

MapR Sandbox for Hadoop

Overview

The MapR Sandbox for Hadoop is a fully-functional single-node cluster that gently introduces business analysts, current and aspiring Hadoop developers, and administrators (database, system, and Hadoop) to the big data promises of Hadoop and its ecosystem. Use the sandbox to experiment with Hadoop technologies using the MapR Control System (MCS) and Hue.

MapR's Sandbox for Hadoop includes a link to tutorials that can help you understand the power and capabilities of Hadoop through the MCS and Hue tools. Hadoop administrators can launch the MCS and use the tutorials to configure, monitor, and manage a cluster. Hadoop developers or analysts, looking to gain an understanding of Hadoop and MapR, can launch the Hue interface and use the tutorials to perform tasks using the applications included in Hue. Explore solutions to your use cases, and run jobs on your data in the MapR Sandbox for Hadoop.

To use the MapR Sandbox for Hadoop, perform the following tasks:

- Verify that the host system meets the prerequisites.
- Install the MapR Sandbox for Hadoop.
- Launch Hue or the MapR Control System.

Prerequisites

The MapR Sandbox for Hadoop runs on VMware Player and VirtualBox, free desktop applications that you can use to run a virtual machine on a Windows, Mac, or Linux PC. Before you install the MapR Sandbox for Hadoop, verify that the host system meets the following prerequisites:

- VMware Player or VirtualBox is installed.
- At least 20 GB free hard disk space, at least 4 physical cores, and 8 GB of RAM is available. Performance increases with more RAM and free hard disk space.
- Uses one of the following 64-bit x86 architectures:
 - A 1.3 GHz or faster AMD CPU with segment-limit support in long mode
 - A 1.3 GHz or faster Intel CPU with VT-x support
- If you have an Intel CPU with VT-x support, verify that VT-x support is enabled in the host system BIOS. The BIOS settings that must be enabled for VT-x support vary depending on the system vendor. See the VMware knowledge base article at <http://kb.vmware.com/kb/1003944> for information about how to determine if VT-x support is enabled.

VM Player Downloads

For Linux, Mac, or Windows, download the free [VMware Player](#) or [VirtualBox](#). Optionally, you can purchase [VMware Fusion](#) for Mac.

VM Player Installation

The following list provides links to the virtual machine installation instructions:

- To install the VMware Player, see the [VMware documentation](#). Use of VMware Player is subject to the VMware Player end user license terms. VMware does not provide support for VMware Player. For self-help resources, see the [VMware Player FAQ](#).
- To install VirtualBox, see the [Oracle VM VirtualBox User Manual](#). By downloading VirtualBox, you agree to the terms and conditions of the respective license.

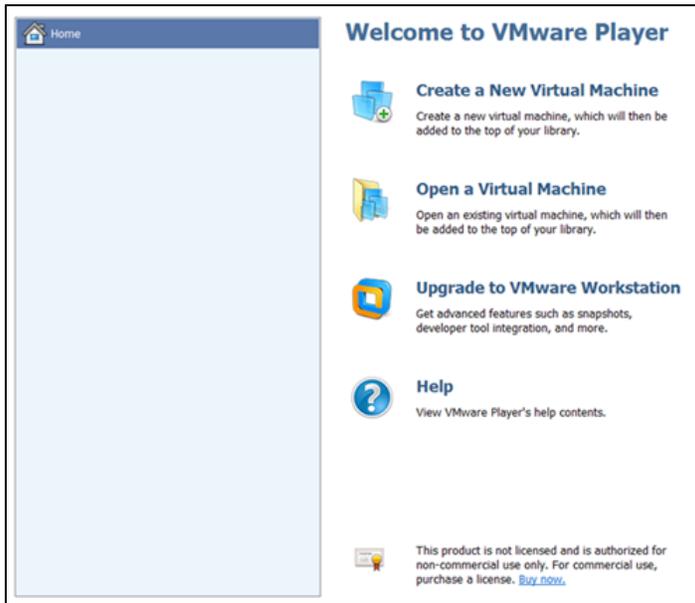
Install the MapR Sandbox for Hadoop

Download the MapR Sandbox for Hadoop file and then install the MapR Sandbox for Hadoop on your virtual machine player.

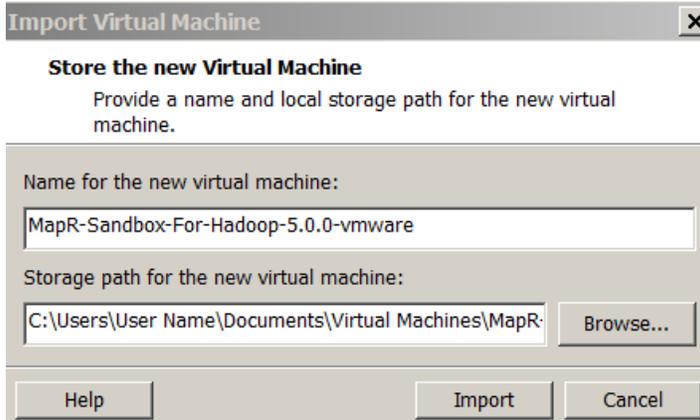
✓ [Installing the Sandbox on VMware Player/VMware Fusion](#)

Complete the following steps to install the MapR Sandbox for Hadoop on VMware Player or VMware Fusion:

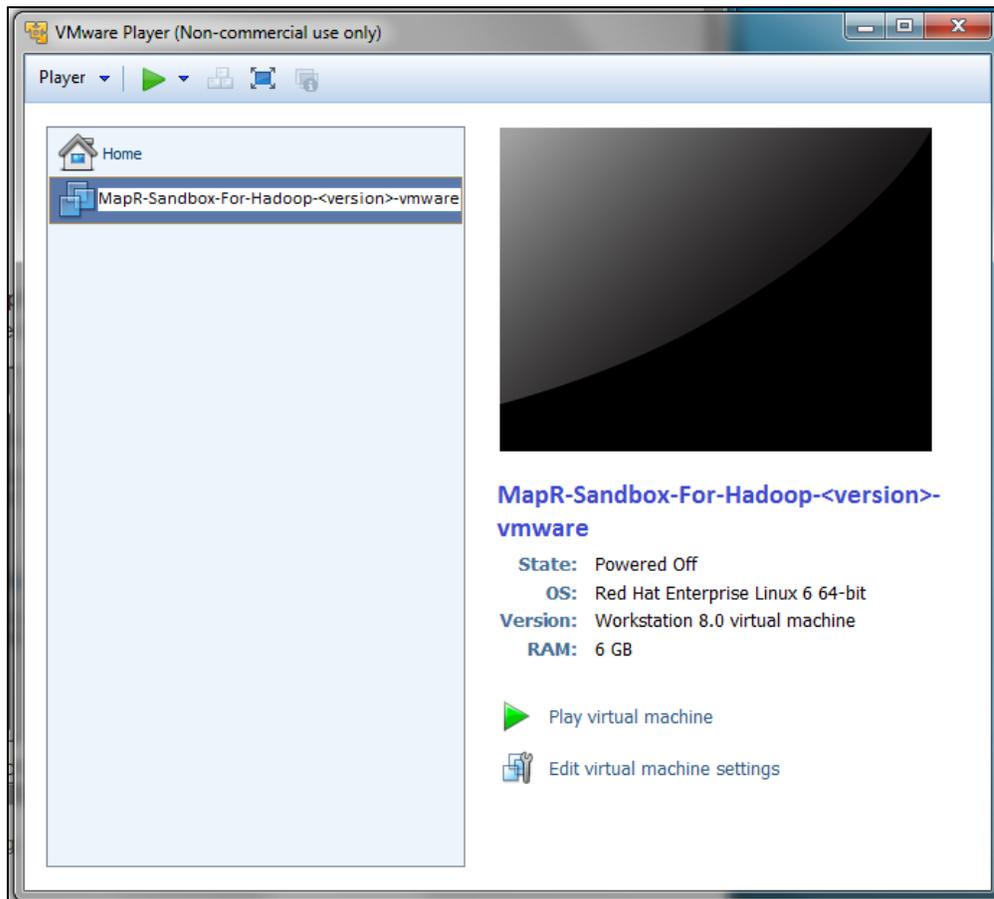
1. Download the MapR Sandbox for Hadoop file to a directory on your machine.
 - a. Navigate to <http://package.mapr.com/releases/>, and select the directory for the latest MapR release.
 - b. Select `sandbox > vmware > MapR-Sandbox-For-Hadoop-<release_version>-vmware.ova`
Example: `http://package.mapr.com/releases/v5.0.0/sandbox/vmware/MapR-Sandbox-For-Hadoop-5.0.0-vmware.ova`
2. Open the virtual machine player, and select the **Open a Virtual Machine** option.
 - ✓ [Tip for VMware Fusion](#)
If you are running VMware Fusion, select **Import**.



3. Navigate to the directory where you downloaded the MapR Sandbox for Hadoop file, and select `MapR-Sandbox-For-Hadoop-<version>-vmware.ova`. The Import Virtual Machine dialog appears.



4. Click **Import**. The virtual machine player imports the sandbox.



5. Select `MapR-Sandbox-For-Hadoop-<version>-vmware`, and click **Play virtual machine**. It takes a few minutes for the MapR services to start. After the MapR services start and installation completes, a screen that displays the following instructional message appears:

Message

MapR-Sandbox-For-Hadoop installation finished successfully.
 Please go to <http://192.168.16.128:8443> to begin your experience.
 Open a browser on your host machine and enter the URL in the browser's address field.
 You can access the host via SSH by `ssh mapr@192.168.16.128`
 The following credentials should be used for the MCS and HUE: `mapr/mapr`
 Log in to this virtual machine: Linux/Windows `<Alt+F2>`, Mac OS X `<Options+F5>`

6. Verify that a DNS entry was created on the host machine for the virtual machine. If not, create the entry.
 - ~ For Linux and Mac, create the entry in `/etc/hosts`.
 - ~ For Windows, create the entry in the `%WINDIR%\system32\drivers\etc\hosts` file.
 Example: `127.0.1.1 <vm_hostname>`
7. You can navigate to the URL provided in the message to experience Hue and the MapR Control System, or you can log in to the sandbox through the command line.
 - a. To navigate to the MapR Sandbox for Hadoop, enter the provided URL in your browser's address bar.
 - b. To log in to the virtual machine and access the command line, enter `Alt+F2` on Windows or `Option+F5` on Mac. When prompted, enter `mapr` as the login and password.

▼ Installing the Sandbox on VirtualBox

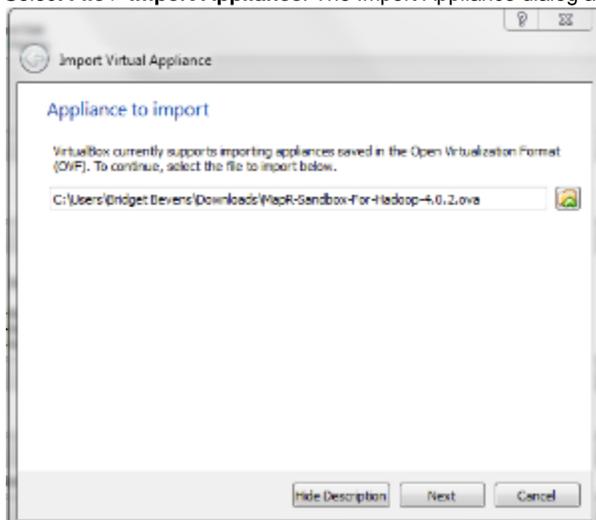
The MapR Sandbox for Hadoop on VirtualBox comes with NAT ports forwarded to `localhost`.

Complete the following steps to install the MapR Sandbox for Hadoop on VirtualBox:

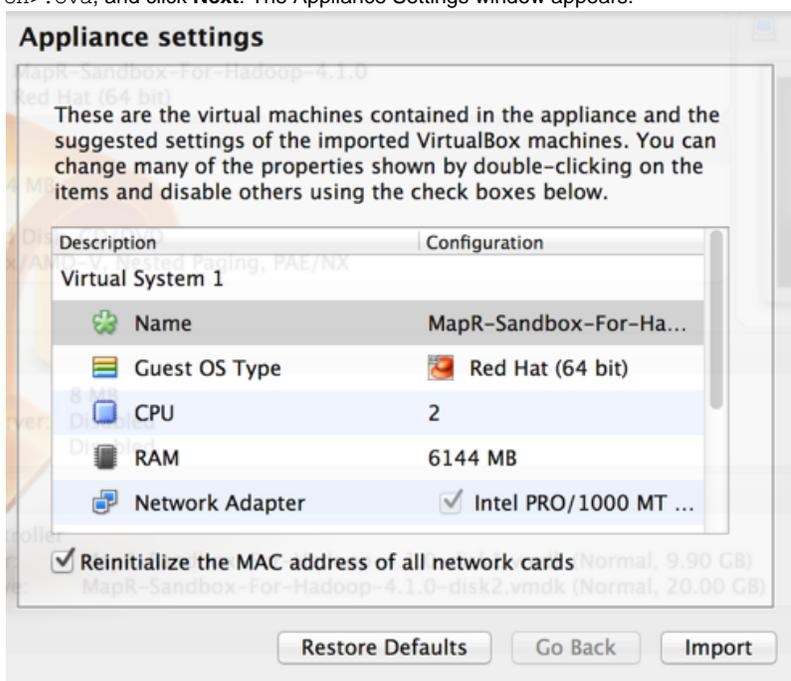
1. Download the MapR Sandbox for Hadoop file to a directory on your machine.
 - a. Navigate to <http://package.mapr.com/releases/>, and select the directory for the latest MapR release.

- b. Select `sandbox > ova > MapR-Sandbox-For-Hadoop-<release_version>.ova`
 Example: `http://package.mapr.com/releases/v5.0.0/sandbox/ova/MapR-Sandbox-For-Hadoop-5.0.0.ov`

2. Open the virtual machine player.
3. Select **File > Import Appliance**. The Import Appliance dialog appears.



4. Navigate to the directory where you downloaded the MapR Sandbox for Hadoop file, select `MapR-Sandbox-For-Hadoop-<version>.ova`, and click **Next**. The Appliance Settings window appears.



5. Select the check box at the bottom of the screen to **Reinitialize the MAC address of all network cards**, then click **Import**. The Import Appliance imports the sandbox.



6. Select the `MapR-Sandbox-For-Hadoop-<version>` in the virtual machine, and click **Start**. It takes a few minutes for the MapR services to start. After the MapR services start and installation completes, a screen that displays the following instructional message appears:

Message

MapR-Sandbox-For-Hadoop installation finished successfully.
Please go to <http://127.0.0.1:8443> to begin your experience.
Open a browser on your host machine and enter the URL in the browser's address field.
You can access the host via SSH by `ssh mapr@localhost -p 2222`
The following credentials should be used for the MCS and HUE: `mapr/mapr`
Log in to this virtual machine: Linux/Windows <Alt+F2>, Mac OS X <Options+F5>

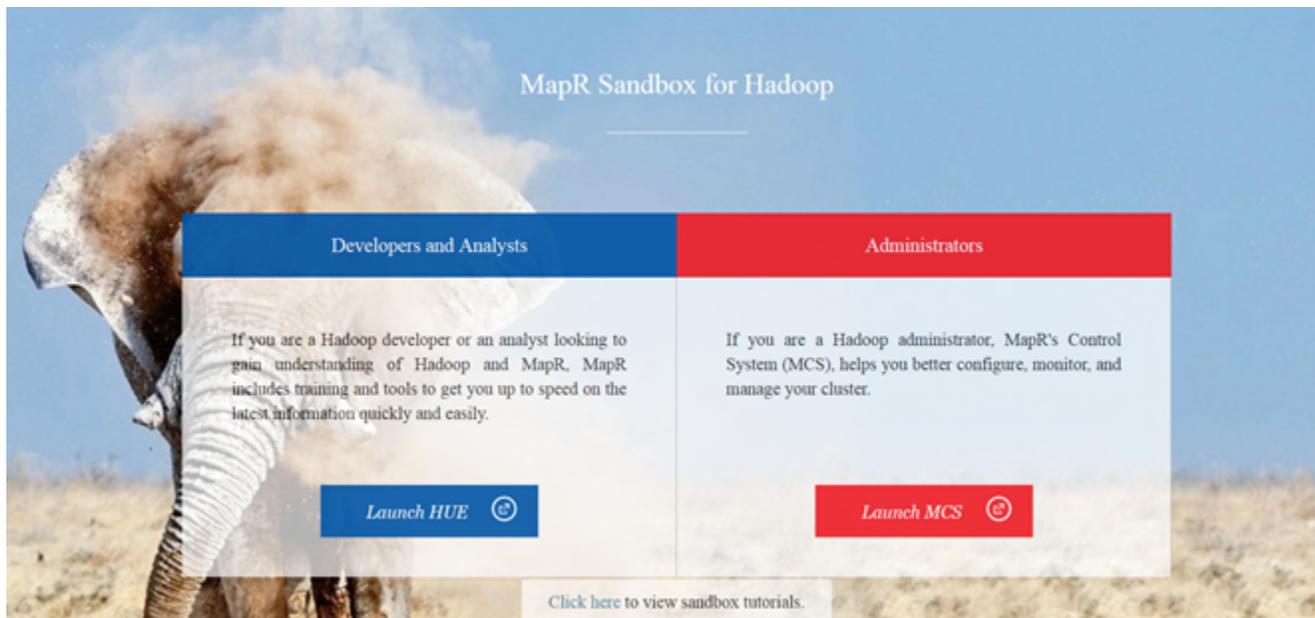


You can use the URL provided in the message or you can use <http://localhost:8443>.

7. Verify that a DNS entry was created on the host machine for the virtual machine. If not, create the entry.
 - ~ For Linux and Mac, create the entry in `/etc/hosts`.
 - ~ For Windows, create the entry in the `%WINDIR%\system32\drivers\etc\hosts` file.Example: `127.0.1.1 <vm_hostname>`
8. You can navigate to the URL to experience Hue and the MapR Control System, or you can log in to the sandbox through the command line.
 - a. To navigate to the MapR Sandbox for Hadoop, enter <http://localhost:8443> in your browser's address bar.
 - b. To log in to the virtual machine and access the command line, enter `Alt+F2` on Windows or `Option+F5` on Mac. When prompted, enter `mapr` as the login and password.

Launch Hue or the MCS

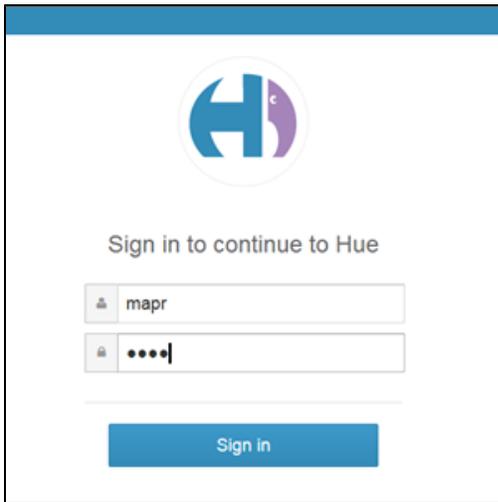
After you install the sandbox and open it in your virtual machine, you can navigate to the URL provided to access the following MapR Sandbox for Hadoop launch page:



The launch page includes launch buttons to access HUE and the MCS. The page also includes a link to [the sandbox tutorials](#).

To access Hue or the MapR Control System (MCS), complete the following steps:

1. Click a launch button. A login window appears.



2. Log in to the application with the username `mapr` and the password `mapr`.



If you launch the MCS, the system populates the username and password fields for you. By default, the MCS username is `root`. You can log in to the MCS as `root` or `mapr`. Or, you can just click **OK** to proceed. When prompted by the End User License Agreement, click **I Agree** to continue.

You have completed the steps required to use the MapR Sandbox for Hadoop. Have fun!