Fully redundant networks

EU MUM 2016

Hello!

I am Mihai Săftoiu

MTCNA, MTCTCE, MTCWE,
MTCUME, MTCRE, MTCINE
MikroTik Trainer

Consulting & Training

A long time ago in a galaxy far, far away....

X-Wing Fighter

How can we make it work?

??

What is network redundancy?

Is redundant network redundancy ***redundant****?*

1. Network redundancy

Equipment redundancy

Path redundancy

Equipment redundancy

Path redundancy

Path redundancy

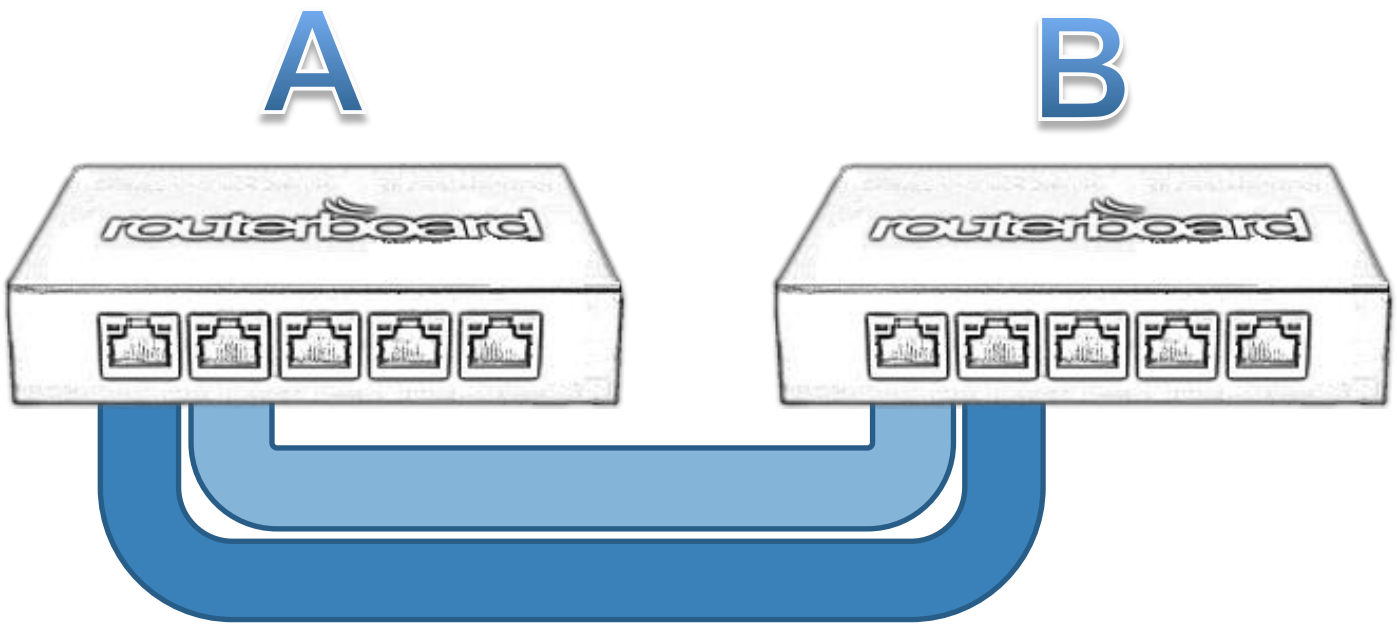

Upstream path redundancy

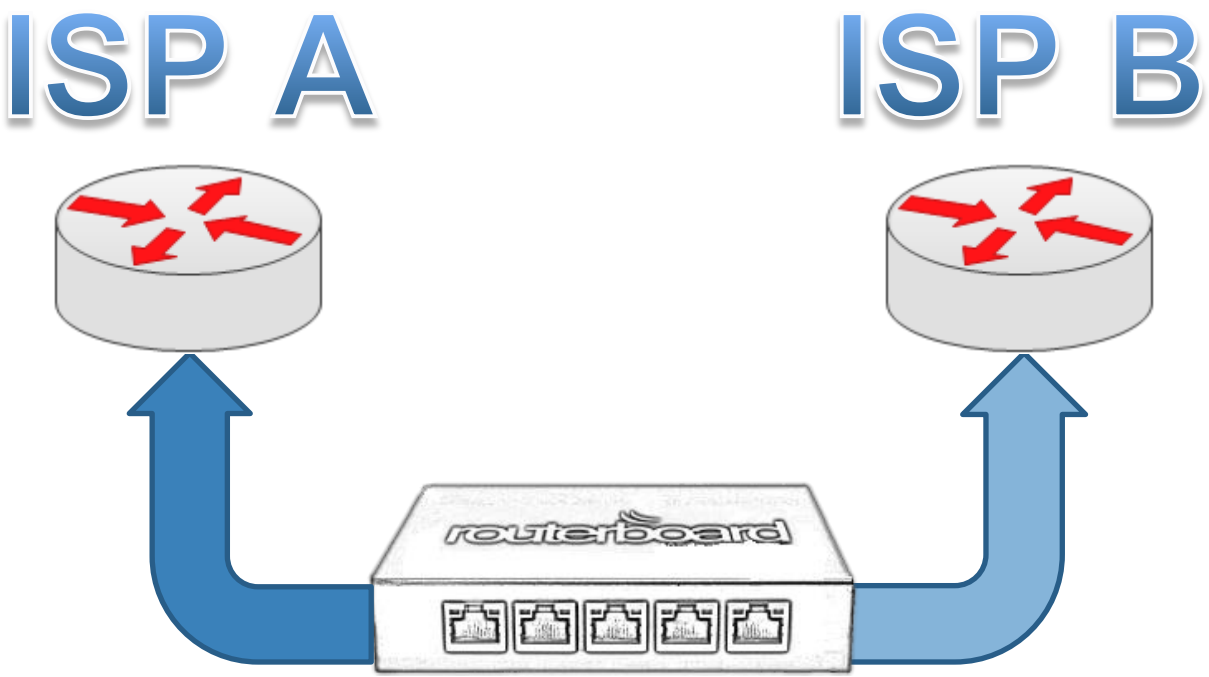

SLAs are based on calculations

Network uptime SLA:

This is not TIER based, as in case of data centers, but mathematically calculated.

Calculation based on 1%

For every single point of failure in the network we take 1% (thumb rule) of monthly downtime into account

2.
Non redundant networks

No equipment redundancy

No path redundancy

EXAMPLE

- 1% means: 30,5 days x 24h / 100 = **7,32h**
- 1% of downtime for device A
- 1% of downtime for device B
- 1% of downtime for the link

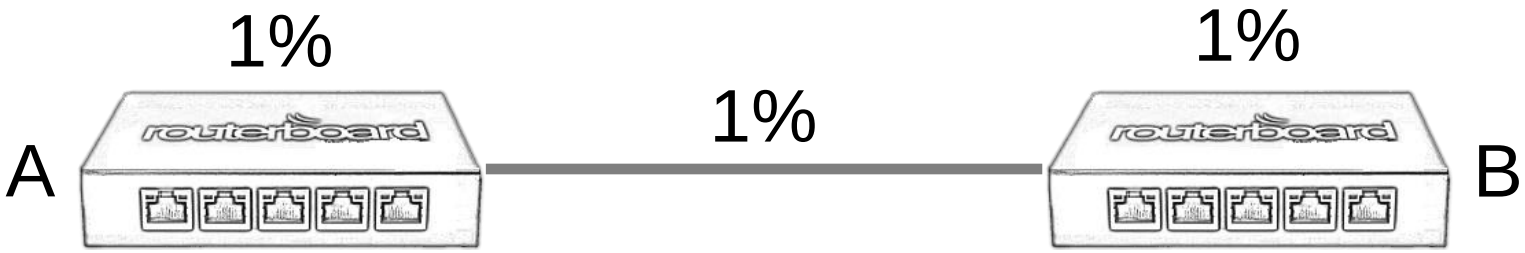

Max network downtime = 3%
(SLA = 97% guaranteed uptime)

3. Multiple upstream networks

No equipment redundancy

Partial upstream path redundancy

WHAT IF WE DO IT LIKE THIS?

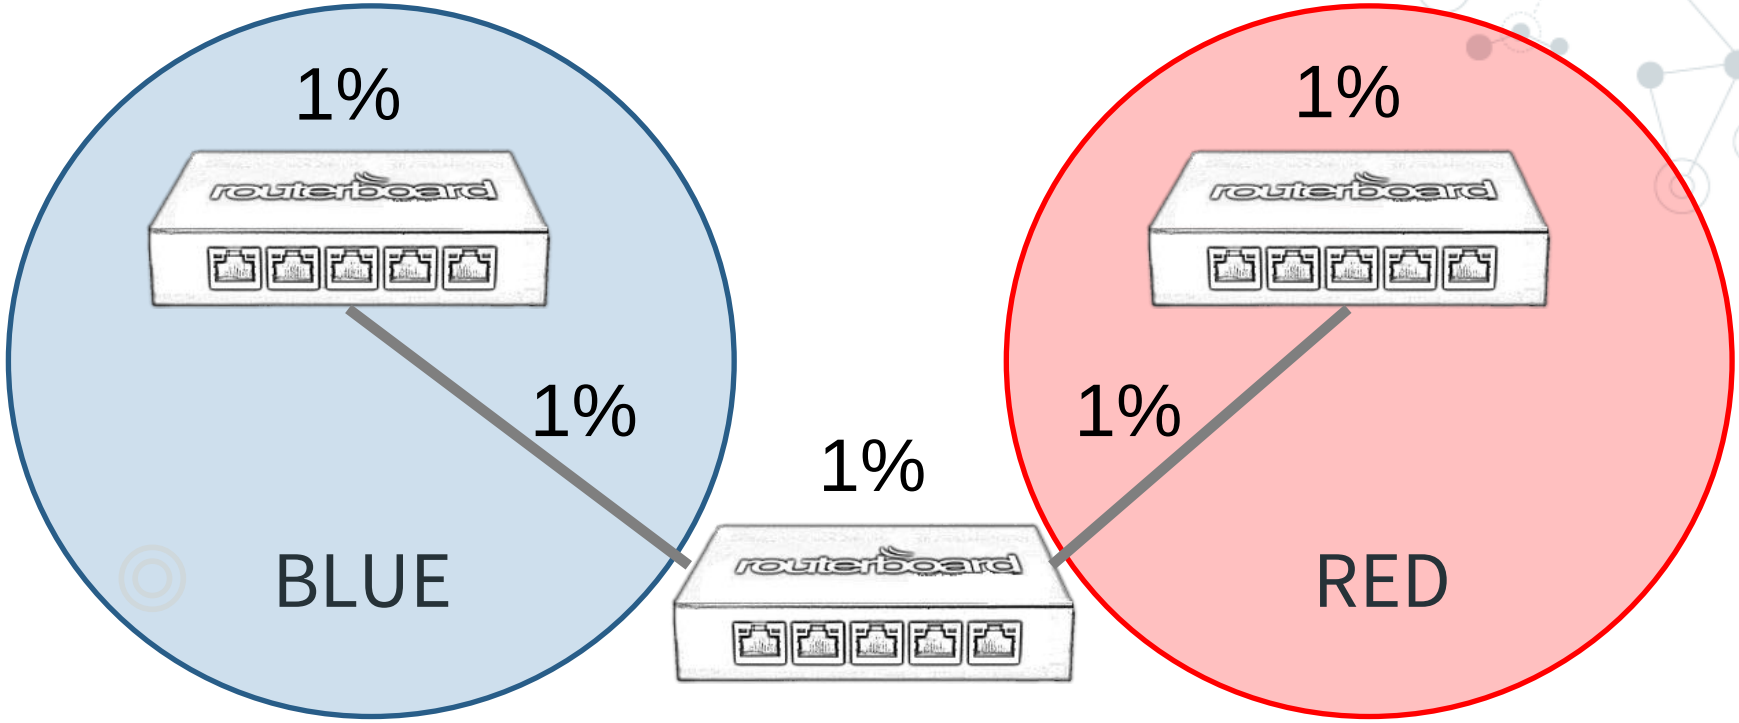

- BLUE max downtime = 2%
- RED max downtime = 2%

CALCULATION

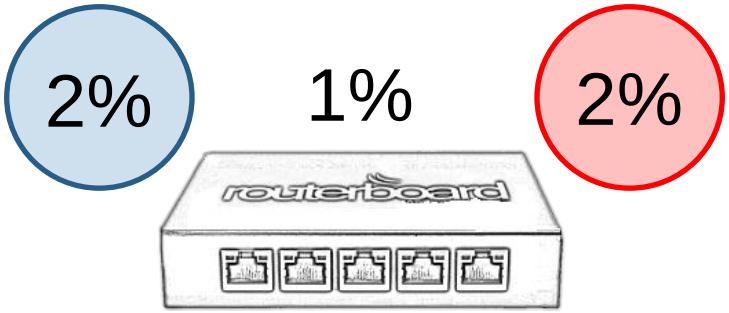

- The chance BLUE and RED will both be down is: 2% x 2% = 0,04%
- The worst case that RB and both ISPs will be down is: 1% + 0,04% = 1,04%

- SLA for the current network: 98,96%

4.

Partial mesh networks

Equipment redundancy

Multiple upstream path redundancy

WITH UP TO 3 POINTS OF FAILURE THIS NETWORK CAN FAIL IN 6 WAYS

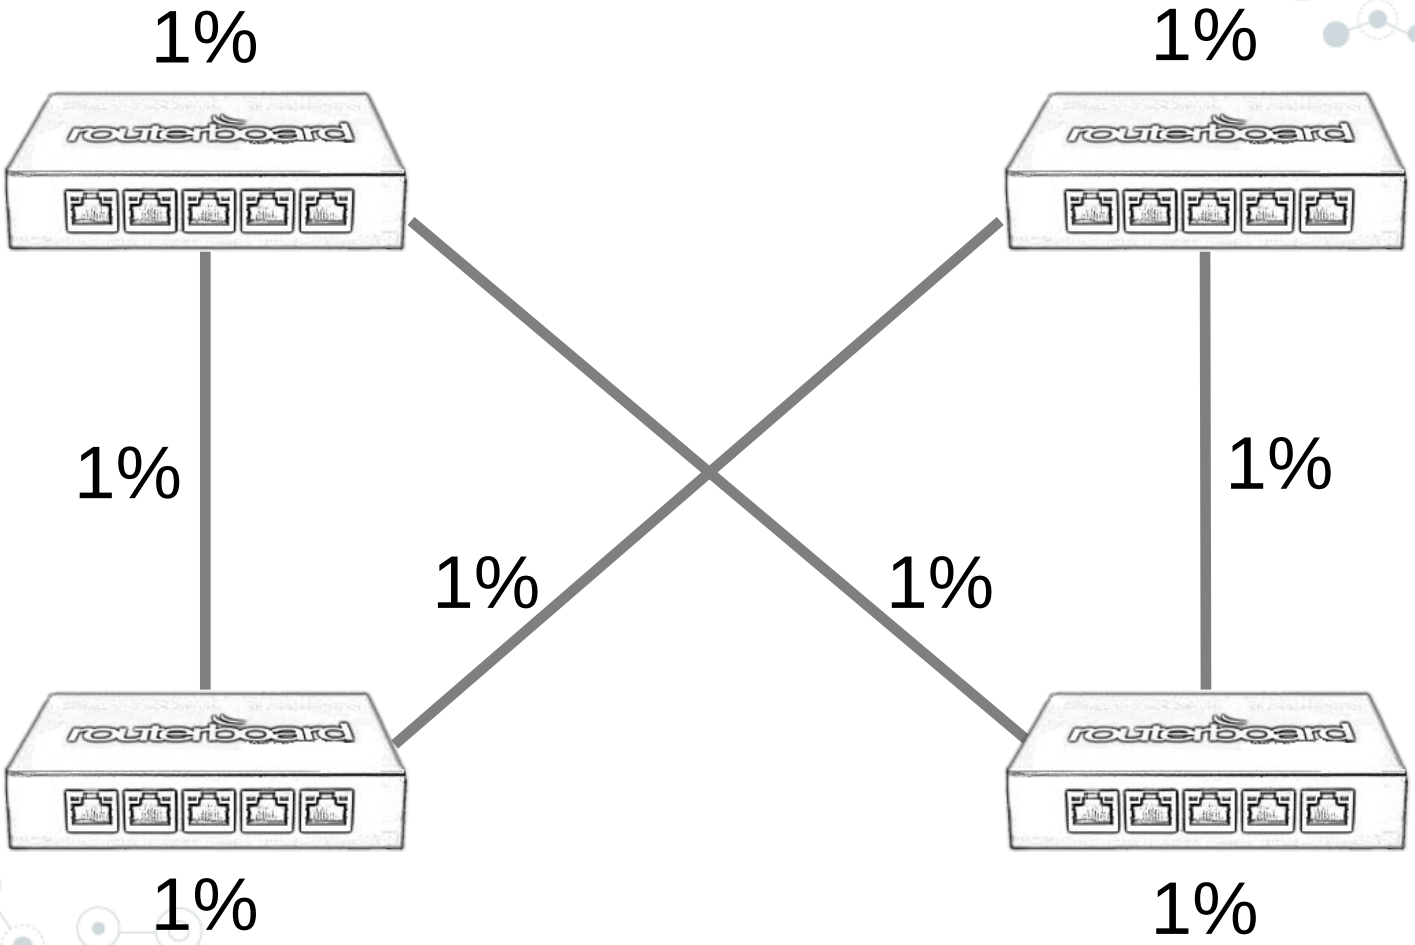

FAIL NO. 1

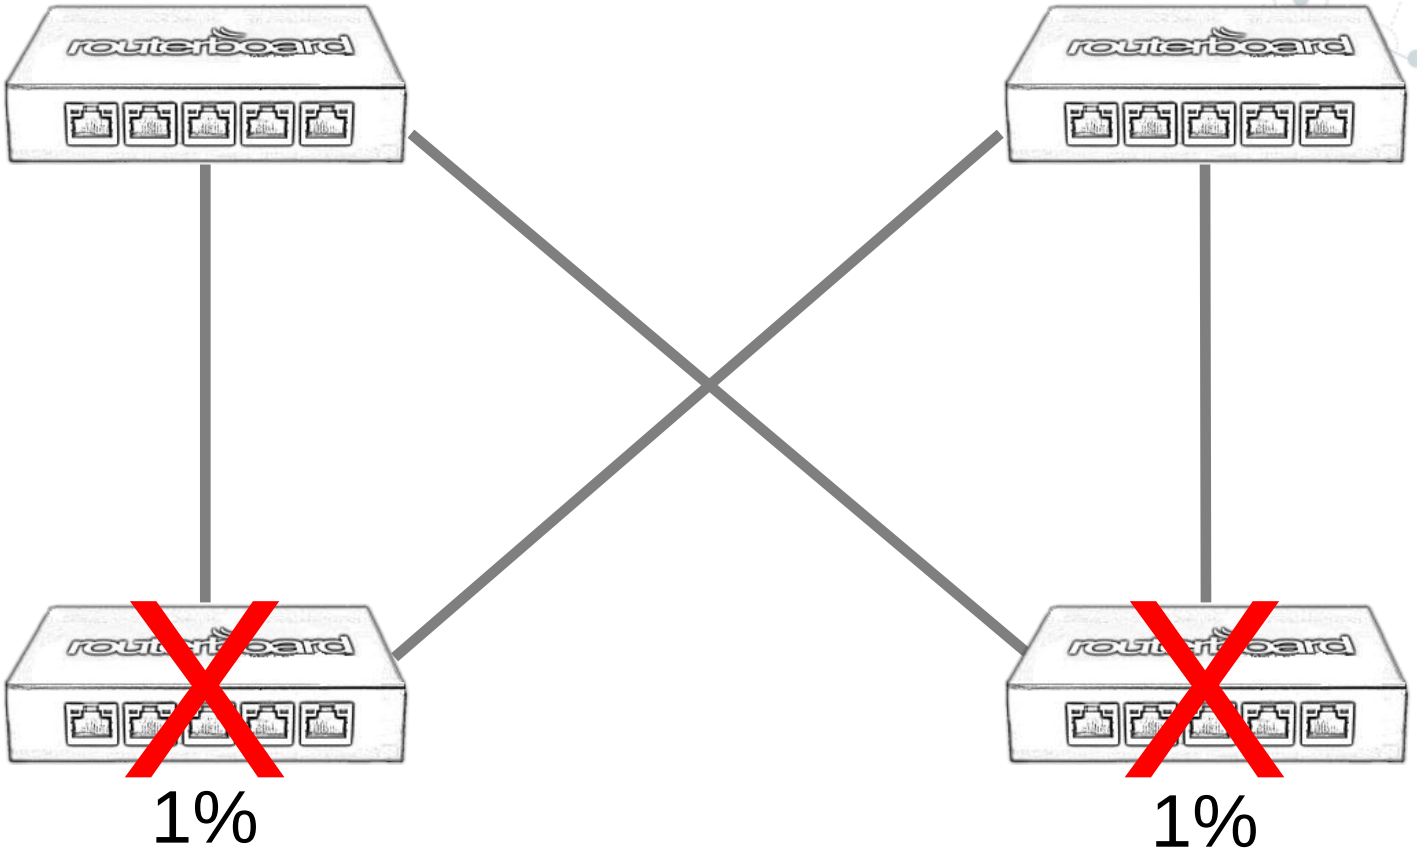

The chance for this to happen:

1% x 1% = 0,01%

FAIL NO. 2

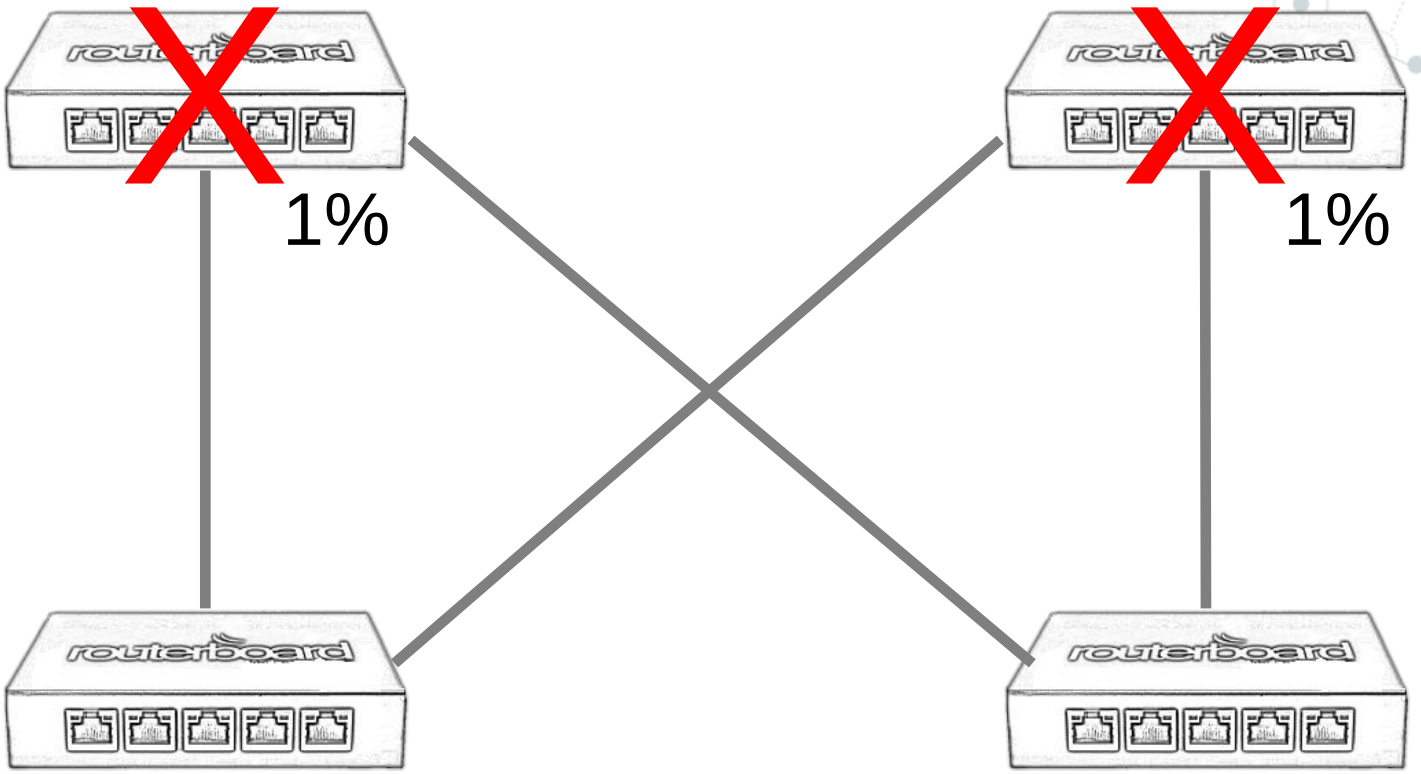

The chance for this to happen:

1% x 1% = 0,01%

FAIL NO. 3

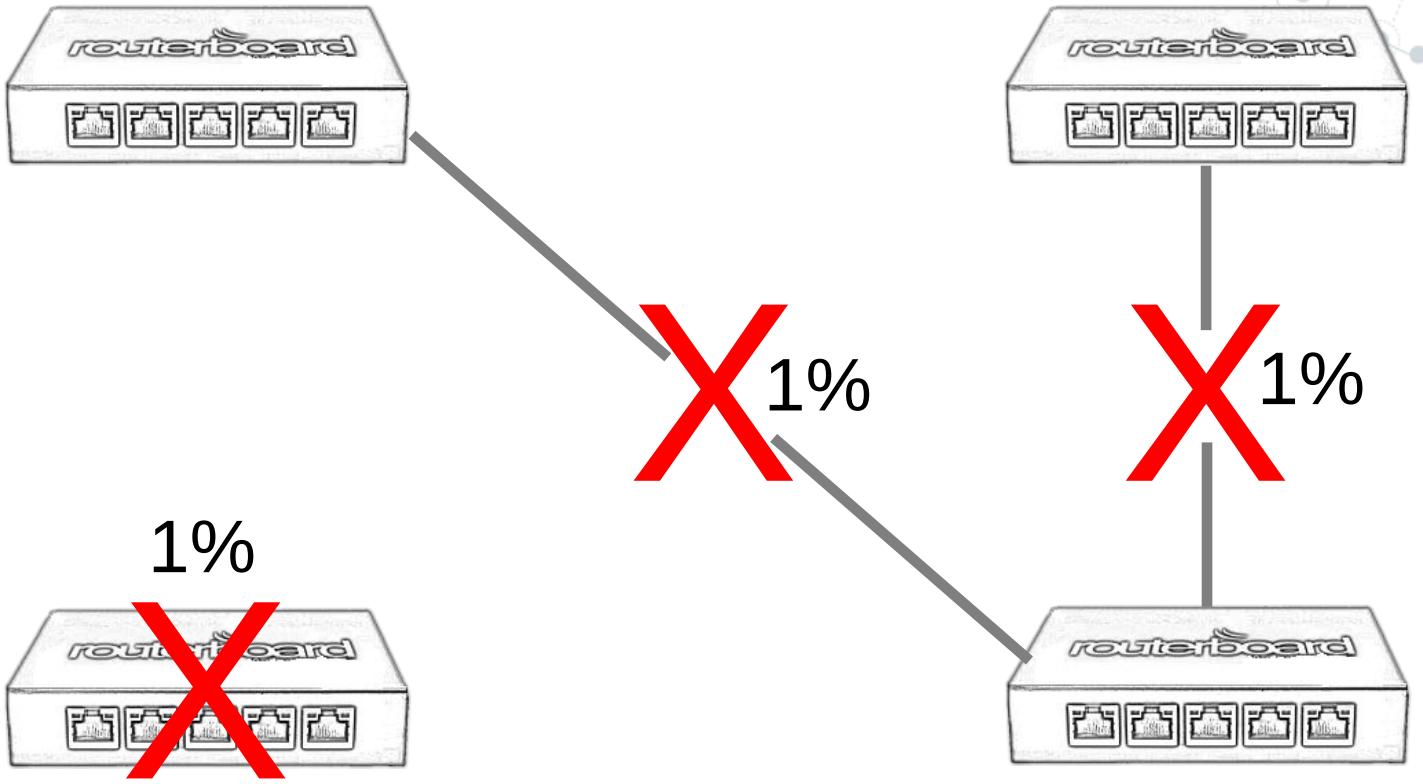

- The chance for this to happen:

  1% x 1% x 1% = 0,0001%

FAIL NO. 4

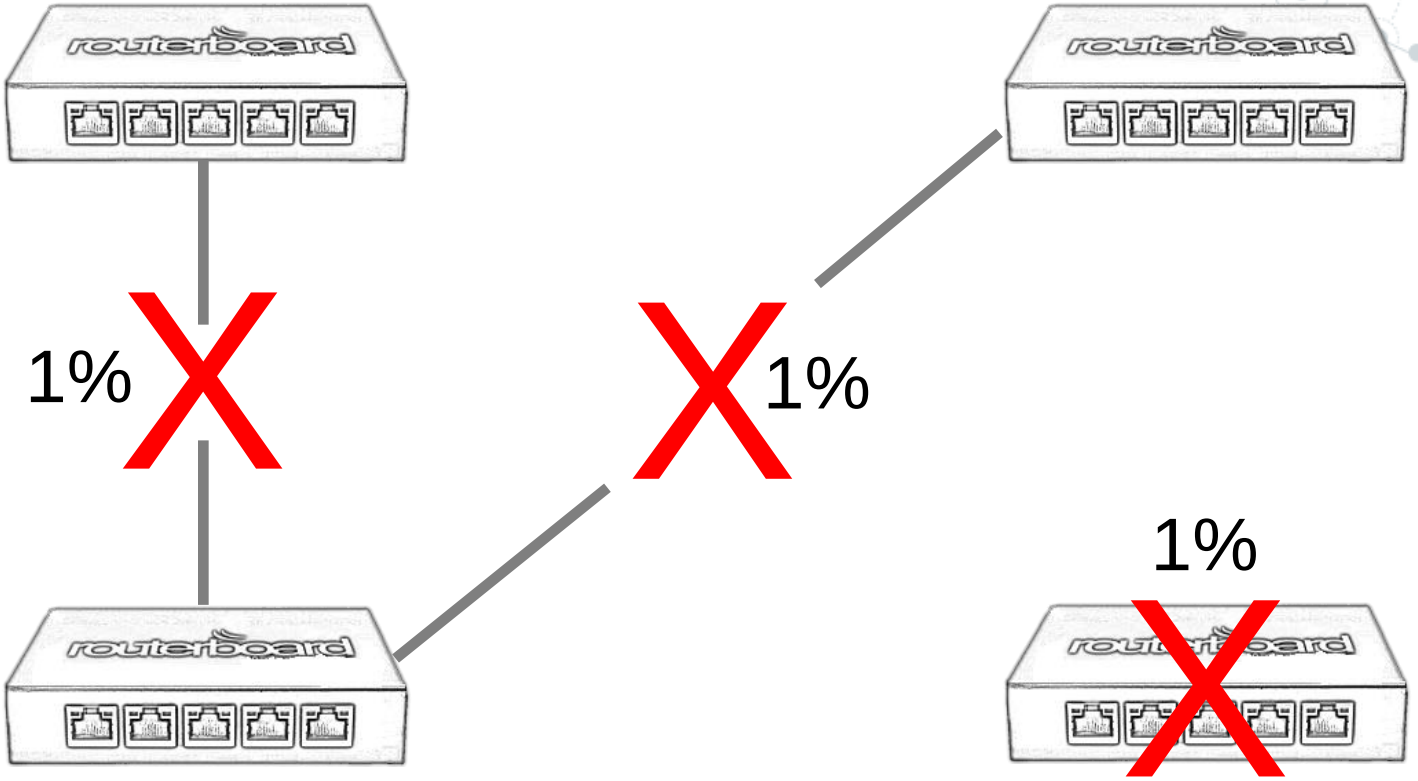

- The chance for this to happen:

  1% x 1% x 1% = 0,0001%

FAIL NO. 5

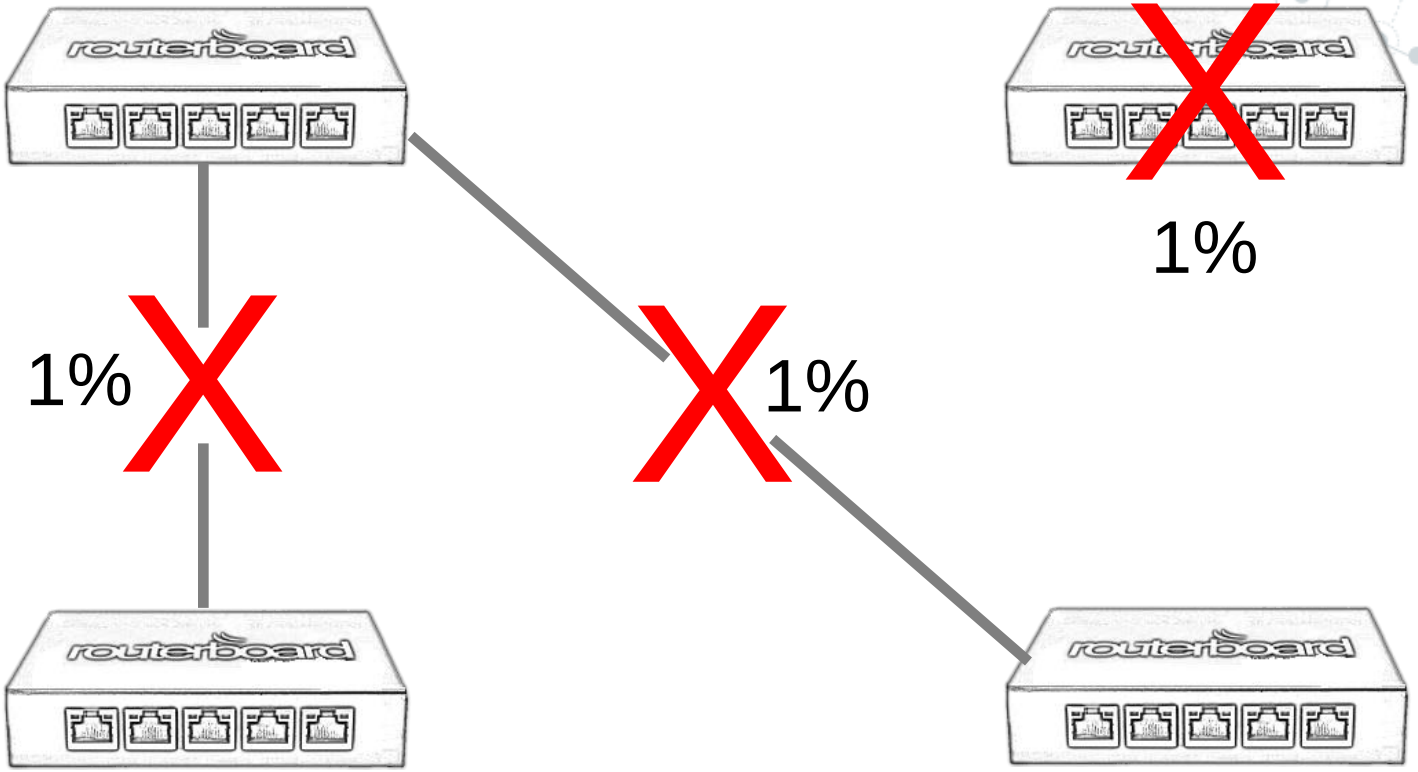

- The chance for this to happen:

  1% x 1% x 1% = 0,0001%

FAIL NO. 6

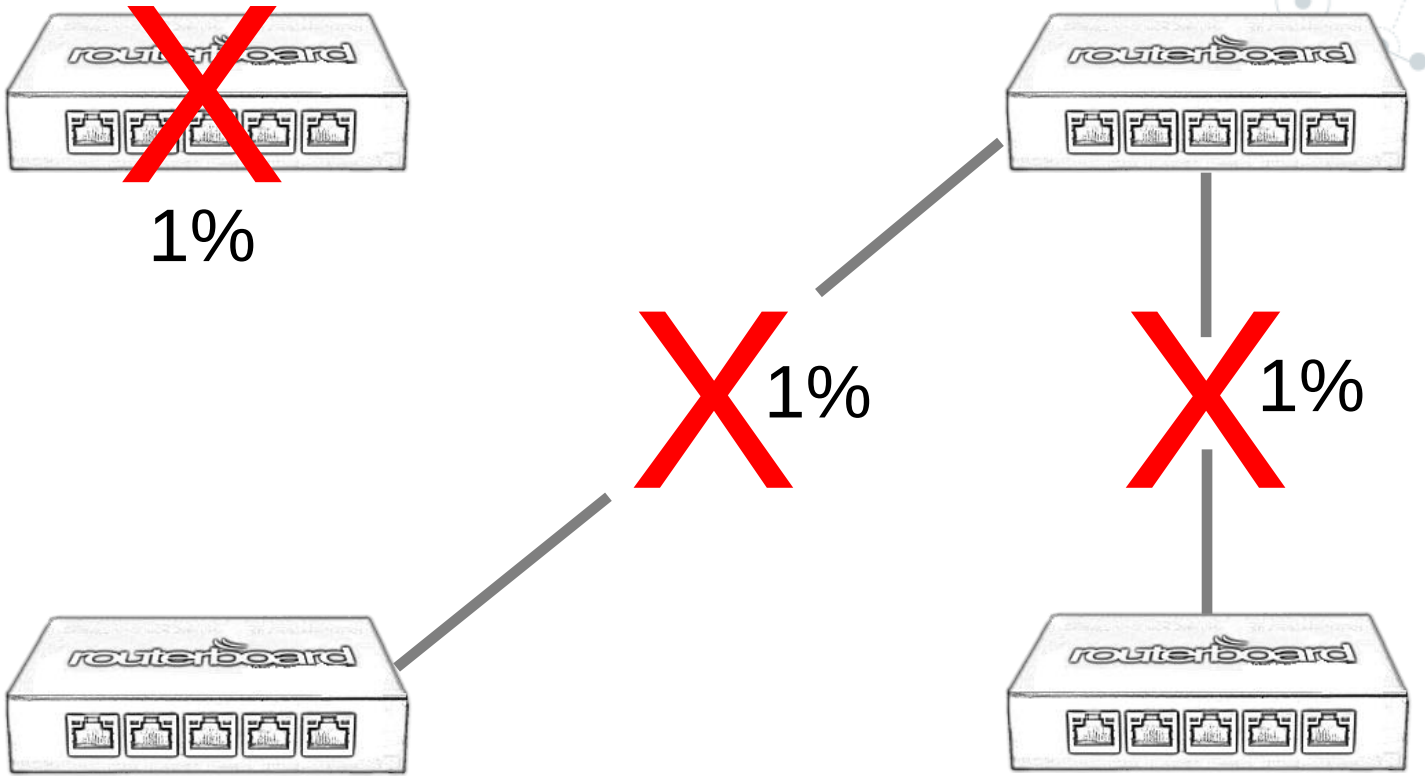

- The chance for this to happen:

  1% x 1% x 1% = 0,0001%

TOTAL FAILURE RATE

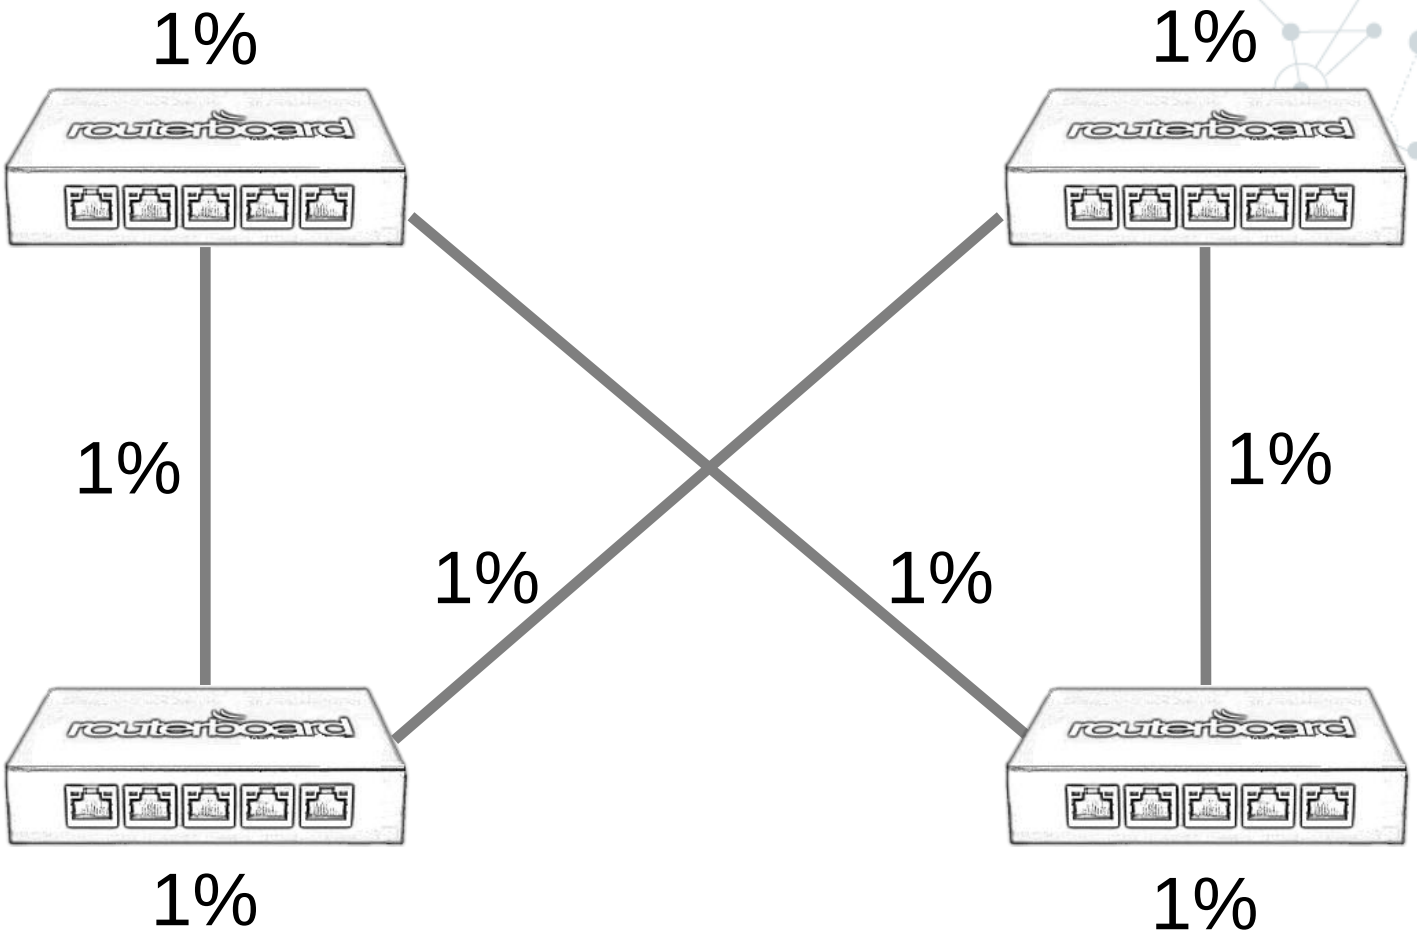

4 x 0,0001% + 2 x 0,01% = 0,0204%

SLA: 99,9796% (99,9%)

5.
Mesh networks

Equipment redundancy

Multiple path redundancy and usptream path redundancy

99,9% MESH NETWORK

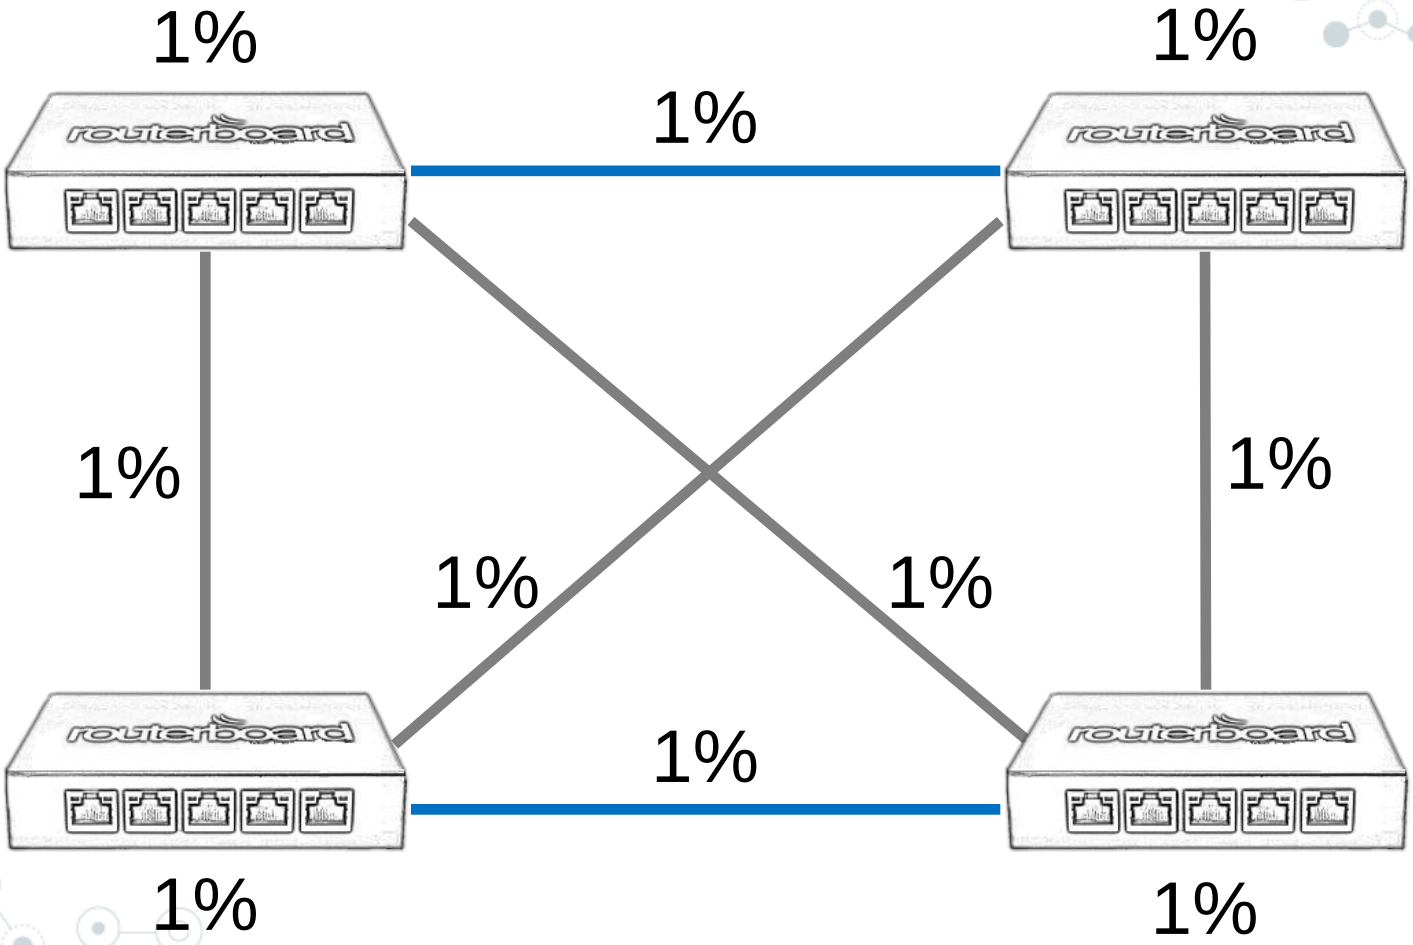

6.
Fully redundant networks

Equipment redundancy

Double path redundancy with
double upstream path redundancy

99,98%
FULLY REDUNDANT NETWORK

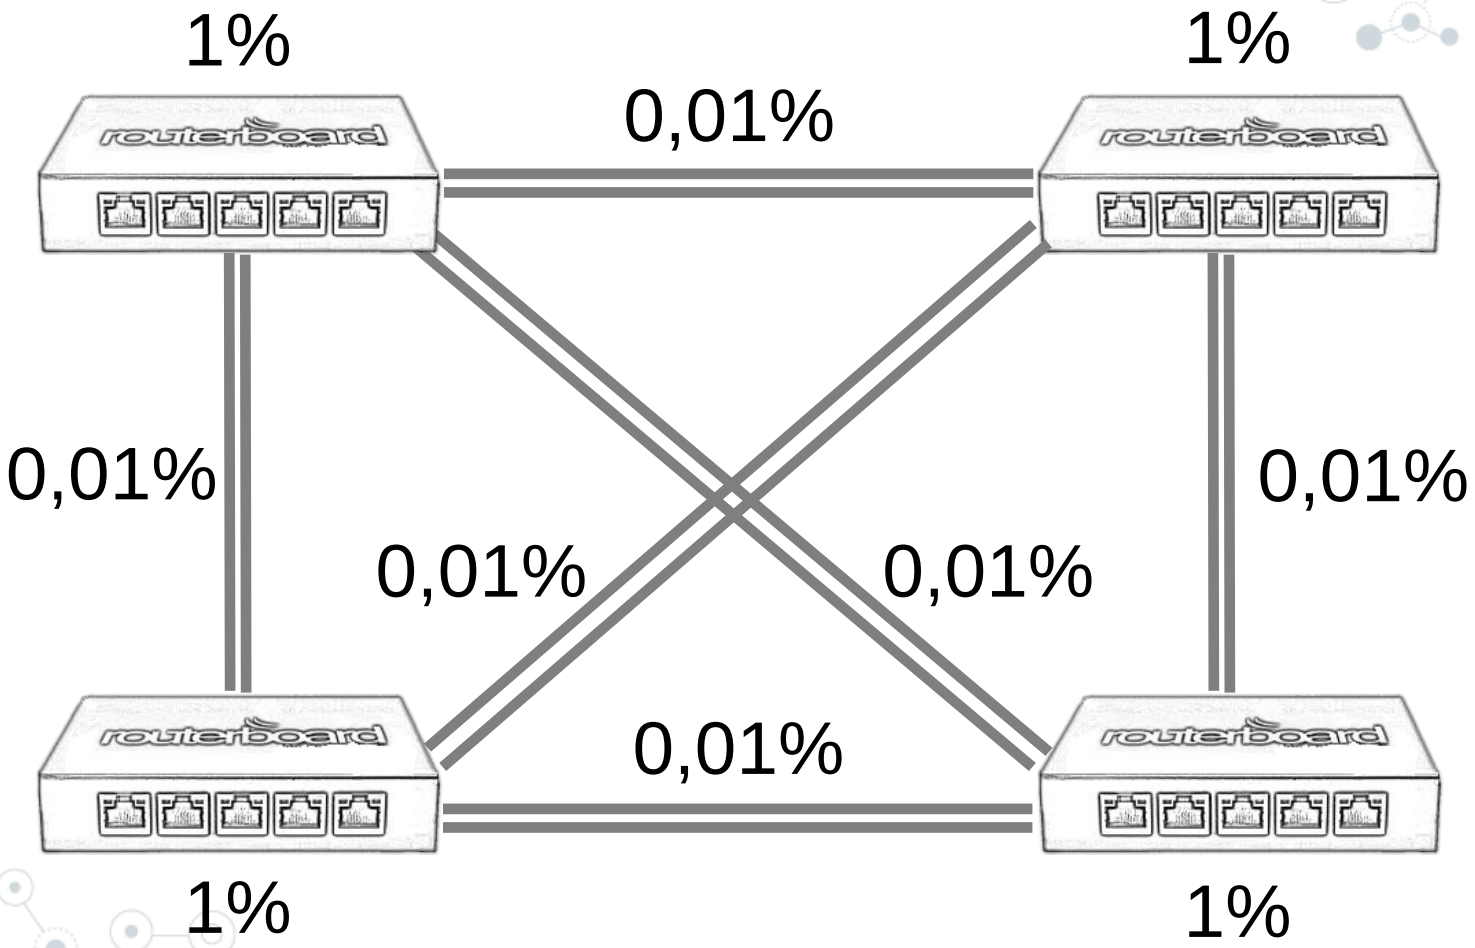

7.

Let's do it on MikroTik

There are different ways to do it

We will use multiple protocols and layers to prove the concept

Use bonding active-backup

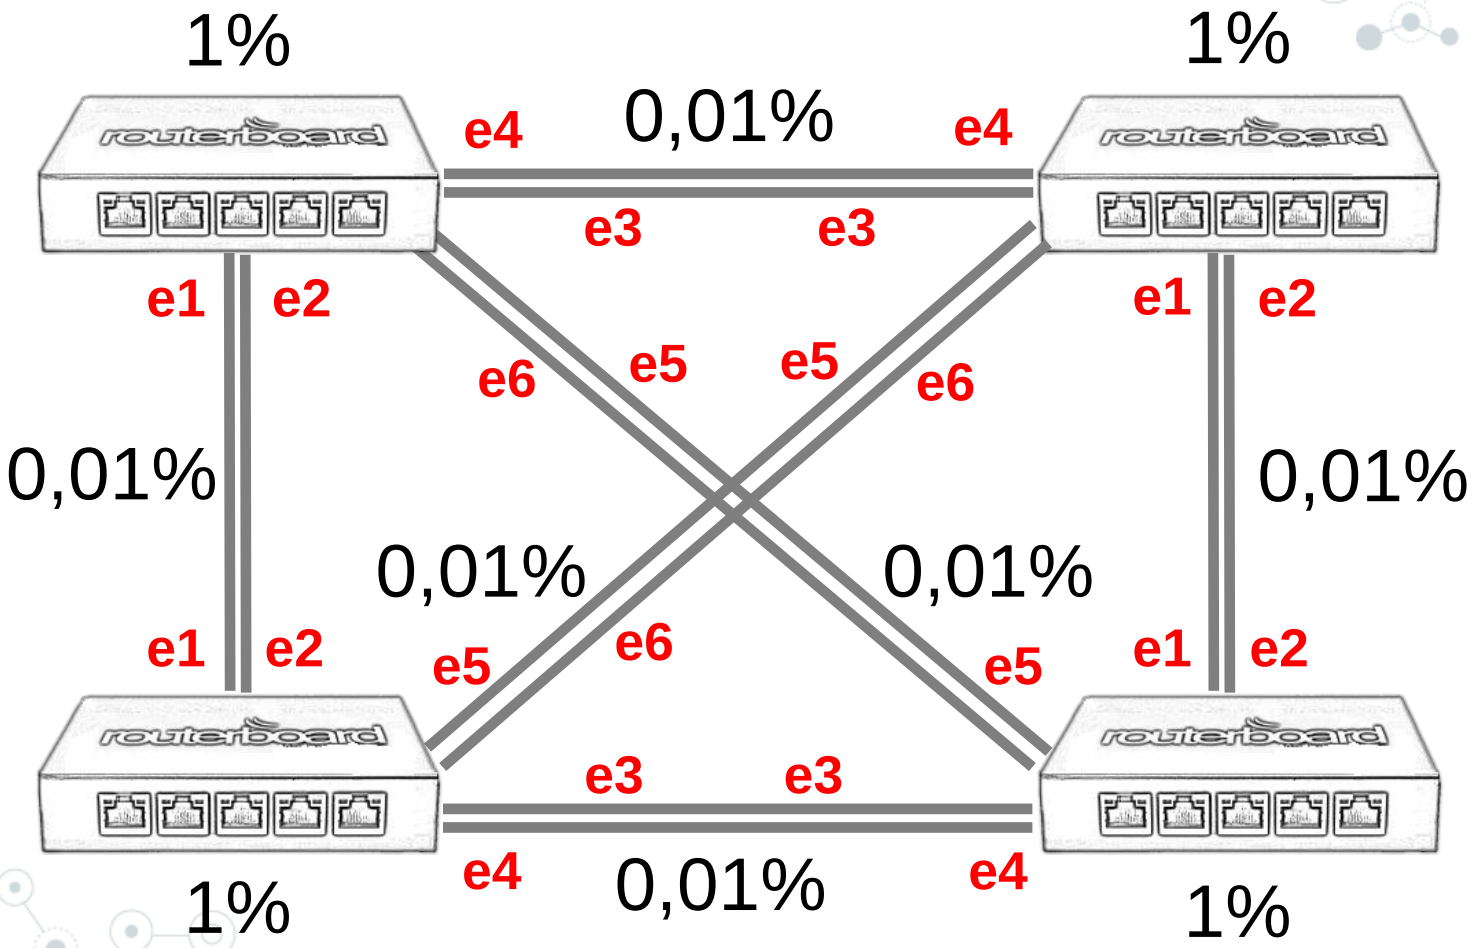

Create the bonding interfaces

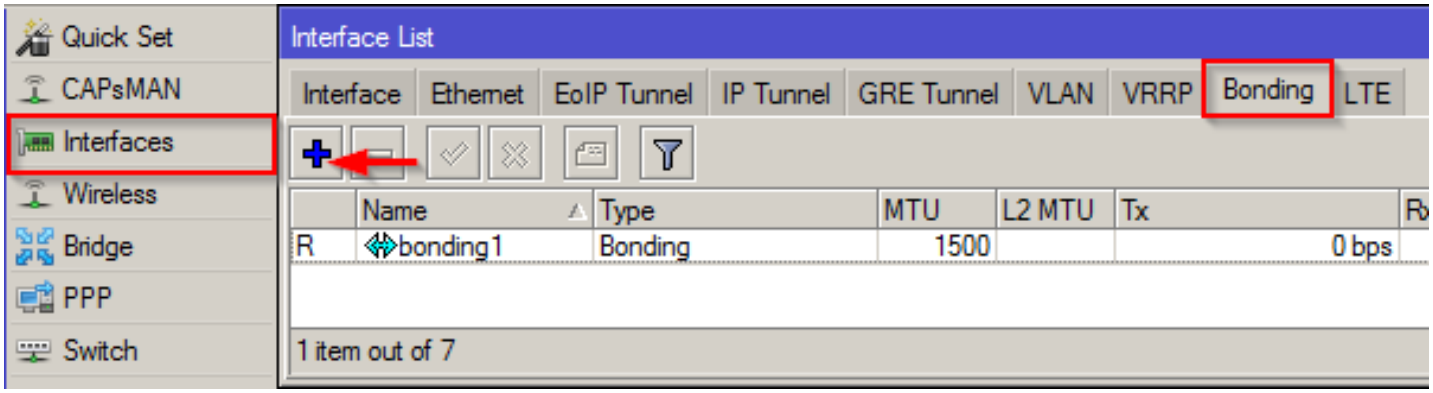

Repeat this process for every bonding link on every router until you get all 3 bonding interfaces on each router

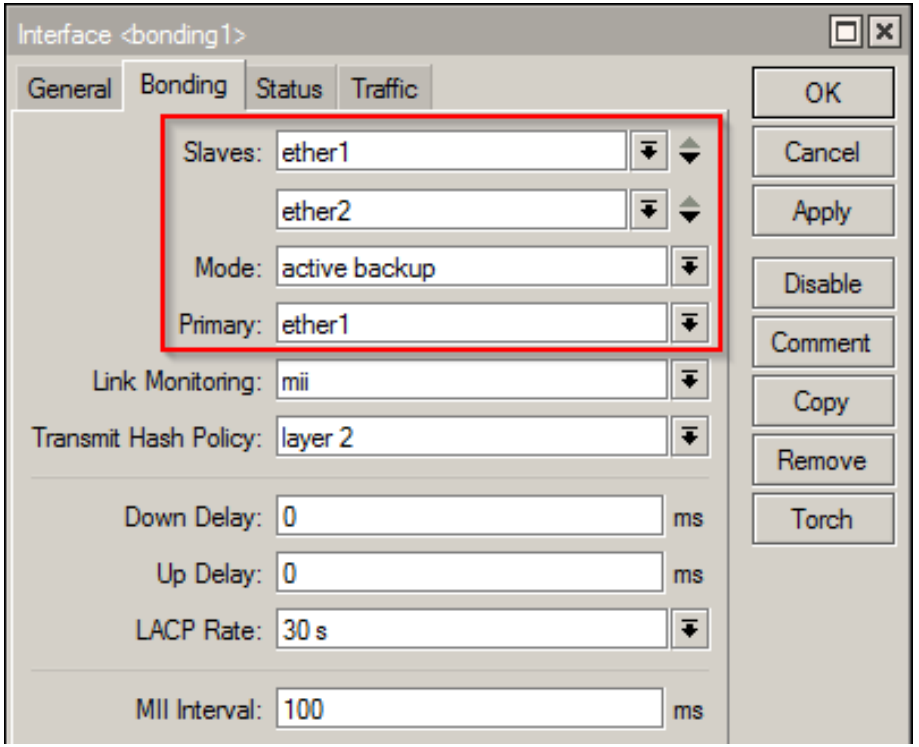

Resulting network

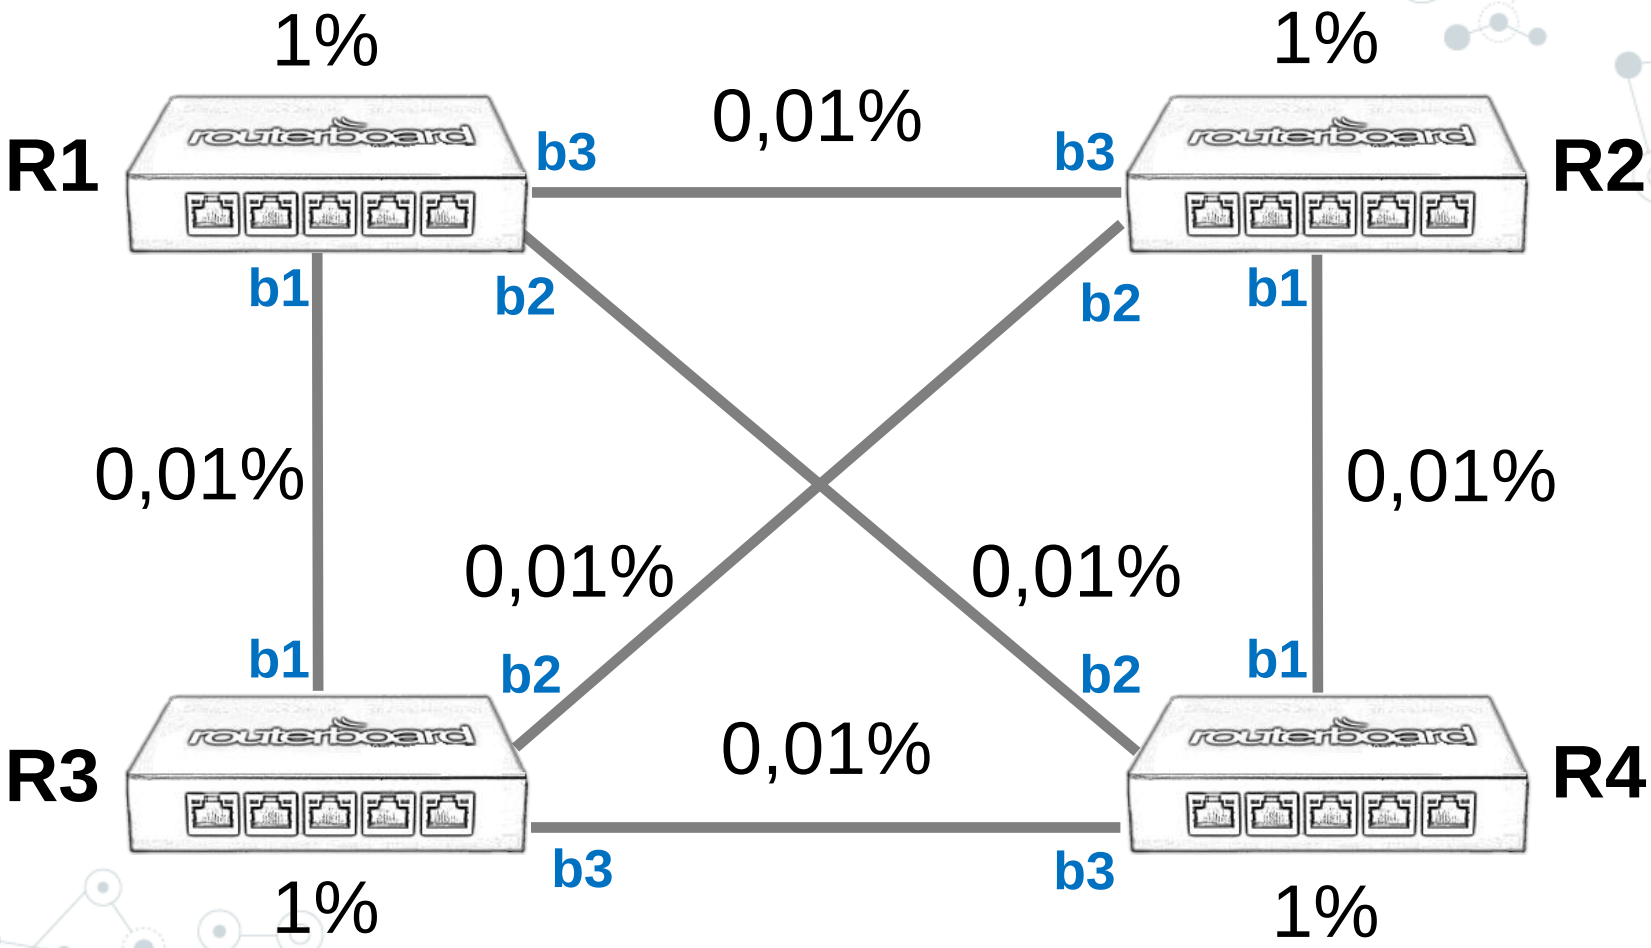

Loopback setup

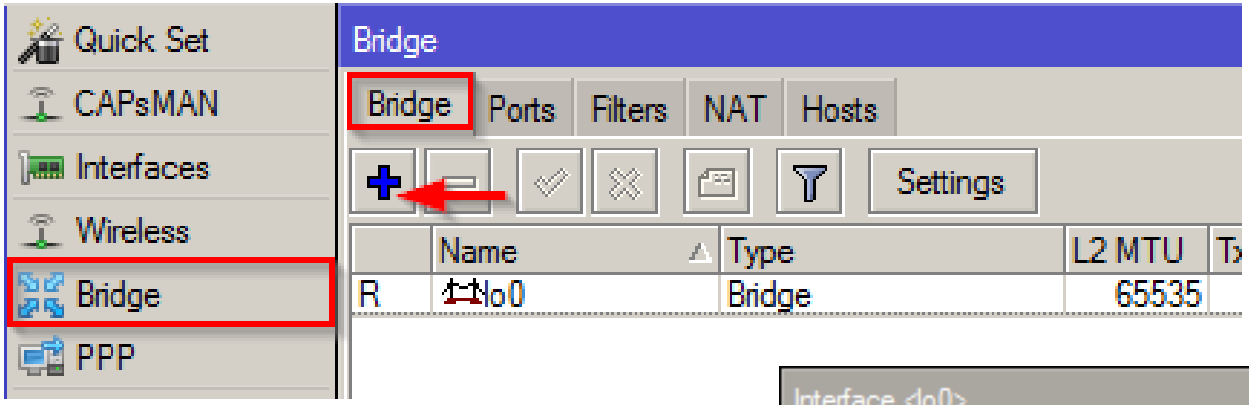

On every router add a loopback device (a bridge with no ports attached)

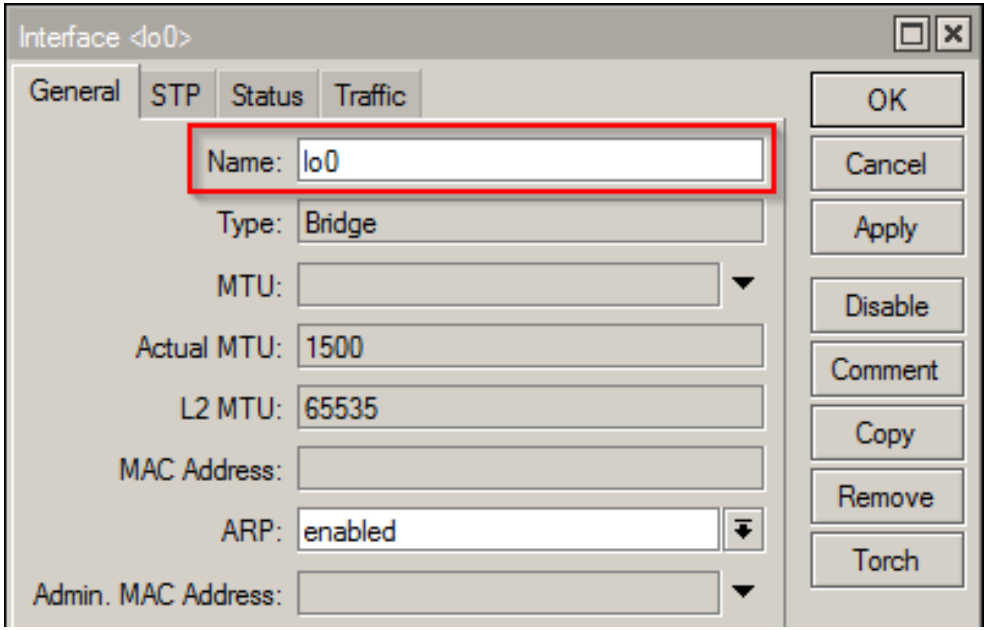

Adding loopback IPs

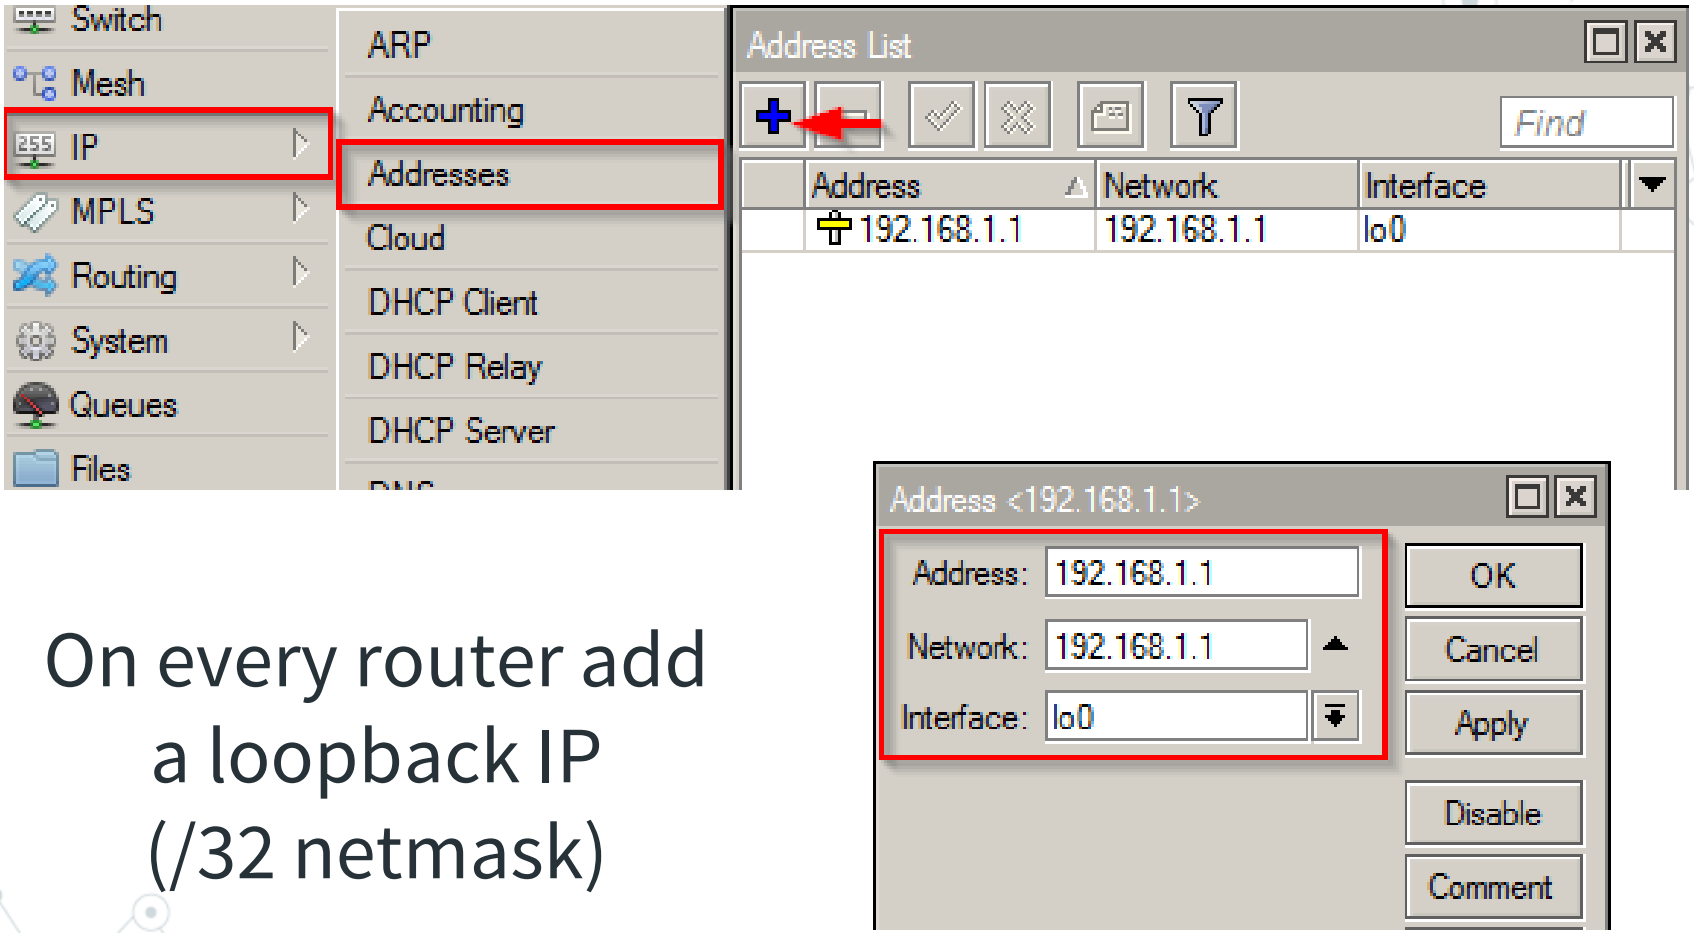

On every router add a loopback IP (/32 netmask)

Interconnections

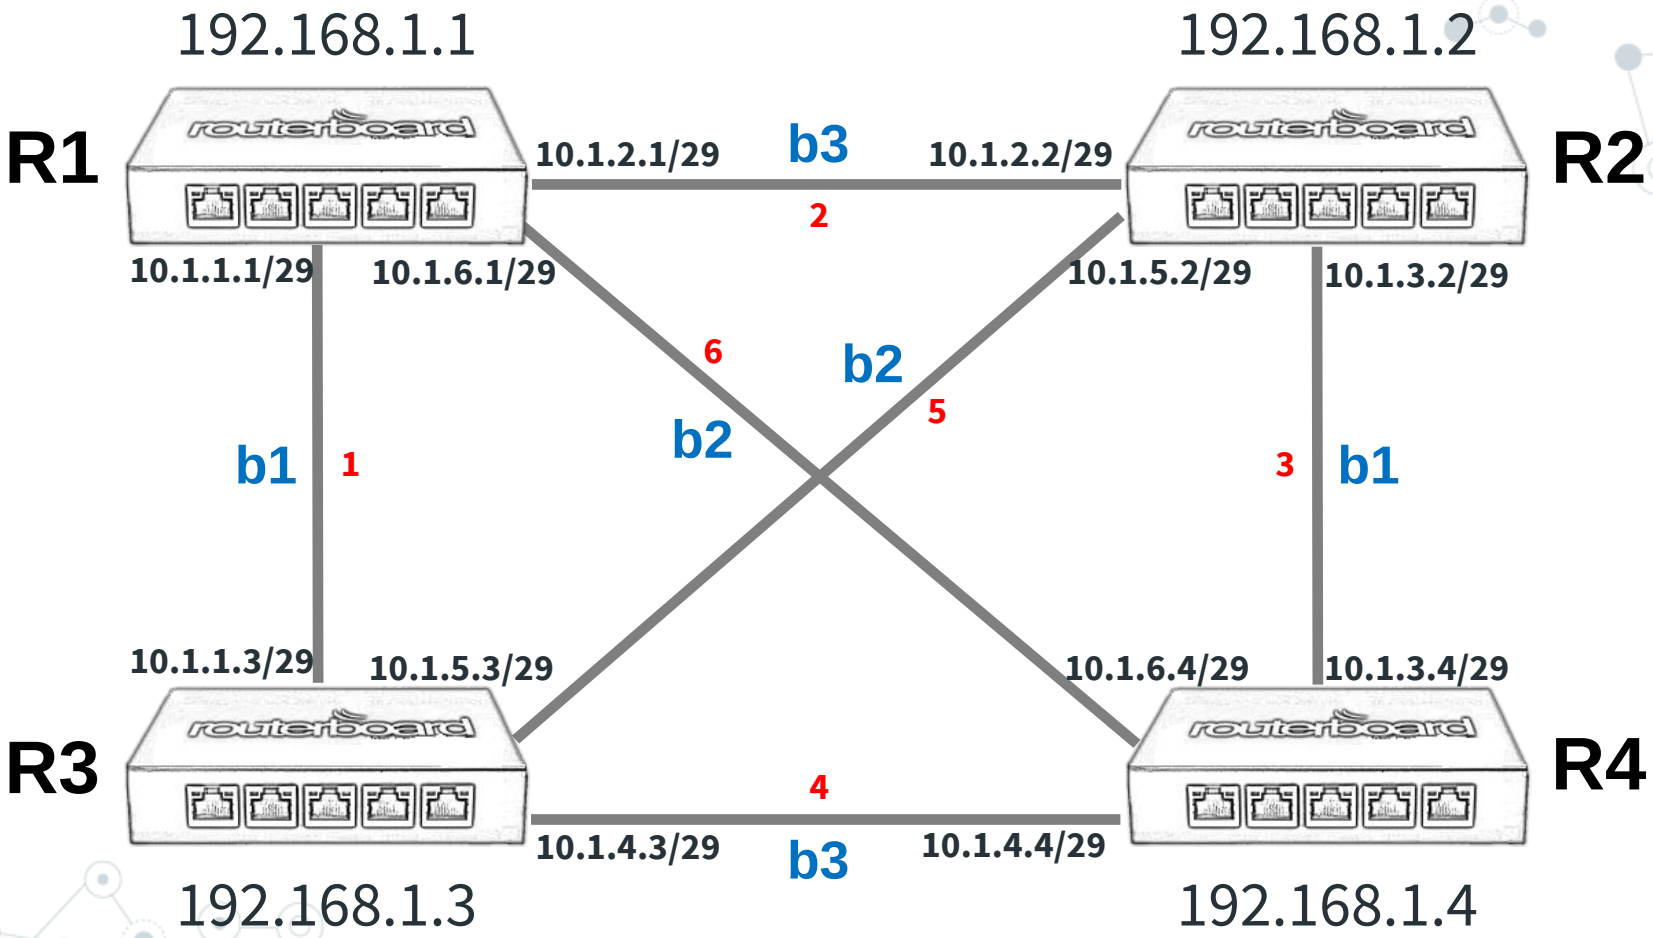

Adding the interconnect networks

Link number

Router number

R1

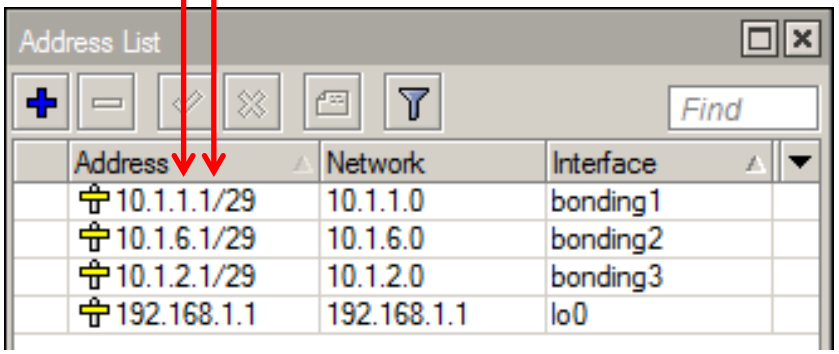

R2

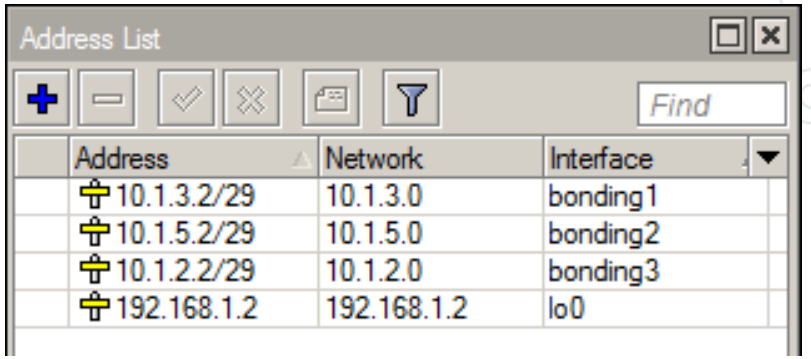

R3

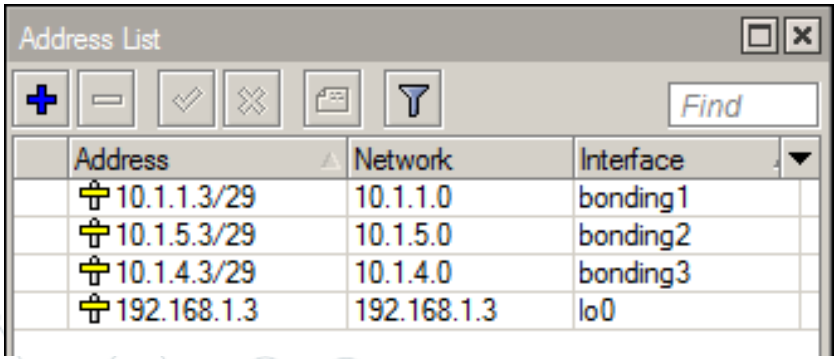

R4

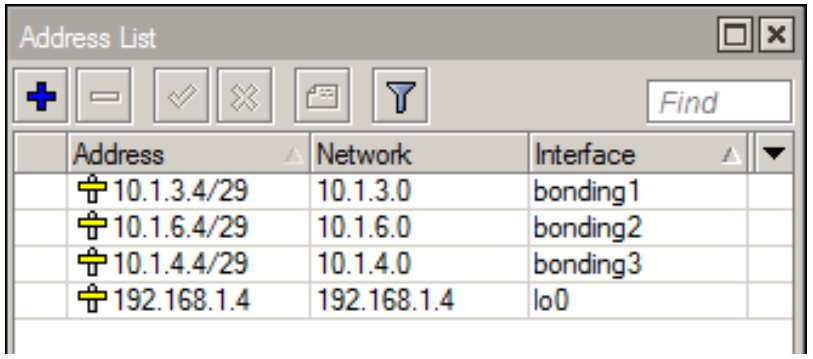

OSPF setup

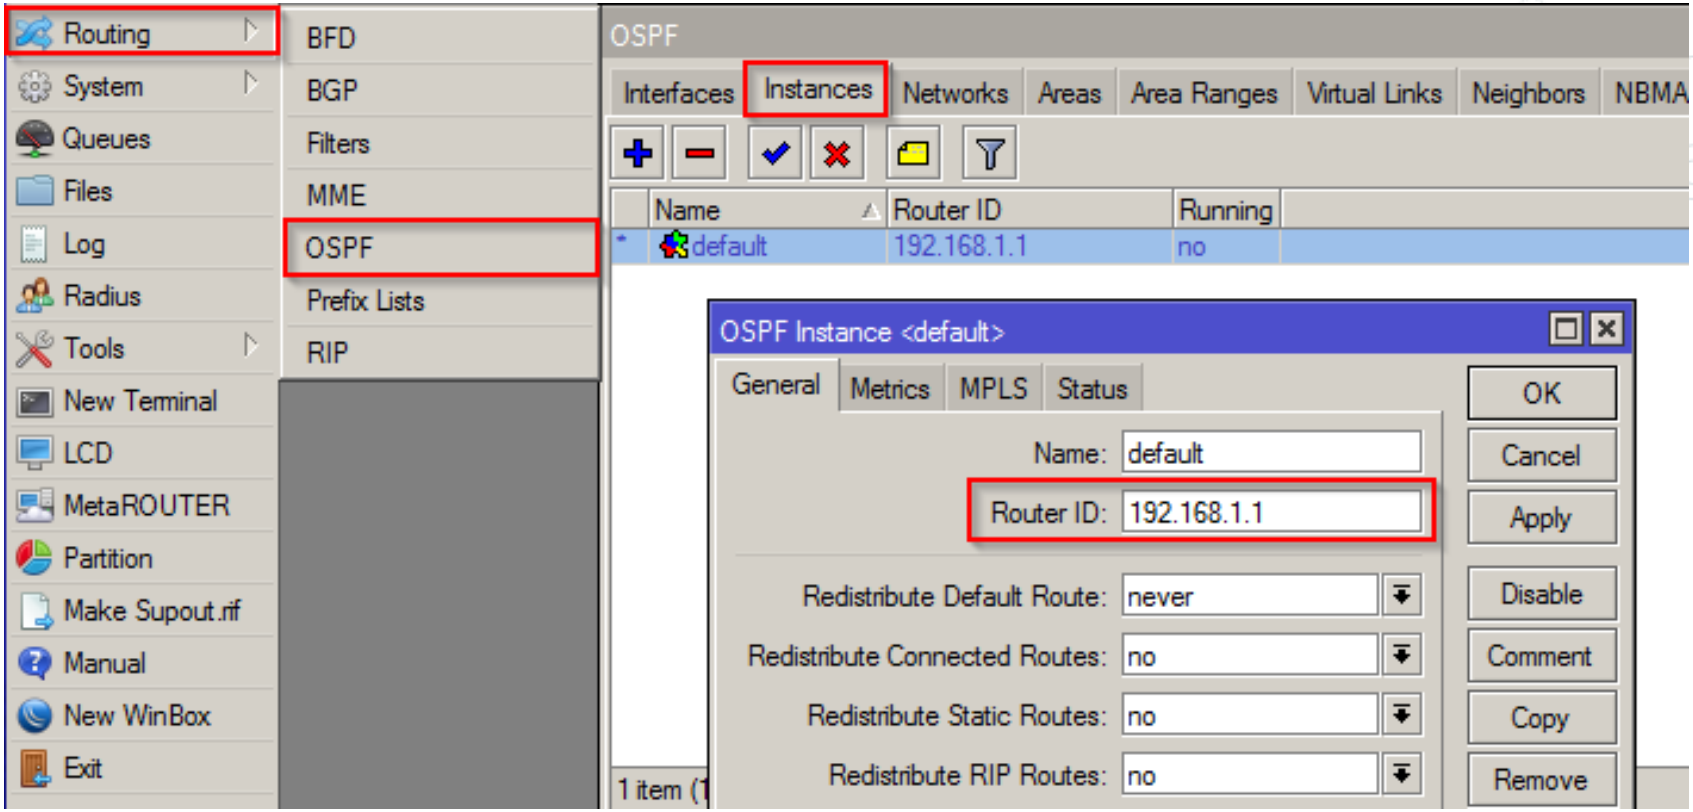

Set the OSPF router ID to the loopback address

OSPF setup

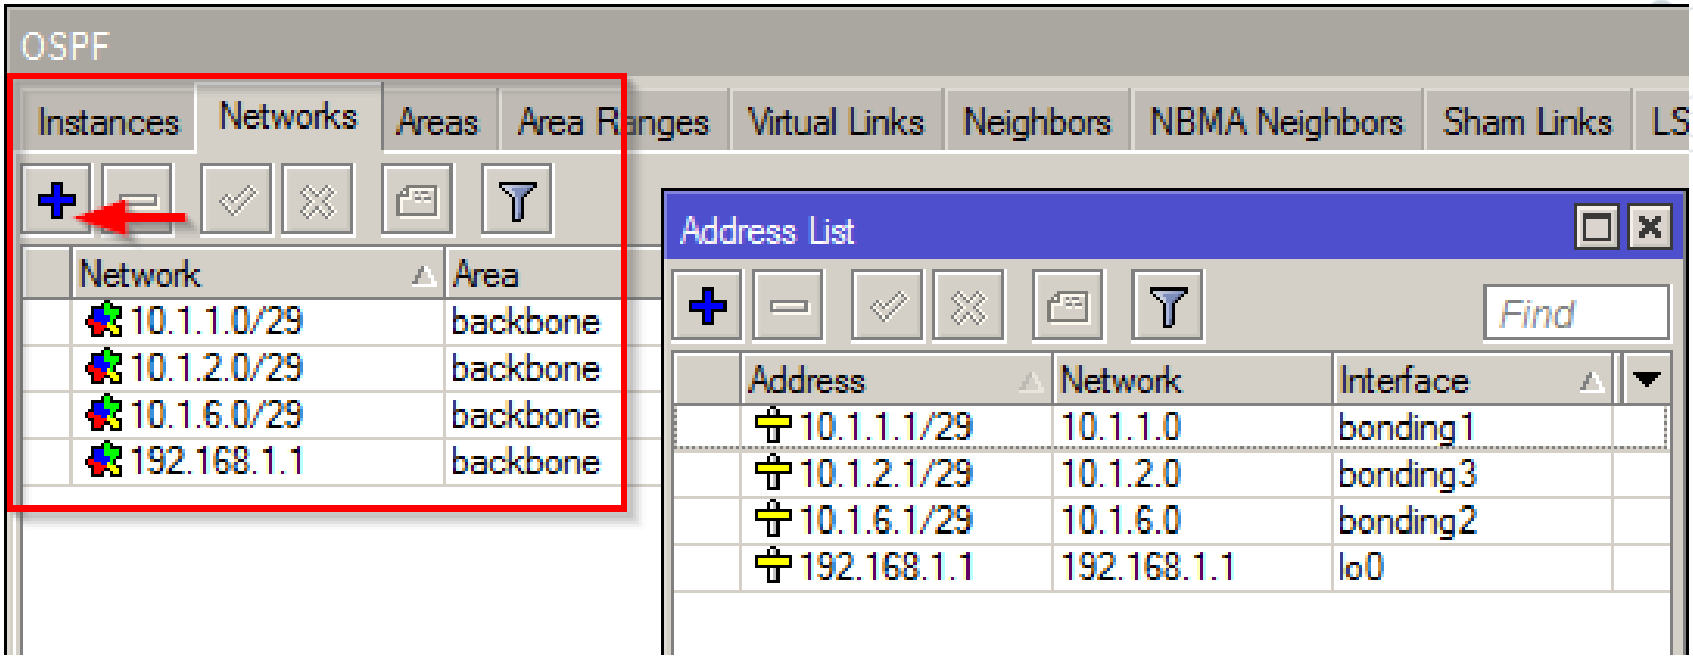

Add all connected networks to OSPF

Adding the interconnect networks

R1

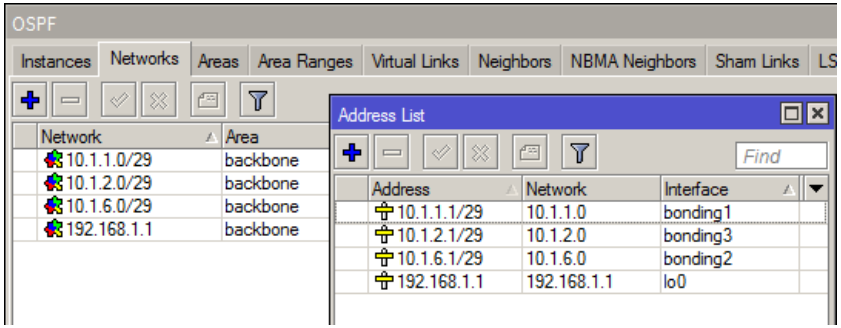

R2

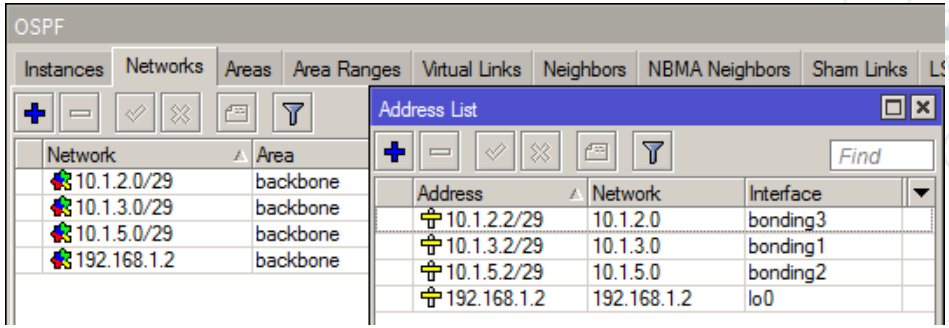

R3

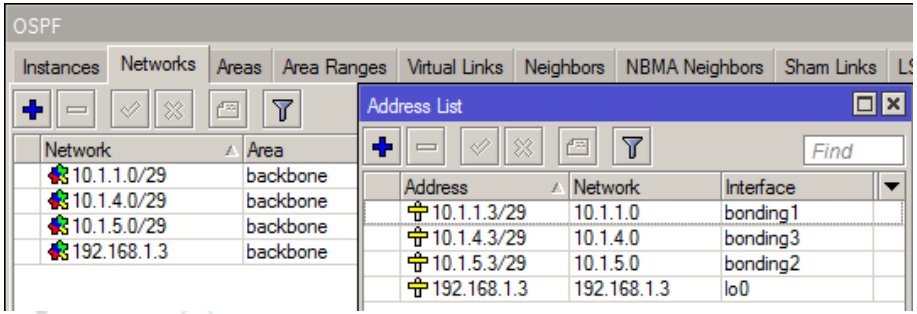

R4

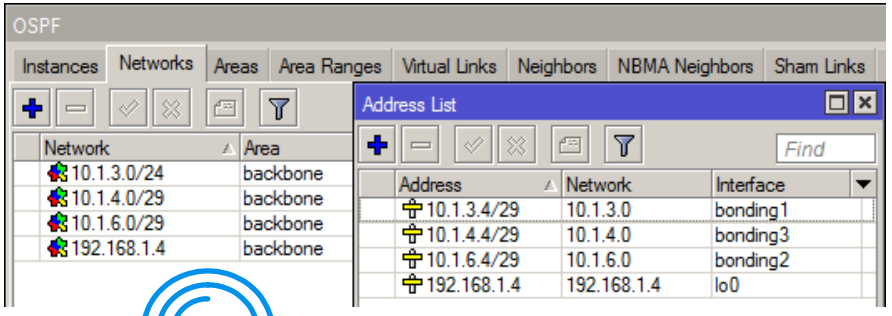

R4 - Routing table result

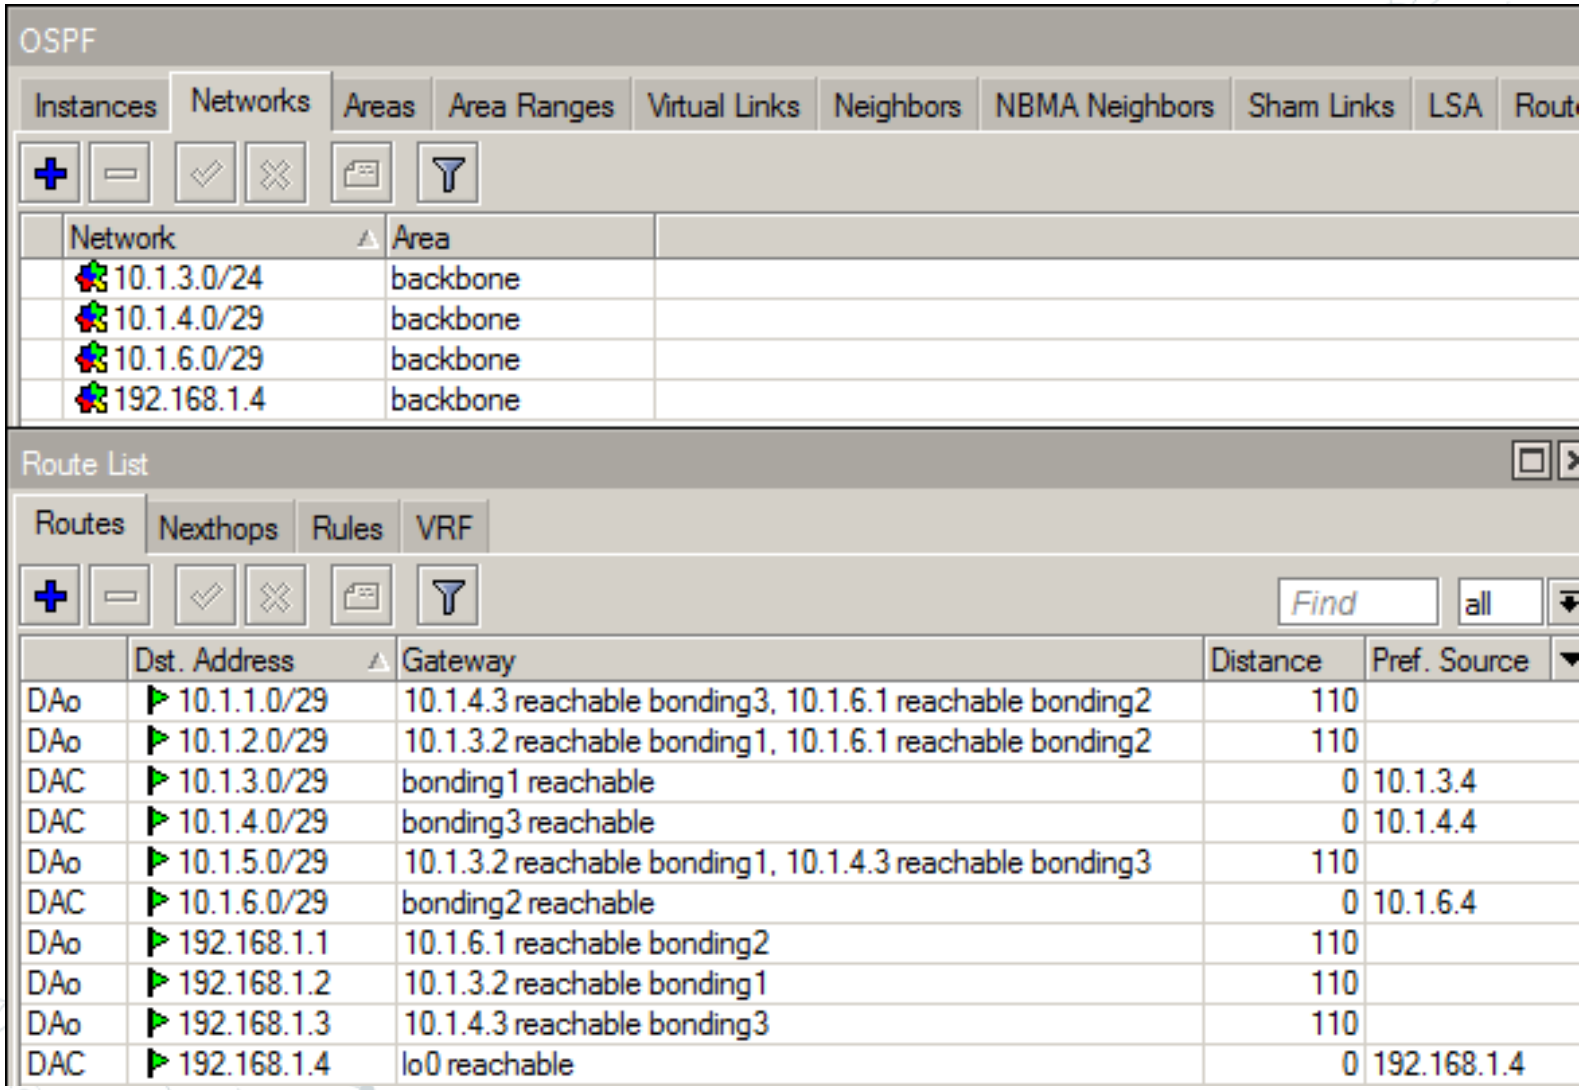

OSPF

Instances | Networks | Areas | Area Ranges | Virtual Links | Neighbors | NBMA Neighbors | Sham Links | LSA | Rout

Network	Area
10.1.3.0/24	backbone
10.1.4.0/29	backbone
10.1.6.0/29	backbone
192.168.1.4	backbone

Route List

Routes | Nexthops | Rules | VRF

Find | all

	Dst. Address	Gateway	Distance	Pref. Source
DAo	10.1.1.0/29	10.1.4.3 reachable bonding3, 10.1.6.1 reachable bonding2	110	
DAo	10.1.2.0/29	10.1.3.2 reachable bonding1, 10.1.6.1 reachable bonding2	110	
DAC	10.1.3.0/29	bonding1 reachable	0	10.1.3.4
DAC	10.1.4.0/29	bonding3 reachable	0	10.1.4.4
DAo	10.1.5.0/29	10.1.3.2 reachable bonding1, 10.1.4.3 reachable bonding3	110	
DAC	10.1.6.0/29	bonding2 reachable	0	10.1.6.4
DAo	192.168.1.1	10.1.6.1 reachable bonding2	110	
DAo	192.168.1.2	10.1.3.2 reachable bonding1	110	
DAo	192.168.1.3	10.1.4.3 reachable bonding3	110	
DAC	192.168.1.4	lo0 reachable	0	192.168.1.4

Providers and customers

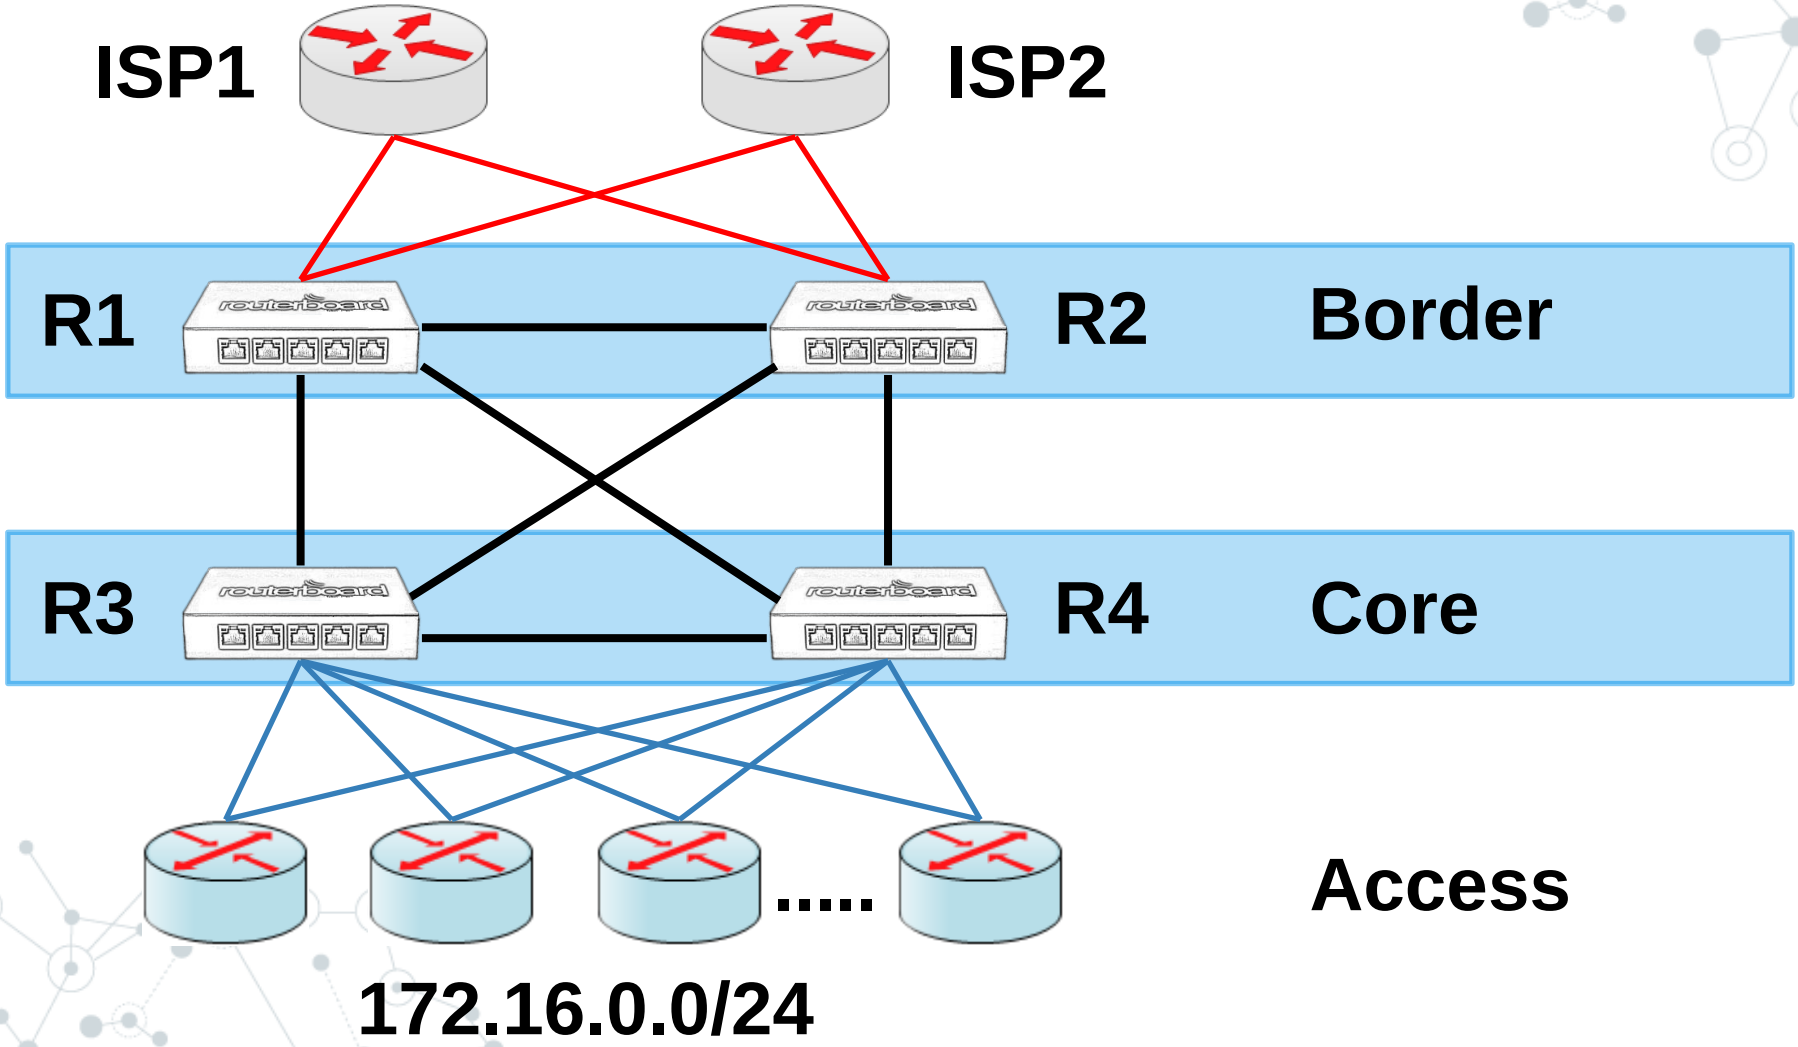

OSPF specifics

The network is now fully redundant but it depends on OSPF for path selection and convergence

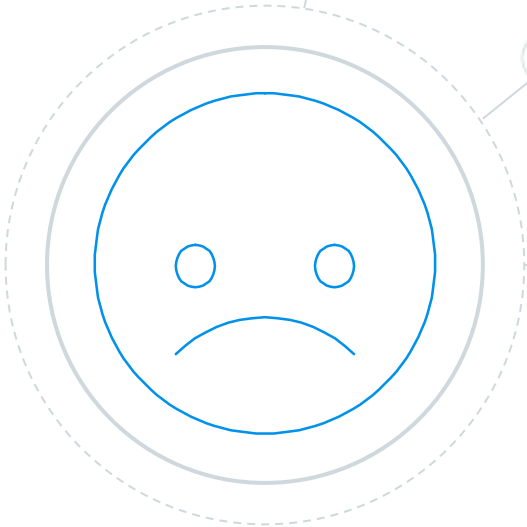

MPLS TE to the rescue

Deploy MPLS TE for faster and better path control

For demo purposes we will turn this network…

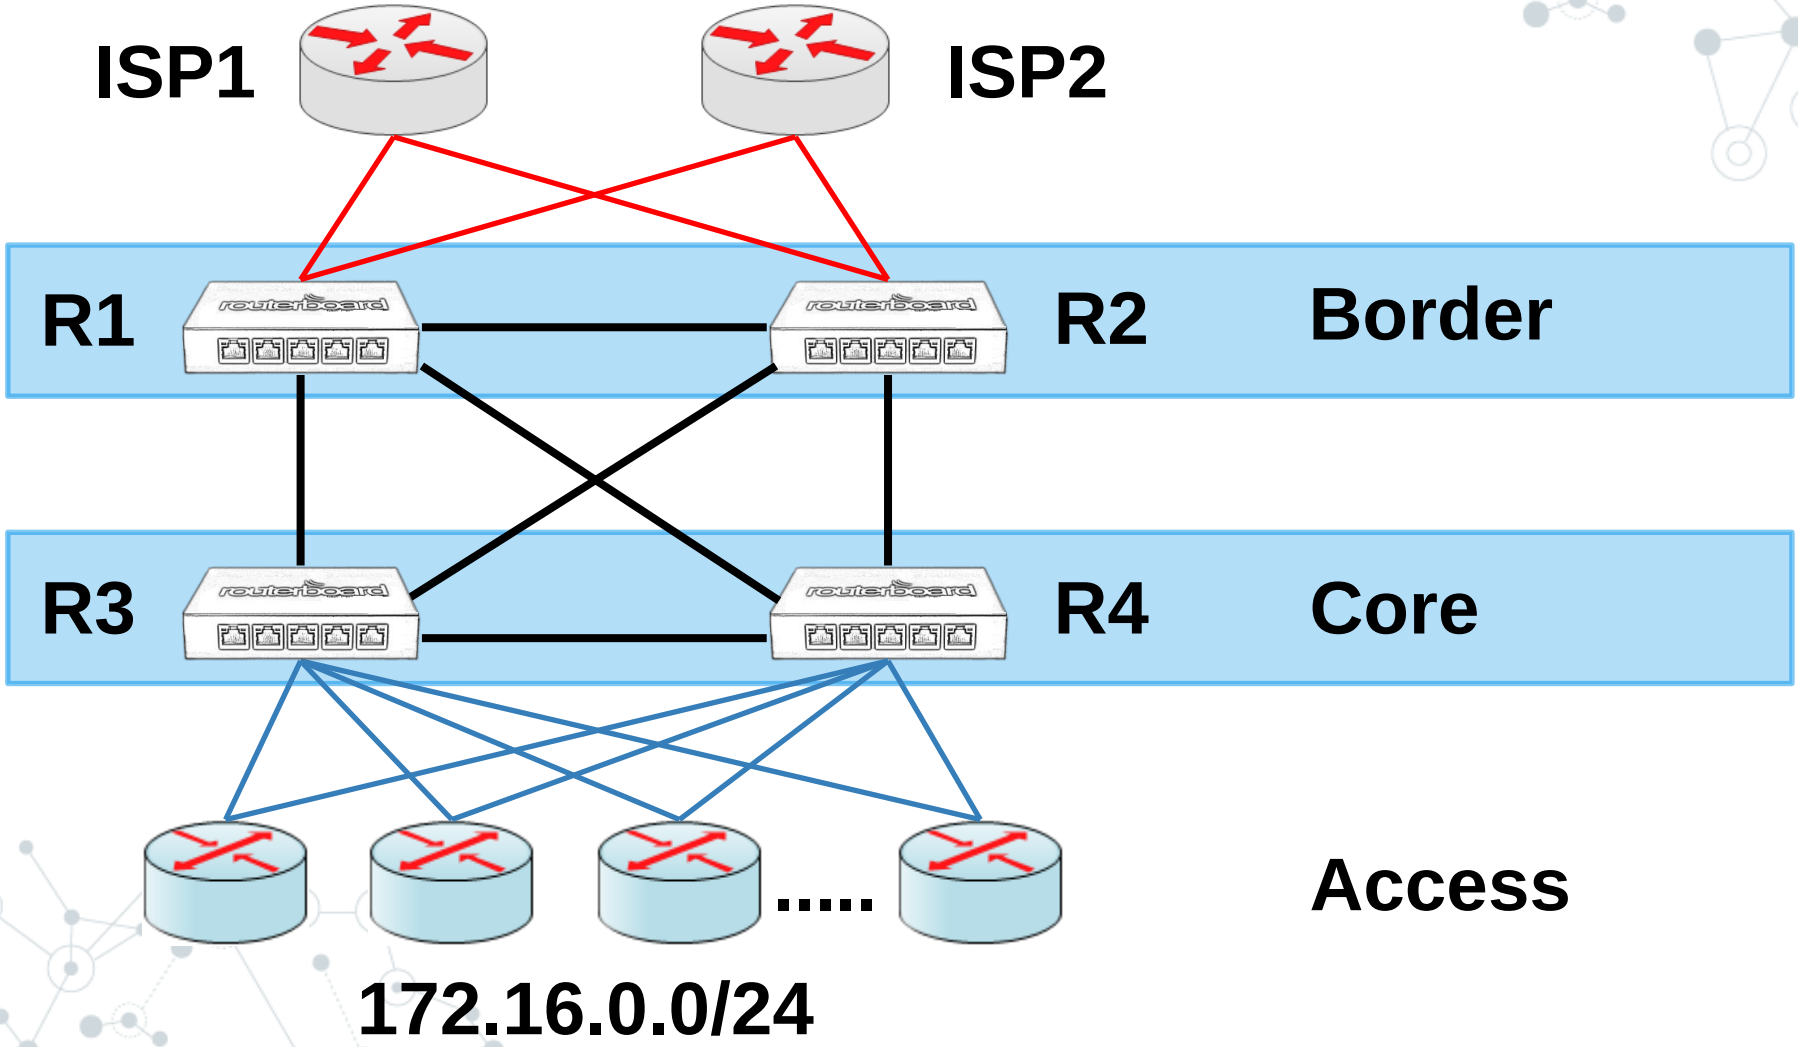

Into this simplified network

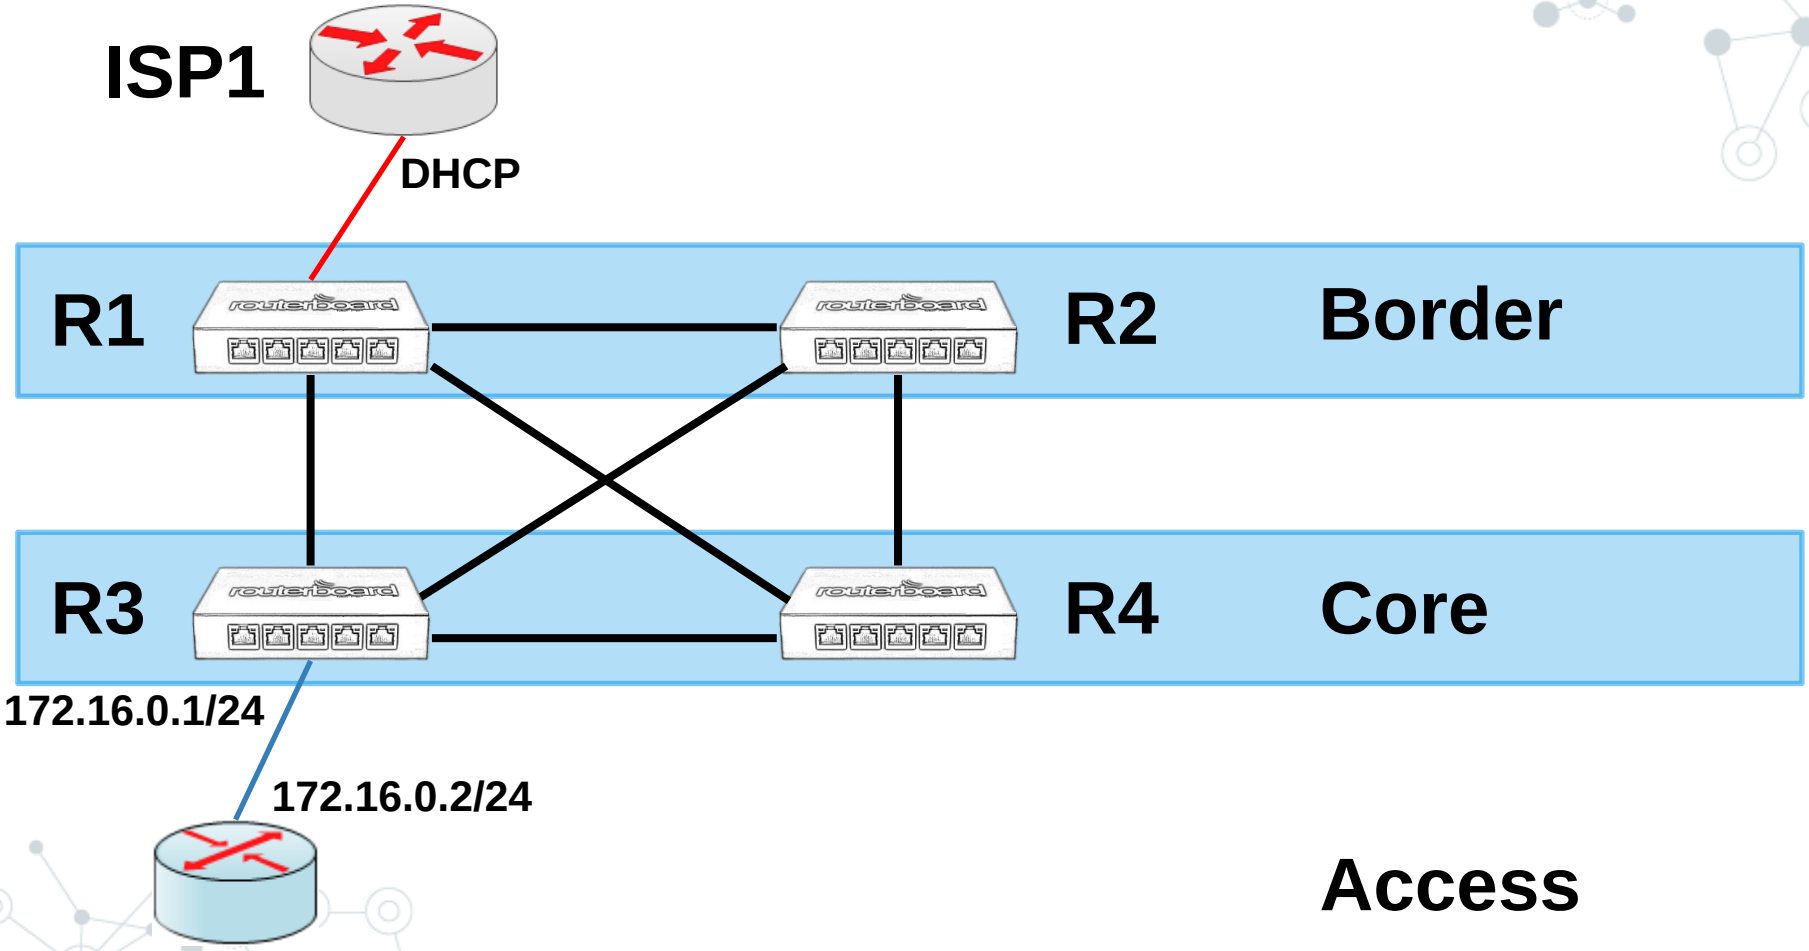

The rest remains the same

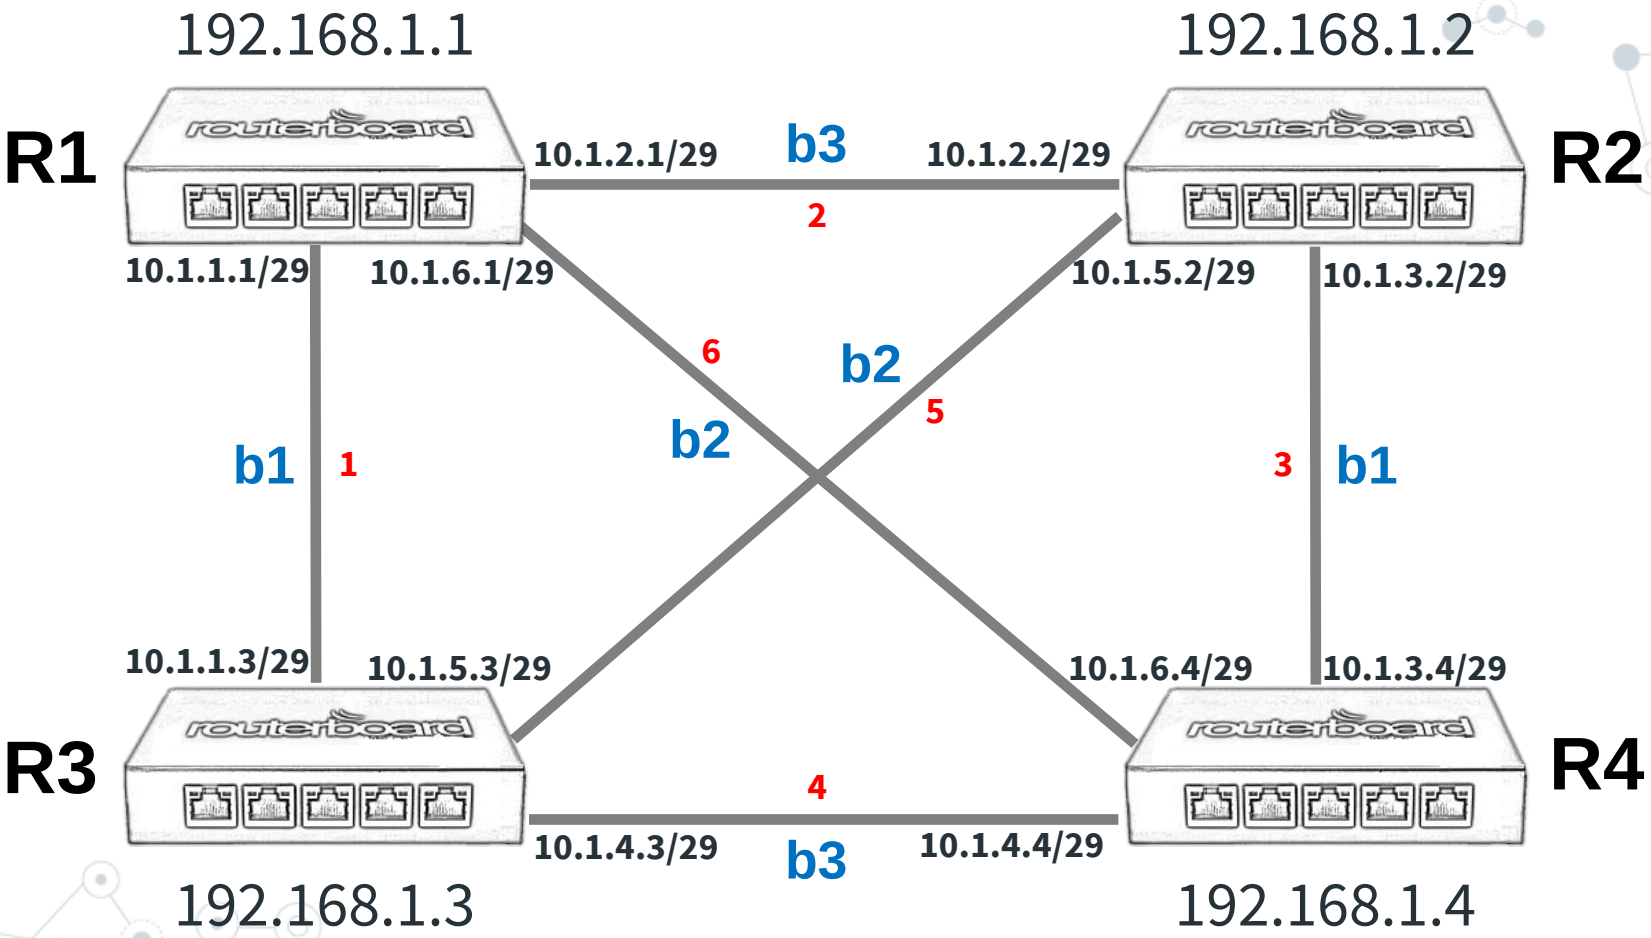

http://mikrotik-academy.com/mum2016

First things first (R1 default route)

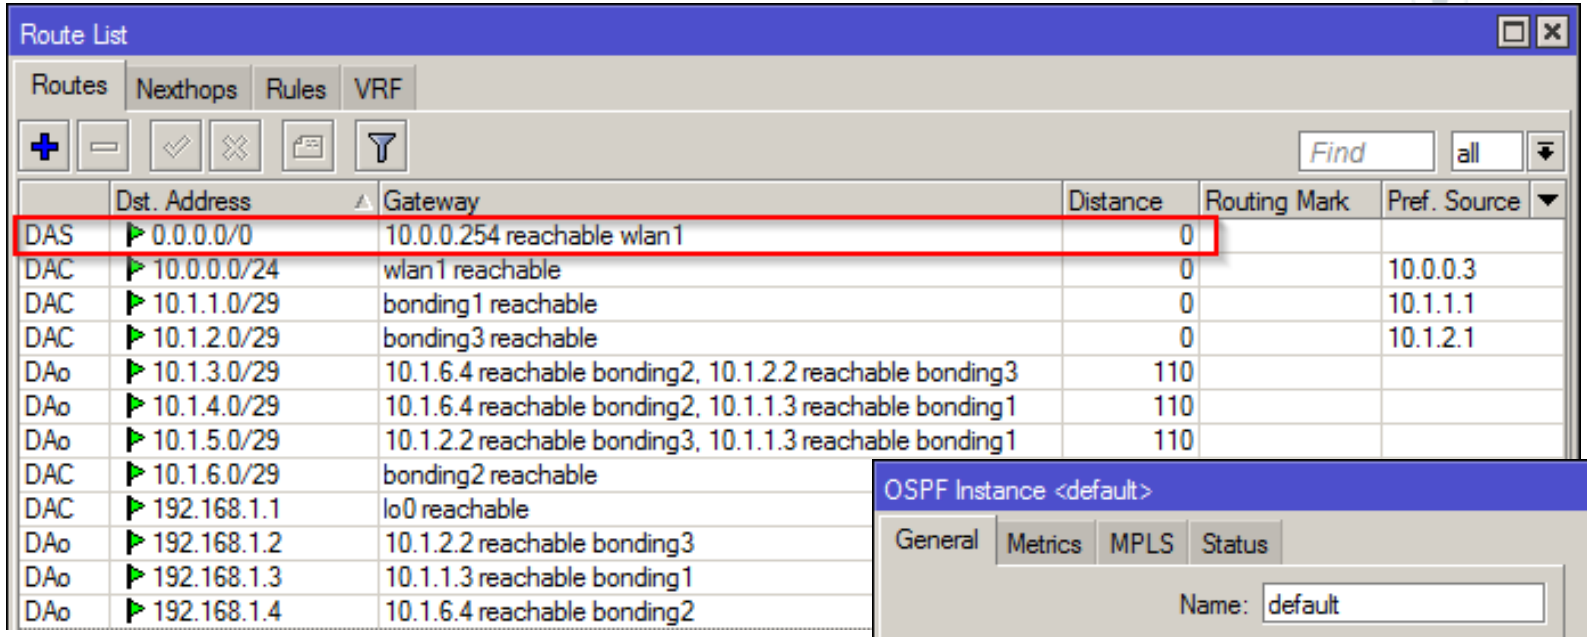

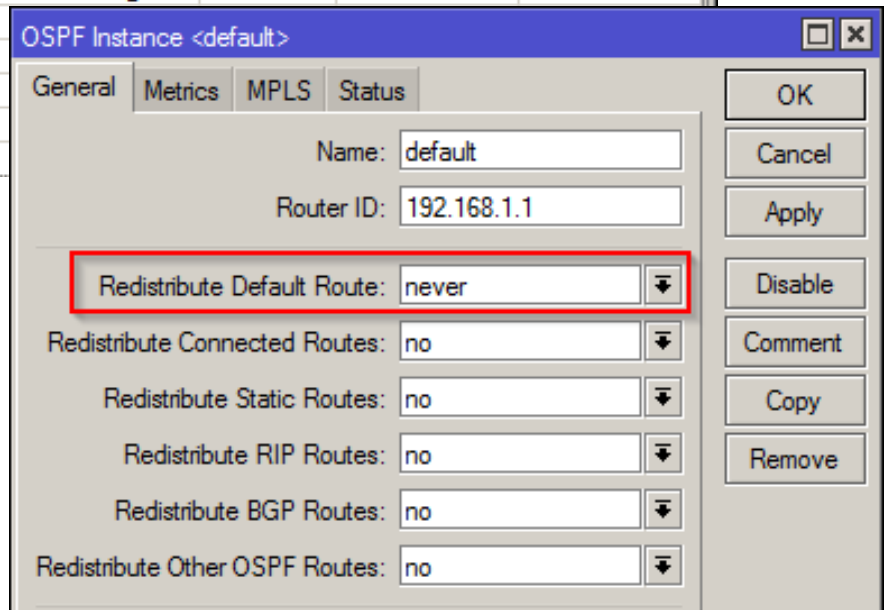

Do the same for **R2** on non-simplified setups

R3 & Customer setup

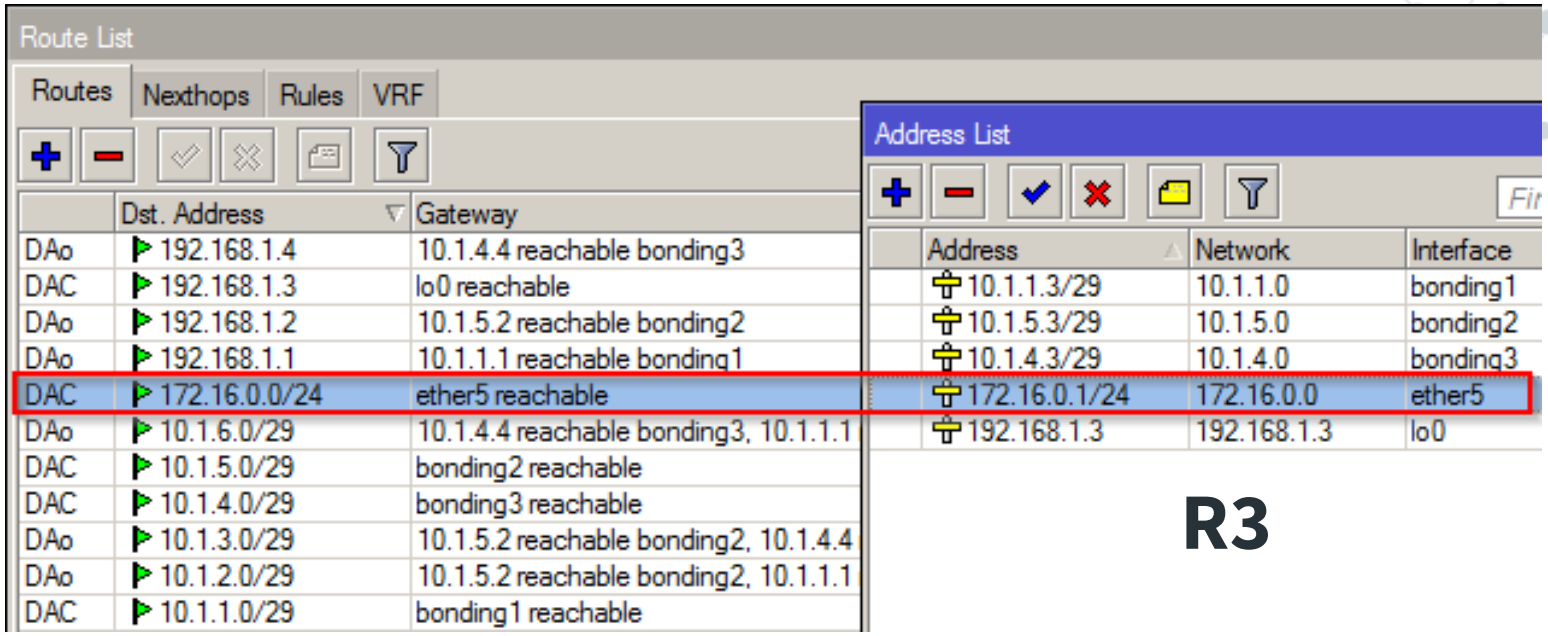

Customer equipment towards **R3**

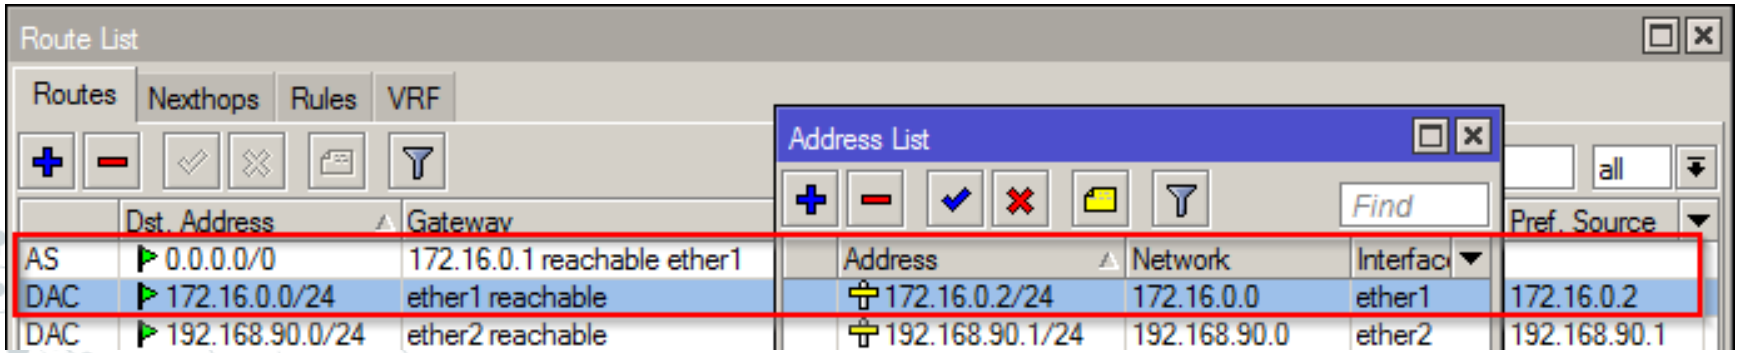

Preliminary MPLS TE over OSPF

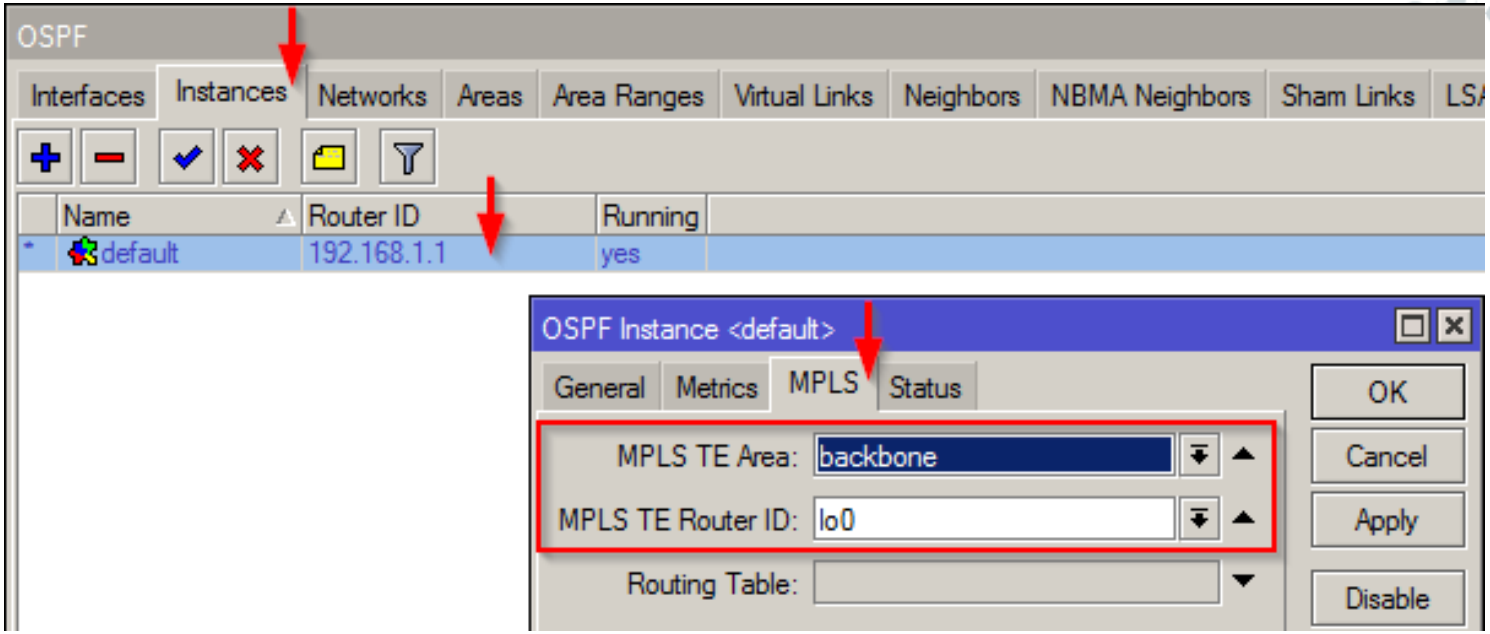

Do this for all 4 routers

Setting up MPLS TE interfaces

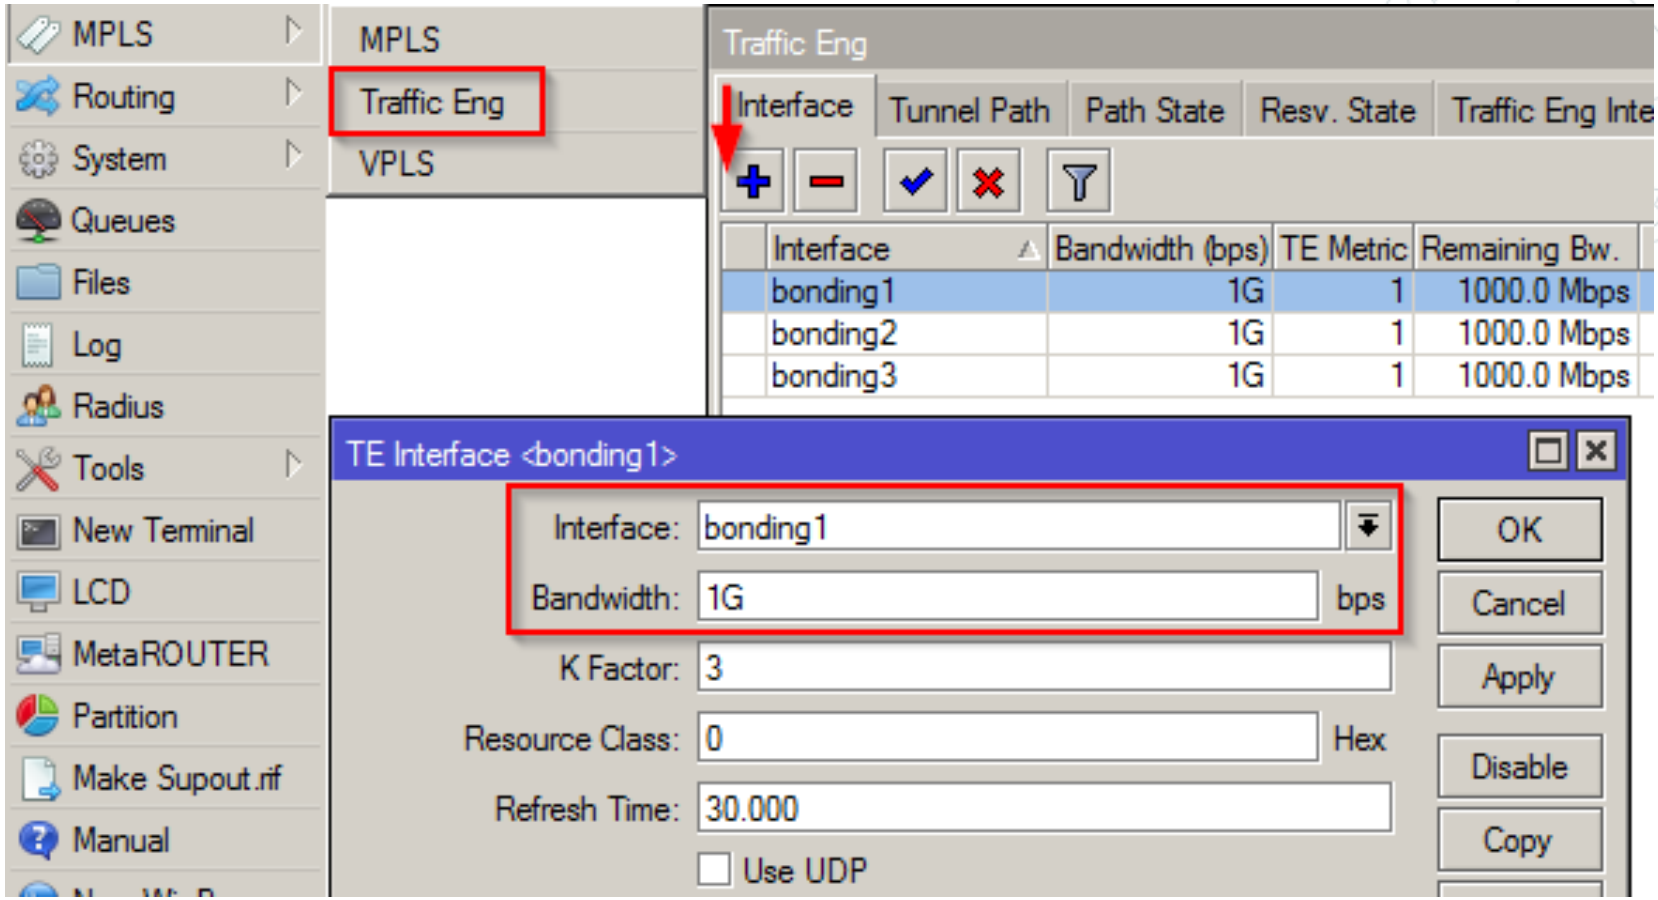

Add each bonding interface to TE on all 4 routers

TE tunnels are unidirectional

We need to set up paths
in both directions

Possible paths from **R1** to **R3**

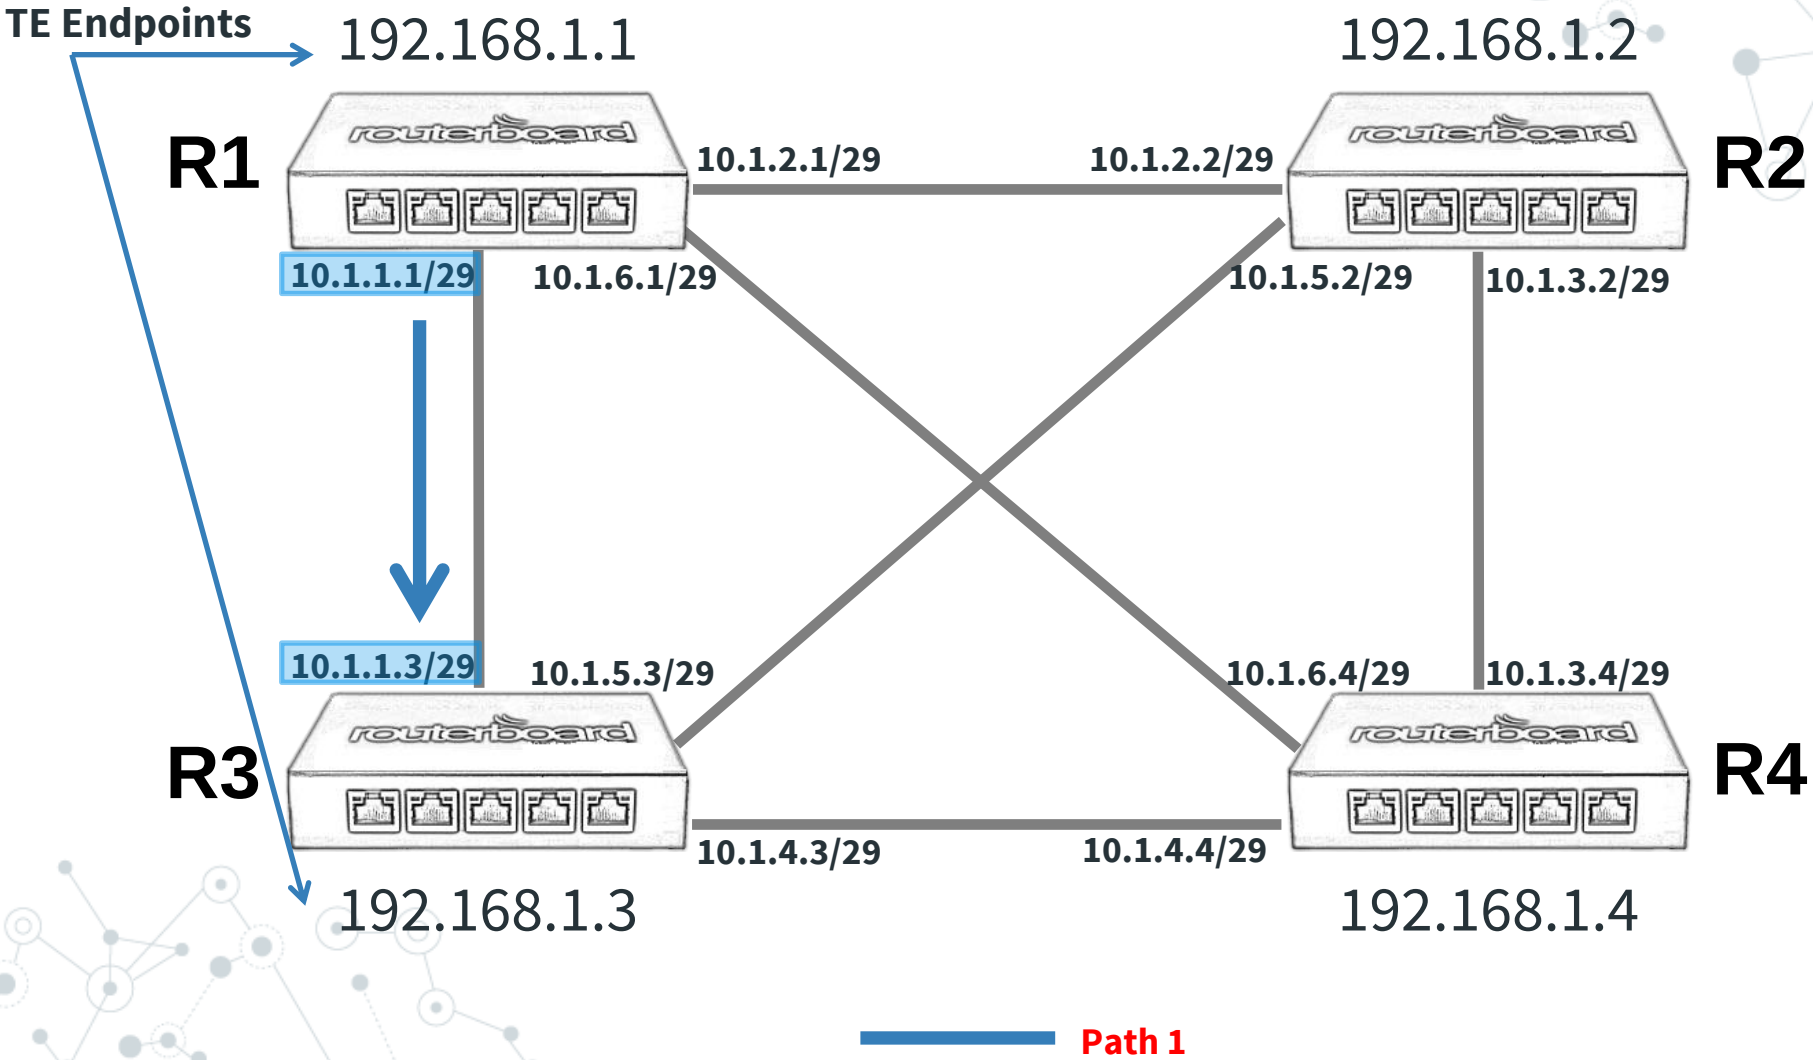

Possible paths from **R1** to **R3**

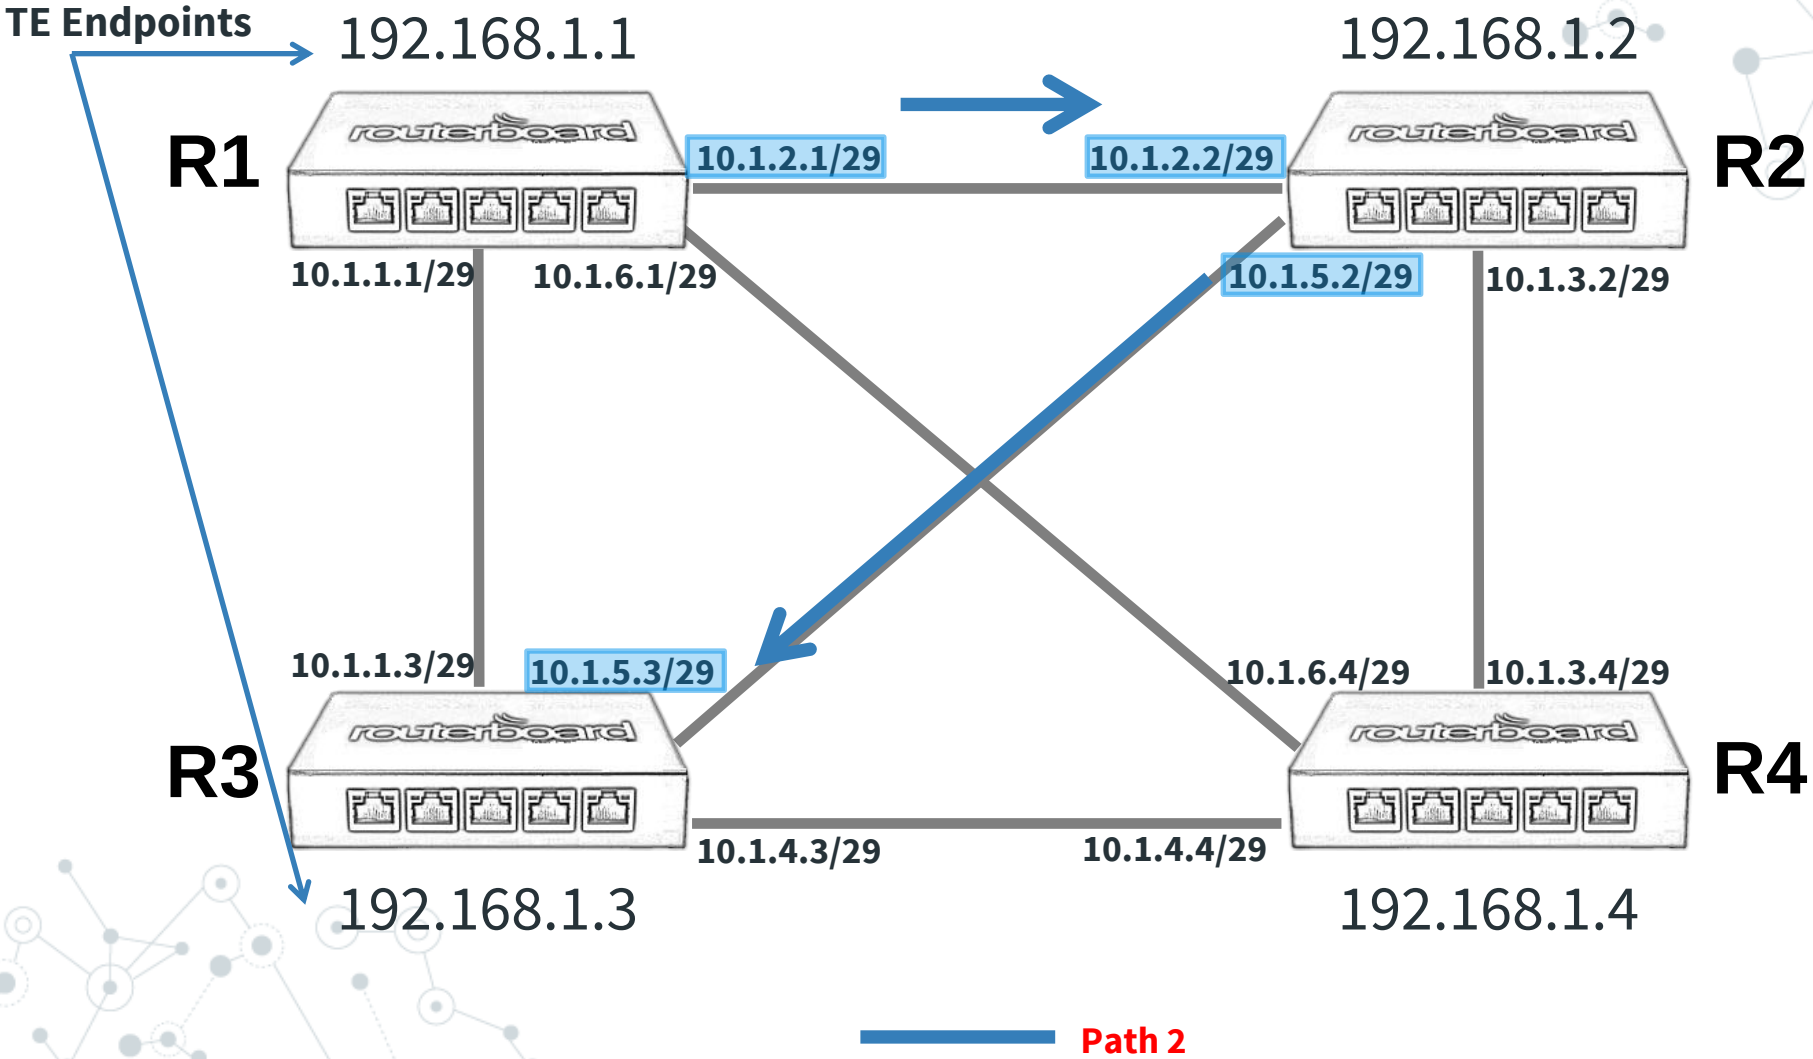

Possible paths from **R1** to **R3**

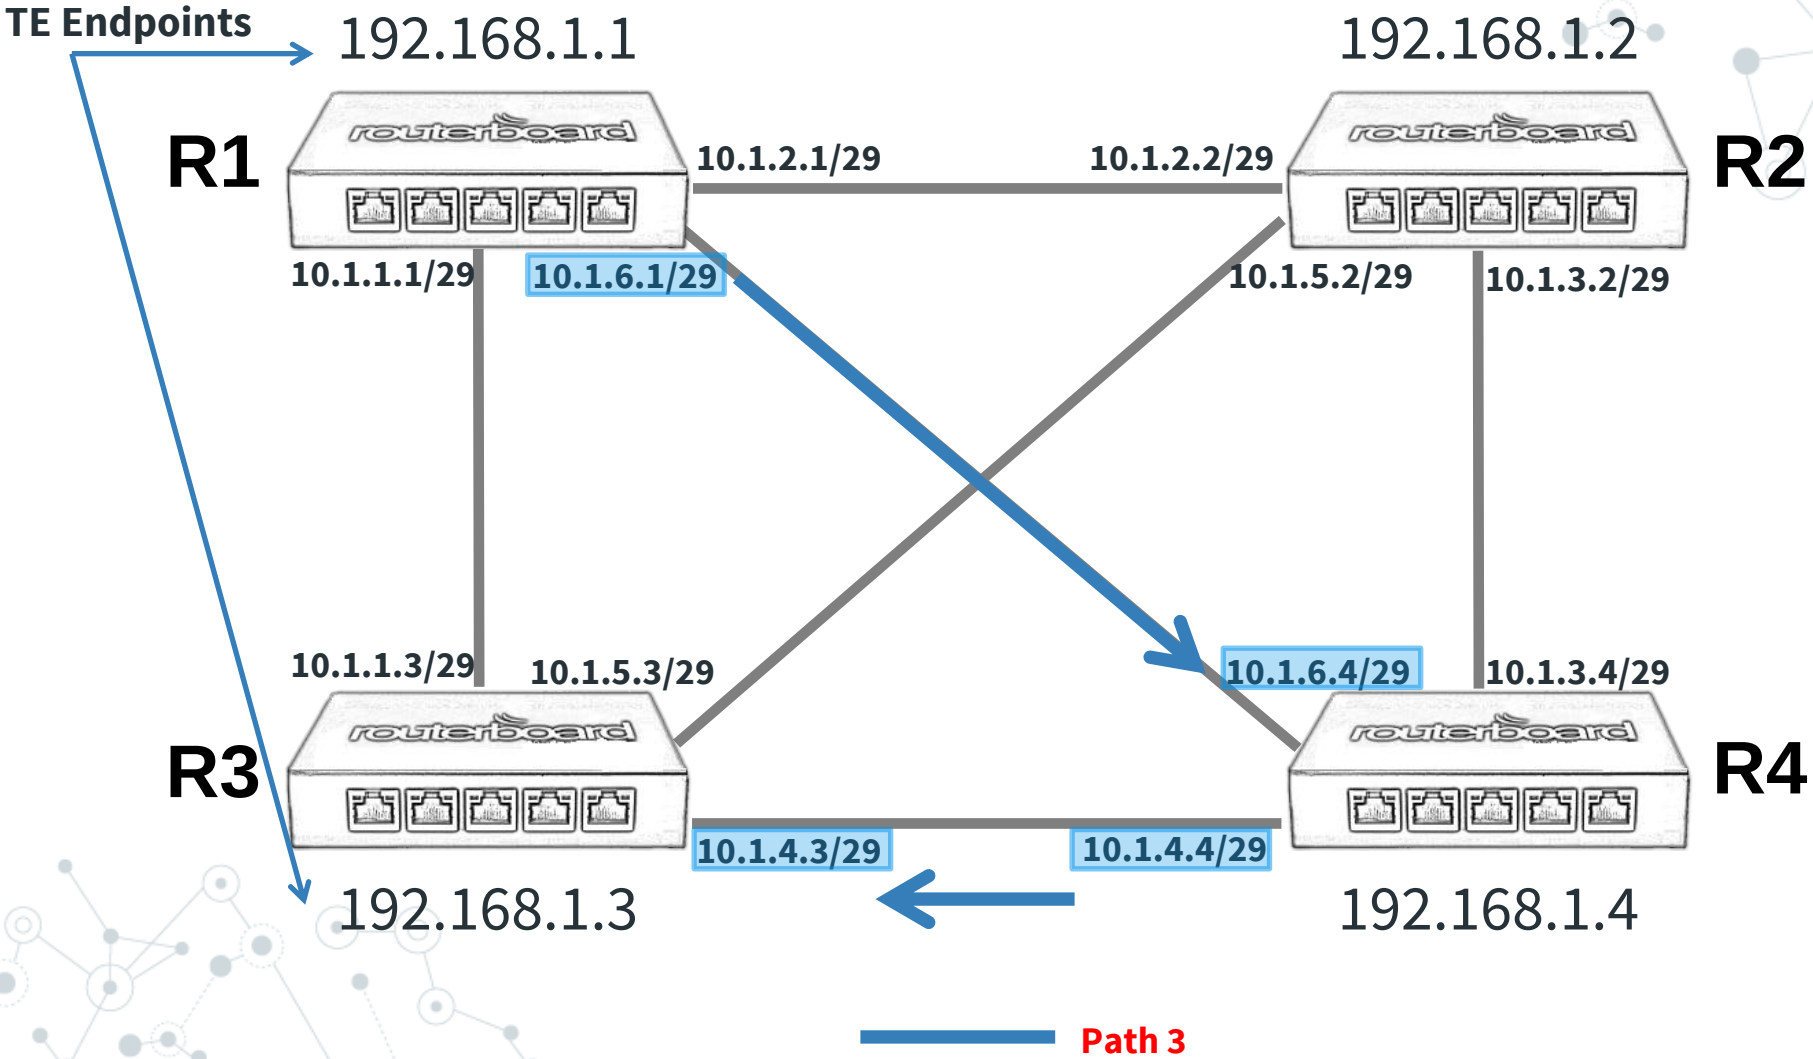

Possible paths from **R1** to **R3**

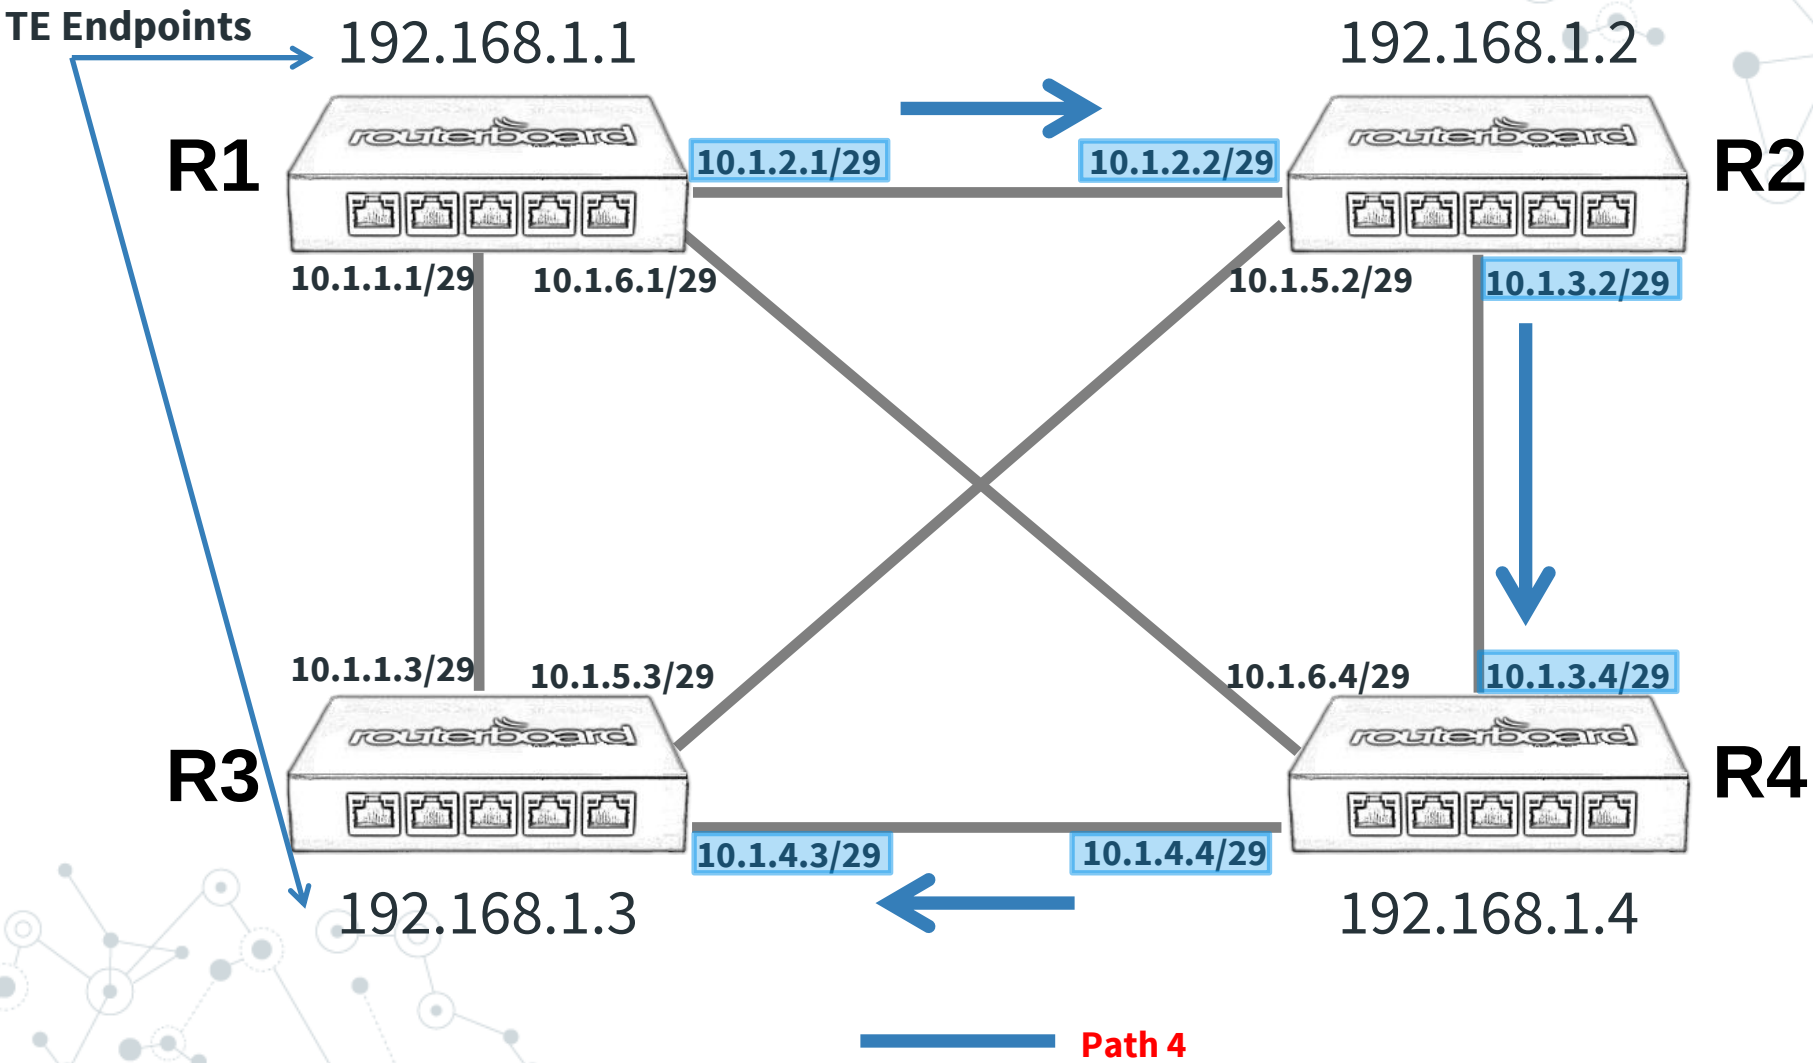

Possible paths from **R1** to **R3**

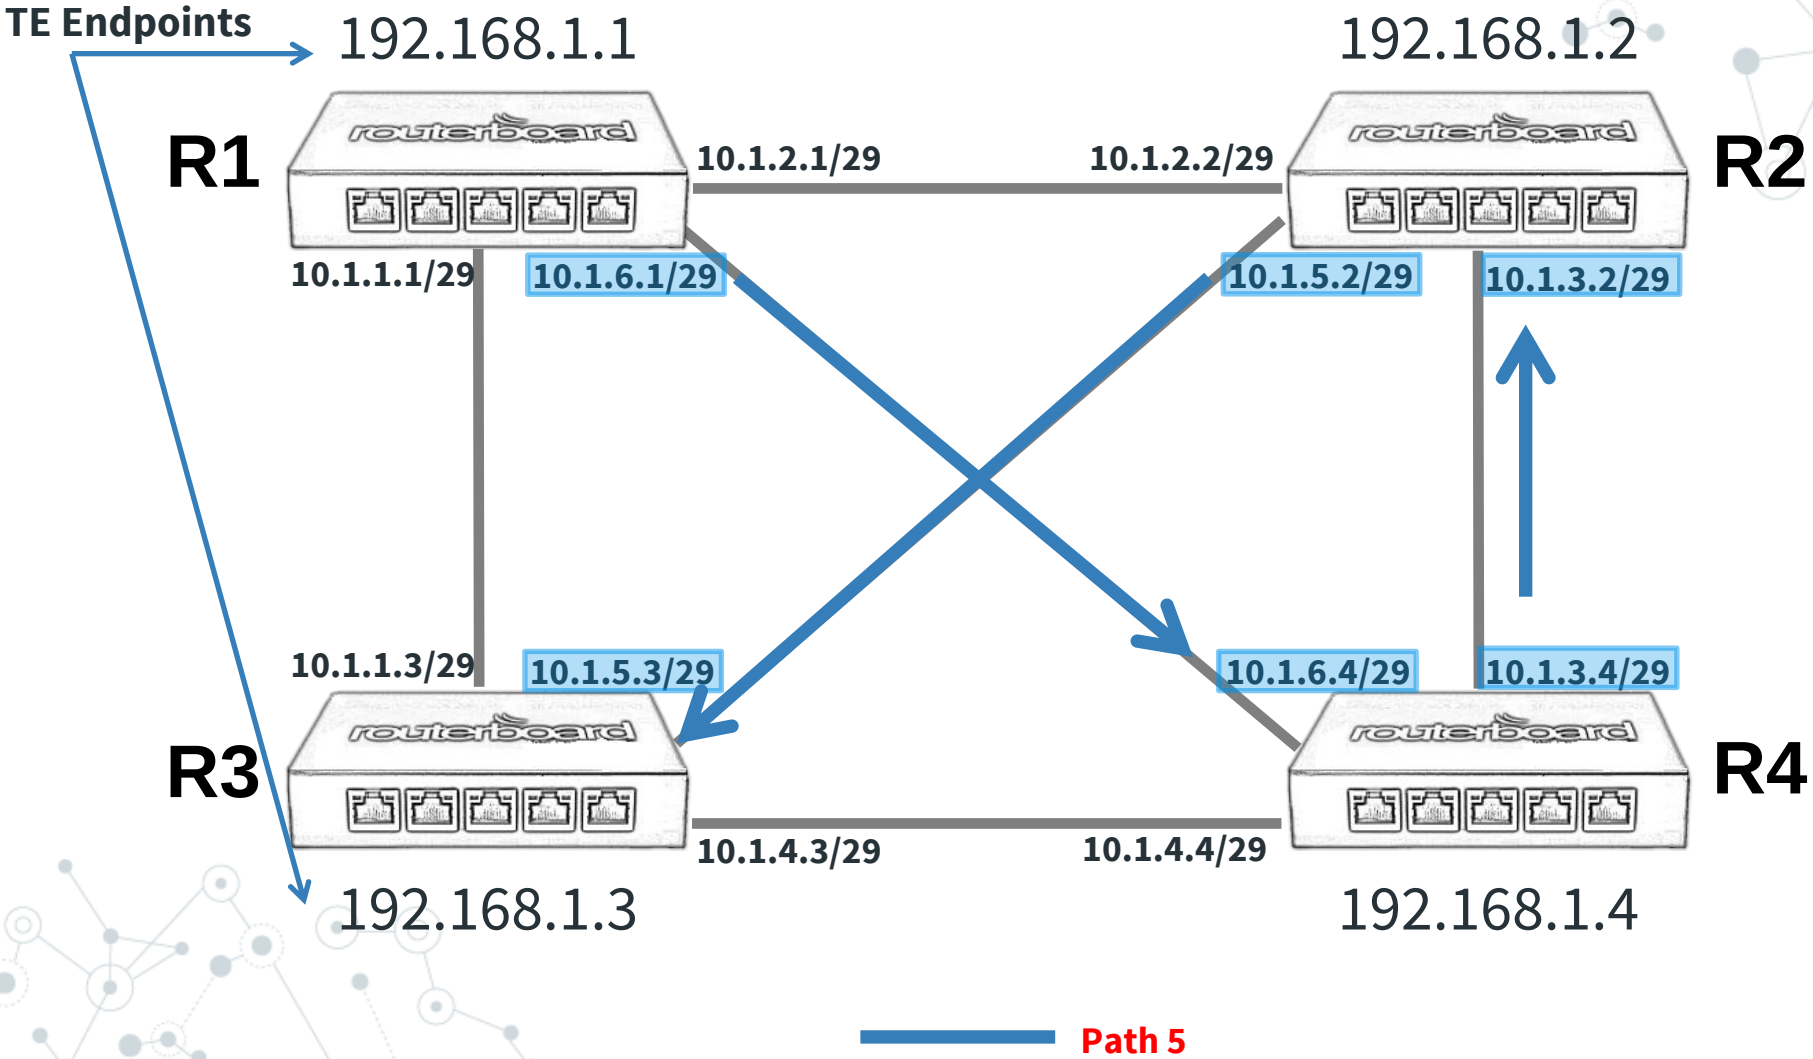

Set up TE unidirectional paths on R1

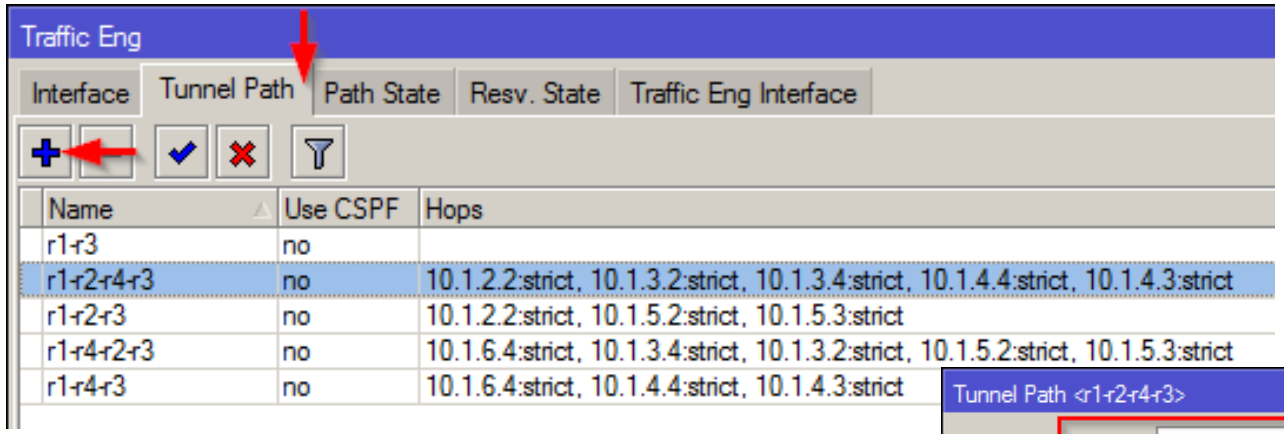

Add all possible paths

Add *strict next-hops* as well as the *strict* local Pref-src towards the next *next-hop*

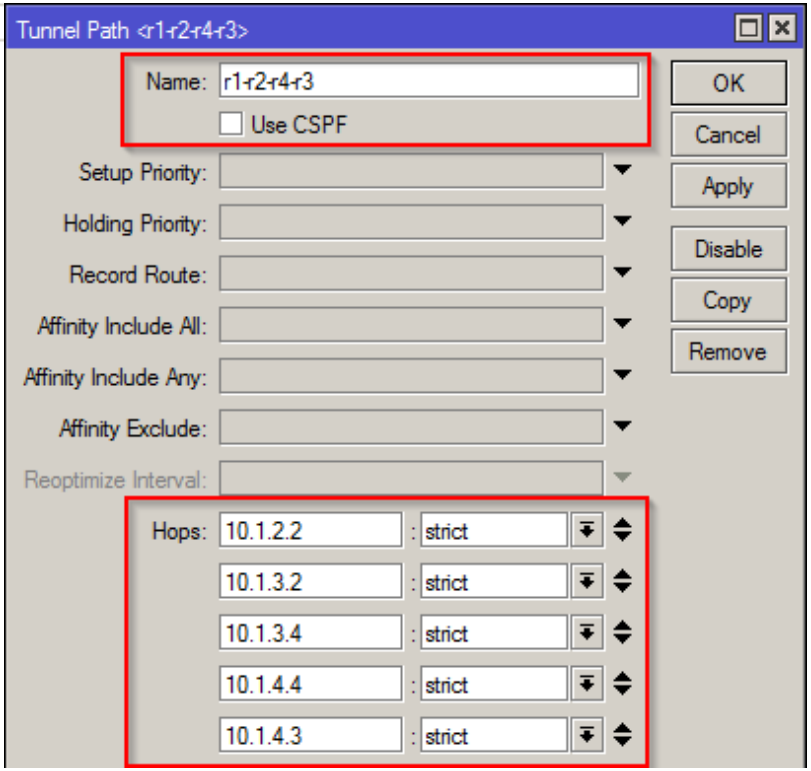

Possible paths from **R3** to **R1**

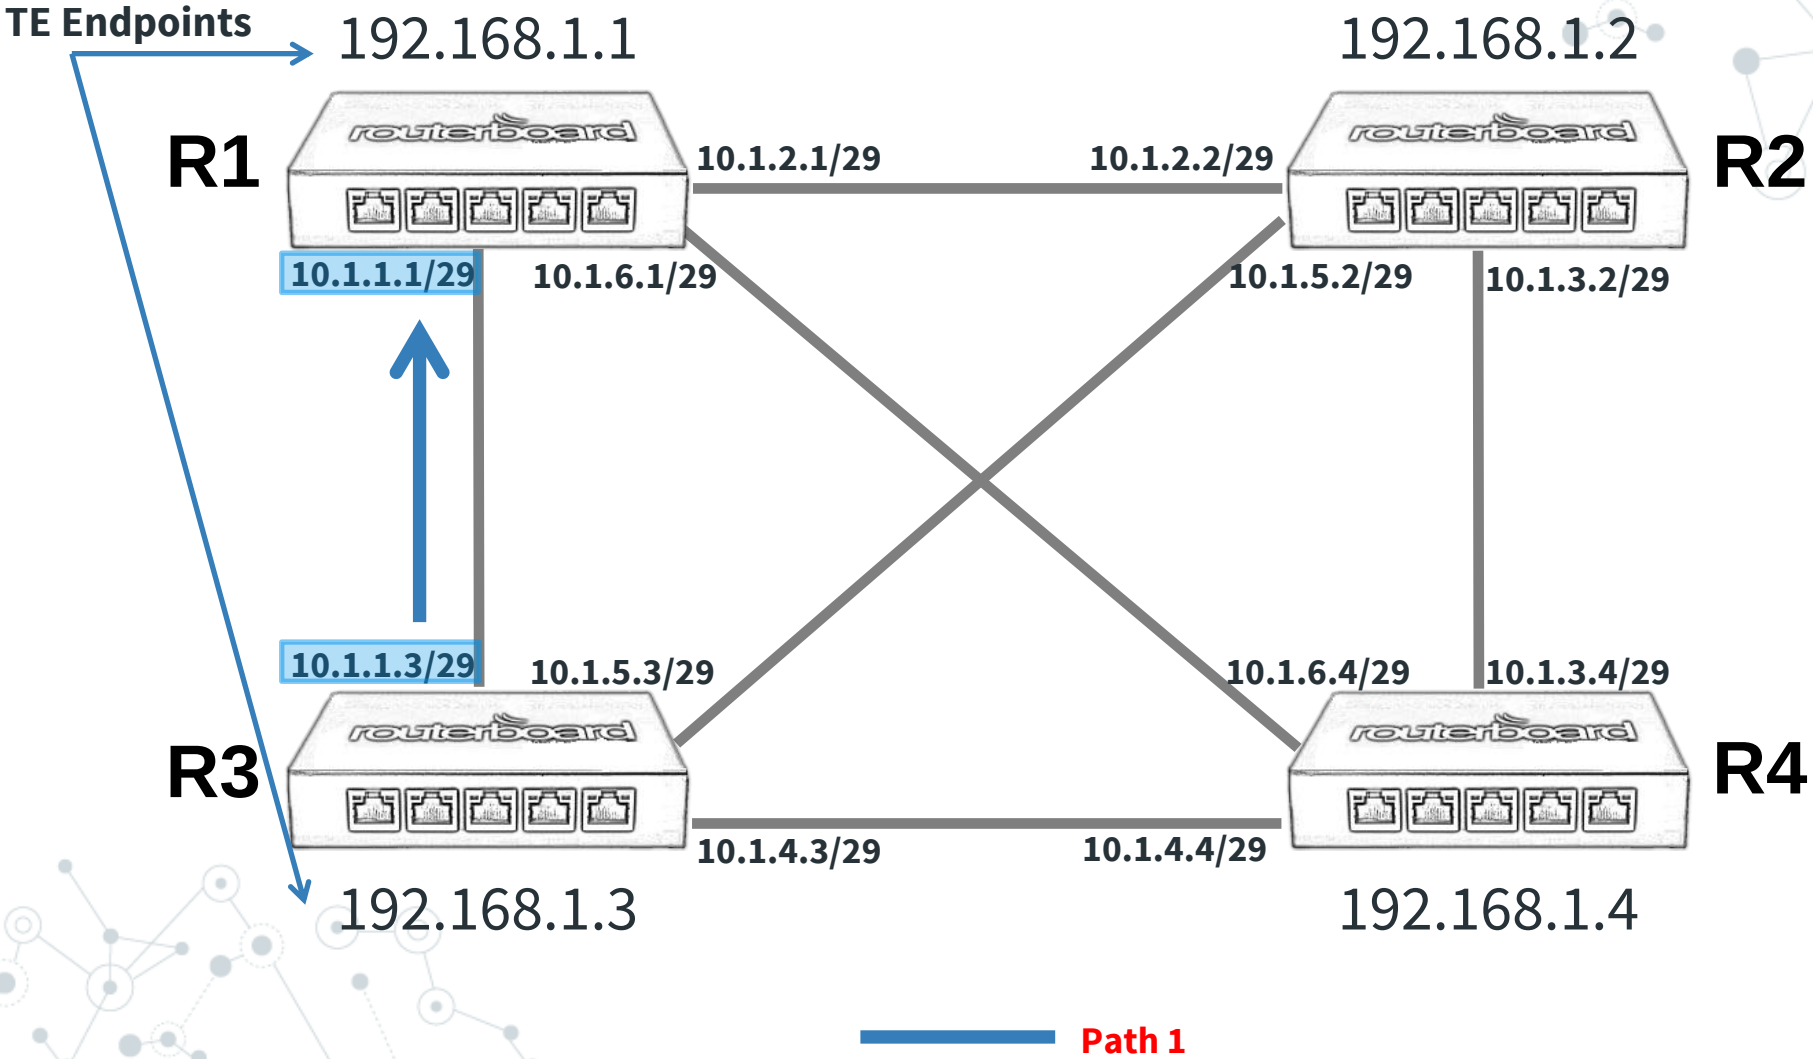

Possible paths from **R3** to **R1**

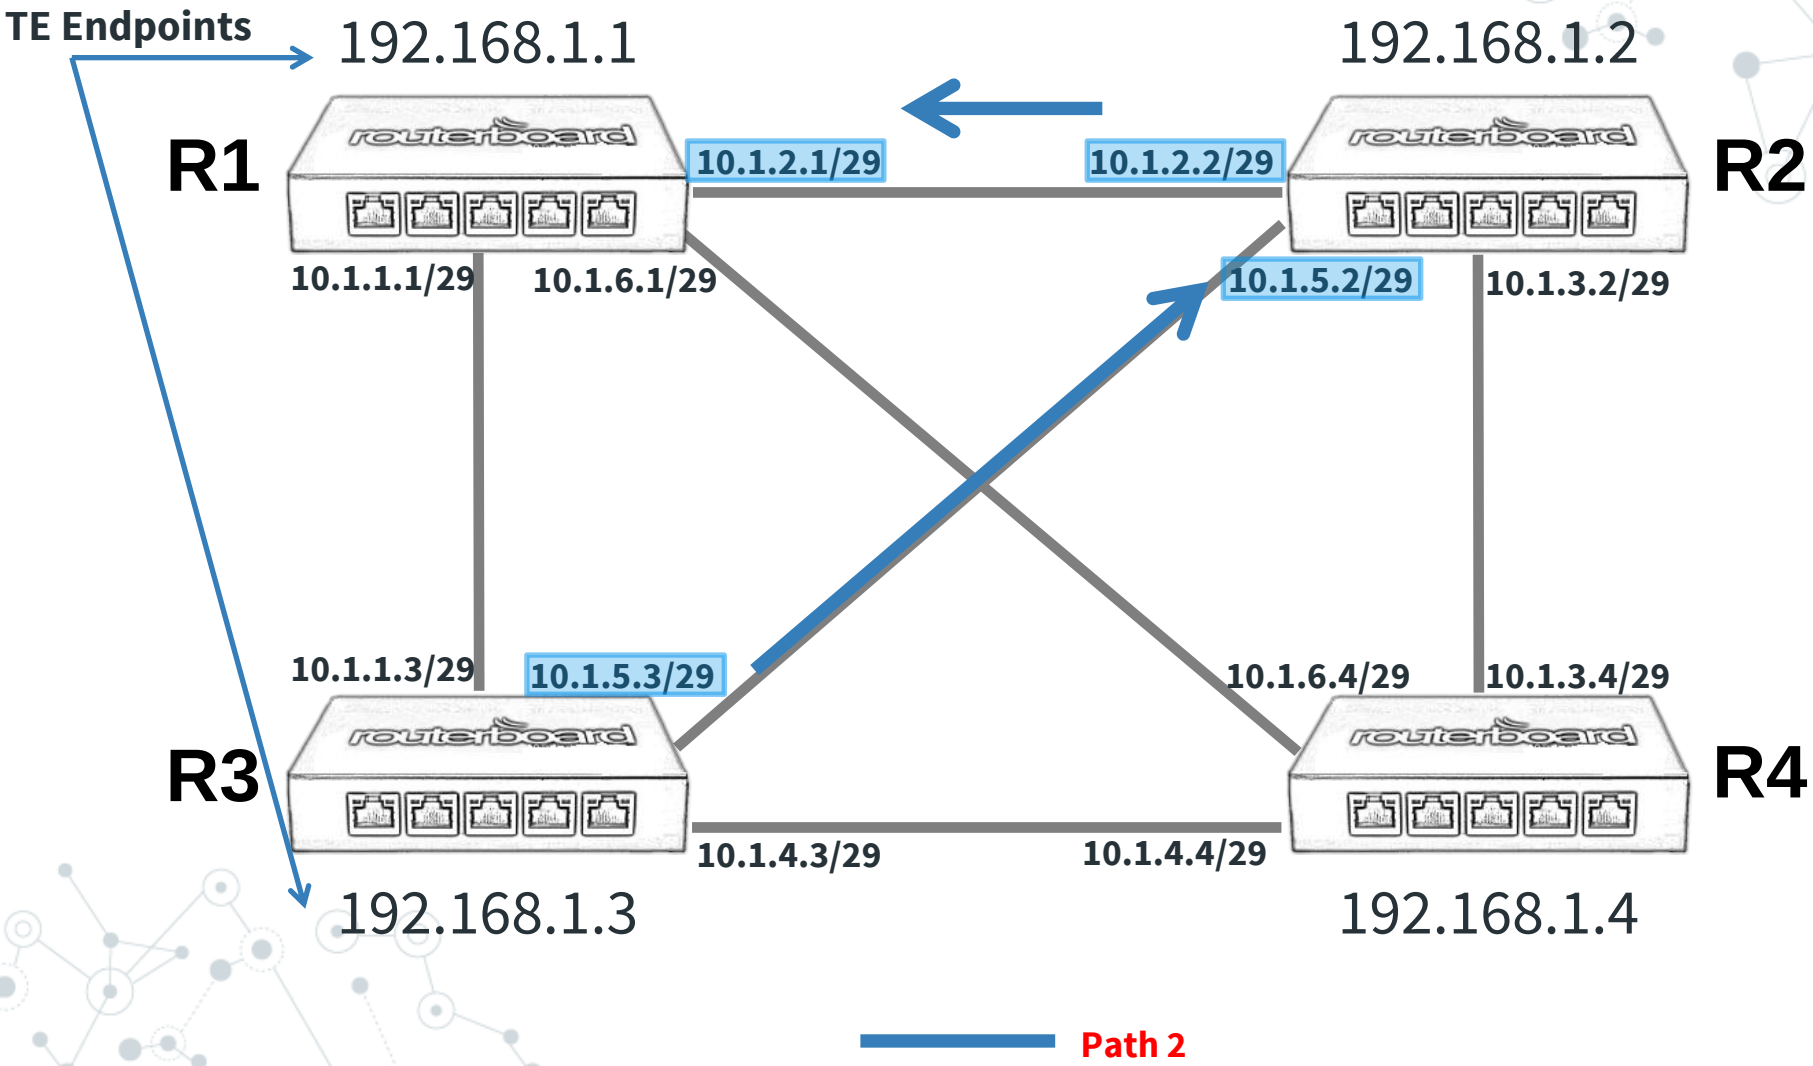

Possible paths from **R3** to **R1**

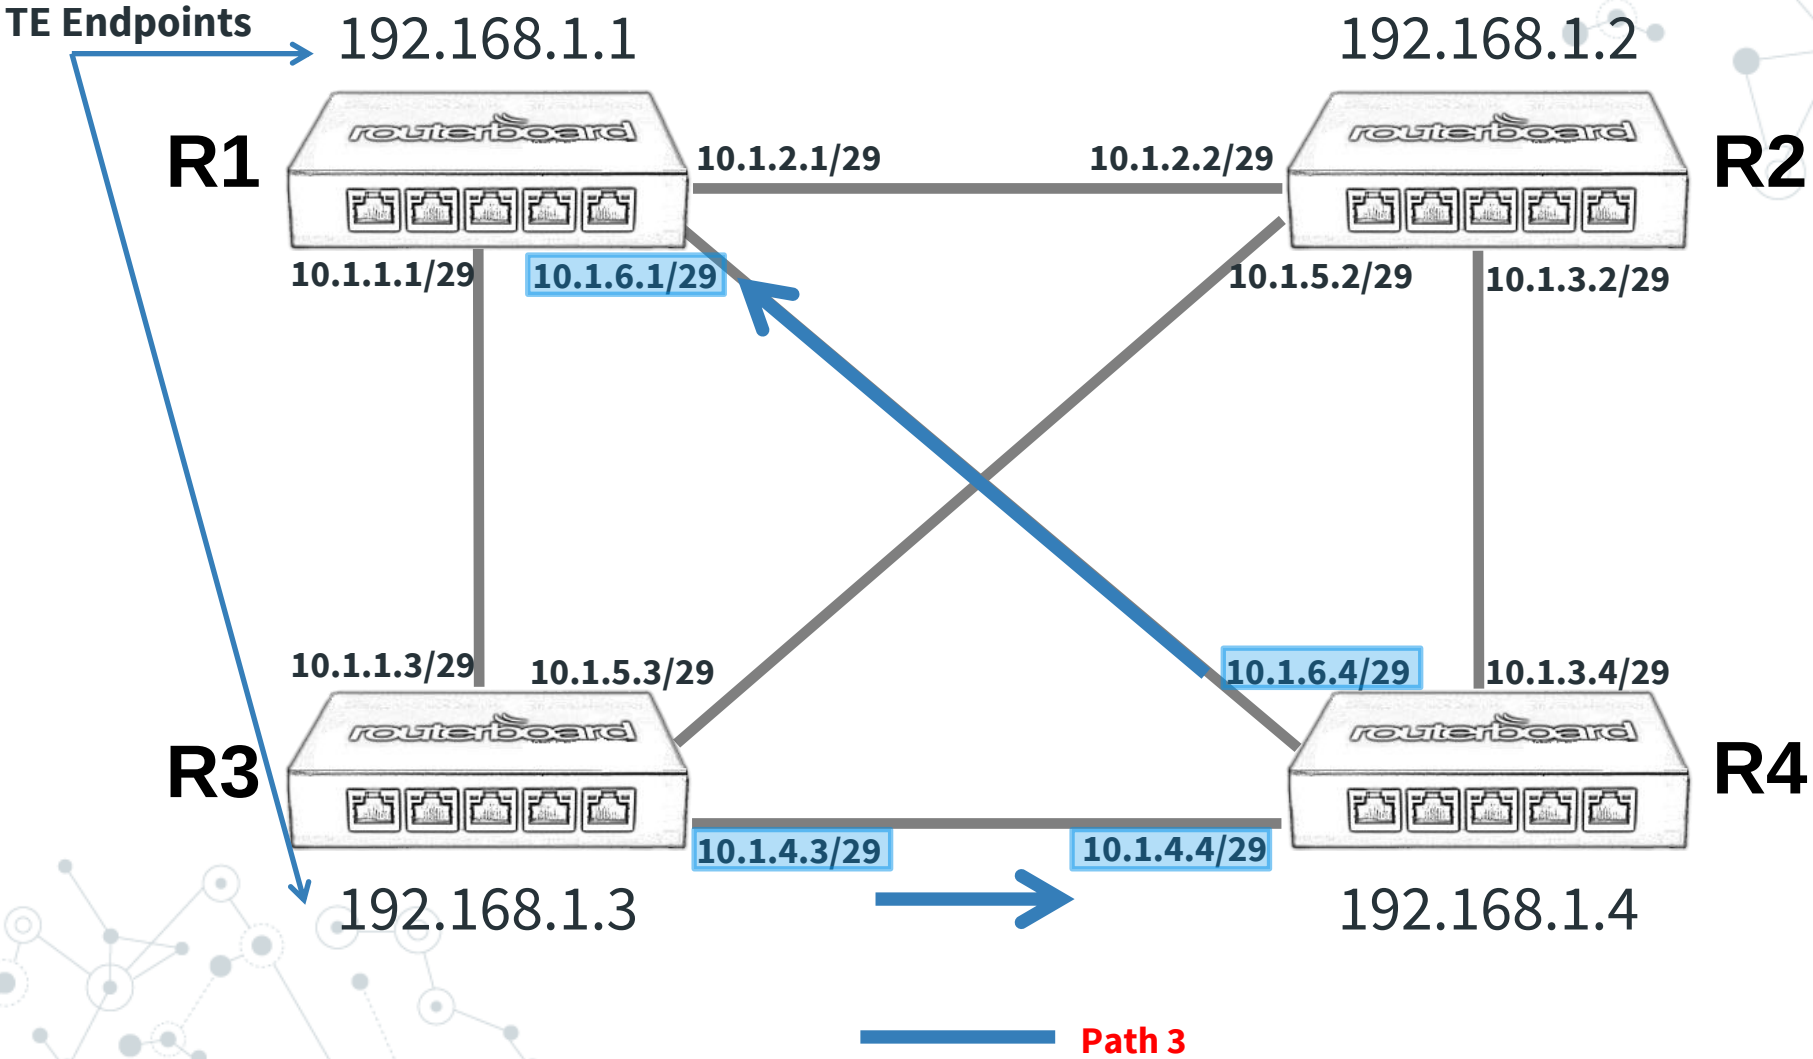

Possible paths from **R3** to **R1**

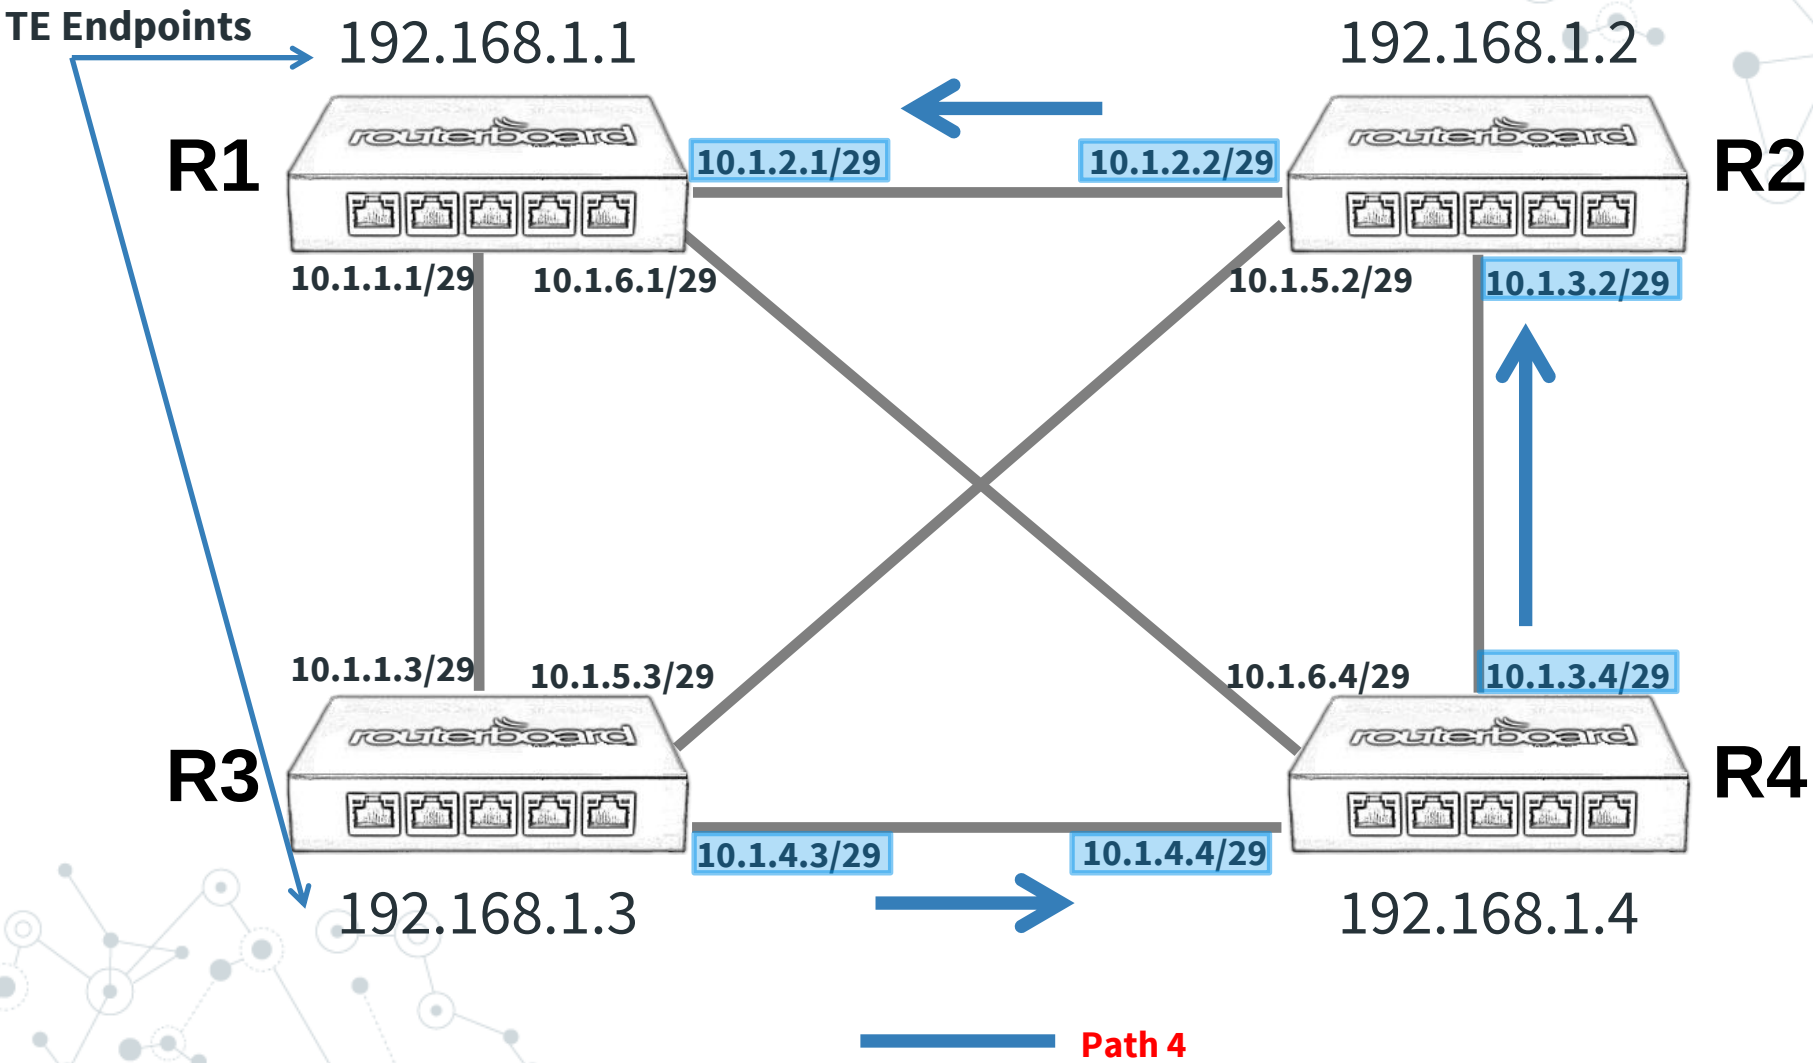

Possible paths from **R3** to **R1**

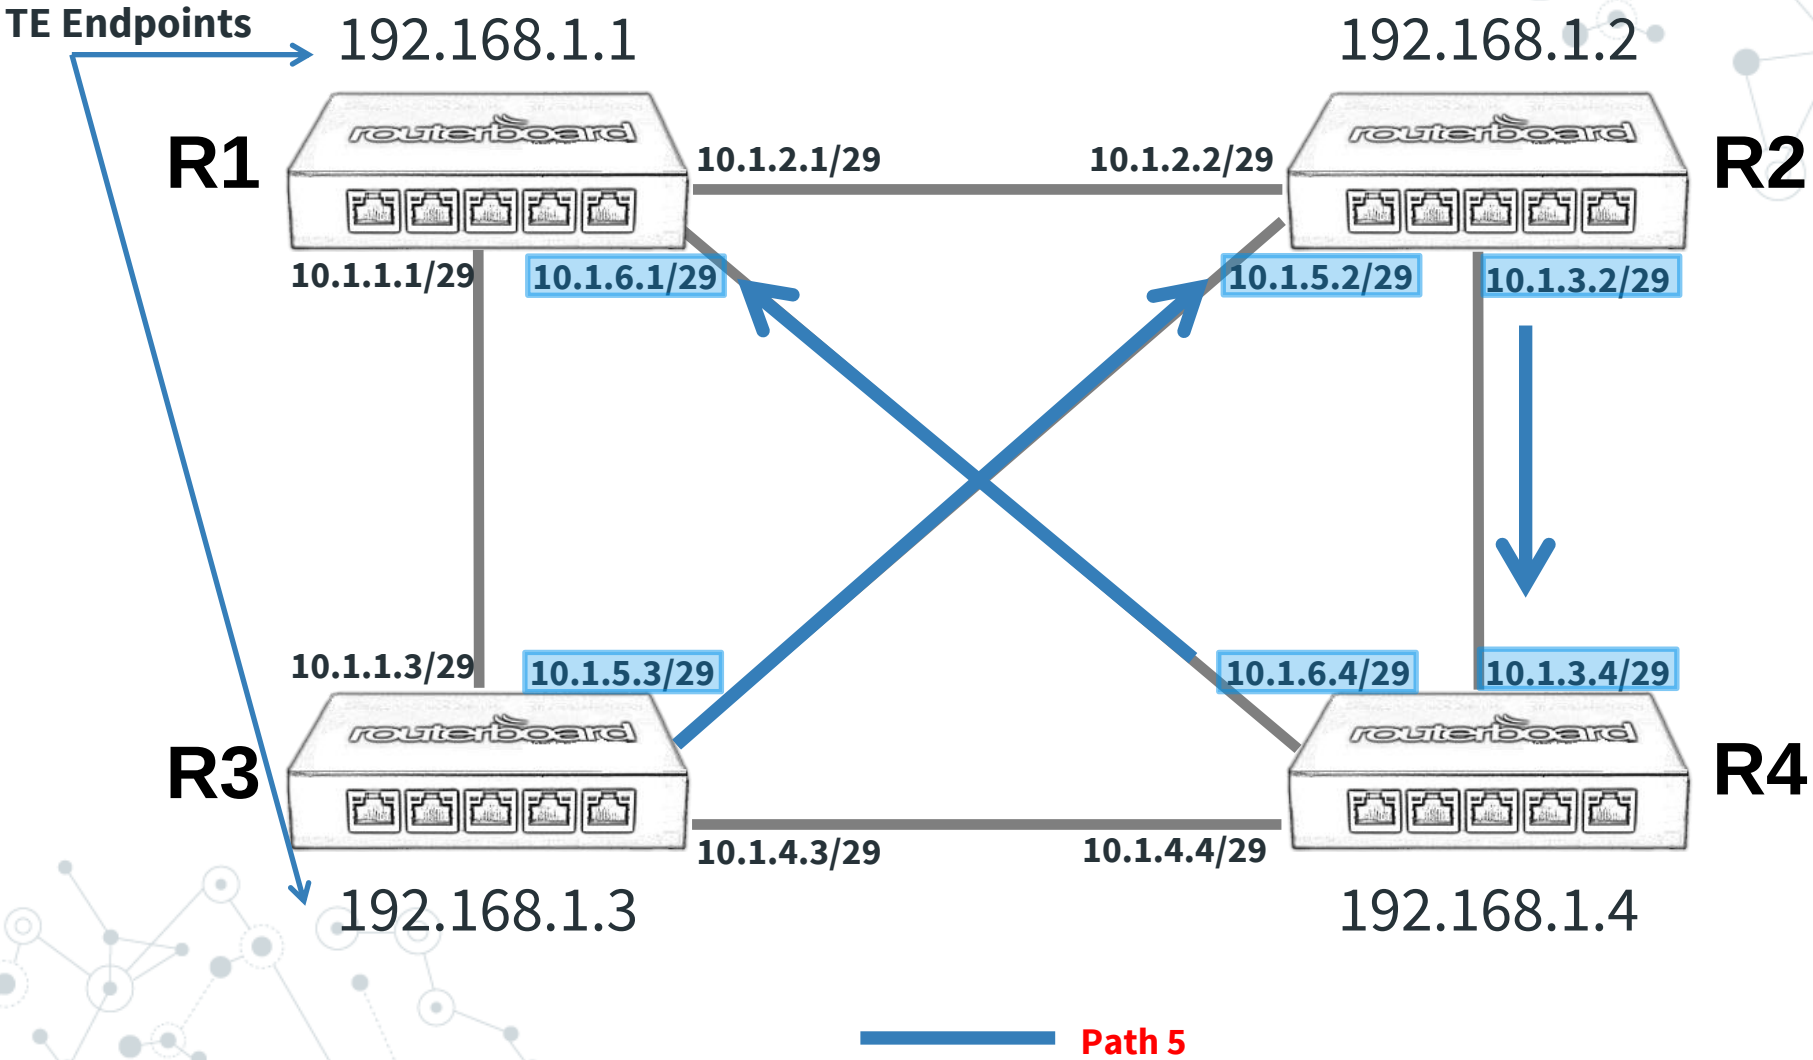

Set up TE unidirectional paths on R3

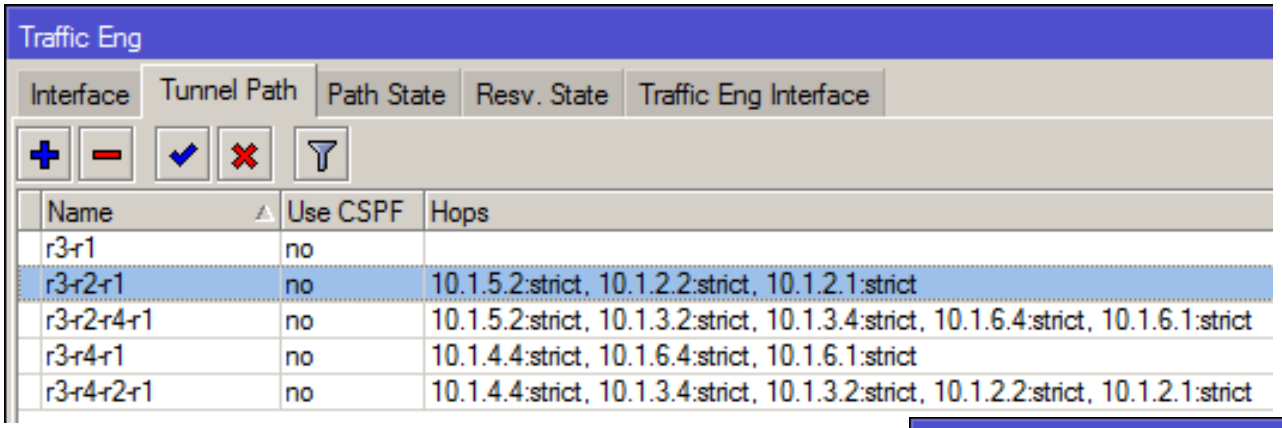

Repeat the exact setup as on **R1**

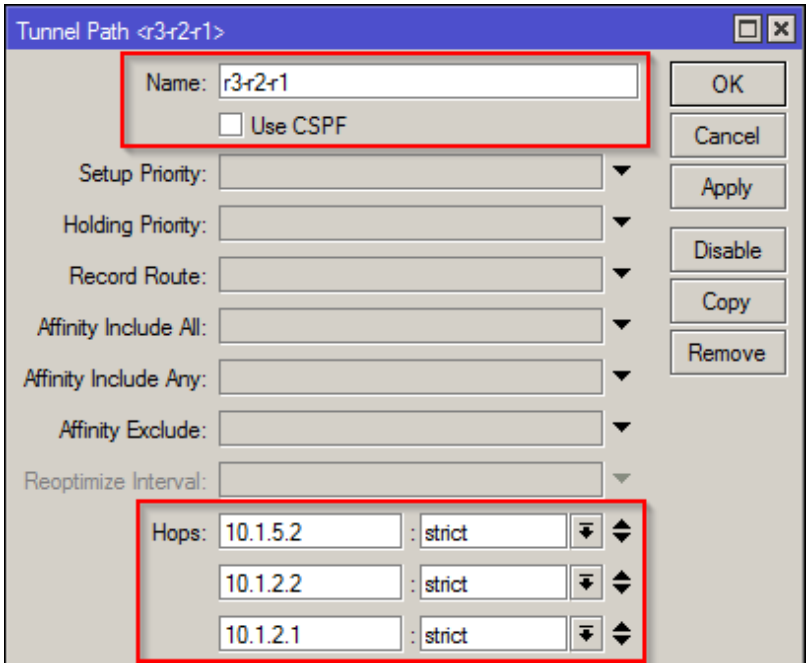

Set up TE tunnel on R1

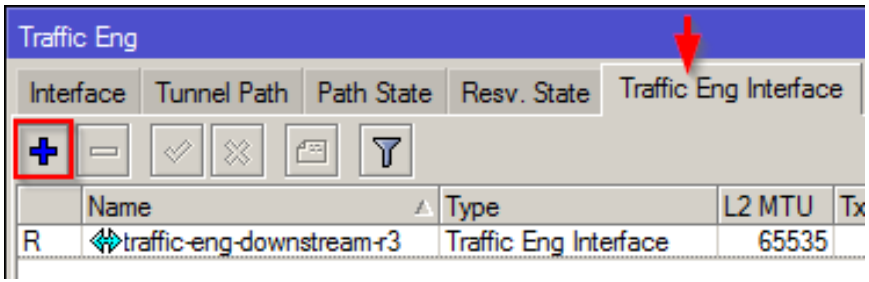

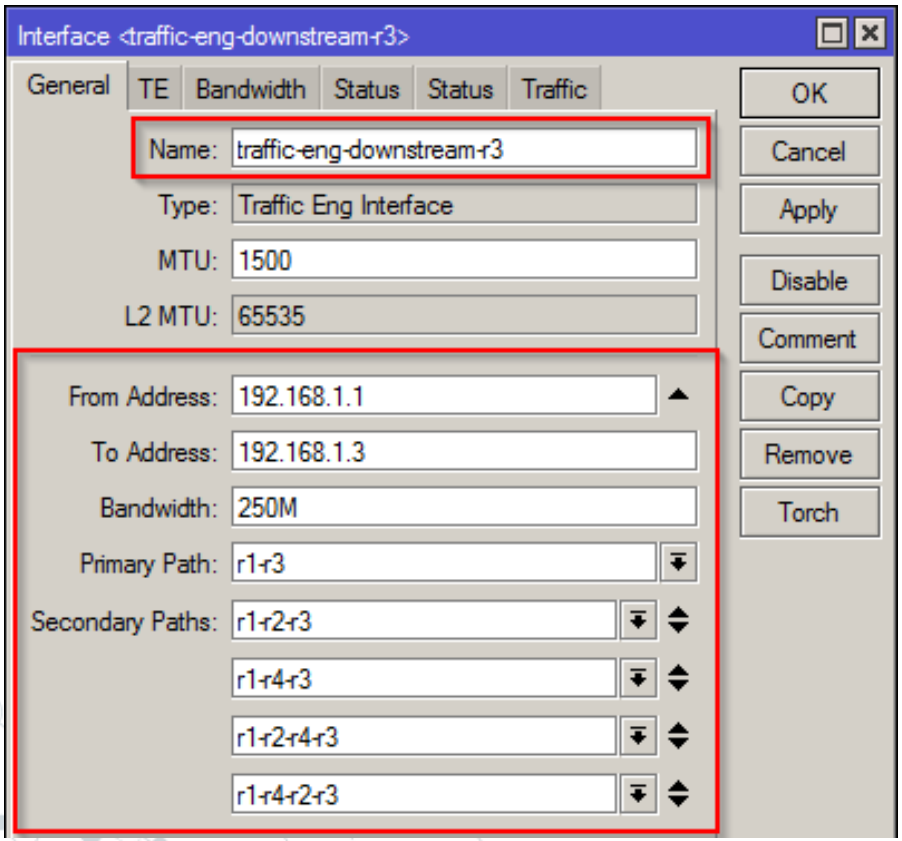

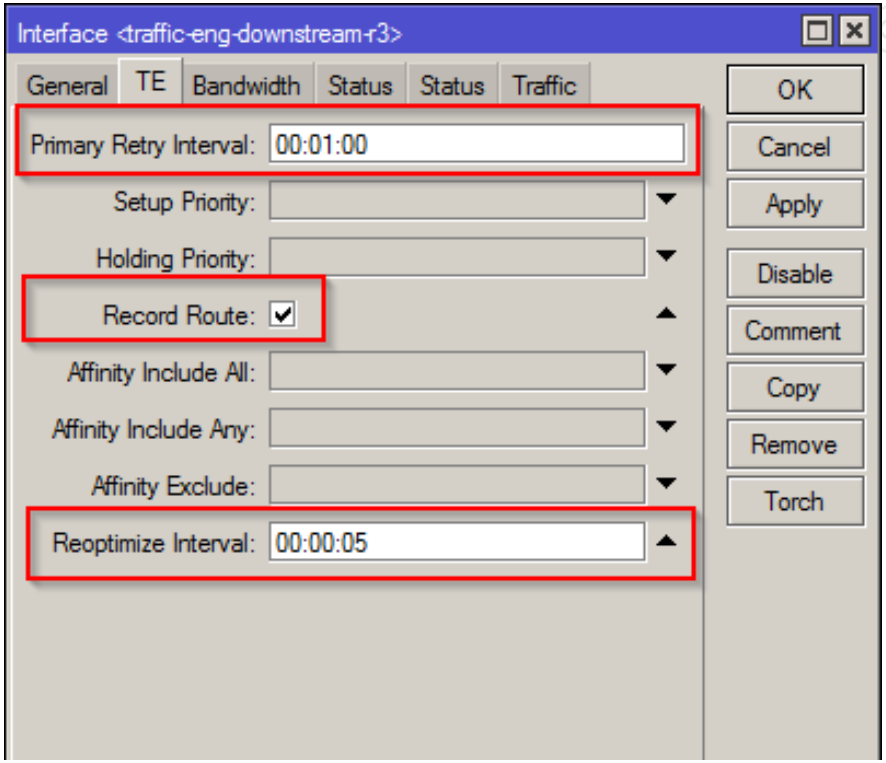

Set up TE tunnel on R3

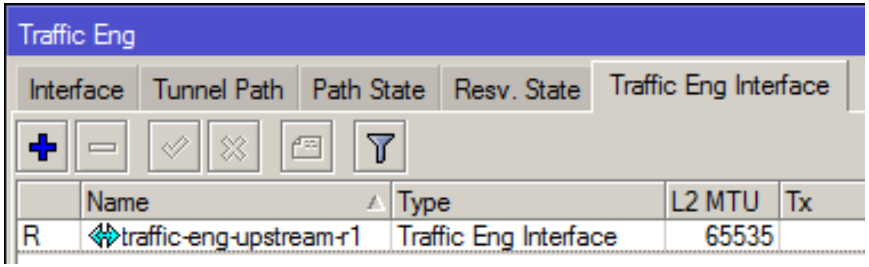

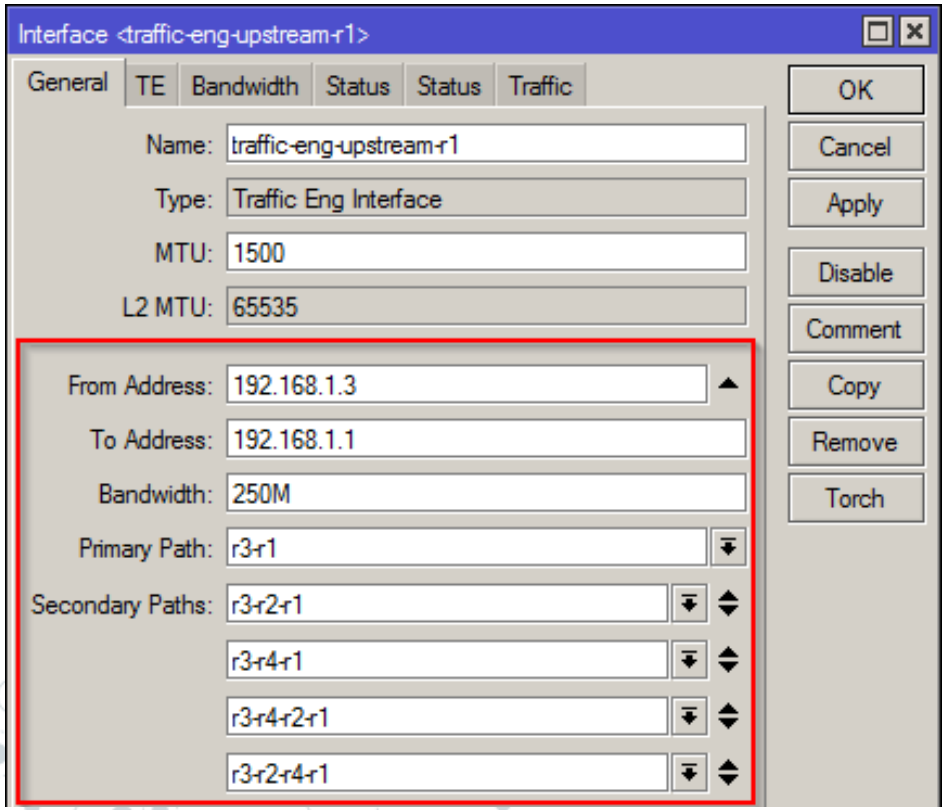

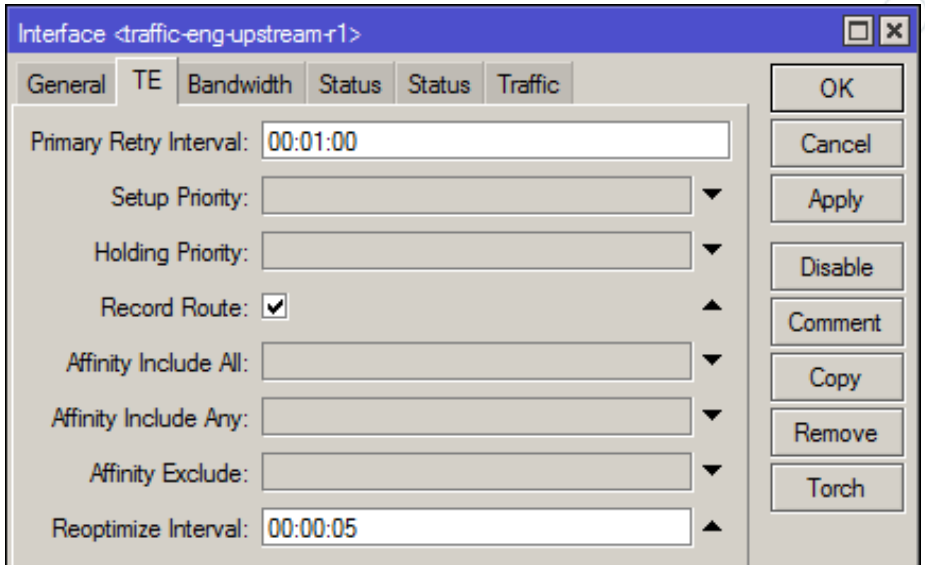

TE effects on R3

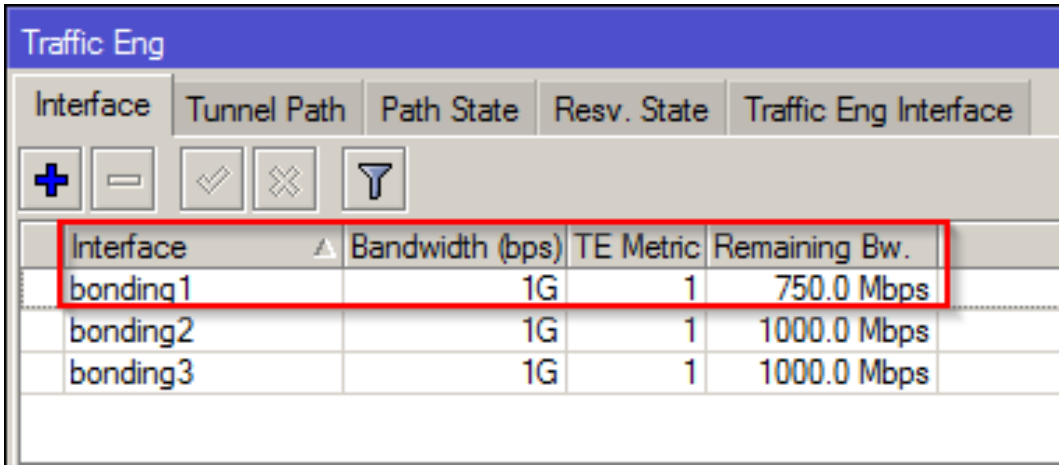

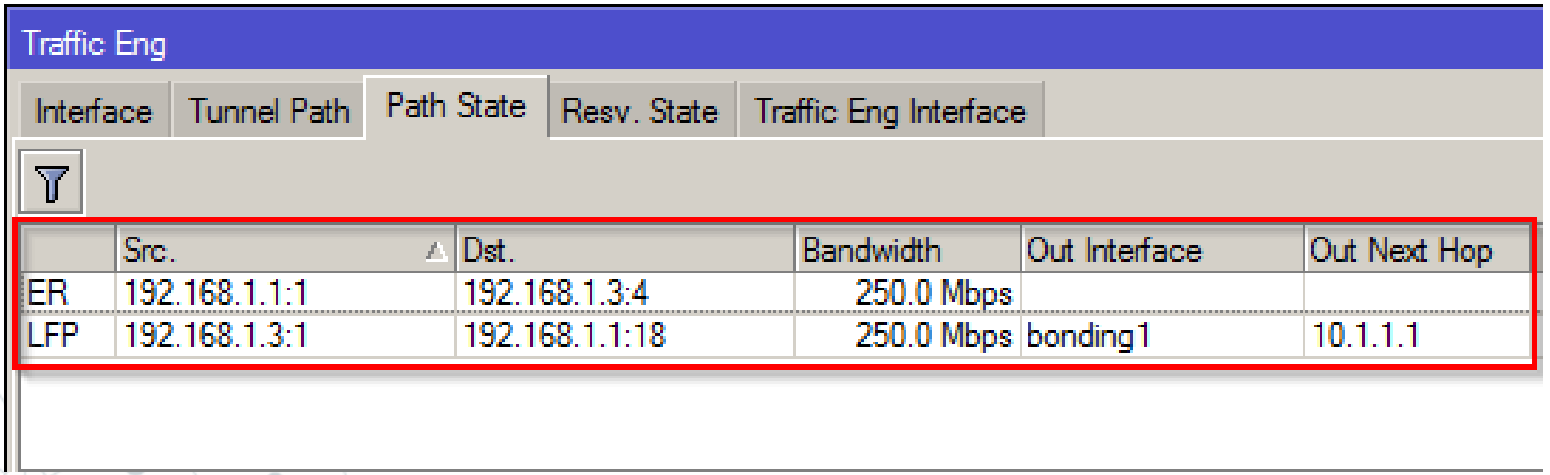

R1 routing (download traffic)

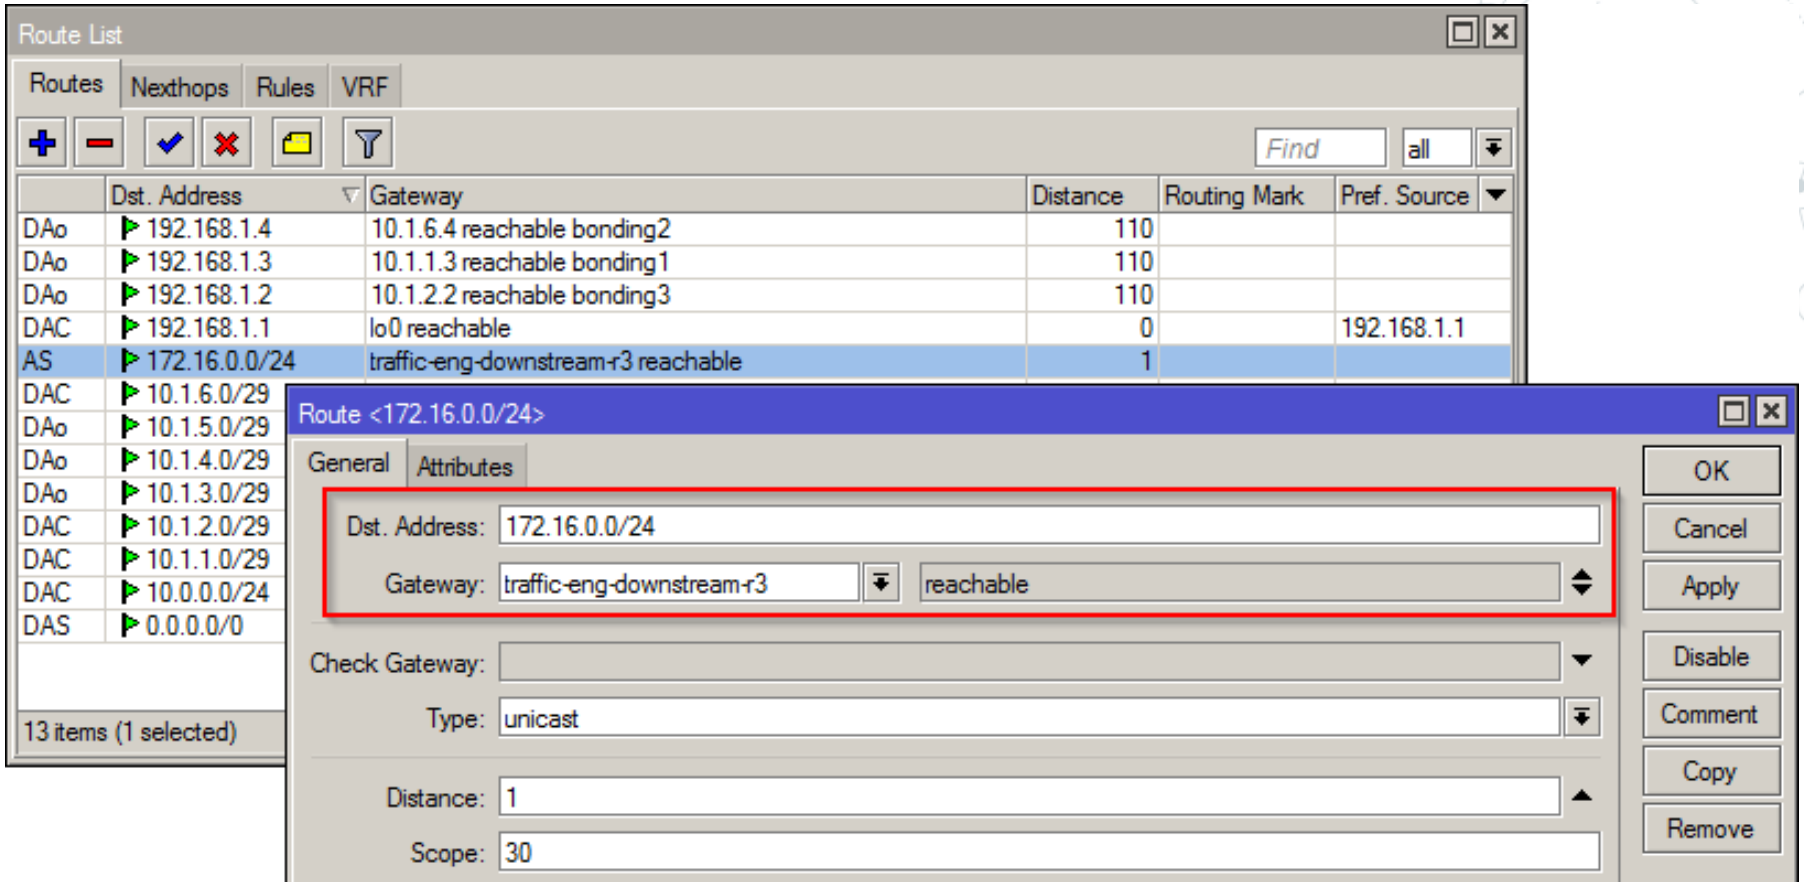

From **R1** push the traffic for 172.16.0.0/24 through the TE tunnel

R3 routing (upload traffic)

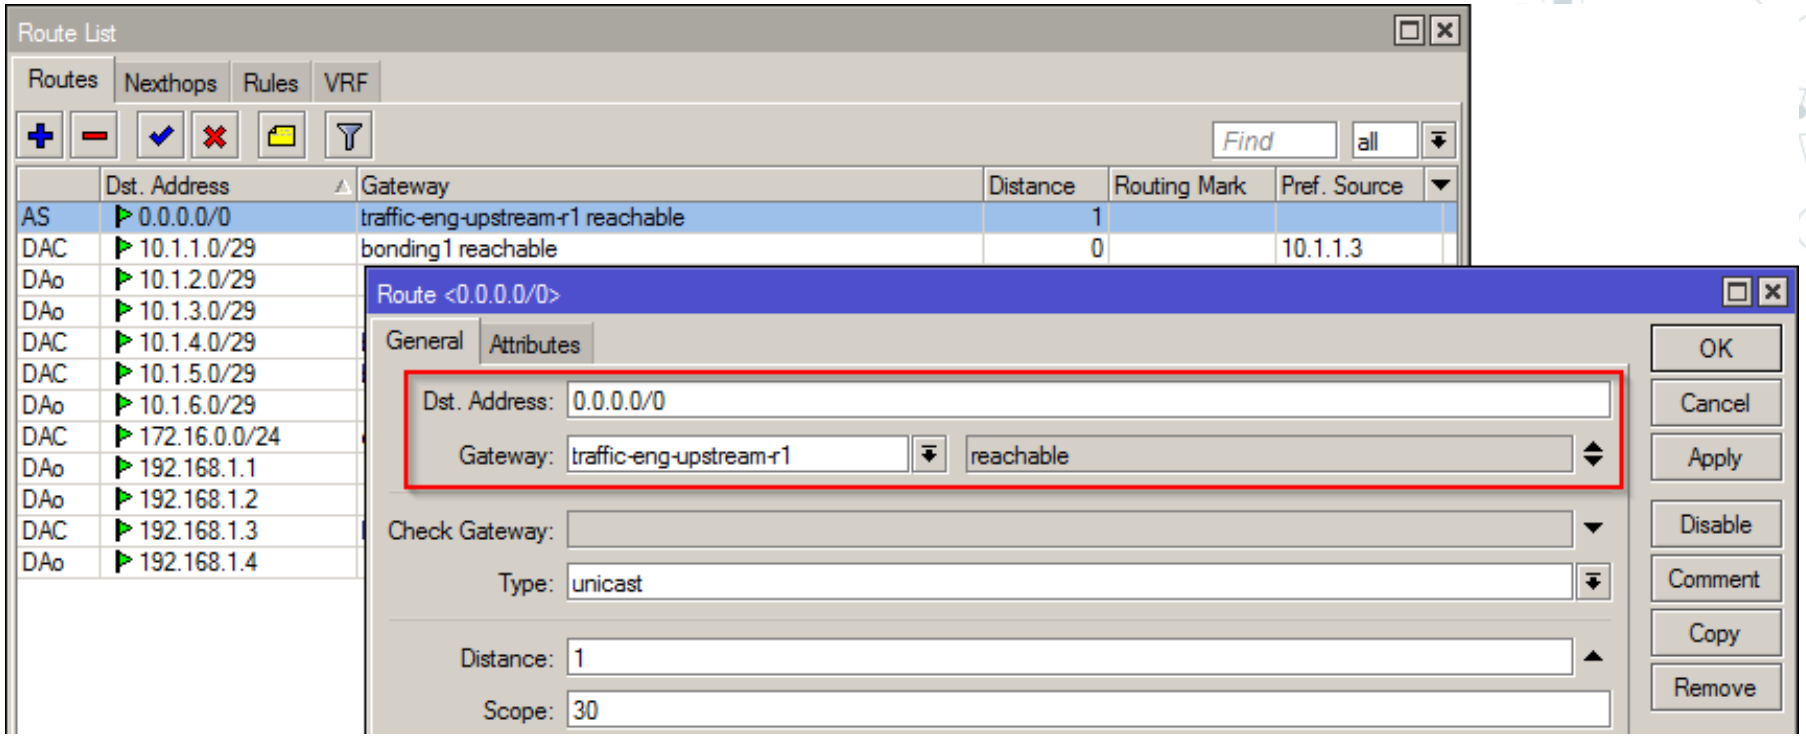

From **R3** push all traffic through the TE tunnel

Final checks and test

- Check if DNS is working on all routers
- Check that MPLS TE opaque LSAs exist
- Check if you need NAT on the border

- Test if it’s working from customer

- Verify that tunnels are pushing traffic correctly (only TX) and that reservations are met on the bonding interfaces that have the active paths attached

Customer test

Terminal

```
[admin@Management] > /tool traceroute mikrotik.com
 # ADDRESS                                  LOSS SENT    LAST     AVG    BEST   WORST
 1 172.16.0.1                                0%   15   0.2ms     0.5     0.2     5.1
 2 192.168.1.1                               0%   15   0.3ms     0.4     0.2     1.5
 3 10.0.0.254                                0%   15   0.6ms     1.3     0.6     4.4
 4 85.204.99.149                             0%   15   5.4ms     6.5     5.4    11.1
 5 81.183.1.104                              0%   15  20.5ms      22    19.1    36.7
 6 80.81.194.165                             0%   15  68.9ms      71    68.3      75
 7 195.13.224.86                             0%   15  70.5ms    71.1    70.5      73
 8 159.148.147.196                           0%   14  68.2ms    69.4    68.2    71.8
-- [Q quit|D dump|C-z pause]
```

R1 & R3 tunnels

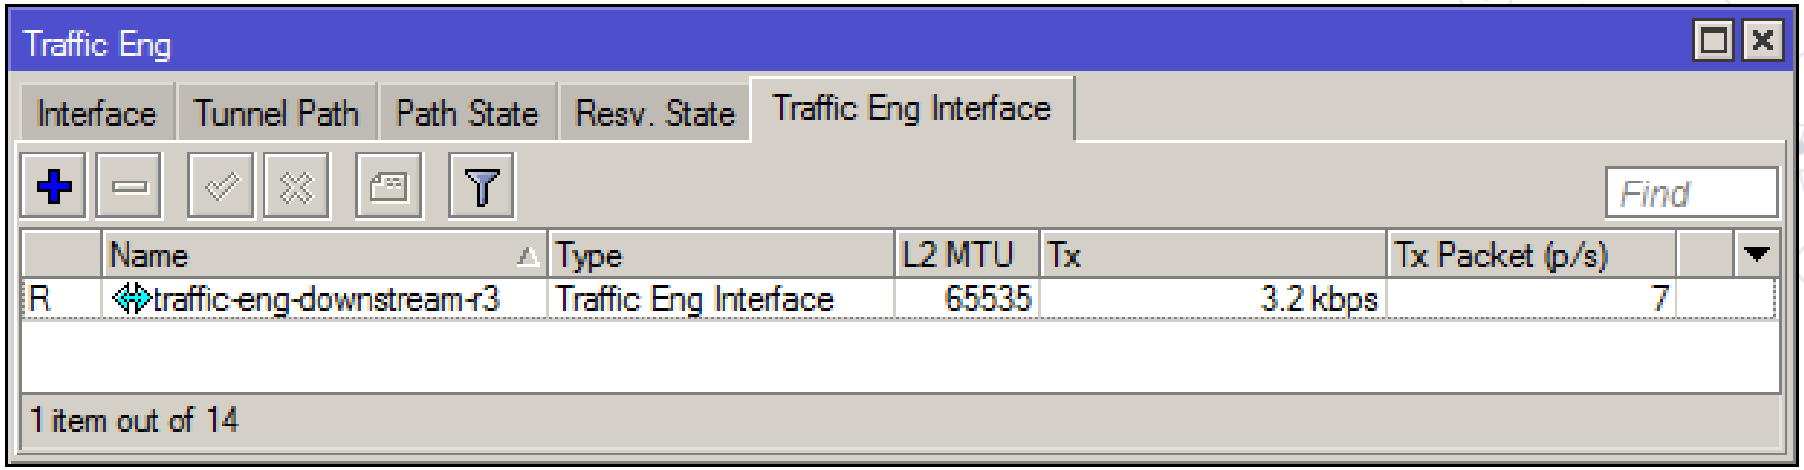

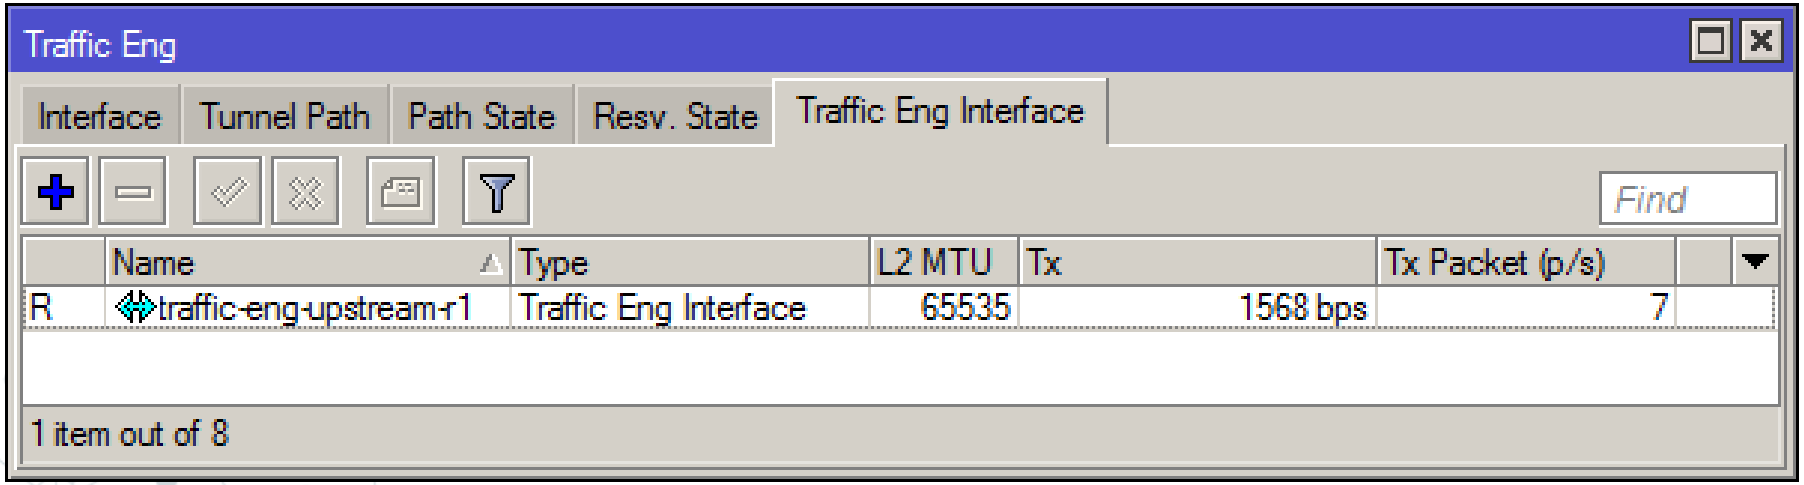

Thank you!

Any questions?

Contact me

mihai.saftoiu@tier.ro

+4 0751-160-169