

Instituto Superior Técnico, Universidade de Lisboa  
**Network and Computer Security**

**Assignment 2**  
**Implementation and Analysis of**  
**a Virtual Computer Network**

VirtualBox version. Revised on 2016-10-12

**Alpha version:** This is an early version and may contain some incorrect features.

## Goals

- Implement and test a virtual computer network.
- Perform a simple TCP/IP packet analysis.

## 1. Introduction

The laboratory assignment uses a virtual network consisting of multiple virtual machines supported by Oracle's VirtualBox virtualization software. VirtualBox works with Windows, Mac and Linux operating systems, although specific instructions may vary slightly.

The physical machine is called the **Host** whereas the virtual machines are called **Guests**. The host will be the computer running VirtualBox. Each guest will use a version of Caixa Mágica Linux with low memory requirements.

You can use the manual command to read the documentation about the Linux commands:

```
$ man ln
```

### 1.1 *Create a virtual machine instance*

Each virtual machine instance will have its own virtual disk.

The virtual disk image can be obtained at: <https://ftp.rnl.tecnico.ulisboa.pt/classes/sirs/vm-SIRS2016-17.vdi.zip>

Create a virtual machine named **vm1** by executing the following steps:

1. Open Virtual Box
2. Create a Virtual Machine:
  - **Machine → New.**
  - Name: vm1
  - Operating system: Linux, version: Linux Red Hat (64 bits)
  - Use default value for memory (512 MB)
  - **Do not add a virtual drive.**
3. Add the virtual drive as IDE:
  - **Machine → Settings → Storage.**
  - Add an IDE Hard Disk
  - Choose existing disk. Add the **.vdi** file downloaded before
4. Test the Virtual Machine:
  - Start. Log in with the user *root* and the password "*Ins3cur3*"
  - If the GUI does not appear, it is necessary to edit the `/etc/inittab` file:
    - i. `vi /etc/inittab`
    - ii. Make sure the last line is like this: `id:5:initdefault:`
    - iii. Finally: `sudo /sbin/shutdown -r now`

By default you can use **<Right CTRL> or <⌘>** to release the mouse pointer from the guest back to the host.

## **1.2 Copy the virtual machine instance**

We will now create a total of 3 virtual machine instances: 1, 2 and 3.

To create a copy, use the Machine - Clone command:

1. Power-off virtual machine 1 (**vm1**)
2. VirtualBox menu: select virtual machine 1
3. Select Machine -> Clone...
4. Change name to vm2. Select "Reinitialize the MAC Address of all network cards", so that the interface MAC address differs on different virtual machines.
5. Select Full Clone that will create several virtual disks. Linked clone also would work but the cloned machine is created from a snapshot of machine 1, leaving a file dependency.

6. Click on Clone

Apply the same procedure to create vm3.

### **1.3 Create the virtual network**

We will now configure the network settings of each Guest.

**VM2** will access both the **physical network** and the virtual network.

- To configure the access to the physical network, go to Network Adapters and Other Hardware configuration:

1. VirtualBox menu: select vm2 -> Settings... -> Network
2. Network menu: select **Adapter 1**. Click on “Enable Network Adapter” to enable the adapter, if the checkbox is not already checked. **Attached to: “Bridged Adapter”**.
3. Name: Select the default network card: on Windows (“Intel(R) Pro...”, “RealTek...”, “Wifi”, “Ethernet”, etc); on Linux (eth0, wlan0, etc); on MacOS (en0: Wi-Fi (AirPort)) , so that the virtual machine virtual is added to the host virtual network.
4. Expand **Advanced**
  - a. Adapter Type: Intel PRO/1000 MT Desktop (82540EM)
  - a. Promiscuous Mode: Allow All
  - b. MAC Address: use default value
  - c. Verify that “cable connected” is enabled.

**Note: VM1 and VM2 will be connected to the same physical network, so you must follow steps 1-4 above for VM1 as well.**

- To configure the **virtual network**:

2. VirtualBox menu: select vm2 -> Settings... -> Network
3. Network menu: select **Adapter 2**. Click on “Enable Network Adapter” to enable the adapter, if the checkbox is not already checked. **Attached to: “Internal Network”**.
4. Name: Select the default network card: on Windows (“Intel(R) Pro...”, “RealTek...”, “Wifi”, “Ethernet”, etc); on Linux (eth0, wlan0, etc); on MacOS (en0: Wi-Fi (AirPort)), so that the virtual machine virtual is added to the host virtual network.

**Note: if the only option is 'intnet', choose that option.**

## 5. Expand **Advanced**

- Adapter Type: Intel PRO/1000 MT Desktop (82540EM)
- Promiscuous Mode: Allow All
- MAC Address: use default value
- Verify that “cable connected” is enabled.

**Note:** VM2 and VM3 will be connected to the same virtual network, so you must follow steps 1-4 above for VM3, but use Adapter 1 for VM3.

### 1.4 **Configure the IP network**

The physical and virtual networks will be different IP subnets.

For the virtual network, we will use the private IP addresses 192.168.1.0/24 (meaning that the subnet mask is 255.255.255.0 – we can have 254 addresses to use - from 192.168.1.1 to 192.168.1.254 – 192.168.1.255 is reserved for broadcast).

The following diagram presents the desired network layout. IP addresses may vary depending on the host network settings and the virtual network settings.




Diagram 1 – Intended network layout. VM1 and 2 are connected in the physical network. VM 2 and 3 are connected in a virtual network. VM2 will operate as a gateway between the two subnets.

On **VM1**, delete any unnecessary `ifcfg-` files in the `/etc/sysconfig/network-scripts/` folder:

```
$ sudo rm /etc/sysconfig/network-scripts/ifcfg-eth-*
```

Then edit the file `/etc/udev/rules.d/70-persistent-net.rules` and change the MAC address in this file by the MAC Address of VM1 (you can consult this in Virtual Box -> Network -> Adapter 1 -> Advanced). Then restart VM1.

Afterwards, create a network interface configuration file named `/etc/sysconfig/network-scripts/ifcfg-eth0` (or `ifcfg-eth-id-` followed by the MAC address):

```
$ sudo touch /etc/sysconfig/network-scripts/ifcfg-eth0
```

Note: since a lot of the files edited are located in the `/etc/sysconfig/network-scripts/` directory, it is advised to navigate to it to make command execution easier:

```
$ cd /etc/sysconfig/network-scripts/
```

Using a basic editor – like `vi` or `mousepad` – set the *IP* and *Netmask* values (you may use the following **example** contents)

```
$ sudo vi /etc/sysconfig/network-scripts/ifcfg-eth0
```

Press the 'i' key on the keyboard to enter editing mode in **vi** and write the following content:

```
BOOTPROTO=none
ONBOOT=yes
DEVICE=eth0
IPADDR=192.168.0.100
NETMASK=255.255.255.0
GATEWAY=192.168.0.10
NM_CONTROLLED=yes
HWADDR=
```

When you finish writing, press 'Esc' to exit editing mode and then write the following command and press enter (this is done still inside **vi**):

```
:wq
```

Create the file `/etc/sysconfig/network-scripts/route-eth0` to define a default gateway address for interface `eth0`. For **example**, write the following lines to this file:

```
default 192.168.0.10 dev eth0
```

### **Routes file format details**

The files: `/etc/sysconfig/network-scripts/route-<interface>` define routing tables. Each line is a routing rule defined as "column1 column2 column3" and finally the interface name.

- **Column 1** is used to provide a **destination**, which may be a network address, a host address or the string `default` to specify the default gateway. Whether the given address is a host or a network is determined by the netmask with the notation `IP_ADDRESS/NETMASK_LENGTH`. Network route means a route to a network, whereas a host route is a route to a single host;
- **Column 2** is either the **IP address of the router** which should be used to reach the destination detailed in the previous column, or the special value `default`, which means that all traffic headed to the destination will be given to the device specified in the third column;
- **Column 3** contains device responsible for that route.

For example, in the following rule, the netmask is indicated using the notation `IP/NetmaskLength`:

```
192.168.0.0/24 via 172.16.1.3 dev eth0
```

Notice that the route to the default gateway does not include the keyword `via`, for example:

```
default 192.168.0.10 dev eth0
```

More information about routing tables can be consulted in the [manual of CentOS](#).

Reload the network interfaces:

```
$ service network restart
```

To check that the configuration is correct:

```
$ ifconfig
```

```
$ ip route
```

Apply the same procedure to configure **VM2** with IP 192.168.0.10, as indicated in **Diagram 1**. To do this, in **VM2** delete any unnecessary `ifcfg-` files in the `/etc/sysconfig/network-scripts/` folder:

```
$ sudo rm /etc/sysconfig/network-scripts/ifcfg-eth-id-*
```

Afterwards, create a network interface configuration file named `/etc/sysconfig/network-scripts/ifcfg-eth0` (or `ifcfg-eth-id-` followed by the mac address):

```
$ sudo touch /etc/sysconfig/network-scripts/ifcfg-eth0
```

Using a basic editor – like `vi` - set the *IP* and *Netmask* values (you may use the following example contents)

```
$ sudo vi /etc/sysconfig/network-scripts/ifcfg-eth0
```

Write the following content and save the changes:

```
BOOTPROTO=none
ONBOOT=yes
DEVICE=eth0
IPADDR=192.168.0.10
NETMASK=255.255.255.0
GATEWAY=192.168.0.1
NM_CONTROLLED=yes
HWADDR=
```

Change the file `/etc/sysconfig/network-scripts/route-eth0` to define a default gateway address for interface *eth0*. For example, change the file contents to the following:

```
default 192.168.0.1 dev eth0
```

Reload the network interfaces:

```
$ service network restart
```

To check that the configuration is correct:

```
$ ifconfig
$ ip route
```

**VM2** should now be able to ping **VM1** (and vice-versa):

```
$ ping 192.168.0.100
```

Configure **VM3**, with IP `192.168.1.1` and a rule in the `/etc/sysconfig/network-scripts/route-eth0` file to use **VM2** as default gateway.

Configure **VM2** with IP `192.168.1.254` on `eth1`.

*Do not forget that VM2 will have two `ifcfg-ethX` files: one for `eth0` (which already exists at this point) and another for `eth1`.*

Define routing rules in `/etc/sysconfig/network-scripts/route-eth0` for **VM2**:

```
default 192.168.0.0/24 dev eth0
192.168.1.0/24 via 192.168.0.100 dev eth0
```

Define routing rules in `/etc/sysconfig/network-scripts/route-eth1` for **VM2**:

```
192.168.1.0/24 via 192.168.1.1 dev eth1
```

After reloading the network configurations, **VM2** should now be able to ping **VM3**:

```
$ ping 192.168.1.1
```

**VM3** should now be able to ping **VM2**:

```
$ ping 192.168.1.254
```

## 1.5 Configure a gateway

Since **VM2** will be the default gateway for **VM3**, *IP* forwarding must be enabled. This will allow **VM3** to communicate with machines outside the subnet 192.168.1.X.

In **VM2** open the file `/etc/sysctl.conf` and change/add the following line to:

```
net.ipv4.ip_forward = 1
```

Restart **VM2**.

Confirm you will get **1** by executing the following command:

```
/sbin/sysctl net.ipv4.ip_forward
```

If you get **0** the execute the following:

```
/sbin/sysctl -w net.ipv4.ip_forward=1
```

## 1.6 Configure NAT (Network Address Translation)

During the previous step, the connectivity test between **VM3** with **VM1** failed. Since NAT was not enabled in **VM2**, the host sent the reply to its gateway, which eventually dropped the packet. Use the `iptables` command (`man iptables`) in **VM2** to correct this behaviour. NAT will do the source and destination mapping.

```
$ sudo iptables -P FORWARD ACCEPT
```

```
$ sudo iptables -F FORWARD
```

```
$ sudo iptables -t nat -F
```

```
$ sudo iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

Test again the connectivity between **VM3** and **VM1** and check that it is working now:

```
$ ping 192.168.0.100
```



## **1.7 Monitor network traffic**

To monitor network traffic, use **VM2** to run **tcpdump** and capture all network traffic. Make sure you can detect ICMP packets originating at **VM3** and destined to **VM1** (using ping). Use tcpdump with options `-X` and `-XX` and identify the *IP* addresses, MAC addresses and protocol in a given packet.

While still running **tcpdump**, open a **telnet** connection between machines **VM1** and **VM3** using user “*fireman*” and password “*Ins3cur3*”. Verify that you can capture both the username and password with tcpdump.

**You have successfully eavesdropped communications... But what is the difference between executing telnet from VM1 to VM3 with and without NAT? Use tcpdump to analyse the output and compare the differences.**