



---

## IBM InfoSphere BigInsights Quick Start Edition: VM Image README

Welcome to the IBM® InfoSphere® BigInsights™ Quick Start Edition, v2.1.2.

The purpose of the InfoSphere BigInsights Quick Start Edition is for experimenting with the features of InfoSphere BigInsights, while being able to use real data and run real applications in the least amount of time. The InfoSphere BigInsights Quick Start Edition puts no data limit on the cluster and there is no time limit on the license. This InfoSphere BigInsights Quick Start Edition uses a non-warranted program license, and is not for production use.

To read more about the InfoSphere BigInsights Quick Start Edition and the features that it includes, see the InfoSphere BigInsights documentation, IBM InfoSphere BigInsights Quick Start Edition.

### System requirements

Before you download, ensure that your system meets the minimum requirements:

#### Hardware requirements

- A minimum of 2 CPUs for the VM image, but 4 CPUs are recommended.
- A minimum of 5 GB of RAM for the VM image, but 8 GB of RAM is recommended.
- A minimum of 30 GB of free space on the hard disk.
- Supported operating systems:
  - Apple Mac- 64-bit
  - Linux - 64-bit
  - Windows 7 - 64-bit

#### Software requirements

- VMware Workstation 9.0+
- VMware Player 5.0 +
- Other virtual machine software that can support VMDK or OVF formats.

#### Virtual machine software configuration

Depending on the hardware capabilities of the host of the VM, you can allocate additional processors and memory for better performance.

For Windows operating systems, you might need to update the BIOS to enable virtualization.

---

## Which version of the Quick Start Edition VM image is the one for you

Before you download a Quick State Edition, determine which package might suit your needs. Remember, the features that you can experiment with are the same for both editions.

The Quick Start Edition VM packages are available in VMware formats (VMDK) and Open Virtualization Formats (OVF).

#### Single-node use

This version of the Quick Start Edition VM image is capable of running as a single-node cluster only. You cannot add additional nodes in this version. Select one of the following packages to download:

- iibi2120\_QuickStart\_Single\_VMware.7z
- iibi2120\_QuickStart\_Single\_OVF.7z

### Multi-node use

This version of the Quick Start Edition VM image is capable of being modified to run with multiple nodes. Select one of the following packages to download:

- iibi2120\_QuickStart\_Cluster\_VMware.7z
- iibi2120\_QuickStart\_Cluster\_OVF.7z

If you are processing large volumes of data, handling this data in a reasonable amount of time requires a distributed cluster that would work in parallel. The multi-node Quick Start VM image allows you to simulate this distributed cluster environment by allowing multiple nodes to be added to the Quick Start master node.

To learn more about how to add additional nodes to the cluster, see [Adding nodes](#).

## Downloading IBM InfoSphere BigInsights Quick Start Edition, v2.1.2 VM Image

The IBM InfoSphere BigInsights Quick Start Edition can be downloaded and run on any of the supported operating systems.

### Procedure

1. Download the **Quick Start Edition V2.1.2 (VM Image)** to your local file system.
2. Extract the appropriate file, as described in “Which version of the Quick Start Edition VM image is the one for you” on page 1. To extract the contents, use any valid extraction software, but **7-Zip** is recommended. Download **7-Zip** from <http://www.7-zip.org/download.html>.
3. When the extraction completes, make sure that you have the `.vmx` or the `.ovf` file.
4. If you downloaded the multi-node version, decide if you want to create multiple nodes with this image before you launch IBM InfoSphere BigInsights. See “Optional: Copying the files to prepare for adding nodes before starting IBM InfoSphere BigInsights” on page 7.

## Starting IBM InfoSphere BigInsights Quick Start Edition, v2.1.2 VM

Start the VM player and then start the InfoSphere BigInsights Quick Start Edition from the desktop icon of the image.

### About this task

Do these steps for each VM image or node that you created. In these steps, you will be asked for a user name and a password on two separate screens the first time you access the image. Use the following table as a reference:

*Table 1. User names and Passwords*

When to use	Username	Password
Log in from the command-line to accept the licenses	root	password
Log in from the Linux SUSE Desktop to access the BigInsights desktop	biadmin	biadmin

## Procedure

1. Open the image in the VM player of your choice.

Option	Description
Double-click the *.vmx or the *.ovf file.	It opens in the VM player.
Launch the VM player.	Select the VM image.

2. After the VM player completes its start-up procedures, you are requested to log in:


Option	Description
bivm login:	Type root and click <b>Enter</b> .

3. Then, you are prompted for the password: .

Option	Description
Password:	Type password and click <b>Enter</b>

4. The next several screens are to ensure that you understand the license agreements. Click **Enter** while the text, **I understand** is highlighted for each screen:
  - Linux Distribution Statement.
  - Novell SLES License Agreement.
  - VMware Tools License Agreement (if you are using the VMware player).
  - IBM InfoSphere BigInsights License Agreement.
  - IBM InfoSphere BigInsights Notice
5. The final License Agreement screen shows the acceptance for each of the licenses. To continue, click **Enter** while the text **I accept** is highlighted.
6. In the next screen, enter the **root** password: password, and then click **Enter** while **OK** is highlighted. The next screen is a confirmation of the password, Type password again and click **Enter** while **OK** is highlighted.
7. On the biadmin screen, enter the **password** for the **biadmin** account, which is biadmin. The next screen is a confirmation of the biadmin account password, Type biadmin again and click **Enter**.
8. In the next screen, type the user name for the **biadmin** account, which is biadmin. Then, click **Log in**.
9. In the next screen, type the password for the **biadmin** account, which is biadmin. Then, click **Log in**.
10. You now see the InfoSphere BigInsights Quick Start Edition desktop. Double-click the **Start**



**BigInsights** icon (  ) to start the BigInsights services.






11. When the script for starting the BigInsights services is complete , a message is displayed on the InfoSphere BigInsights desktop: Start BigInsights script completed. Click **OK** to close the message.
12. Now you are ready to explore the Quick Start Edition.
13. If you want to view any of the videos available on the VM image, install Adobe Flash updates.

---

## Using IBM InfoSphere BigInsights Quick Start Edition, V2.1.2

You can use the this InfoSphere BigInsights Quick Start Edition to explore the features of IBM InfoSphere BigInsights by using real data and running real applications.

## Procedure

1. Double-click the **Start BigInsights** desktop icon () to start the BigInsights services. You can double-click **Stop BigInsights** to stop the services.
2. Double-click the Web console desktop icon () to open the InfoSphere BigInsights Console. As an administrator, you can use the Console to inspect the overall health of the system and to complete basic functions such as starting and stopping specific servers and components. If you have the multi-node version of the VM image, you can add nodes to the cluster. As a user, you can interact with files in the distributed file system and manage and run applications.
3. From the BigInsights desktop, you can learn about BigInsights by double-clicking the **BigInsights Tutorials** icon () and performing the tutorials.
4. You can develop applications with the Eclipse client, by double-clicking the Eclipse icon ().  
**Note:** The Eclipse client that is part of this VM image is already enabled for InfoSphere BigInsights. To create BigInsights project in Eclipse, double-click the Eclipse icon to open the InfoSphere BigInsights Tools for Eclipse. This client tool comes pre-loaded with some sample projects that you can use to validate your work in the tutorials. In the Eclipse Project Explorer, right-click any of the sample projects and select **Open Project** to see the contents.
5. To work with the Big R component, double-click the **Install Big R** icon (), which begins running a script that installs the Big R components. Then, to use Big R with R, you must download an R client.
6. To see the complete BigInsights documentation, double-click the **Knowledge Center** icon.
7. Access the **BigInsights Forum**, and the **BigInsights Video** by double-clicking their respective desktop icons.
8. Open the **BigInsights Shell** folder to access these functions:
  - Big SQL Shell
  - HBase Shell
  - Hive Shell
  - Jaql Shell
  - Pig Shell
  - Terminal window
9. If you have problems while using the VM image, double-click the **Reset VM** icon on your BigInsights desktop. For other issues, double-click the **BigInsights Forum** icon.

---

## Related tasks

The following tasks are for your reference. Depending on the environment on which you install the IBM InfoSphere BigInsights Quick Start Edition, and the scenarios that you want to run, you might find some of these tasks necessary.

- “Installing Adobe Flash” on page 5
- “Optional: Accessing the InfoSphere BigInsights Console from a host machine” on page 5
- “Optional: Updating the default VM configuration” on page 6

- “Optional: Updating the Windows BIOS” on page 6
- “Optional: Copying the files to prepare for adding nodes before starting IBM InfoSphere BigInsights” on page 7
- “Optional: Configuring your additional nodes after starting IBM InfoSphere BigInsights” on page 7

## Installing Adobe Flash

If you want to view any of the videos that are included on the desktop of the VM image, you must have an ability to play these videos. If you do not have Adobe Flash, download and install the latest updates for Adobe Flash:

### Procedure

1. Double-Click the BigInsights WebConsole icon to open a browser to the web console. Then add another tab to open a new browser window.
2. From your VM image browser, access the Flash update site: [http://support.mozilla.org/en-US/kb/install-flash-plugin-view-videos-animations-games#w\\_installing-the-flash-plugin-manually](http://support.mozilla.org/en-US/kb/install-flash-plugin-view-videos-animations-games#w_installing-the-flash-plugin-manually). This site contains instructions for installing the software. It also contains a troubleshooting link. The troubleshooting link contains information that helps you to configure Adobe Flash for your specific environment.
3. Follow the instructions on the page. When prompted for your super user (sudo user) password, type password.

## Optional: Accessing the InfoSphere BigInsights Console from a host machine

From the VM image, you can access the InfoSphere BigInsights Console from your own host machine if necessary.

### Procedure

You should never modify the host name properties of the VM image. Instead, if you need to access the console from your own host machine, do the following steps:

Option	Description
Windows 7 - 64-bit environments	<ol style="list-style-type: none"> <li>1. In the VM image, open a Linux command line and type <code>su root</code>. The password is password.</li> <li>2. Run the following command to determine the IP address: <code>/sbin/ifconfig</code></li> <li>3. Edit the <code>c:\Windows\System32\drivers\etc\hosts</code> file.</li> <li>4. Add a new line with the following information: <code>&lt;ip address from step 2&gt; hostname</code></li> </ol>

Option	Description
<b>Apple Mac- 64-bit environments</b>	<ol style="list-style-type: none"> <li>1. In the VM image, open a Linux command line and type <code>su root</code>. The password is <code>password</code>.</li> <li>2. Run the following command to determine the IP address: <code>/sbin/ifconfig</code></li> <li>3. From a Linux command line, type the following command as the root user: <code>sudo nano /private/etc/hosts</code>  If the root password is requested, type <code>password</code>.</li> <li>4. Add a new line with the following information: <code>&lt;ip address from step 2&gt; hostname</code></li> </ol>
<b>Linux environments</b>	<ol style="list-style-type: none"> <li>1. In the VM image, open a Linux command line and type <code>su root</code>. The password is <code>password</code>.</li> <li>2. Run the following command to determine the IP address: <code>/sbin/ifconfig</code></li> <li>3. From a Linux command line, type the following command as the root user: <code>vi /etc/hosts</code>  If the root password is requested, type <code>password</code>.</li> <li>4. Add a new line with the following information: <code>&lt;ip address from step 2&gt; hostname</code></li> </ol>

## Optional: Updating the default VM configuration

You can change the settings in the VM software package before you start the InfoSphere BigInsights VM.

### About this task

In some VM Players, you might need to do further configuration than what is provided with the default image. Click the **Edit virtual machine** setting or the appropriate interface controls in the VM Player that you use.

If you change the settings, after you have started InfoSphere BigInsights, follow these steps:

### Procedure

1. Double-click the Terminal Shell icon. Open a terminal window.
2. Change the directory in the VM command-line to `$BIGINSIGHTS_HOME/bin`
3. Run **`syncconf.sh`**, so that InfoSphere BigInsights recognizes your new configuration
4. Click the **Stop BigInsights** icon and then click **Start BigInsights**
5. Type `exit` to leave the terminal screen.

## Optional: Updating the Windows BIOS

For Windows operating systems, you might need to update the BIOS to enable virtualization.

## About this task

Read the BIOS configuration documentation for your particular operating system. The following are general guidelines:

### Procedure

1. Shutdown, then reboot your system.
2. Access the BIOS of your system before it completely reboots.
3. In the Configuration section, click CPU.
4. Enable both of the virtualization settings: Virtualization Technology and VT-d.
5. Save the setting.
6. In the Configuration section, click Intel AMT and make sure that it is also enabled. The Console Type should be VT100+.
7. Save and exit. The system continues to reboot.
8. Then, Shut down, and Restart.

## Optional: Copying the files to prepare for adding nodes before starting IBM InfoSphere BigInsights

After you extract the `iibi2120_QuickStart_VM_Cluster` file, decide if you want to prepare your VM image BigInsights environment for multiple nodes.

### Procedure

1. For each additional IBM InfoSphere BigInsights Quick Start Edition VM node that you want to use, copy the `iibi2120_QuickStart_Cluster` folder into another destination on your current machine, or onto another machine.

**Tip:** You might want to make only one copy or one additional node at this point, because in a later step you are going to clean the second node to make it smaller. The *cleaned* node is the node that you should copy from for planning additional nodes.

Each copy represents a separate node on this cluster. The additional nodes can be on the same physical machine or other machines.

2. Complete the tasks to start BigInsights, and then configure your additional nodes.

## Optional: Configuring your additional nodes after starting IBM InfoSphere BigInsights

You can add multiple nodes to your IBM InfoSphere BigInsights Quick Start Edition, to explore a production level type of environment.

## About this task

You now have one node, or VM image, and one copy of the node or VM image. Each node has been started and you see the InfoSphere BigInsights desktop on each node.

For convenience in referencing each node, refer to the original `BIQuickStartVM2.1.2.Cluster` as **MasterNode**. Refer to the copied VM image as **ChildNode**.

To learn more about how to add additional nodes to the cluster, see Adding nodes.

## Procedure

1. On the **ChildNode** VM image that you made in “Optional: Copying the files to prepare for adding nodes before starting IBM InfoSphere BigInsights” on page 7, log in as the **biadmin** user and complete the following steps from the desktop of the VM image of **ChildNode**:
  - a. Double-click the **Clean Local BI** icon.

This operation uninstalls the existing InfoSphere BigInsights on this node and resets the VM image so that you can use it later to add as a Data Node in the cluster. It also removes the **Start BigInsights, Stop BigInsights, BigInsights WebConsole, BigInsights Shell, and the Clean Local BI** icons.
  - b. If you are planning to create additional nodes, copy the resulting `<copied_node>\BIQuickStartVM2.1.2.Cluster` folder that contains your **ChildNode** to other destinations. This VM image is reduced and therefore easier to copy.
2. On the **MasterNode**, or management VM image node, double-click the **Start BigInsights** icon.

A Terminal window opens on the desktop that shows you the progress of starting BigInsights. When it is completed, you are returned to the desktop.
3. As a prerequisite for adding new nodes in the cluster, all nodes in the cluster should be able to successfully communicate with each other by hostname. You can do that by modifying `/etc/hosts` (run the command `sudo vi /etc/hosts` on the command line) on each node of the cluster to make it contain the host name and IP address information of all of the nodes. The hostname of the master node cannot be changed. It should always be *bivm*. The hostname of the child node must be different from the hostname of the master node. You can get the IP address of the node by entering `ipconfig` on the command line.

For example, if you have three nodes, one is the master node *bivm*, the other two are child nodes. Through running the `ipconfig` command, you know these three nodes respectively have IP addresses of "192.168.0.100", "192.168.0.101" and "192.168.0.102". If you want to name the child nodes `binode1` and `binode2`, then the `etc/hosts` file on the management node and the child nodes might look like the following list:

```
191.168.0.100 bivm
192.168.0.101 binode1
192.168.0.102 binode2
```

To quickly check if it is ready, you can ping them through the host name by each other. For example, log on `binode1`, run command `ping binode2` and `ping bivm`. If they both succeed, then it is working.
4. On the **MasterNode**, double-click the **InfoSphere BigInsights Web Console** icon, and make sure that you are in the **Welcome** page.
  - a. From the Tasks pane of the Welcome page, click **Add or remove a node**.
  - b. Click **Add Nodes**. In the Add Nodes window, select the service that you want to add and the IP address of the node on which you want to start it. Complete the remainder of the dialog and click **Save**.