


Configuration réseau, «sniffing» & virtualisation & VLANs

■ ■ ■ Plateforme virtuelle et VLANs

On va reprendre le réseau du TP n°1, où au lieu de 3 swiches, on va n'en utiliser qu'un seul :

- ☐ 3 réseaux correspondant sur un seul switch ;
- ☐ deux routeurs r1 et r2 ;
- ☐ deux hôtes, h1 et h2, chacun connecté dans le réseau s respectif.

switch/réseau	adresse
s1	192.168.10.0/24
s2	192.168.20.0/24
internet	10.0.0.0/24



■ ■ ■ Utilisation d'un seul switch pour tout connecter

On modifie le fichier «build_architecture » :

```
# créer le switch
# ovs-vsctl add-br internet
# ovs-vsctl add-br s1
# ovs-vsctl add-br s2
ovs-vsctl add-br switch
```

On remplace les 3 swiches par un seul appelé « switch ».

Puis on modifie les connexions :

```
ovs-vsctl add-port switch s1-h1
ovs-vsctl add-port switch s1-h3
ovs-vsctl add-port switch s1-r1
ovs-vsctl add-port switch s2-h2
ovs-vsctl add-port switch s2-r2

ovs-vsctl add-port switch internet-r1
ovs-vsctl add-port switch internet-r2
```

On «ping» h2 depuis h1 :

```
xterm
pef@palomino:~/NET_LAB_VLAN$ [h1] ping -c 1 192.168.20.1
PING 192.168.20.1 (192.168.20.1) 56(84) bytes of data.
64 bytes from 192.168.20.1: icmp_seq=1 ttl=62 time=3.05 ms

--- 192.168.20.1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 3.046/3.046/3.046/0.000 ms
```

Ça marche !

Mais comment ? On passe par r1 !

```
xterm
pef@palomino:~$ [h1] sudo tcpdump -lnvveX -i h1-eth0 not ip6
tcpdump: listening on h1-eth0, link-type EN10MB (Ethernet), snapshot length 262144
bytes
23:28:59.015816 d2:cc:44:e8:3c:6d > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Request who-has 192.168.10.254 tell
192.168.10.1, length 28
    0x0000: 0001 0800 0604 0001 d2cc 44e8 3c6d c0a8 .....D.<m..
    0x0010: 0a01 0000 0000 0000 c0a8 0afe .....
23:28:59.016617 66:f1:76:90:a9:21 > d2:cc:44:e8:3c:6d, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Reply 192.168.10.254 is-at
66:f1:76:90:a9:21, length 28
    0x0000: 0001 0800 0604 0002 66f1 7690 a921 c0a8 .....f.v...!..
    0x0010: 0afe d2cc 44e8 3c6d c0a8 0a01 ....D.<m....
23:28:59.016632 d2:cc:44:e8:3c:6d > 66:f1:76:90:a9:21, ethertype IPv4 (0x0800),
length 98: (tos 0x0, ttl 64, id 25348, offset 0, flags [DF], proto ICMP (1), length
84)
    192.168.10.1 > 192.168.20.1: ICMP echo request, id 21920, seq 1, length 64
    0x0000: 4500 0054 6304 4000 4001 3852 c0a8 0a01 E..Tc.@.@.8R....
    0x0010: c0a8 1401 0800 5280 55a0 0001 ab68 4965 .....R.U....hIe
    0x0020: 0000 0000 9c3d 0000 0000 0000 1011 1213 .....=.....
    0x0030: 1415 1617 1819 1a1b 1c1d 1e1f 2021 2223 .....!"#
    0x0040: 2425 2627 2829 2a2b 2c2d 2e2f 3031 3233 $%&'()*+,-./0123
    0x0050: 3435 3637 .....4567
23:28:59.016704 f2:01:9e:99:31:ac > d2:cc:44:e8:3c:6d, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Reply 192.168.10.254 is-at
f2:01:9e:99:31:ac, length 28
    0x0000: 0001 0800 0604 0002 f201 9e99 31ac c0a8 .....1...
    0x0010: 0afe d2cc 44e8 3c6d c0a8 0a01 ....D.<m....
23:28:59.017087 f2:01:9e:99:31:ac > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Request who-has 10.0.0.2 tell 10.0.0.1,
length 28
    0x0000: 0001 0800 0604 0001 f201 9e99 31ac 0a00 .....1...
    0x0010: 0001 0000 0000 0000 0a00 0002 .....
23:28:59.017743 b6:08:5d:75:74:a4 > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Request who-has 192.168.20.1 tell
192.168.20.254, length 28
    0x0000: 0001 0800 0604 0001 b608 5d75 74a4 c0a8 .....]ut...
    0x0010: 14fe 0000 0000 0000 c0a8 1401 .....
23:28:59.018818 66:f1:76:90:a9:21 > d2:cc:44:e8:3c:6d, ethertype IPv4 (0x0800),
length 98: (tos 0x0, ttl 62, id 51432, offset 0, flags [none], proto ICMP (1),
length 84)
    192.168.20.1 > 192.168.10.1: ICMP echo reply, id 21920, seq 1, length 64
    0x0000: 4500 0054 c8e8 0000 3e01 146e c0a8 1401 E..T....>..n....
    0x0010: c0a8 0a01 0000 5a80 55a0 0001 ab68 4965 .....Z.U....hIe
    0x0020: 0000 0000 9c3d 0000 0000 0000 1011 1213 .....=.....
    0x0030: 1415 1617 1819 1a1b 1c1d 1e1f 2021 2223 .....!"#
    0x0040: 2425 2627 2829 2a2b 2c2d 2e2f 3031 3233 $%&'()*+,-./0123
    0x0050: 3435 3637 .....4567
23:29:04.164547 66:f1:76:90:a9:21 > d2:cc:44:e8:3c:6d, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Request who-has 192.168.10.1 tell
192.168.10.254, length 28
    0x0000: 0001 0800 0604 0001 66f1 7690 a921 c0a8 .....f.v...!..
    0x0010: 0afe 0000 0000 0000 c0a8 0a01 .....
23:29:04.164568 d2:cc:44:e8:3c:6d > 66:f1:76:90:a9:21, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Reply 192.168.10.1 is-at
d2:cc:44:e8:3c:6d, length 28
    0x0000: 0001 0800 0604 0002 d2cc 44e8 3c6d c0a8 .....D.<m..
    0x0010: 0a01 66f1 7690 a921 c0a8 0afe ..f.v...!.....
```

Ok $\Rightarrow$ comme dans un réseau bien fait : on fait de l'ARP pour obtenir l'@MAC du routeur, puis on passe par lui pour envoyer le paquet « *ICMP request* » et c'est lui qui nous retourne le paquet « *ICMP reply* » !

Mais c'est quoi ces messages **ARP supplémentaires** ??

▷ un « *echo reply* » de la part de la seconde interface de r1 ??

C'est comme si r1 avait reçu le broadcast de la requête ARP de h1 qui le concerne sur ces deux interfaces

$\Rightarrow$ c'est vrai le broadcast est "floodé" sur tous les ports du switch.

▷ une requête ARP pour 10.0.0.2 $\Rightarrow$ c'est r1 qui cherche r2

▷ une requête ARP pour 192.168.20.1 $\Rightarrow$ c'est r2 qui cherche h2

Toutes les requêtes ARP transmises en broadcast sont reçues par h1 !!!

$\Rightarrow$ pas d'isolation !!!

Un seul switch n'apporte aucune isolation de niveau 2 $\Rightarrow$ il faut des VLANs !

Pour connaître les adresses MAC respectives et interpréter correctement la capture de tcpdump :

```
xterm
pef@palomino:~$ [r1] ip 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
        valid_lft forever preferred_lft forever
80: r1-eth0@if79: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
group default qlen 1000
    link/ether 66:f1:76:90:a9:21 brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 192.168.10.254/24 scope global r1-eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::64f1:76ff:fe90:a921/64 scope link
        valid_lft forever preferred_lft forever
82: r1-eth1@if81: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
group default qlen 1000
    link/ether f2:01:9e:99:31:ac brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 10.0.0.1/24 scope global r1-eth1
        valid_lft forever preferred_lft forever
    inet6 fe80::f001:9eff:fe99:31ac/64 scope link
        valid_lft forever preferred_lft forever
```

```
xterm
pef@palomino:~/NET_LAB_VLAN$ [h1] ip 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
        valid_lft forever preferred_lft forever
74: h1-eth0@if73: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
group default qlen 1000
    link/ether d2:cc:44:e8:3c:6d brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 192.168.10.1/24 scope global h1-eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::d0cc:44ff:fee8:3c6d/64 scope link
        valid_lft forever preferred_lft forever
```

Mais, si sur h1, je me configure dans le réseau de h2 ?

```
xterm
pef@palomino:~/NET_LAB_VLAN$ [h1] sudo ip a add 192.168.20.42/24 dev h1-eth0
pef@palomino:~/NET_LAB_VLAN$ [h1] ip 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
        valid_lft forever preferred_lft forever
74: h1-eth0@if73: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
group default qlen 1000
    link/ether d2:cc:44:e8:3c:6d brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 192.168.10.1/24 scope global h1-eth0
        valid_lft forever preferred_lft forever
    inet 192.168.20.42/24 scope global h1-eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::d0cc:44ff:fee8:3c6d/64 scope link
        valid_lft forever preferred_lft forever
pef@palomino:~/NET_LAB_VLAN$ [h1] ping -c 1 192.168.20.1
PING 192.168.20.1 (192.168.20.1) 56(84) bytes of data.
64 bytes from 192.168.20.1: icmp_seq=1 ttl=64 time=1.20 ms

--- 192.168.20.1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 1.200/1.200/1.200/0.000 ms
```

On regarde comment ça marche :

```
xterm
23:29:21.963787 d2:cc:44:e8:3c:6d > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Request who-has 192.168.20.1 tell
192.168.20.42, length 28
    0x0000: 0001 0800 0604 0001 d2cc 44e8 3c6d c0a8 .....D.<m..
    0x0010: 142a 0000 0000 0000 c0a8 1401 .....*.....
23:29:21.964593 5e:6d:ee:ba:fc:7b > d2:cc:44:e8:3c:6d, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Reply 192.168.20.1 is-at
5e:6d:ee:ba:fc:7b, length 28
    0x0000: 0001 0800 0604 0002 5e6d eeba fc7b c0a8 .....^m...{..
    0x0010: 1401 d2cc 44e8 3c6d c0a8 142a .....D.<m...*
23:29:21.964608 d2:cc:44:e8:3c:6d > 5e:6d:ee:ba:fc:7b, ethertype IPv4 (0x0800),
length 98: (tos 0x0, ttl 64, id 2476, offset 0, flags [DF], proto ICMP (1), length
84)
    192.168.20.42 > 192.168.20.1: ICMP echo request, id 52693, seq 1, length 64
    0x0000: 4500 0054 09ac 4000 4001 8781 c0a8 142a E..T..@.@.....*
    0x0010: c0a8 1401 0800 c3d3 cdd5 0001 c168 4965 .....hIe
    0x0020: 0000 0000 8eb4 0e00 0000 0000 1011 1213 .....
    0x0030: 1415 1617 1819 1a1b 1c1d 1e1f 2021 2223 .....!"#
    0x0040: 2425 2627 2829 2a2b 2c2d 2e2f 3031 3233 $%&'()*+,-./0123
    0x0050: 3435 3637 .....4567
23:29:21.964926 5e:6d:ee:ba:fc:7b > d2:cc:44:e8:3c:6d, ethertype IPv4 (0x0800),
length 98: (tos 0x0, ttl 64, id 18751, offset 0, flags [none], proto ICMP (1),
length 84)
    192.168.20.1 > 192.168.20.42: ICMP echo reply, id 52693, seq 1, length 64
    0x0000: 4500 0054 493f 0000 4001 87ee c0a8 1401 E..TI?...@.....
    0x0010: c0a8 142a 0000 cbd3 cdd5 0001 c168 4965 ...*.....hIe
    0x0020: 0000 0000 8eb4 0e00 0000 0000 1011 1213 .....
    0x0030: 1415 1617 1819 1a1b 1c1d 1e1f 2021 2223 .....!"#
    0x0040: 2425 2627 2829 2a2b 2c2d 2e2f 3031 3233 $%&'()*+,-./0123
    0x0050: 3435 3637 .....4567
```

Pour connaître les adresses MAC respectives et interpréter correctement la capture de tcpdump :

```
xterm
pef@palomino:~$ [h2] ip 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
        valid_lft forever preferred_lft forever
76: h2-eth0@if75: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
group default qlen 1000
    link/ether 5e:6d:ee:ba:fc:7b brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 192.168.20.1/24 scope global h2-eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::5c6d:eeff:feba:fc7b/64 scope link
        valid_lft forever preferred_lft forever
```

⇒ h1 ne fait plus de routage !!!

⇒ On va directement de h1 vers h2 !!!

⇒ On a bien **aucune isolation !!!**

■ ■ ■ Première configuration de VLAN : un port ⇒ un VLAN

On modifie le fichier «build_architecture »:

```
ovs-vsctl add-port switch s1-h1 tag=100
ovs-vsctl add-port switch s1-h3 tag=100
ovs-vsctl add-port switch s1-r1 tag=100
ovs-vsctl add-port switch s2-h2 tag=200
ovs-vsctl add-port switch s2-r2 tag=200

ovs-vsctl add-port switch internet-r1 tag=300
ovs-vsctl add-port switch internet-r2 tag=300
```

▷ Le VLAN 100 correspond au réseau 192.168.10.0/24

▷ Le VLAN 200 correspond au réseau 192.168.20.0/24

▷ Le VLAN 300 correspond au réseau 10.0.0.0/24

Est-ce que l'on a bien de l'isolation ? On va essayer de rejoindre le réseau 10.0.0.0/24 :

```
xterm
pef@palomino:~/NET_LAB_VLAN$ [h1] ip 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
        valid_lft forever preferred_lft forever
58: h1-eth0@if57: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
    link/ether d2:cc:44:e8:3c:6d brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 192.168.10.1/24 scope global h1-eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::d0cc:44ff:fee8:3c6d/64 scope link
        valid_lft forever preferred_lft forever
pef@palomino:~/NET_LAB_VLAN$ [h1] sudo ip a add 10.0.0.42/24 dev h1-eth0
pef@palomino:~/NET_LAB_VLAN$ [h1] ping -c 1 10.0.0.1
PING 10.0.0.1 (10.0.0.1) 56(84) bytes of data.

--- 10.0.0.1 ping statistics ---
1 packets transmitted, 0 received, 100% packet loss, time 0ms
```

Ouf l'isolation est bonne !

Mais est-ce vraiment le cas ?

```
xterm
pef@palomino:~$ [h1] sudo tcpdump -lnvveX -i h1-eth0 not ip6
tcpdump: listening on h1-eth0, link-type EN10MB (Ethernet), snapshot length 262144 bytes
23:15:31.778517 d2:cc:44:e8:3c:6d > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806), length 42: Ethernet (len 6), IPv4 (len 4), Request who-has 10.0.0.1 tell 10.0.0.42, length 28
    0x0000: 0001 0800 0604 0001 d2cc 44e8 3c6d 0a00 .....D.<m..
    0x0010: 002a 0000 0000 0000 0a00 0001 .....*.....
23:15:31.779141 66:f1:76:90:a9:21 > d2:cc:44:e8:3c:6d, ethertype ARP (0x0806), length 42: Ethernet (len 6), IPv4 (len 4), Reply 10.0.0.1 is-at 66:f1:76:90:a9:21, length 28
    0x0000: 0001 0800 0604 0002 66f1 7690 a921 0a00 .....f.v...!..
    0x0010: 0001 d2cc 44e8 3c6d 0a00 002a ....D.<m....*
23:15:31.779154 d2:cc:44:e8:3c:6d > 66:f1:76:90:a9:21, ethertype IPv4 (0x0800), length 98: (tos 0x0, ttl 64, id 50459, offset 0, flags [DF], proto ICMP (1), length 84)
    10.0.0.42 > 10.0.0.1: ICMP echo request, id 62504, seq 1, length 64
    0x0000: 4500 0054 c51b 4000 4001 6163 0a00 002a E..T..@.@.ac...*
    0x0010: 0a00 0001 0800 8357 f428 0001 8365 4965 .....W.(...eIe
    0x0020: 0000 0000 e9e0 0b00 0000 0000 1011 1213 .....
    0x0030: 1415 1617 1819 1a1b 1c1d 1e1f 2021 2223 .....!"#
    0x0040: 2425 2627 2829 2a2b 2c2d 2e2f 3031 3233 $%&'()*+,-./0123
    0x0050: 3435 3637 .....4567
```

Horreur !!! la requête ARP a été transmise vers r1 et une réponse reçue par h1 !!!

Regardons tout de même les adresses MACs utilisées :

```
xterm
pef@palomino:~$ [r1] ip 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
        valid_lft forever preferred_lft forever
64: r1-eth0@if63: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
    link/ether 66:f1:76:90:a9:21 brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 192.168.10.254/24 scope global r1-eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::64f1:76ff:fe90:a921/64 scope link
        valid_lft forever preferred_lft forever
66: r1-eth1@if65: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
    link/ether f2:01:9e:99:31:ac brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 10.0.0.1/24 scope global r1-eth1
        valid_lft forever preferred_lft forever
    inet6 fe80::f001:9eff:fe99:31ac/64 scope link
        valid_lft forever preferred_lft forever
```

Mais qu'elle est l'adresse MAC donnée en réponse ? Celle de l'interface de r1 connectée au VLAN de h1, c-à-d 66:f1:76:90:a9:21...

Est-ce normal ? Oui le routeur dispose de **deux adresses IPs**, chacune associée à une de ses deux interfaces réseaux, donc il répond **mais n'accepte pas un paquet provenant de la mauvaise interface** (celle qui n'est pas raccordée au réseau d'entrée du paquet de h1).

⇒ **l'isolation est bonne, mais pourrait être meilleure mais r1 est un routeur.**

Et si on recommence, en essayant vers h2 :

```
xterm
pef@palomino:~/NET_LAB_VLAN$ [h1] ip 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
        valid_lft forever preferred_lft forever
58: h1-eth0@if57: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
    link/ether d2:cc:44:e8:3c:6d brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 192.168.10.1/24 scope global h1-eth0
        valid_lft forever preferred_lft forever
    inet 10.0.0.42/24 scope global h1-eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::d0cc:44ff:fee8:3c6d/64 scope link
        valid_lft forever preferred_lft forever
pef@palomino:~/NET_LAB_VLAN$ [h1] sudo ip a add 192.168.20.42/24 dev h1-eth0
pef@palomino:~/NET_LAB_VLAN$ [h1] ip 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
        valid_lft forever preferred_lft forever
58: h1-eth0@if57: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
    link/ether d2:cc:44:e8:3c:6d brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 192.168.10.1/24 scope global h1-eth0
        valid_lft forever preferred_lft forever
    inet 192.168.20.42/24 scope global h1-eth0
        valid_lft forever preferred_lft forever
    inet 10.0.0.42/24 scope global h1-eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::d0cc:44ff:fee8:3c6d/64 scope link
        valid_lft forever preferred_lft forever
pef@palomino:~/NET_LAB_VLAN$ [h1] ip r
default via 192.168.10.254 dev h1-eth0
10.0.0.0/24 dev h1-eth0 proto kernel scope link src 10.0.0.42
192.168.10.0/24 dev h1-eth0 proto kernel scope link src 192.168.10.1
192.168.20.0/24 dev h1-eth0 proto kernel scope link src 192.168.20.42
pef@palomino:~/NET_LAB_VLAN$ [h1] ping -c 1 192.168.20.1
PING 192.168.20.1 (192.168.20.1) 56(84) bytes of data.
From 192.168.20.42 icmp_seq=1 Destination Host Unreachable

--- 192.168.20.1 ping statistics ---
1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms
```

Rien ne passe et si on observe le trafic :

```
xterm
23:20:56.718269 d2:cc:44:e8:3c:6d > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806), length 42: Ethernet (len 6), IPv4 (len 4), Request who-has 192.168.20.1 tell 192.168.20.42, length 28
    0x0000: 0001 0800 0604 0001 d2cc 44e8 3c6d c0a8 .....D.<m..
    0x0010: 142a 0000 0000 0000 c0a8 1401 .....*.....
23:20:57.732774 d2:cc:44:e8:3c:6d > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806), length 42: Ethernet (len 6), IPv4 (len 4), Request who-has 192.168.20.1 tell 192.168.20.42, length 28
    0x0000: 0001 0800 0604 0001 d2cc 44e8 3c6d c0a8 .....D.<m..
    0x0010: 142a 0000 0000 0000 c0a8 1401 .....*.....
23:20:58.756571 d2:cc:44:e8:3c:6d > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806), length 42: Ethernet (len 6), IPv4 (len 4), Request who-has 192.168.20.1 tell 192.168.20.42, length 28
    0x0000: 0001 0800 0604 0001 d2cc 44e8 3c6d c0a8 .....D.<m..
    0x0010: 142a 0000 0000 0000 c0a8 1401 .....*.....
```

⇒ **pas de réponse ARP depuis h2 : l'isolation est bonne !!!**

Essayons d'aller d'un VLAN à un autre en ajoutant une encapsulation VLAN sur h1, pour rejoindre h2 :

```
xterm
sudo ip l add link h1-eth0 name vers_vlan200 type vlan id 200
sudo ip l set vers_vlan200 up
sudo ip a add 192.168.20.42/24 dev vers_vlan200
```

on supprimera avant, l'adresse 192.168.20.42 de l'interface h1-eth0.

On vérifie la configuration :

```
xterm
pef@palomino:~/NET_LAB_VLAN$ [h1] ip l
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT
group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: vers_vlan200@h1-eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue
state UP mode DEFAULT group default qlen 1000
    link/ether be:f3:b3:20:e0:96 brd ff:ff:ff:ff:ff:ff
10: h1-eth0@if9: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
mode DEFAULT group default qlen 1000
    link/ether be:f3:b3:20:e0:96 brd ff:ff:ff:ff:ff:ff link-netnsid 0

pef@palomino:~/NET_LAB_VLAN$ [h1] ip 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
        valid_lft forever preferred_lft forever
2: vers_vlan200@h1-eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue
state UP group default qlen 1000
    link/ether be:f3:b3:20:e0:96 brd ff:ff:ff:ff:ff:ff
    inet 192.168.20.42/24 scope global vers_vlan200
        valid_lft forever preferred_lft forever
    inet6 fe80::bcf3:b3ff:fe20:e096/64 scope link
        valid_lft forever preferred_lft forever
10: h1-eth0@if9: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
group default qlen 1000
    link/ether be:f3:b3:20:e0:96 brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 192.168.10.1/24 scope global h1-eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::bcf3:b3ff:fe20:e096/64 scope link
        valid_lft forever preferred_lft forever

pef@palomino:~/NET_LAB_VLAN$ [h1] ip r
default via 192.168.10.254 dev h1-eth0
192.168.10.0/24 dev h1-eth0 proto kernel scope link src 192.168.10.1
192.168.20.0/24 dev vers_vlan200 proto kernel scope link src 192.168.20.42
```

Si on essaye de 'ping' vers la machine h2 :

```
xterm
pef@palomino:~/NET_LAB_VLAN$ [h1] ping -c 1 192.168.20.1
PING 192.168.20.1 (192.168.20.1) 56(84) bytes of data.
From 192.168.20.42 icmp_seq=1 Destination Host Unreachable

--- 192.168.20.1 ping statistics ---
1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms
```

⇒ **ça ne marche pas !**

Si on sniffe sur l'interface h1-eth0 :

```
xterm
pef@palomino:~$ [h1] sudo tcpdump -lnvveX -i h1-eth0 not ip6
tcpdump: listening on h1-eth0, link-type EN10MB (Ethernet), snapshot length 262144
bytes
22:24:32.311172 be:f3:b3:20:e0:96 > ff:ff:ff:ff:ff:ff, ethertype 802.1Q (0x8100),
length 46: vlan 200, p 0, ethertype ARP (0x0806), Ethernet (len 6), IPv4 (len 4),
Request who-has 192.168.20.1 tell 192.168.20.42, length 28
    0x0000: 0001 0800 0604 0001 bef3 b320 e096 c0a8 .....
    0x0010: 142a 0000 0000 0000 c0a8 1401 .....
22:24:33.316526 be:f3:b3:20:e0:96 > ff:ff:ff:ff:ff:ff, ethertype 802.1Q (0x8100),
length 46: vlan 200, p 0, ethertype ARP (0x0806), Ethernet (len 6), IPv4 (len 4),
Request who-has 192.168.20.1 tell 192.168.20.42, length 28
    0x0000: 0001 0800 0604 0001 bef3 b320 e096 c0a8 .....
    0x0010: 142a 0000 0000 0000 c0a8 1401 .....

```

On voit bien l'étiquette VLAN ajoutée à la trame ARP, mais il n'y a pas de réponse

⇒ le paquet ne rejoint pas le VLAN 200...

⇒ **Notre isolation fonctionne.**

■ ■ ■ Seconde version de configuration VLAN : on va autoriser des paquets taggés VLAN

On modifie la connexion de h1 sur le « switch » dans le fichier « build_architecture » :

```
ovs-vsctl add-port switch s1-h1 tag=100 vlan_mode=native-tagged
```

On configure sur h1, l'encapsulation VLAN pour rejoindre h2 :

```
xterm
pef@palomino:~/NET_LAB_VLAN$ [h1] sudo ip l add link h1-eth0 name vers_vlan200 type
vlan id 200
pef@palomino:~/NET_LAB_VLAN$ [h1] sudo ip l set vers_vlan200 up
pef@palomino:~/NET_LAB_VLAN$ [h1] sudo ip a add 192.168.20.42/24 dev vers_vlan200
pef@palomino:~/NET_LAB_VLAN$ [h1] ping -c 1 192.168.20.1
PING 192.168.20.1 (192.168.20.1) 56(84) bytes of data.
64 bytes from 192.168.20.1: icmp_seq=1 ttl=64 time=1.40 ms

--- 192.168.20.1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 1.396/1.396/1.396/0.000 ms
```

Dans une autre fenêtre sur h1 :

```
xterm
pef@palomino:~$ [h1] sudo tcpdump -lnvveX -i h1-eth0 not ip6
tcpdump: listening on h1-eth0, link-type EN10MB (Ethernet), snapshot length 262144
bytes
22:52:21.391105 d2:cc:44:e8:3c:6d > ff:ff:ff:ff:ff:ff, ethertype 802.1Q (0x8100),
length 46: vlan 200, p 0, ethertype ARP (0x0806), Ethernet (len 6), IPv4 (len 4),
Request who-has 192.168.20.1 tell 192.168.20.42, length 28
0x0000: 0001 0800 0604 0001 d2cc 44e8 3c6d c0a8 .....D.<m..
0x0010: 142a 0000 0000 0000 c0a8 1401 .....*.....
22:52:21.391919 5e:6d:ee:ba:fc:7b > d2:cc:44:e8:3c:6d, ethertype 802.1Q (0x8100),
length 46: vlan 200, p 0, ethertype ARP (0x0806), Ethernet (len 6), IPv4 (len 4),
Reply 192.168.20.1 is-at 5e:6d:ee:ba:fc:7b, length 28
0x0000: 0001 0800 0604 0002 5e6d ee ba fc7b c0a8 .....^m...{..
0x0010: 1401 d2cc 44e8 3c6d c0a8 142a .....D.<m...*
22:52:21.391933 d2:cc:44:e8:3c:6d > 5e:6d:ee:ba:fc:7b, ethertype 802.1Q (0x8100),
length 102: vlan 200, p 0, ethertype IPv4 (0x0800), (tos 0x0, ttl 64, id 42702,
offset 0, flags [DF], proto ICMP (1), length 84)
192.168.20.42 > 192.168.20.1: ICMP echo request, id 46535, seq 1, length 64
0x0000: 4500 0054 a6ce 4000 4001 ea5e c0a8 142a E..T..@.@..^...*
0x0010: c0a8 1401 0800 93a7 b5c7 0001 1560 4965 .....`Ie
0x0020: 0000 0000 8bf7 0500 0000 0000 1011 1213 .....
0x0030: 1415 1617 1819 1a1b 1c1d 1e1f 2021 2223 .....!"#
0x0040: 2425 2627 2829 2a2b 2c2d 2e2f 3031 3233 $%&'()*+,-./0123
0x0050: 3435 3637 .....4567
22:52:21.392447 5e:6d:ee:ba:fc:7b > d2:cc:44:e8:3c:6d, ethertype 802.1Q (0x8100),
length 102: vlan 200, p 0, ethertype IPv4 (0x0800), (tos 0x0, ttl 64, id 7758,
offset 0, flags [none], proto ICMP (1), length 84)
192.168.20.1 > 192.168.20.42: ICMP echo reply, id 46535, seq 1, length 64
0x0000: 4500 0054 1e4e 0000 4001 b2df c0a8 1401 E..T.N..@.....
0x0010: c0a8 142a 0000 9ba7 b5c7 0001 1560 4965 .....*.....`Ie
0x0020: 0000 0000 8bf7 0500 0000 0000 1011 1213 .....
0x0030: 1415 1617 1819 1a1b 1c1d 1e1f 2021 2223 .....!"#
0x0040: 2425 2627 2829 2a2b 2c2d 2e2f 3031 3233 $%&'()*+,-./0123
0x0050: 3435 3637 .....4567
```

Ici, le trafic est encapsulé dans des trames VLANs ⇒ ça marche !

Par contre, si on essaye de "ping" une patte externe du routeur :

```
xterm
pef@palomino:~/NET_LAB_VLAN$ [h1] ping -c 1 10.0.0.1
PING 10.0.0.1 (10.0.0.1) 56(84) bytes of data.
From 192.168.10.1 icmp_seq=1 Destination Host Unreachable

--- 10.0.0.1 ping statistics ---
1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms

22:56:01.873244 d2:cc:44:e8:3c:6d > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Request who-has 192.168.10.254 tell
192.168.10.1, length 28
0x0000: 0001 0800 0604 0001 d2cc 44e8 3c6d c0a8 .....D.<m..
0x0010: 0a01 0000 0000 0000 c0a8 0afe .....
22:56:01.873855 66:f1:76:90:a9:21 > d2:cc:44:e8:3c:6d, ethertype 802.1Q (0x8100),
length 46: vlan 100, p 0, ethertype ARP (0x0806), Ethernet (len 6), IPv4 (len 4),
Reply 192.168.10.254 is-at 66:f1:76:90:a9:21, length 28
0x0000: 0001 0800 0604 0002 66f1 7690 a921 c0a8 .....f.v...!..
0x0010: 0afe d2cc 44e8 3c6d c0a8 0a01 .....D.<m....
```

La machine h1 n'interprète pas la trame avec l'encapsulation VLAN

Ici, le port se comporte en "trunk", mais les paquets sortant vers h1 (ceux du routeur) reçoivent une étiquette VLAN, ce qui perturbe h1 !

■ ■ ■ Troisième version de configuration VLAN : on va autoriser des paquets taggés VLAN

On modifie la connexion de h1 sur le « switch » dans le fichier « build_architecture » :

```
ovs-vsctl add-port switch s1-h1 tag=100 vlan_mode=native-untagged
```

On configure l'encapsulation VLAN sur h1 pour rejoindre h2 :

```
xterm
pef@palomino:~/NET_LAB_VLAN$ [h1] sudo ip l add link h1-eth0 name vers_vlan200 type
vlan id 200
pef@palomino:~/NET_LAB_VLAN$ [h1] sudo ip l set vers_vlan200 up
pef@palomino:~/NET_LAB_VLAN$ [h1] sudo ip a add 192.168.20.42/24 dev vers_vlan200
pef@palomino:~/NET_LAB_VLAN$ [h1] ping -c 1 192.168.20.1
PING 192.168.20.1 (192.168.20.1) 56(84) bytes of data.
64 bytes from 192.168.20.1: icmp_seq=1 ttl=64 time=1.18 ms

--- 192.168.20.1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 1.176/1.176/1.176/0.000 ms
pef@palomino:~/NET_LAB_VLAN$ [h1] ping -c 1 10.0.0.1
PING 10.0.0.1 (10.0.0.1) 56(84) bytes of data.
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=1.09 ms

--- 10.0.0.1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 1.090/1.090/1.090/0.000 ms
```

Tout fonctionne !

La capture du trafic dans une autre fenêtre :

```

xterm
pef@palomino:~$ [h1] sudo tcpdump -lnvveX -i h1-eth0 not ip6
tcpdump: listening on h1-eth0, link-type EN10MB (Ethernet), snapshot length 262144
bytes
23:06:00.241368 d2:cc:44:e8:3c:6d > ff:ff:ff:ff:ff:ff, ethertype 802.1Q (0x8100),
length 46: vlan 200, p 0, ethertype ARP (0x0806), Ethernet (len 6), IPv4 (len 4),
Request who-has 192.168.20.1 tell 192.168.20.42, length 28
0x0000: 0001 0800 0604 0001 d2cc 44e8 3c6d c0a8 .....D.<m..
0x0010: 142a 0000 0000 0000 c0a8 1401 .....*.....
23:06:00.242007 5e:6d:ee:ba:fc:7b > d2:cc:44:e8:3c:6d, ethertype 802.1Q (0x8100),
length 46: vlan 200, p 0, ethertype ARP (0x0806), Ethernet (len 6), IPv4 (len 4),
Reply 192.168.20.1 is-at 5e:6d:ee:ba:fc:7b, length 28
0x0000: 0001 0800 0604 0002 5e6d ee8a fc7b c0a8 .....^m...{..
0x0010: 1401 d2cc 44e8 3c6d c0a8 142a .....D.<m...*
23:06:00.242023 d2:cc:44:e8:3c:6d > 5e:6d:ee:ba:fc:7b, ethertype 802.1Q (0x8100),
length 102: vlan 200, p 0, ethertype IPv4 (0x0800), (tos 0x0, ttl 64, id 23291,
offset 0, flags [DF], proto ICMP (1), length 84)
192.168.20.42 > 192.168.20.1: ICMP echo request, id 10437, seq 1, length 64
0x0000: 4500 0054 5afb 4000 4001 3632 c0a8 142a E..TZ.@.62...*
0x0010: c0a8 1401 0800 d1ef 28c5 0001 4863 4965 .....(...HcIe
0x0020: 0000 0000 a9ae 0300 0000 0000 1011 1213 .....
0x0030: 1415 1617 1819 1a1b 1c1d 1e1f 2021 2223 .....!"#
0x0040: 2425 2627 2829 2a2b 2c2d 2e2f 3031 3233 $%&'()*+,-./0123
0x0050: 3435 3637 4567
23:06:00.242497 5e:6d:ee:ba:fc:7b > d2:cc:44:e8:3c:6d, ethertype 802.1Q (0x8100),
length 102: vlan 200, p 0, ethertype IPv4 (0x0800), (tos 0x0, ttl 64, id 21369,
offset 0, flags [none], proto ICMP (1), length 84)
192.168.20.1 > 192.168.20.42: ICMP echo reply, id 10437, seq 1, length 64
0x0000: 4500 0054 5379 0000 4001 7db4 c0a8 1401 E..TSy..@.).....
0x0010: c0a8 142a 0000 d9ef 28c5 0001 4863 4965 ...*....(...HcIe
0x0020: 0000 0000 a9ae 0300 0000 0000 1011 1213 .....
0x0030: 1415 1617 1819 1a1b 1c1d 1e1f 2021 2223 .....!"#
0x0040: 2425 2627 2829 2a2b 2c2d 2e2f 3031 3233 $%&'()*+,-./0123
0x0050: 3435 3637 4567
23:06:05.348631 5e:6d:ee:ba:fc:7b > d2:cc:44:e8:3c:6d, ethertype 802.1Q (0x8100),
length 46: vlan 200, p 0, ethertype ARP (0x0806), Ethernet (len 6), IPv4 (len 4),
Request who-has 192.168.20.42 tell 192.168.20.1, length 28
0x0000: 0001 0800 0604 0001 5e6d ee8a fc7b c0a8 .....^m...{..
0x0010: 1401 0000 0000 0000 c0a8 142a .....*
23:06:05.348656 d2:cc:44:e8:3c:6d > 5e:6d:ee:ba:fc:7b, ethertype 802.1Q (0x8100),
length 46: vlan 200, p 0, ethertype ARP (0x0806), Ethernet (len 6), IPv4 (len 4),
Reply 192.168.20.42 is-at d2:cc:44:e8:3c:6d, length 28
0x0000: 0001 0800 0604 0002 d2cc 44e8 3c6d c0a8 .....D.<m..
0x0010: 142a 5e6d ee8a fc7b c0a8 1401 .....*^m...{....
23:06:09.253030 d2:cc:44:e8:3c:6d > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Request who-has 192.168.10.254 tell
192.168.10.1, length 28
0x0000: 0001 0800 0604 0001 d2cc 44e8 3c6d c0a8 .....D.<m..
0x0010: 0a01 0000 0000 0000 c0a8 0afe .....
23:06:09.253631 66:f1:76:90:a9:21 > d2:cc:44:e8:3c:6d, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Reply 192.168.10.254 is-at
66:f1:76:90:a9:21, length 28
0x0000: 0001 0800 0604 0002 66f1 7690 a921 c0a8 .....f.v...!..
0x0010: 0afe d2cc 44e8 3c6d c0a8 0a01 .....D.<m....
23:06:09.253644 d2:cc:44:e8:3c:6d > 66:f1:76:90:a9:21, ethertype IPv4 (0x0800),
length 98: (tos 0x0, ttl 64, id 6244, offset 0, flags [DF], proto ICMP (1), length
84)
192.168.10.1 > 10.0.0.1: ICMP echo request, id 52196, seq 1, length 64
0x0000: 4500 0054 1864 4000 4001 4d9b c0a8 0a01 E..T.d@.@.M.....
0x0010: 0a00 0001 0800 9ea2 cbe4 0001 5163 4965 .....QcIe
0x0020: 0000 0000 38dc 0300 0000 0000 1011 1213 ....8.....
0x0030: 1415 1617 1819 1a1b 1c1d 1e1f 2021 2223 .....!"#
0x0040: 2425 2627 2829 2a2b 2c2d 2e2f 3031 3233 $%&'()*+,-./0123
0x0050: 3435 3637 4567
23:06:09.254074 66:f1:76:90:a9:21 > d2:cc:44:e8:3c:6d, ethertype IPv4 (0x0800),
length 98: (tos 0x0, ttl 64, id 761, offset 0, flags [none], proto ICMP (1), length
84)
10.0.0.1 > 192.168.10.1: ICMP echo reply, id 52196, seq 1, length 64
0x0000: 4500 0054 02f9 0000 4001 a306 0a00 0001 E..T....@.....
0x0010: c0a8 0a01 0000 9ea2 cbe4 0001 5163 4965 .....QcIe
0x0020: 0000 0000 38dc 0300 0000 0000 1011 1213 ....8.....
0x0030: 1415 1617 1819 1a1b 1c1d 1e1f 2021 2223 .....!"#
0x0040: 2425 2627 2829 2a2b 2c2d 2e2f 3031 3233 $%&'()*+,-./0123
0x0050: 3435 3637 4567
23:06:14.309074 66:f1:76:90:a9:21 > d2:cc:44:e8:3c:6d, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Request who-has 192.168.10.1 tell
192.168.10.254, length 28
0x0000: 0001 0800 0604 0001 66f1 7690 a921 c0a8 .....f.v...!..
0x0010: 0afe 0000 0000 0000 c0a8 0a01 .....
23:06:14.309096 d2:cc:44:e8:3c:6d > 66:f1:76:90:a9:21, ethertype ARP (0x0806),
length 42: Ethernet (len 6), IPv4 (len 4), Reply 192.168.10.1 is-at
d2:cc:44:e8:3c:6d, length 28
0x0000: 0001 0800 0604 0002 d2cc 44e8 3c6d c0a8 .....D.<m..
0x0010: 0a01 66f1 7690 a921 c0a8 0afe ...f.v...!.....

```

Mais peut-on atteindre le réseau des routeurs depuis h1 ?

On configure une nouvelle encapsulation VLAN pour rejoindre le réseau de routeurs 10.0.0.0/24 :

```
xterm
pef@palomino:~/NET_LAB_VLAN$ [h1] sudo ip l add link h1-eth0 name vers_vlan300 type
vlan id 300
pef@palomino:~/NET_LAB_VLAN$ [h1] sudo ip l set vers_vlan300 up
pef@palomino:~/NET_LAB_VLAN$ [h1] sudo ip a add 10.0.0.42/24 dev vers_vlan300
pef@palomino:~/NET_LAB_VLAN$ [h1] ping -c 1 10.0.0.1
PING 10.0.0.1 (10.0.0.1) 56(84) bytes of data.
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=1.28 ms

--- 10.0.0.1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 1.282/1.282/1.282/0.000 ms
```

Bien sûr !

La capture de trafic :

```
xterm
23:10:42.374270 d2:cc:44:e8:3c:6d > ff:ff:ff:ff:ff:ff, ethertype 802.1Q (0x8100),
length 46: vlan 300, p 0, ethertype ARP (0x0806), Ethernet (len 6), IPv4 (len 4),
Request who-has 10.0.0.1 tell 10.0.0.42, length 28
0x0000: 0001 0800 0604 0001 d2cc 44e8 3c6d 0a00 .....D.<m..
0x0010: 002a 0000 0000 0000 0a00 0001 .....*.....
23:10:42.374974 ca:10:a1:18:e0:b2 > d2:cc:44:e8:3c:6d, ethertype 802.1Q (0x8100),
length 46: vlan 300, p 0, ethertype ARP (0x0806), Ethernet (len 6), IPv4 (len 4),
Reply 10.0.0.1 is-at ca:10:a1:18:e0:b2, length 28
0x0000: 0001 0800 0604 0002 ca10 a118 e0b2 0a00 .....
0x0010: 0001 d2cc 44e8 3c6d 0a00 002a ....D.<m....*
23:10:42.374990 d2:cc:44:e8:3c:6d > ca:10:a1:18:e0:b2, ethertype 802.1Q (0x8100),
length 102: vlan 300, p 0, ethertype IPv4 (0x0800), (tos 0x0, ttl 64, id 28012,
offset 0, flags [DF], proto ICMP (1), length 84)
10.0.0.42 > 10.0.0.1: ICMP echo request, id 10637, seq 1, length 64
0x0000: 4500 0054 6d6c 4000 4001 b912 0a00 002a E..Tm1@.@.....*
0x0010: 0a00 0001 0800 961f 298d 0001 6264 4965 .....)...bdIe
0x0020: 0000 0000 c8b5 0500 0000 0000 1011 1213 .....
0x0030: 1415 1617 1819 1a1b 1c1d 1e1f 2021 2223 .....!"#
0x0040: 2425 2627 2829 2a2b 2c2d 2e2f 3031 3233 $%&'()*+,-./0123
0x0050: 3435 3637 4567
23:10:42.375498 ca:10:a1:18:e0:b2 > d2:cc:44:e8:3c:6d, ethertype 802.1Q (0x8100),
length 102: vlan 300, p 0, ethertype IPv4 (0x0800), (tos 0x0, ttl 64, id 28065,
offset 0, flags [none], proto ICMP (1), length 84)
10.0.0.1 > 10.0.0.42: ICMP echo reply, id 10637, seq 1, length 64
0x0000: 4500 0054 6da1 0000 4001 f8dd 0a00 0001 E..Tm...@.....
0x0010: 0a00 002a 0000 9e1f 298d 0001 6264 4965 ...*....)...bdIe
0x0020: 0000 0000 c8b5 0500 0000 0000 1011 1213 .....
0x0030: 1415 1617 1819 1a1b 1c1d 1e1f 2021 2223 .....!"#
0x0040: 2425 2627 2829 2a2b 2c2d 2e2f 3031 3233 $%&'()*+,-./0123
0x0050: 3435 3637 4567
23:10:47.460684 ca:10:a1:18:e0:b2 > d2:cc:44:e8:3c:6d, ethertype 802.1Q (0x8100),
length 46: vlan 300, p 0, ethertype ARP (0x0806), Ethernet (len 6), IPv4 (len 4),
Request who-has 10.0.0.42 tell 10.0.0.1, length 28
0x0000: 0001 0800 0604 0001 ca10 a118 e0b2 0a00 .....
0x0010: 0001 0000 0000 0000 0a00 002a .....*
23:10:47.460716 d2:cc:44:e8:3c:6d > ca:10:a1:18:e0:b2, ethertype 802.1Q (0x8100),
length 46: vlan 300, p 0, ethertype ARP (0x0806), Ethernet (len 6), IPv4 (len 4),
Reply 10.0.0.42 is-at d2:cc:44:e8:3c:6d, length 28
0x0000: 0001 0800 0604 0002 d2cc 44e8 3c6d 0a00 .....D.<m..
0x0010: 002a ca10 a118 e0b2 0a00 0001 .....*.....
```

On a perdu toute isolation !!!

L'utilisateur peut faire du « VLAN Hoping » et sauter d'un VLAN à un autre.

⇒ Il faut indiquer au switch les étiquettes VLAN autorisées.

Pour connaître les adresses MAC respectives et interpréter correctement la capture de tcpdump :

```
xterm
pef@palomino:~/NET_LAB_VLAN$ [h1] ip l
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT
group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: vers_vlan200@h1-eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue
state UP mode DEFAULT group default qlen 1000
    link/ether d2:cc:44:e8:3c:6d brd ff:ff:ff:ff:ff:ff
3: vers_vlan300@h1-eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue
state UP mode DEFAULT group default qlen 1000
    link/ether d2:cc:44:e8:3c:6d brd ff:ff:ff:ff:ff:ff
42: h1-eth0@if41: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
mode DEFAULT group default qlen 1000
    link/ether d2:cc:44:e8:3c:6d brd ff:ff:ff:ff:ff:ff link-netnsid 0
pef@palomino:~/NET_LAB_VLAN$ [h1] ip 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
        valid_lft forever preferred_lft forever
2: vers_vlan200@h1-eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue
state UP group default qlen 1000
    link/ether d2:cc:44:e8:3c:6d brd ff:ff:ff:ff:ff:ff
    inet 192.168.20.42/24 scope global vers_vlan200
        valid_lft forever preferred_lft forever
    inet6 fe80::d0cc:44ff:fee8:3c6d/64 scope link
        valid_lft forever preferred_lft forever
3: vers_vlan300@h1-eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue
state UP group default qlen 1000
    link/ether d2:cc:44:e8:3c:6d brd ff:ff:ff:ff:ff:ff
    inet 10.0.0.42/24 scope global vers_vlan300
        valid_lft forever preferred_lft forever
    inet6 fe80::d0cc:44ff:fee8:3c6d/64 scope link
        valid_lft forever preferred_lft forever
42: h1-eth0@if41: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
group default qlen 1000
    link/ether d2:cc:44:e8:3c:6d brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 192.168.10.1/24 scope global h1-eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::d0cc:44ff:fee8:3c6d/64 scope link
        valid_lft forever preferred_lft forever
pef@palomino:~/NET_LAB_VLAN$ [h1] ip r
default via 192.168.10.254 dev h1-eth0
10.0.0.0/24 dev vers_vlan300 proto kernel scope link src 10.0.0.42
192.168.10.0/24 dev h1-eth0 proto kernel scope link src 192.168.10.1
192.168.20.0/24 dev vers_vlan200 proto kernel scope link src 192.168.20.42
```