OK ]
Starting resource manager [ OK ]
```

10. Eventually you will see the screen below indicating Hortonworks Sandbox is fully booted and displays login instructions.

```
Sandbox_HDP_2.2_VMware
Starting Virtual Printing daemon: done
Hortonworks Sandbox with HDP 2.2
http://hortonworks.com

To initiate your Hortonworks Sandbox session,
please open a browser and enter this address
in the browser's address field:
http://192.168.235.130/

You can access SSH by $ ssh root@192.168.235.130

Log in to this virtual machine: Linux/Windows <Alt+F5>, Mac OS X <Ctrl+Alt+F5>
```


NOTE: Because what is being displayed is a conceptually separate machine, control of the mouse and the keyboard must be passed back and forth between the host and the VM. This is particularly useful when the VM has a GUI. If you accidentally let the console “capture” your mouse or keyboard, you can release them back to the host machine by pressing the Ctrl and Command keys. A reminder appears in the upper right corner of the console window if you forget.

11. Use one of the supported browsers mentioned in the prerequisites section of this document within your host machine. Enter the URL displayed on the console. By default it should be <http://127.0.0.1:8888>. The registration form should appear:



The screenshot shows a registration page for the Hortonworks HDP 2.2 Sandbox. The page has a green header with the Hortonworks logo on the left and "HDP 2.2 Sandbox" on the right. Below the header, the text reads "Welcome to Sandbox!" followed by a paragraph explaining the sandbox's purpose. Below this is a registration form with the following fields: "First Name:", "Last Name:", "Email Address:", "Company Name:", and "Phone Number:". There is also a checkbox for "I agree to [Terms of Use](#)". A "Submit" button is located at the bottom right of the form.

NOTE: You just need to do this once. The next time you reboot Sandbox, you will not be prompted for this registration page.


12. Once you fill it out, clicking the submit button brings you to the Welcome page. That's it. Read over the information and links to get started in exploring HDP with the Hortonworks Sandbox!



Welcome to the Hortonworks Sandbox with HDP 2.2

Hortonworks Data Platform Version 2.2 represents yet [another major step forward](#) for Hadoop as the foundation of a Modern Data Architecture. This release incorporates the last six months of innovation across Apache Hadoop and its related projects.

Every component is updated and we have added some key technologies and capabilities to HDP 2.2. If you would like to experience our new features, try some of the "what's New" tutorials below.



Dive right in

Get started using Sandbox with your own datasets, and connect it to your existing tools and applications.

To use the web interface to explore HDP 2.2 Sandbox navigate to <http://192.168.235.133:8000>. The username and password is hue and 1111 respectively.

To SSH into the VM and explore HDP 2.2 Sandbox from the command line

```
ssh root@192.168.235.133 -p 2222;
```

The password is `hadoop`

Tutorials to get you started: