

Virtual Machine Guide Advanced Databases

2015/2016

This document describes the virtual machine running Centos 6.4. that is intended for students enrolled at the Advanced databases course. The machine contains software required for the purposes of trying out exercises published alongside lectures and implementing the course's project assignments. Should any project require additional software, it will be discussed in the instructions for that particular project.

CentOS 6.4 x86 minimal with no graphical interface is installed on the virtual machine and set to 512MB of RAM.

1. VirtualBox and virtual machine setup

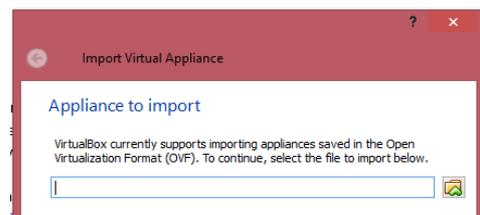
To start (use) the virtual machine, you have to install VirtualBox , which can be downloaded from: <https://www.virtualbox.org/wiki/Downloads>.

After successful installation, you need to run VirtualBox and select File- > Import Appliance

Then choose the file:

kent.zpr.fer.hr/CentOS-6.4-i386-minimal_15_16.ova (2,4 GB)
acquired from:

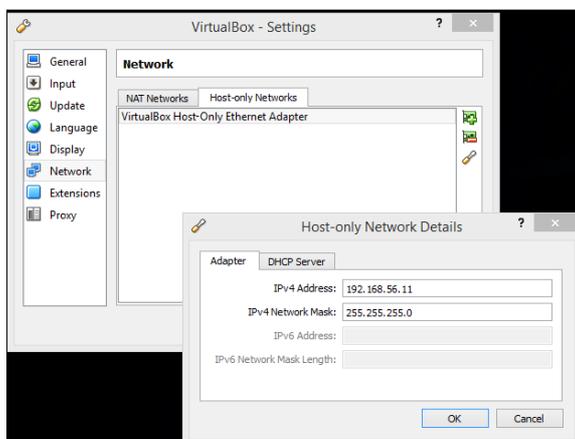
kent.zpr.fer.hr/CentOS-6.4-i386-minimal_15_16.ova.zip



Note that this creates another working copy of the virtual machine in the VBOX format, which is probably stored at: C:\Users\USERNAME\VirtualBox VMs. In other words all subsequent changes on the virtual machine are applied to the working copy, not the initial .ova file.

Once the VM is loaded, set:

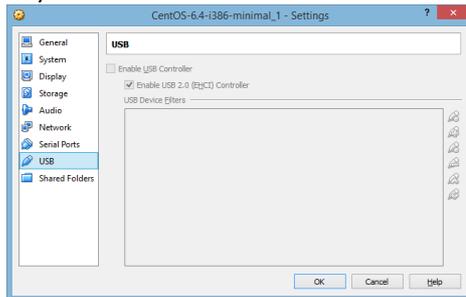
- a) File->Preferences->Network, tab Host-Only Networks, Edit



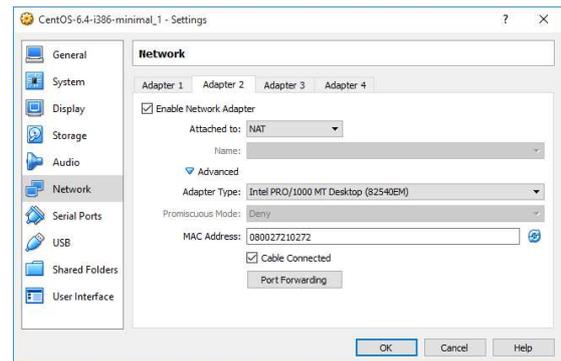
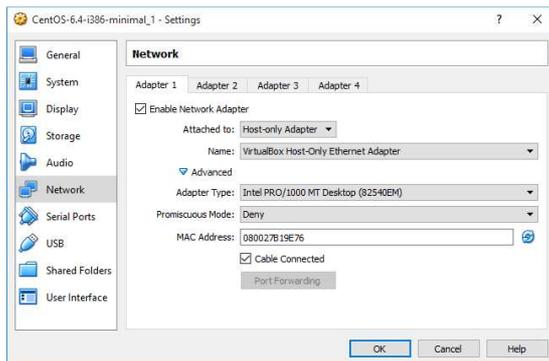
Also, make sure that DHCP is turned off (adjacent tab) .

b) Then select the properties of the virtual machine (right the click > Settings), and :

b1) Disable USB:



b2) Make sure that the Network is set as shown below (for Adapter 1 and Adapter 2):



Leave MAC addresses as on pictures above.
This is done only once.

Start the virtual machine.



2. Connecting to the VM

Virtual PC is set to fixed local IP address **192.168.56.12** and is configured so that it is **not possible** to access an outside network, but only the host computer (your computer).

The computer has the following user accounts:

Username/password = **root/reverse**

Username/password = **posgres/reverse**

Although the virtual machine can be used via in the VirtualBox window, we recommend that you work using a terminal emulator e.g. Putty. Putty can be downloaded from:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>

To connect via Putty to 192.168.56.12 use the given root user credentials.
It is recommended to save the connection settings for future sessions.

If you should have to enable the outer network connection, e.g. to download and install some additional software on the VM, you can enable the connection to the Internet following these steps:

- entering the following command from a command prompt start the network interface `eth1` (previously configured as Adapter 2 that runs in NAT mode):

```
ifup eth1
```

- make sure that the connection to the Internet through the interface eth1 is successfully enabled:

```

CentOS-6.4-i386-minimal_1 [Running] - Oracle VM VirtualBox
File Machine View Input Devices Help
root@localhost ~# ifup eth1
Determining IP information for eth1... done.
root@localhost ~#
root@localhost ~#
root@localhost ~# ping www.fer.hr
PING www.fer.hr (161.53.72.119) 56(84) bytes of data:
64 bytes from skynet.cc.fer.hr (161.53.72.119): icmp_seq=1 ttl=62 time=1.07 ms
64 bytes from skynet.cc.fer.hr (161.53.72.119): icmp_seq=2 ttl=62 time=2.63 ms
64 bytes from skynet.cc.fer.hr (161.53.72.119): icmp_seq=3 ttl=62 time=2.23 ms
64 bytes from skynet.cc.fer.hr (161.53.72.119): icmp_seq=4 ttl=62 time=1.63 ms
^C
--- www.fer.hr ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3494ms
rtt min/avg/max/mdev = 1.070/1.891/2.633/0.595 ms
root@localhost ~# _

```

- If you still can not access the Internet, try the same procedure but in the settings of the virtual machine Adapter 2 set to Bridged Adapter instead of NAT (Figure 1b)

Use the ifconfig command to find out the IP address assigned to the eth1 interface. Connecting to address 192.168.56.12 can still be as before (via existing interfaces eth0).

When the connection to the Internet is no longer required, you can switch back to the host-only mode by disabling network adapter eth1. Enter the following command from a command prompt:

```
ifdown eth1
```

3. Installed software

- **PostgreSQL 9.4 + PostGIS 2.93 + temporal extension**
- **Development Tools, Development Libraries**
- **MongoDB 3.0.6.**
- **Apache Spark 1.5.1**
- **Riak + Erlang**
- **Jena Fuseki Server**
- **Fuseki Server is a SPARQL HTTP server that enables the SPARQL queries over the RDF graph data. This installation already contains the Linked Movie database RDF graph.**

More information at: http://jena.apache.org/documentation/serving_data/

Fuseki is installed in /usr/fuseki-server/ folder the and **is not started** by default.

To start it, you have to give the following commands:

```
cd /usr/fuseki-server/
./fuseki-server --loc=LMDB/ /LMDB
```

or, if you also want to enable the SPARQL UPDATE commands:

```
./fuseki-server --update --loc=LMDB/ /LMDB
```

After that, when the Fuseki is running, you can access it via: http://IP_ADDRESS:3030/

- **Apache+PHP**
Apache+PHP is not required in any of the projects, it is installed for you convenience.
- **FTP**
FTP is enabled to facilitate the file exchange between the host and guest computer. Also, if you find Vi editor on the guest machine difficult to use, you can bypass it by using some other more friendly text editor on the host computer that can save files over the FTP, such as Notepad++, UltraEdit, Sublime Text,
- **locate command**
locate command searches the entire disk and locates the assigned file.

For instance, try this:

```
> locate test.php
```

locate finds (returns):

```
/var/www/html/test.php
```

You may try opening that file from the host computer to see what it does:

<http://192.168.56.12/test.php>

Should you change/add new files, you must run the following command to update the locate command's database, in order to be able to find those new files:

```
> locate /etc/cron.daily/mlocate.cron
```

- **curl**

For more advanced curl commands you may opt to use Linux's curl implementation as it is better than Window's curl. For instance, you can use the multi-line command mode (if you end a line with backslash, more about this in the NoSQL project assignment). Also, if you finish the line with the "@-" characters, curl will continue to read the console input stream until you issue CTRL+D.