

Each question worth 2 points:

Tip: When not logged in as root, you can still use many of the network commands but you must preface them with `/sbin/` because you don't have root's environment. For example, typing `ifconfig` will result in a "command not found" error. Use `/sbin/ifconfig` instead.

The first three questions are all related. They refer to the same computer and NIC drivers.

1. **xxx** Use Putty or ssh to login as cis192 on **172.30.1.18**. This computer may have one or more Ethernet controllers (NICs). Determine the number of NICs installed. For every NIC found, identify the NIC vendor and model:

Using `/sbin/lspci` or `/sbin/lspci | grep Eth` commands

```
[cis192@mystery ~]$ /sbin/lspci | grep Eth
00:03.0 Ethernet controller: Intel Corporation 82540EM Gigabit Ethernet Controller
(rev 02)
00:08.0 Ethernet controller: Intel Corporation 82543GC Gigabit Ethernet Controller
(Copper) (rev 02)
00:09.0 Ethernet controller: Advanced Micro Devices [AMD] 79c970 [PCnet32 LANCE]
(rev 40)
[cis192@mystery ~]$
```

```
[cis192@mystery ~]$ /sbin/ifconfig
eth0      Link encap:Ethernet  HWaddr 08:00:27:E7:9B:20
          inet addr:172.30.1.18  Bcast:172.30.1.255  Mask:255.255.255.0
          < snipped >

eth1      Link encap:Ethernet  HWaddr 08:00:27:C1:6B:A5
          inet addr:10.205.163.107  Bcast:10.205.167.255  Mask:255.255.248.0
          < snipped >

eth1:1    Link encap:Ethernet  HWaddr 08:00:27:C1:6B:A5
          inet addr:10.205.163.200  Bcast:10.205.167.255  Mask:255.255.248.0
          < snipped >

eth1:2    Link encap:Ethernet  HWaddr 08:00:27:C1:6B:A5
          inet addr:10.205.163.201  Bcast:10.205.167.255  Mask:255.255.248.0
          < snipped >

eth1:3    Link encap:Ethernet  HWaddr 08:00:27:C1:6B:A5
          inet addr:10.205.163.202  Bcast:10.205.167.255  Mask:255.255.248.0
          < snipped >

eth2      Link encap:Ethernet  HWaddr 08:00:27:C6:2A:D8
          inet addr:192.168.45.200  Bcast:192.168.45.255  Mask:255.255.255.0
          < snipped >

[cis192@mystery ~]$
```

Number of NICs: **3**

NIC 1 vendor and model: **Intel 82540EM Gigabit**

NIC 2 vendor and model (if present): **AMD 79c970**

NIC 3 vendor and model (if present): Intel 82543GC Gigabit

2. **xxxx** Based on the Linux Ethernet Howto at <http://tldp.org/HOWTO/Ethernet-HOWTO.html> or googling "*vendor model linux driver*", what driver should be used for NIC 1, where is this driver located on the computer, and how would you check to see if this specific driver (not all drivers) has been loaded?

```
Status: Supported, Driver Name: e100, or eeepro100

The e100 driver was supplied by intel, and the eeepro100 driver is the original driver by Dona
listed in the driver are i82557, i82558, i82559, i82801, and about 25 other PCI IDs. For dri

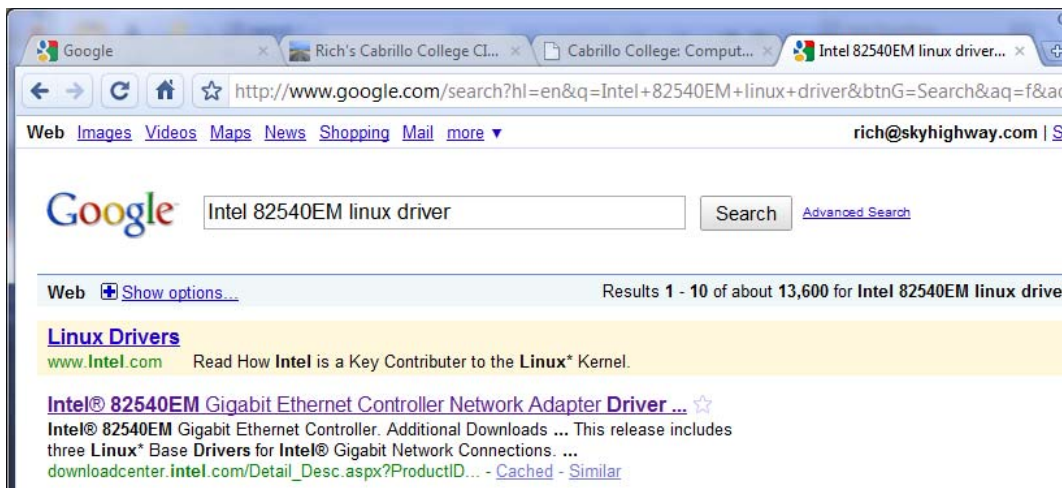
EEPro-100B Page

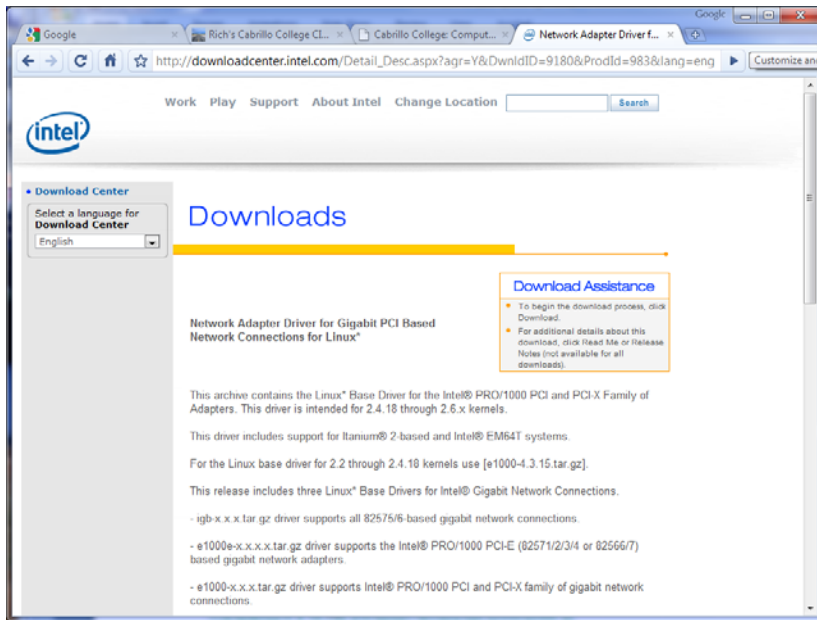
E1000 Gigabit

Status: Supported, Driver Name: e1000

4.25. Vignettes
```

<http://tldp.org/HOWTO/Ethernet-HOWTO-4.html#ss4.24>





Driver name: **e1000**

Absolute pathname of the driver:

/lib/modules/2.6.18-164.el5/kernel/drivers/net/e1000/e1000.ko

```
[cis192@mystery ~]$ ls /lib/modules/2.6.18-164.el5/kernel/drivers/net/e1000/e1000.ko
/lib/modules/2.6.18-164.el5/kernel/drivers/net/e1000/e1000.ko
[cis192@mystery ~]$
```

Command to show specific loaded driver: **lsmod | grep e1000**

```
[cis192@mystery ~]$ /sbin/lsmod | grep e1000
e1000                115797  0
[cis192@mystery ~]$
```

(hint: use this command to verify you have indeed identified the correct driver name)

3. NIC drivers are dynamic kernel modules. To unload and reload the driver (above) what specific commands would you use?


To unload the driver: **rmmmod e1000**

To load the driver: **modprobe e1000**

4. **xx** Given an IP address of 15.87.147.132 and a network mask of 255.224.0.0, what is the network (using /xx notation) and broadcast address?

Network: **15.64.0.0/11**

Broadcast address: **15.95.255.255**



5. **xxxxxxx** On Linux **client kate**, what complete command adds a static route to enable sending packets to hosts on VMnet4?

route add -net 192.168.30.0 netmask 255.255.255.0 gw 172.30.4.107

6. Given the following:

```


root@lilly:~# ping google.com
PING google.com (74.125.19.104) 56(84) bytes of data:
64 bytes from nuq04s01-in-f104.1e100.net (74.125.19.104): icmp_seq=1 ttl=55 time=23.0 ms
64 bytes from nuq04s01-in-f104.1e100.net (74.125.19.104): icmp_seq=2 ttl=55 time=21.1 ms
^C
--- google.com ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 2006ms
rtt min/avg/max/mdev = 16.929/20.373/23.071/2.565 ms
root@lilly:~# arp -n
Address                  HWtype  HWaddress          Flags Mask          Iface
16.112.94.254            ether    00:0c:42:3f:01:33  C                   eth0
root@lilly:~#

```

What is the likely default gateway IP address on Lilly?

16.112.94.254

- 7.



xxxxxxx What commands would you use to configure the eth0 IP addresses on Critter that would stay in effect till the next reboot?

ifconfig eth0 10.205.163.107 netmask 255.255.248.0

ifconfig eth0:1 10.205.163.200 netmask 255.255.248.0

ifconfig eth0:2 10.205.163.201 netmask 255.255.248.0

8. **x** Which of the following are routing protocols?

OSPF

RIP

Ethernet

IP

ARP

RIP and OSPF

9. **xxxxxxx** Using this routing table:

```
[root@elrond ~]# route -n
Kernel IP routing table
Destination      Gateway          Genmask          Flags Metric Ref    Use Iface
192.168.2.8      0.0.0.0          255.255.255.252 U        0      0      0 eth1
192.168.2.0      0.0.0.0          255.255.255.252 U        0      0      0 eth0
192.168.2.4      192.168.2.2     255.255.255.252 UG       2      0      0 eth0
172.30.4.0        0.0.0.0          255.255.255.0   U        0      0      0 eth2
10.10.10.0        192.168.2.2     255.255.255.0   UG       2      0      0 eth0
169.254.0.0       0.0.0.0          255.255.0.0     U        0      0      0 eth2
0.0.0.0           172.30.4.1       0.0.0.0          UG       0      0      0 eth2
[root@elrond ~]#
```

What would happen with an outgoing packet having an destination IP address of 192.168.2.6? Applying the genmask 255.255.255.252 indicates to that IP address yields network 192.168.2.4. There is a route for that network in the routing table. The outgoing packet will be forwarded out eth0 to the gateway at 192.168.2.2.

10. **xxxxxxx** Given the following ping test result, what is the likely problem and how would you fix it?

```
root@frodo:~# ping 172.30.1.222
PING 172.30.1.222 (172.30.1.222) 56(84) bytes of data:
From 172.30.1.151 icmp_seq=2 Destination Host Unreachable
From 172.30.1.151 icmp_seq=3 Destination Host Unreachable
From 172.30.1.151 icmp_seq=4 Destination Host Unreachable
^C
--- 172.30.1.222 ping statistics ---
6 packets transmitted, 0 received, +3 errors, 100% packet loss, time 5006ms
, pipe 3
root@frodo:~#
```

The destination host either does not exist or is powered off. Power it on or correct the IP address to an existing device.

11. **xxxxxxx** At what OSI layer are source and destination ports used? Layer 4 (Transport)

12. **xxxx** Two students are working in the CIS lab at different stations. They both are using the bridged Frodo VM and with correct IP, mask, DNS, and default gateway settings handed out by the DHCP server. Using their Frodo, they both can ping the lab router address 172.30.4.1 as shown below. However only Student 2 can ping Internet addresses.

```
Student 1
[root@frodo:~# ping -c1 172.30.4.1
PING 172.30.4.1 (172.30.4.1) 56(84) bytes of data.
64 bytes from 172.30.4.1: icmp_seq=1 ttl=64 time=1.32 ms

--- 172.30.4.1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 1.329/1.329/1.329/0.000 ms

[root@frodo:~# arp -n
Address                HWtype  HWaddress          Flags Mask          Iface
172.30.4.1             ether   00:0c:29:bf:e4:f9  C                   eth0

Student 2
root@frodo:~# ping -c1 172.30.4.1
PING 172.30.4.1 (172.30.4.1) 56(84) bytes of data.
64 bytes from 172.30.4.1: icmp_seq=1 ttl=64 time=2.41 ms

--- 172.30.4.1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 2.415/2.415/2.415/0.000 ms

root@frodo:~# arp -n
Address                HWtype  HWaddress          Flags Mask          Iface
172.30.4.1             ether   00:B0:64:53:42:01  C                   eth0
root@frodo:~#
```

You only have the troubleshooting information above to diagnose the problem. What is happening on the Lab network and what is the solution to fix it?

The arp tables on each system show a different MAC address for the lab router at 172.30.4.1. This indicates a duplicate IP situation where another interface on the lab network has been configured with the same IP address as the lab router!

The real lab router is a Cisco 2621 router. Googling “mac vendor 00-0c-29” shows student 1 has an ARP cache entry for 172.30.4.1 matching a VMware NIC interface! Student 1 cannot get to Internet using this 172.30.4.1 device as a gateway. Googling “mac vendor 00-B0-64” shows student 2 has a good ARP cache entry with 172.30.4.1 matching a Cisco interface. Student 2 is using the real Cisco router at 172.30.4.1 and can get out to the Internet without any problems.

To fix, the VM with an interface cabled to the lab network and configured with the same address as the lab router must be located and reconfigured with a unique IP address.

13. **xxxx** With respect to RFC **2131**, what is this RFC about and what is the **fifth** word in section **3.1.1**?

RFC **2131** is about: **DHCP**

5th word in section **3.1.1**: **DHCPDISCOVER**

14. **x** Given:

```
[root@elrond ~]# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:0C:29:4E:21:9B
          inet addr:192.168.2.1  Bcast:192.168.2.3  Mask:255.255.255.252
          inet6 addr: fe80::20c:29ff:fe4e:219b/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:28589 errors:0 dropped:0 overruns:0 frame:0
          TX packets:28949 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:2622712 (2.5 MiB)  TX bytes:3198137 (3.0 MiB)
          Interrupt:177 Base address:0x1400

[root@elrond ~]#
```

What is the **IPv6** and MAC address for this interface?

IPv6: **fe80::20c:29ff:fe4e:219b**

MAC: **00:0C:29:4E:21:9B**

15. **xxxxx** Given this routing table

```
[root@legolas ~]# route -n
Kernel IP routing table
Destination      Gateway         Genmask         Flags Metric Ref    Use Iface
192.168.2.8      192.168.2.1    255.255.255.252 UG    2     0      0 eth0
192.168.2.0      0.0.0.0        255.255.255.252 U      0     0      0 eth0
192.168.2.4      0.0.0.0        255.255.255.252 U      0     0      0 eth1
172.30.4.0       192.168.2.1    255.255.255.0  UG    2     0      0 eth0
10.10.10.0       0.0.0.0        255.255.255.0  U      0     0      0 eth2
169.254.0.0      0.0.0.0        255.255.0.0    U      0     0      0 eth2
0.0.0.0          192.168.2.1    0.0.0.0        UG    0     0      0 eth0

[root@legolas ~]#
```

What networks (use the /xx notation) is this Linux router directly connected to?


192.168.2.0 /30

192.168.2.4/30

10.10.10.0/24

169.254.0.0 /16

16. Extra Credit
(2 points)



In the network above, **hugo** fails to ping **kate** getting:

```
root@hugo:~# ping kate
PING kate (172.30.4.222) 56(84) bytes of data:
^C
--- kate ping statistics ---
12 packets transmitted, 0 received, 100% packet loss, time 11014ms
root@hugo:~#
```

The routing tables are as follows:

```
root@hugo:~# route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
192.168.30.0 0.0.0.0 255.255.255.0 U 0 0 0 eth0
169.254.0.0 0.0.0.0 255.255.0.0 U 1000 0 0 eth0
0.0.0.0 192.168.30.1 0.0.0.0 UG 100 0 0 eth0
root@hugo:~#
```

```
[root@r2 ~]# route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
192.168.2.0 0.0.0.0 255.255.255.0 U 0 0 0 eth0
192.168.30.0 0.0.0.0 255.255.255.0 U 0 0 0 eth1
169.254.0.0 0.0.0.0 255.255.0.0 U 0 0 0 eth1
0.0.0.0 192.168.2.107 0.0.0.0 UG 0 0 0 eth0
[root@r2 ~]#
```

```
[root@r1 ~]# route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
172.30.4.0 0.0.0.0 255.255.255.0 U 0 0 0 eth0
192.168.2.0 0.0.0.0 255.255.255.0 U 0 0 0 eth1
192.168.30.0 192.168.2.2 255.255.255.0 UG 0 0 0 eth1
169.254.0.0 0.0.0.0 255.255.0.0 U 0 0 0 eth1
0.0.0.0 172.30.4.1 0.0.0.0 UG 0 0 0 eth0
[root@r1 ~]#
```


```
root@kate:~# route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
172.30.4.0 0.0.0.0 255.255.255.0 U 0 0 0 eth0
192.168.3.0 172.30.4.107 255.255.255.0 UG 0 0 0 eth0
192.168.2.0 172.30.4.107 255.255.255.0 UG 0 0 0 eth0
169.254.0.0 0.0.0.0 255.255.0.0 U 1000 0 0 eth0
0.0.0.0 172.30.4.1 0.0.0.0 UG 100 0 0 eth0
root@kate:~#
```

What is happening and how do you fix it?

kate needs a static route to the 192.168.30.0/24 network. Add it with:
route add -net 192.168.30.0 netmask 255.255.255.0 gw 172.30.4.107

17. Extra Credit

(2 points)



Assuming Critter is a member of the Red Hat family of Linux distributions, how would you configure the eth0 IP addresses so they would stay in effect **after** the next reboot?

Create the following files:

```
==> /etc/sysconfig/network-scripts/ifcfg-eth0 <==
```

```
DEVICE=eth0  
ONBOOT=yes  
BOOTPROTO=static  
IPADDR=10.205.163.107  
NETMASK=255.255.248.0
```

```
==> /etc/sysconfig/network-scripts/ifcfg-eth0:1 <==
```

```
DEVICE=eth0:1  
ONBOOT=yes  
BOOTPROTO=static  
IPADDR=10.205.163.200  
NETMASK=255.255.248.0
```

```
==> /etc/sysconfig/network-scripts/ifcfg-eth0:2 <==
```

```
DEVICE=eth0:2  
ONBOOT=yes  
BOOTPROTO=static  
IPADDR=10.205.163.201  
NETMASK=255.255.248.0
```