

TP1 – Routage et translation d'adresses

Sommaire

<i>1 - Rappels.....</i>	<i>2</i>
<i>2 - Configuration réseau du serveur DS1.....</i>	<i>3</i>
<i>3 - Ajout de l'interface enp0s8.....</i>	<i>7</i>
<i>4 - Transformation du serveur en routeur.....</i>	<i>10</i>
<i>5 - Configuration du poste client Ubuntu.....</i>	<i>11</i>
<i>6 - Configuration du NAT sur le serveur DS1.....</i>	<i>15</i>

2

1 - Rappels

→ On met à jour tout les paquets :

```
root@DEB12Server: ~#apt-get update
```

→ On renomme la machine virtuelle en **DS1** en modifiant les fichiers `/etc/hostname` et `/etc/hosts` puis on redémarre la machine avec **reboot** :

```
GNU nano 7.2  
DS1
```

```
GNU nano 7.2  
127.0.0.1    localhost  
127.0.1.1    DEB12Server.sio-exupery.local DS1  
  
# The following lines are desirable for IPv6 capable hosts  
::1         localhost ip6-localhost ip6-loopback  
ff02::1    ip6-allnodes  
ff02::2    ip6-allrouters
```

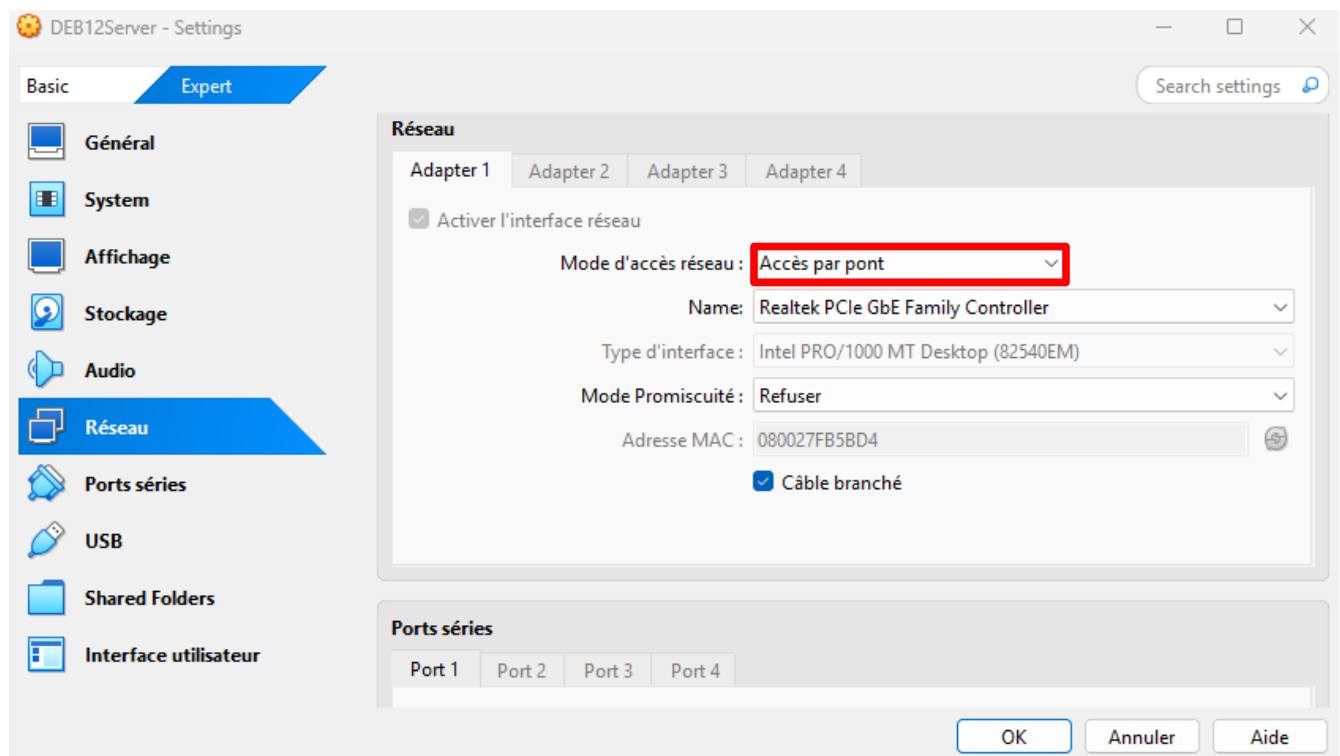
```
root@DS1: ~#reboot
```

2 - Configuration réseau du serveur DS1

→ On vérifie la **configuration** réseau actuelle en mode d'accès **réseau NAT** :

```
root@DS1: ~#ip address
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host noprefixroute
        valid_lft forever preferred_lft forever
2: enp0s3: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc fq_codel state DOWN group default qlen 1000
    link/ether 08:00:27:fb:5b:d4 brd ff:ff:ff:ff:ff:ff
    inet 172.17.110.11/16 brd 172.17.255.255 scope global dynamic enp0s3
        valid_lft 86138sec preferred_lft 86138sec
    inet6 fe80::a00:27ff:fe5b:5bd4/64 scope link
        valid_lft forever preferred_lft forever
```

→ On modifie le mode d'accès réseau par un accès **par pont** :



4

→ On **désactive** la carte réseau enp0s3 avec la commande **ifdown enp0s3** :

```
root@DS1: ~#ifdown enp0s3
Killed old client process
Internet Systems Consortium DHCP Client 4.4.3-P1
Copyright 2004-2022 Internet Systems Consortium.
All rights reserved.
For info, please visit https://www.isc.org/software/dhcp/

Listening on LPF/enp0s3/08:00:27:fb:5b:d4
Sending on      LPF/enp0s3/08:00:27:fb:5b:d4
Sending on      Socket/fallback
DHCPRELEASE of 172.17.110.11 on enp0s3 to 172.17.244.1 port 67
```

→ On passe d'une configuration IP en DHCP à une configuration en **IP fixe** en modifiant le fichier **/etc/network/interfaces** :

```
iface enp0s3 inet static
address 172.17.101.212
netmask 255.255.0.0
network 172.17.0.0
broadcast 172.17.255.255_
gateway 172.17.250.2
dns-nameservers 172.17.254.1
```

→ On réactive la carte réseau avec **ifup enp0s3** puis on vérifie la configuration IP :

```
root@DS1: ~#ifup enp0s3
root@DS1: ~#ip -c a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host noprefixroute
        valid_lft forever preferred_lft forever
2: enp0s3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
    link/ether 08:00:27:fb:5b:d4 brd ff:ff:ff:ff:ff:ff
    inet 172.17.101.212/16 brd 172.17.255.255 scope global enp0s3
        valid_lft forever preferred_lft forever
    inet6 fe80::a00:27ff:fe5b:d4/64 scope link
        valid_lft forever preferred_lft forever
```

5

→ On affiche le contenu du fichier `/etc/resolv.conf` et on vérifie la présence de l'adresse IP du serveur DNS :

```
root@DS1: ~#cat /etc/resolv.conf
domain prince.local
search prince.local
nameserver 172.17.254.1
nameserver 172.17.244.1
nameserver 80.10.246.2
nameserver 8.8.8.8
```

→ On visualise la **table de routage** de DS1 et on vérifie la prise en compte de la **passerelle par défaut (GW)** qu'on a entrée :

```
root@DS1: ~#ip route
default via 172.17.250.2 dev enp0s3 onlink
172.17.0.0/16 dev enp0s3 proto kernel scope link s
rc 172.17.101.212
```

→ On **ping** la **passerelle** ainsi que le **serveur DNS** :

```
root@DS1: ~#ping 172.17.250.2
PING 172.17.250.2 (172.17.250.2) 56(84) bytes of data.
64 bytes from 172.17.250.2: icmp_seq=1 ttl=255 time=1.16 ms
64 bytes from 172.17.250.2: icmp_seq=2 ttl=255 time=0.970 ms
64 bytes from 172.17.250.2: icmp_seq=3 ttl=255 time=1.03 ms
64 bytes from 172.17.250.2: icmp_seq=4 ttl=255 time=1.46 ms
^C
--- 172.17.250.2 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3010ms
rtt min/avg/max/mdev = 0.970/1.153/1.458/0.188 ms
```

```

root@DS1: ~#ping -c3 172.17.254.1
PING 172.17.254.1 (172.17.254.1) 56(84) bytes of data.
64 bytes from 172.17.254.1: icmp_seq=1 ttl=128 time=1.05 ms
64 bytes from 172.17.254.1: icmp_seq=2 ttl=128 time=1.21 ms
64 bytes from 172.17.254.1: icmp_seq=3 ttl=128 time=0.980 ms

--- 172.17.254.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2011ms
rtt min/avg/max/mdev = 0.980/1.078/1.210/0.096 ms

```

* Tout les test pings fonctionnent

→ On vérifie l'accès à Internet puis la résolution DNS avec des tests pings :

```

root@DS1: ~#ping -c 3 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_seq=1 ttl=117 time=31.4 ms
64 bytes from 8.8.8.8: icmp_seq=2 ttl=117 time=32.7 ms
64 bytes from 8.8.8.8: icmp_seq=3 ttl=117 time=33.7 ms

--- 8.8.8.8 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2016ms
rtt min/avg/max/mdev = 31.423/32.610/33.673/0.922 ms

```

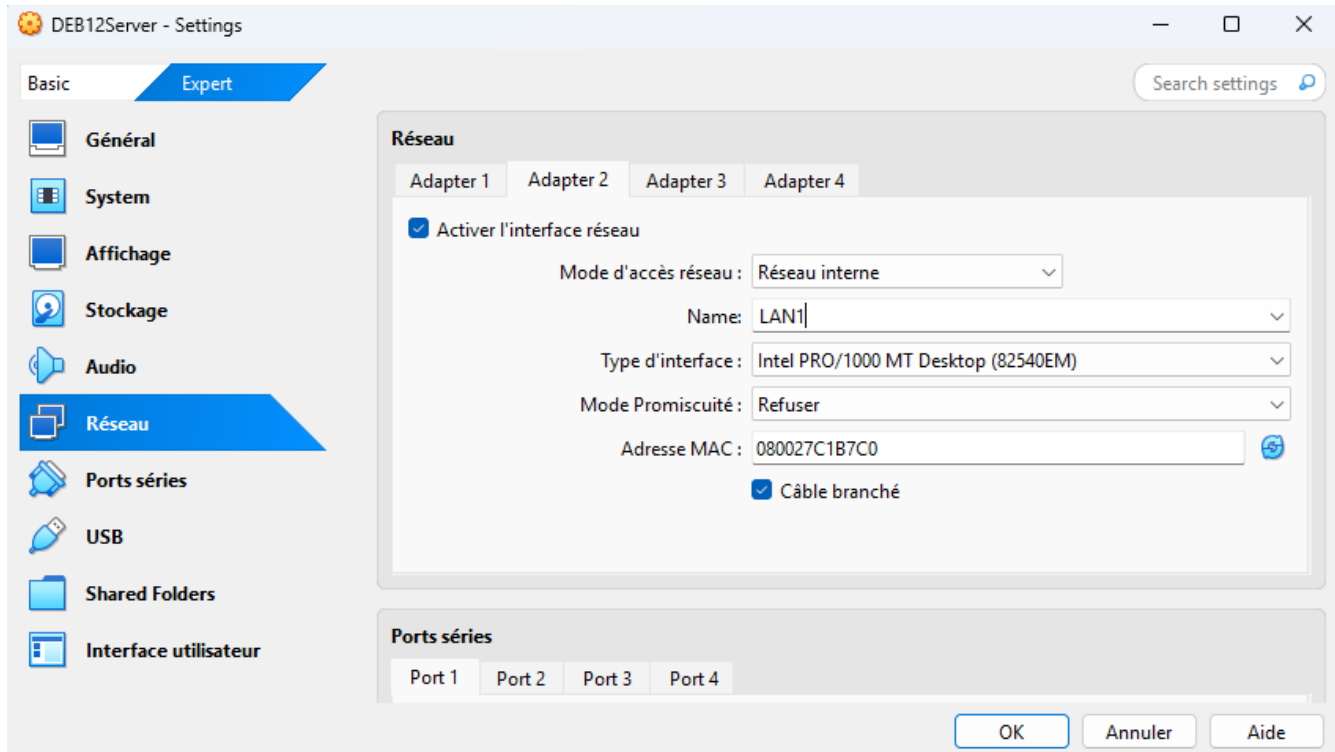
```

root@DS1: ~#ping www.ac-nice.fr
PING cs234.wpc.alphacdn.net (93.184.221.161) 56(84) bytes of data.
64 bytes from 93.184.221.161 (93.184.221.161): icmp_seq=1 ttl=56 time=33.7 ms
64 bytes from 93.184.221.161 (93.184.221.161): icmp_seq=2 ttl=56 time=32.3 ms
64 bytes from 93.184.221.161 (93.184.221.161): icmp_seq=3 ttl=56 time=32.3 ms
^C
--- cs234.wpc.alphacdn.net ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2007ms
rtt min/avg/max/mdev = 32.272/32.778/33.724/0.669 ms

```

3 - Ajout de l'interface enp0s8

→ On arrête la machine virtuelle pour ajouter une deuxième carte réseau à celle-ci configurée en mode **Réseau Interne** :



→ On vérifie que la carte réseau a été prise en compte :

```
root@DS1: ~#ip -c a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
   inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
   inet6 ::1/128 scope host noprefixroute
       valid_lft forever preferred_lft forever
2: enp0s3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
   link/ether 08:00:27:fb:5b:d4 brd ff:ff:ff:ff:ff:ff
   inet 172.17.101.212/16 brd 172.17.255.255 scope global enp0s3
       valid_lft forever preferred_lft forever
   inet6 fe80::a00:27ff:fe5b:d4/64 scope link
       valid_lft forever preferred_lft forever
3: enp0s8: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN group default qlen 1000
   link/ether 08:00:27:c1:b7:c0 brd ff:ff:ff:ff:ff:ff
```

→ On ajoute l'interface d'**enp0s8** dans le fichier **/etc/network/interfaces** :

```
allow-hotplug enp0s8
iface enp0s8 inet static
address 192.168.4.254
netmask 255.255.255.0
network 192.168.4.0
broadcast 192.168.4.255
```

→ On **active** la carte et on vérifie sa bonne configuration IP :

```
root@DS1: ~#ifup enp0s8
root@DS1: ~#ip -c a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
   inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
   inet6 ::1/128 scope host noprefixroute
       valid_lft forever preferred_lft forever
2: enp0s3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
   link/ether 08:00:27:fb:5b:d4 brd ff:ff:ff:ff:ff:ff
   inet 172.17.101.212/16 brd 172.17.255.255 scope global enp0s3
       valid_lft forever preferred_lft forever
   inet6 fe80::a00:27ff:fe5b:bd4/64 scope link
       valid_lft forever preferred_lft forever
3: enp0s8: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
   link/ether 08:00:27:c1:b7:c0 brd ff:ff:ff:ff:ff:ff
   inet 192.168.4.254/24 brd 192.168.4.255 scope global enp0s8
       valid_lft forever preferred_lft forever
   inet6 fe80::a00:27ff:fec1:b7c0/64 scope link
       valid_lft forever preferred_lft forever
```

→ On **ping** les 2 interfaces de la machine pour vérifier leur bonne configuration :

```
root@DS1: ~#ping -c 3 192.168.4.254
PING 192.168.4.254 (192.168.4.254) 56(84) bytes of data.
64 bytes from 192.168.4.254: icmp_seq=1 ttl=64 time=0.038 ms
64 bytes from 192.168.4.254: icmp_seq=2 ttl=64 time=0.038 ms
64 bytes from 192.168.4.254: icmp_seq=3 ttl=64 time=0.037 ms

--- 192.168.4.254 ping statistics ---
3 packets transmitted, 3 received 0% packet loss, time 2027ms
rtt min/avg/max/mdev = 0.037/0.037/0.038/0.000 ms
root@DS1: ~#ping -c 3 172.17.101.212
PING 172.17.101.212 (172.17.101.212) 56(84) bytes of data.
64 bytes from 172.17.101.212: icmp_seq=1 ttl=64 time=0.024 ms
64 bytes from 172.17.101.212: icmp_seq=2 ttl=64 time=0.033 ms
64 bytes from 172.17.101.212: icmp_seq=3 ttl=64 time=0.044 ms

--- 172.17.101.212 ping statistics ---
3 packets transmitted, 3 received 0% packet loss, time 2027ms
rtt min/avg/max/mdev = 0.024/0.033/0.044/0.008 ms
```

→ On affiche la nouvelle **table de routage** de DS1 :

```
root@DS1: ~#ip route
default via 172.17.250.2 dev enp0s3 onlink
172.17.0.0/16 dev enp0s3 proto kernel scope link src 172.17.101.212
192.168.4.0/24 dev enp0s8 proto kernel scope link src 192.168.4.254
```

4 - Transformation du serveur en routeur

→ On change la valeur dans le fichier `ip_forward` en 0 à 1 pour activer le routage :

```
root@DS1: ~#echo 1 > /proc/sys/net/ipv4/ip_forward
root@DS1: ~#cat /proc/sys/net/ipv4/ip_forward
1
```

→ On décommente la ligne `net.ipv4.ip_forward=1` dans le fichier `/etc/sysctl.conf` pour que le routage se mette en place automatiquement à chaque démarrage de la machine :

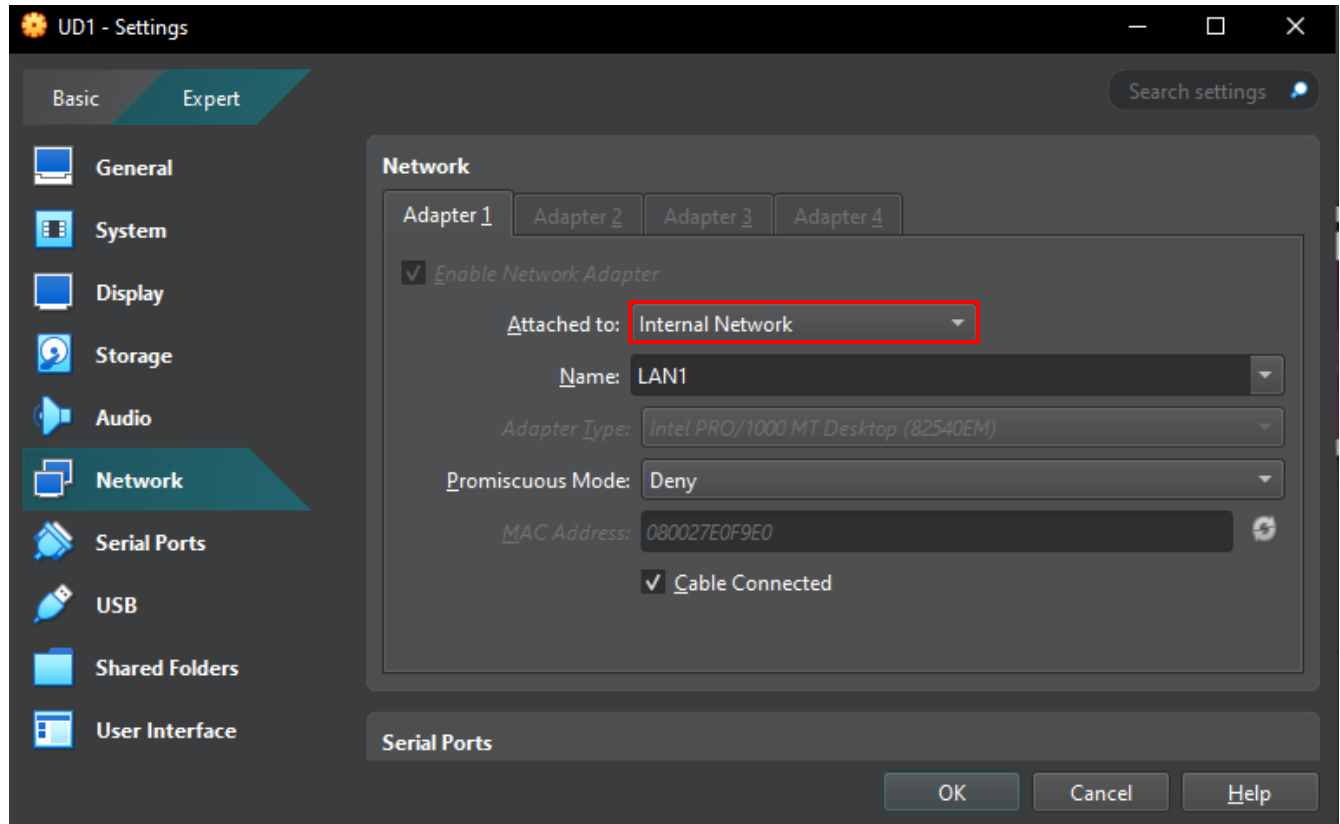
```
# Uncomment the next line to enable packet forwarding for IPv4
net.ipv4.ip_forward=1
```

→ On redémarre la machine avec `reboot` pour que le fichier modifié soit chargé par la machine et on vérifie si c'est le cas :

```
root@DS1: ~#cat /proc/sys/net/ipv4/ip_forward
1
```

5 - Configuration du poste client Ubuntu

→ On règle la machine **UD1** en mode d'accès **LAN** :



→ On modifie la **configuration IP** :

Editing enp0s3

Connection name: enp0s3

General | Ethernet | 802.1X Security | DCB | Proxy | **IPv4 Settings** | IPv6 Settings

Method: Manual

Addresses

Address	Netmask	Gateway	
192.168.4.1	24	192.168.4.254	<input type="button" value="Add"/>
			<input type="button" value="Delete"/>

DNS servers: 192.168.1.254

Search domains:

DHCP client ID:

Require IPv4 addressing for this connection to complete

```
ud1@ud1-VirtualBox: ~  
To run a command as administrator (user "root"), use "sudo <command>".  
See "man sudo_root" for details.  
  
ud1@ud1-VirtualBox:~$ ip -c a  
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000  
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00  
    inet 127.0.0.1/8 scope host lo  
        valid_lft forever preferred_lft forever  
    inet6 ::1/128 scope host noprefixroute  
        valid_lft forever preferred_lft forever  
2: enp0s3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000  
    link/ether 08:00:27:e0:f9:e0 brd ff:ff:ff:ff:ff:ff  
    inet 192.168.4.1/24 brd 192.168.4.255 scope global noprefixroute enp0s3  
        valid_lft forever preferred_lft forever  
    inet6 fe80::a00:27ff:fee0:f9e0/64 scope link  
        valid_lft forever preferred_lft forever  
ud1@ud1-VirtualBox:~$
```

→ On consulte la **table de routage** :

```
ud1@ud1-VirtualBox:~$ ip route  
default via 192.168.4.254 dev enp0s3 proto static metric 20100  
192.168.4.0/24 dev enp0s3 proto kernel scope link src 192.168.4.1 metric 100
```

→ On ping les 2 interfaces réseau de **DS1** :

```

ud1@ud1-VirtualBox:~$ ping -c 3 192.168.1.47
PING 192.168.1.47 (192.168.1.47) 56(84) bytes of data.
64 bytes from 192.168.1.47: icmp_seq=1 ttl=64 time=0.209 ms
64 bytes from 192.168.1.47: icmp_seq=2 ttl=64 time=0.207 ms
64 bytes from 192.168.1.47: icmp_seq=3 ttl=64 time=0.210 ms

--- 192.168.1.47 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2994ms
rtt min/avg/max/mdev = 0.207/0.208/0.210/0.001 ms
ud1@ud1-VirtualBox:~$ ping -c 3 192.168.4.254
PING 192.168.4.254 (192.168.4.254) 56(84) bytes of data.
64 bytes from 192.168.4.254: icmp_seq=1 ttl=64 time=0.239 ms
64 bytes from 192.168.4.254: icmp_seq=2 ttl=64 time=0.227 ms
64 bytes from 192.168.4.254: icmp_seq=3 ttl=64 time=0.205 ms

--- 192.168.4.254 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 4189ms
rtt min/avg/max/mdev = 0.205/0.223/0.239/0.014 ms

```

* Les pings aboutissent.

→ On ping maintenant la passerelle par défaut :

```

ud1@ud1-VirtualBox:~$ ping -c 3 192.168.1.254
PING 192.168.1.254 (192.168.1.254) 56(84) bytes of data.
From 192.168.4.254 icmp_seq=1 Destination Host Unreachable
From 192.168.4.254 icmp_seq=2 Destination Host Unreachable
From 192.168.4.254 icmp_seq=3 Destination Host Unreachable

--- 192.168.1.254 ping statistics ---
3 packets transmitted, 0 received, +3 errors, 100% packet loss, time 2627ms
pipe 3

```

* On constate que UD1 ne communique pas avec l'extérieur car la translation d'adresses n'est pas configurée donc le routeur ne connaît pas encore UD1.

6 - Configuration du NAT sur le serveur DS1

→ On installe le paquet **iptables** :

```
root@DS1: ~#apt-get install iptables
```

→ On met en place l'IP Masquerading :

```
root@DS1: ~#iptables -t nat -A POSTROUTING -o enp0s3 -s 192.168.4.0/24 -j MASQUERADE
```

→ On installe le paquet **iptables-persistent** pour que la **translation d'adresse NAT** se fasse à chaque démarrage :

```
root@DS1: ~#apt-get install iptables-persistent_
```

→ On redémarre la machine puis on vérifie que la nouvelle règle soit bien prise en compte au démarrage :

```
root@DS1: ~#iptables -t nat -L -v
Chain PREROUTING (policy ACCEPT 0 packets, 0 bytes)
 pkts bytes target      prot opt in     out     source
Chain INPUT (policy ACCEPT 0 packets, 0 bytes)
 pkts bytes target      prot opt in     out     source
Chain OUTPUT (policy ACCEPT 0 packets, 0 bytes)
 pkts bytes target      prot opt in     out     source
Chain POSTROUTING (policy ACCEPT 0 packets, 0 bytes)
 0      0 MASQUERADE  all  --  any    enp0s3  192.168.4.0/24  anywhere
```

→ On effectue à nouveau des **test pings** sur le routeur à partir de la machine **UD1** :

```
ud1@ud1-VirtualBox:~$ ping -c 3 192.168.1.254
PING 192.168.1.254 (192.168.1.254) 56(84) bytes of data.
64 bytes from 192.168.1.254: icmp_seq=1 ttl=63 time=2.04 ms
64 bytes from 192.168.1.254: icmp_seq=2 ttl=63 time=2.97 ms
64 bytes from 192.168.1.254: icmp_seq=3 ttl=63 time=2.27 ms

--- 192.168.1.254 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2003ms
rtt min/avg/max/mdev = 2.043/2.426/2.970/0.394 ms
ud1@ud1-VirtualBox:~$ |
```

→ On installe le paquet **tcpdump** sur la machine **DS1** :

```
root@DS1: ~#apt-get install tcpdump
```

→ On capture les **trâmes ICMP** des **tests pings** à partir de **UD1** en utilisant **tcpdump** :

```
root@DS1: ~#tcpdump -i enp0s3 icmp -n
tcpdump: verbose output suppressed, use -v[v]... for full protocol decode
listening on enp0s3, link-type EN10MB (Ethernet), snapshot length 262144 bytes
01:51:18.630441 IP 192.168.1.80 > 192.168.1.254: ICMP echo request, id 8286, seq 7, length 64
01:51:18.633511 IP 192.168.1.254 > 192.168.1.80: ICMP echo reply, id 8286, seq 7, length 64

root@DS1: ~#tcpdump -i enp0s8 icmp -n
tcpdump: verbose output suppressed, use -v[v]... for full protocol decode
listening on enp0s8, link-type EN10MB (Ethernet), snapshot length 262144 bytes
19:34:36.426710 IP 192.168.4.1 > 192.168.1.80: ICMP echo request, id 4655, seq 1, length 64
19:34:36.426744 IP 192.168.1.80 > 192.168.4.1: ICMP echo reply, id 4655, seq 1, length 64
```

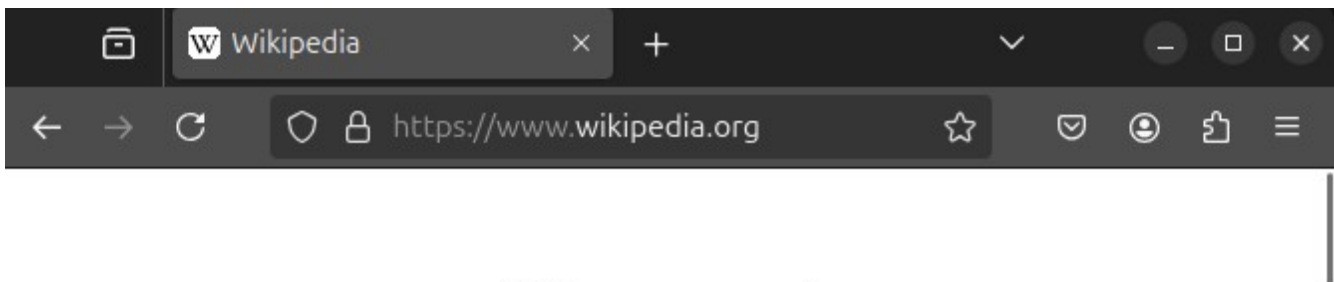
→ On teste la **translation** et la **résolution DNS** depuis **UD1** en effectuant des **tests pings** :

```
ud1@ud1-VirtualBox: /mnt
ud1@ud1-VirtualBox:/mnt$ ping -c 3 www.ac-nice.fr
PING www.ac-nice.fr.cdn.cloudflare.net (141.101.90.107) 56(84)
bytes of data.
64 bytes from 141.101.90.107: icmp_seq=1 ttl=55 time=6.52 ms
64 bytes from 141.101.90.107: icmp_seq=2 ttl=55 time=7.37 ms
64 bytes from 141.101.90.107: icmp_seq=3 ttl=55 time=7.32 ms

--- www.ac-nice.fr.cdn.cloudflare.net ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2011ms
rtt min/avg/max/mdev = 6.515/7.067/7.365/0.390 ms
```

* **Tout fonctionne correctement.**

→ On lance le navigateur pour vérifier l'**accès à Internet** :



* **L'accès à Internet fonctionne donc aussi.**