MUM

зелёнаяточка

смотри. слушай. говори

Интеграция MikroTik в существующую сеть MPLS, построенную на Cisco.

MUM2019
Rostov-on-Don

Владимир Кузнецов
ООО «Зеленая точка Липецк»

Об авторе

- Кузнецов Владимир Петрович
- Россия, Липецк

- Ведущий инженер службы СПД.

  ООО «Зеленая точка Липецк»

Для чего применяется MPLS

- Организация L2VPN
- Организация L3VPN

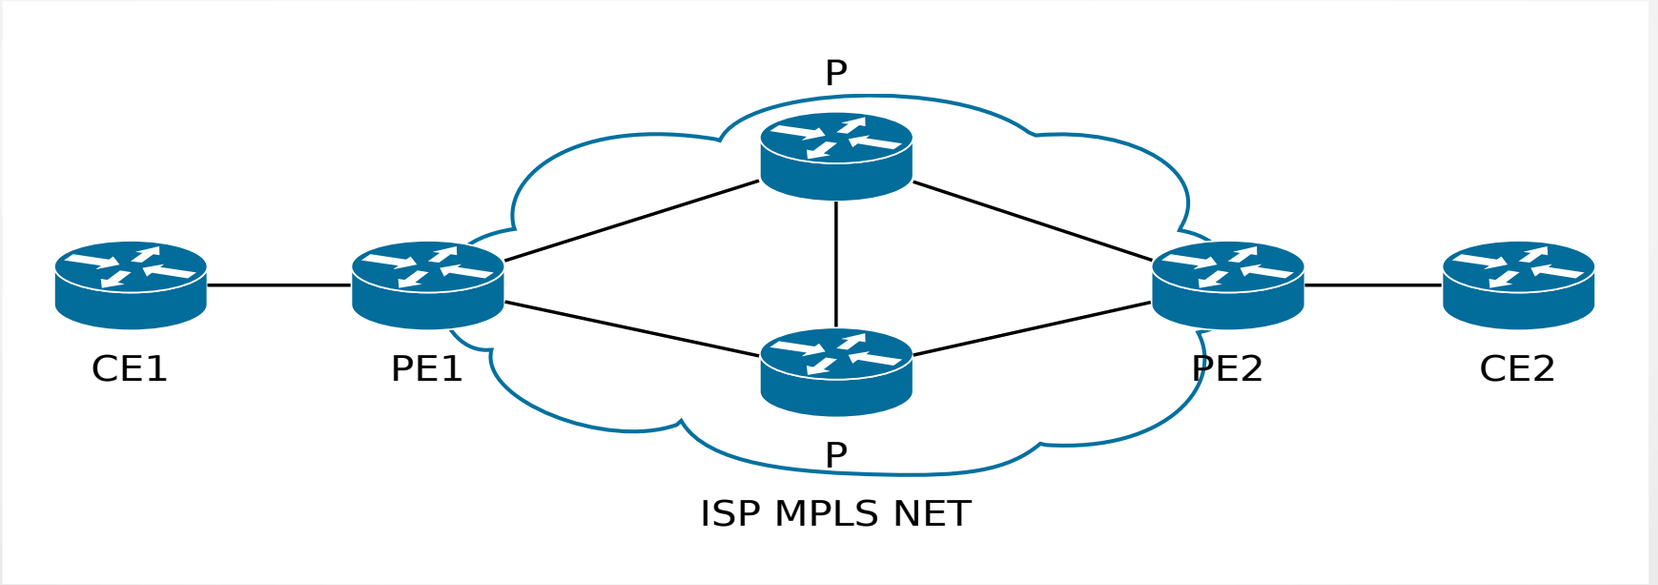

MPLS. Packet flow diagram

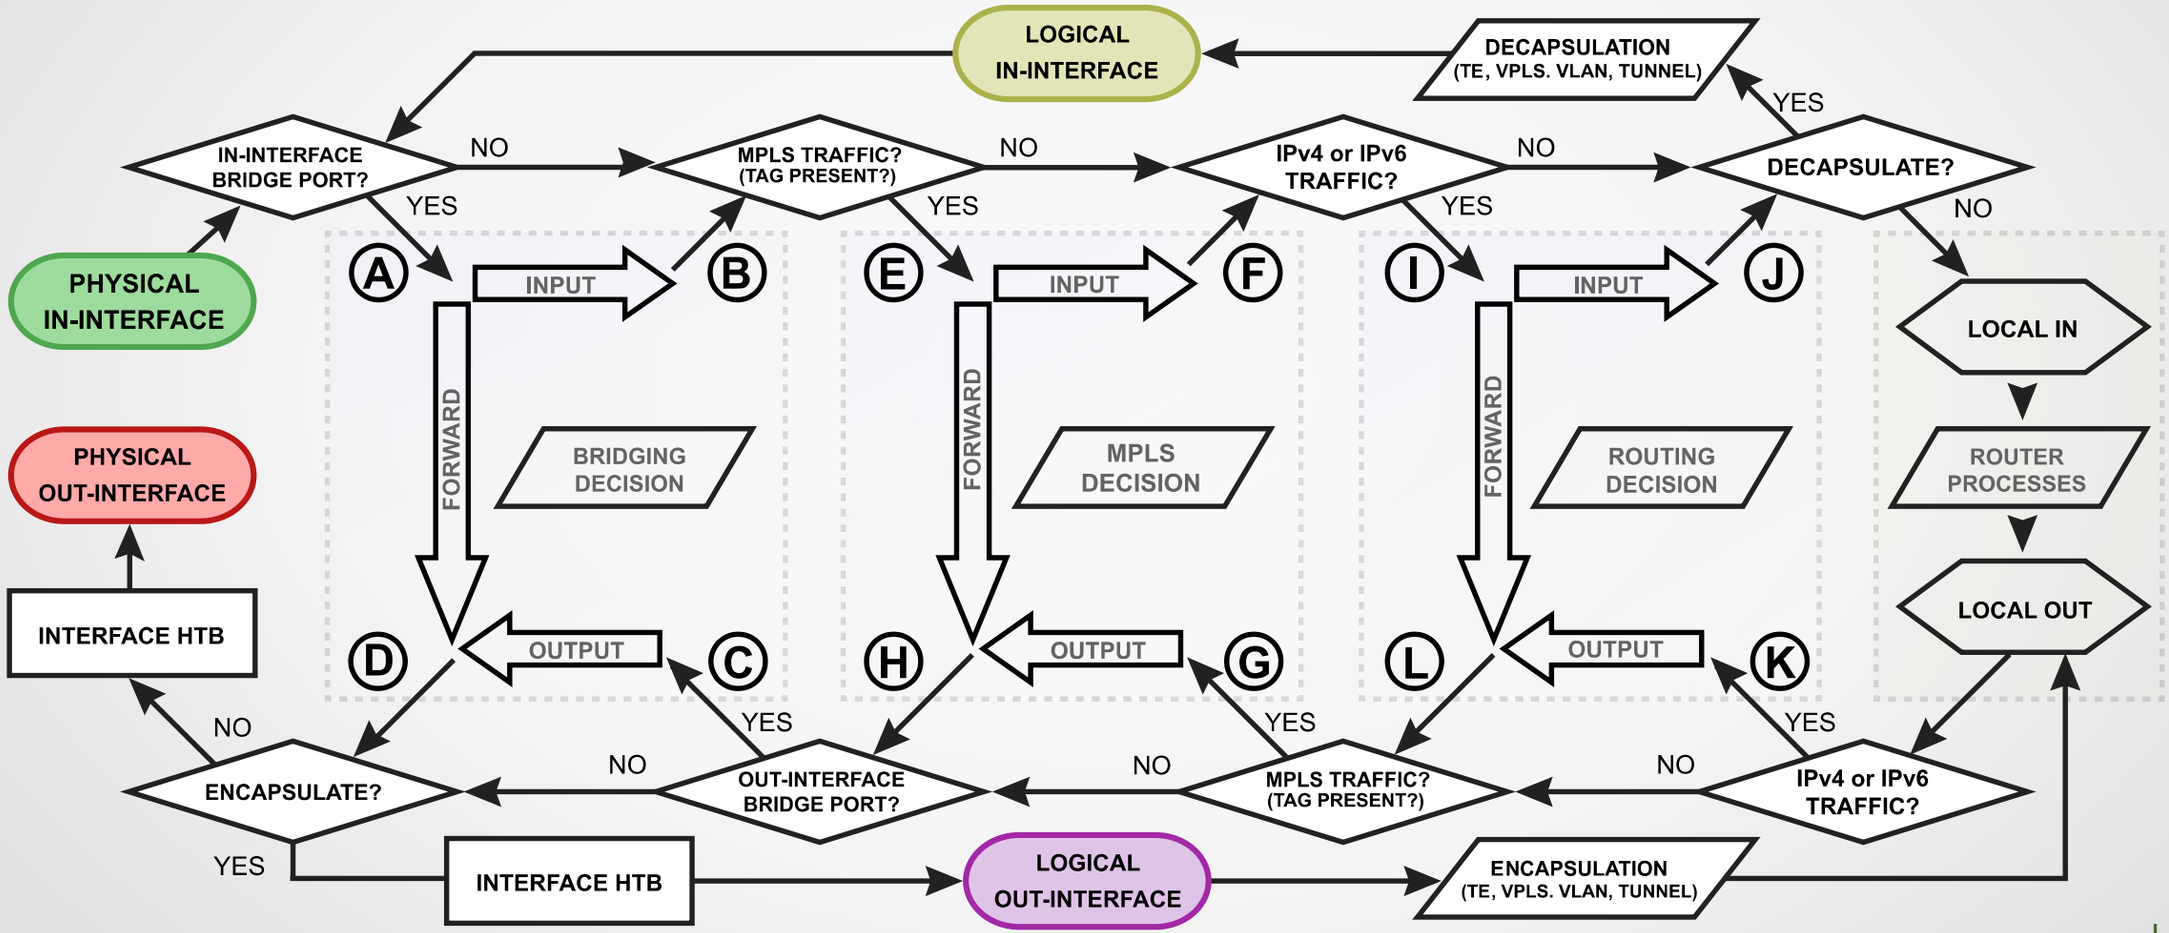

MPLS. Packet flow diagram

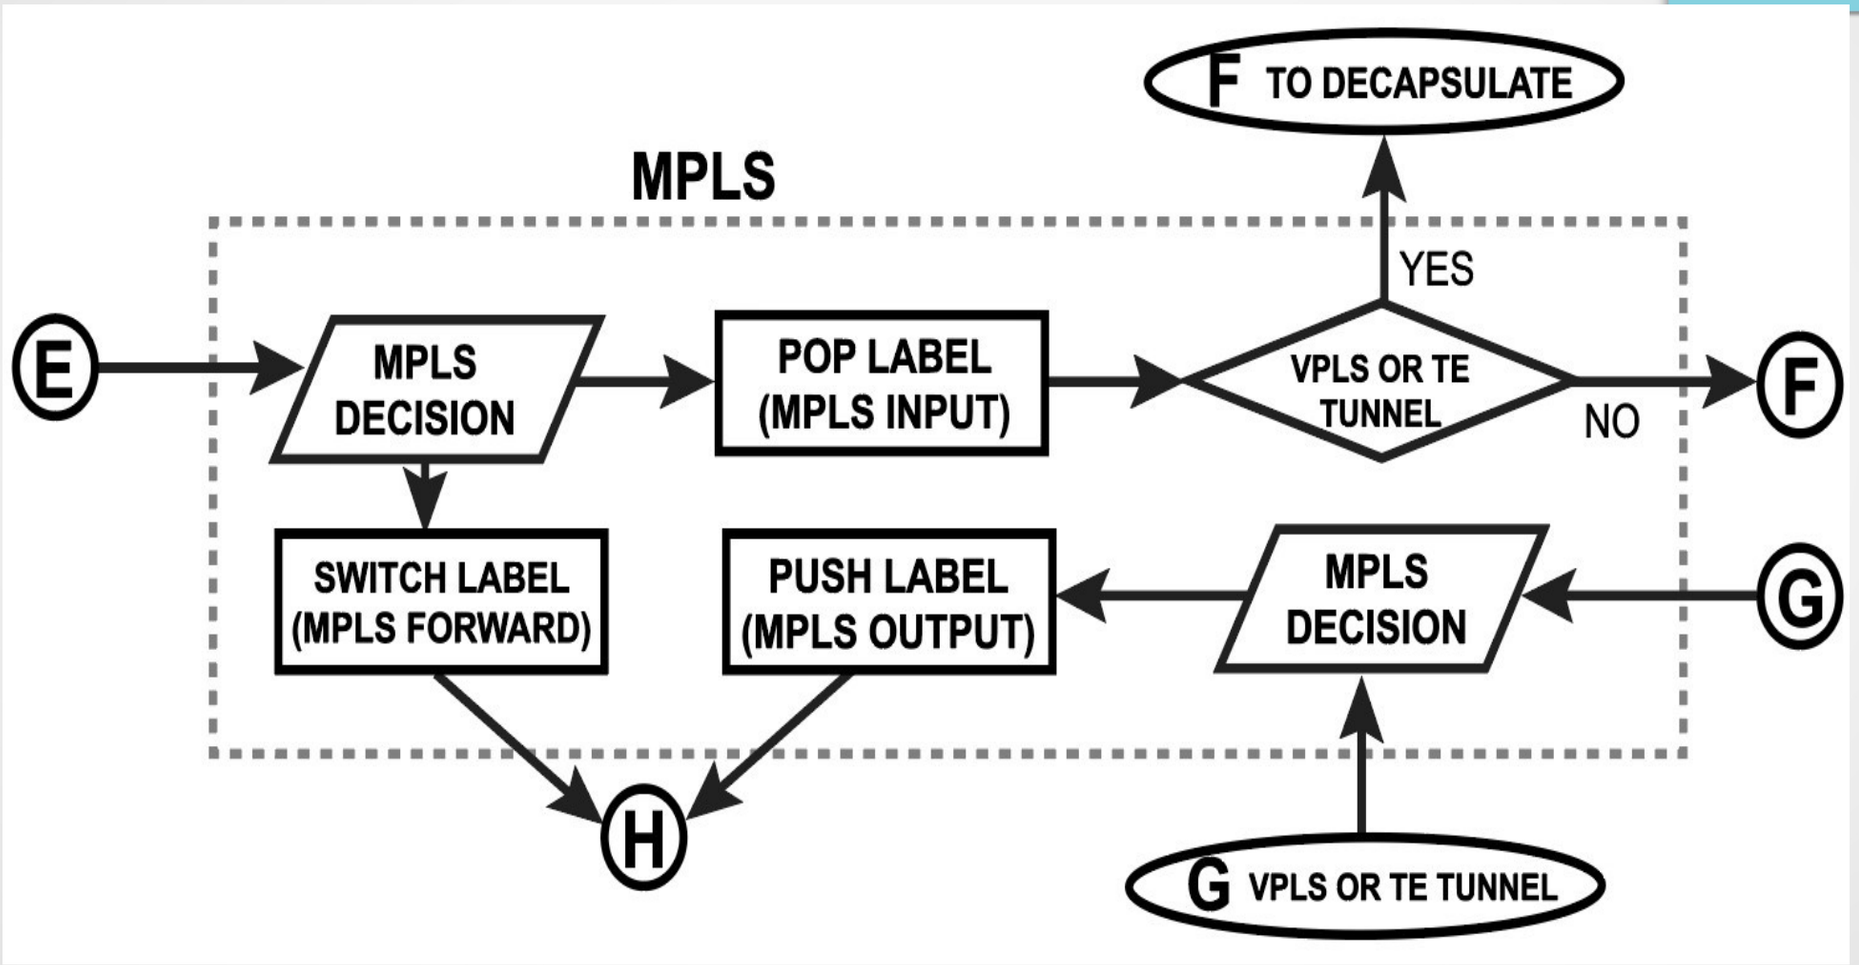

MPLS. LDP

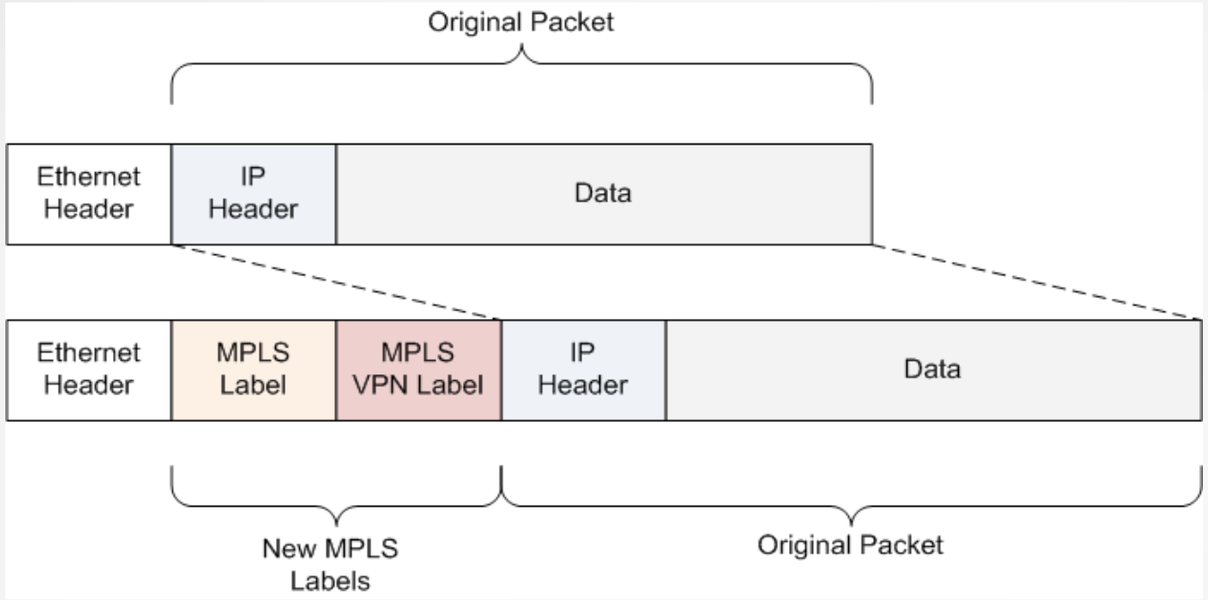

- MPLS - multiprotocol label switching — многопротокольная коммутация по меткам.
- Метки назначит нам LDP.
- MPLS не заменяет IP-маршрутизацию, а работает поверх неё.
- Поэтому надо обеспечить маршрутизацию. IGP.

Модель опорной сети

- R2,R3,R4 маршрутизаторы Cisco C7201
- CHR1, CHR5 маршрутизаторы MikroTik

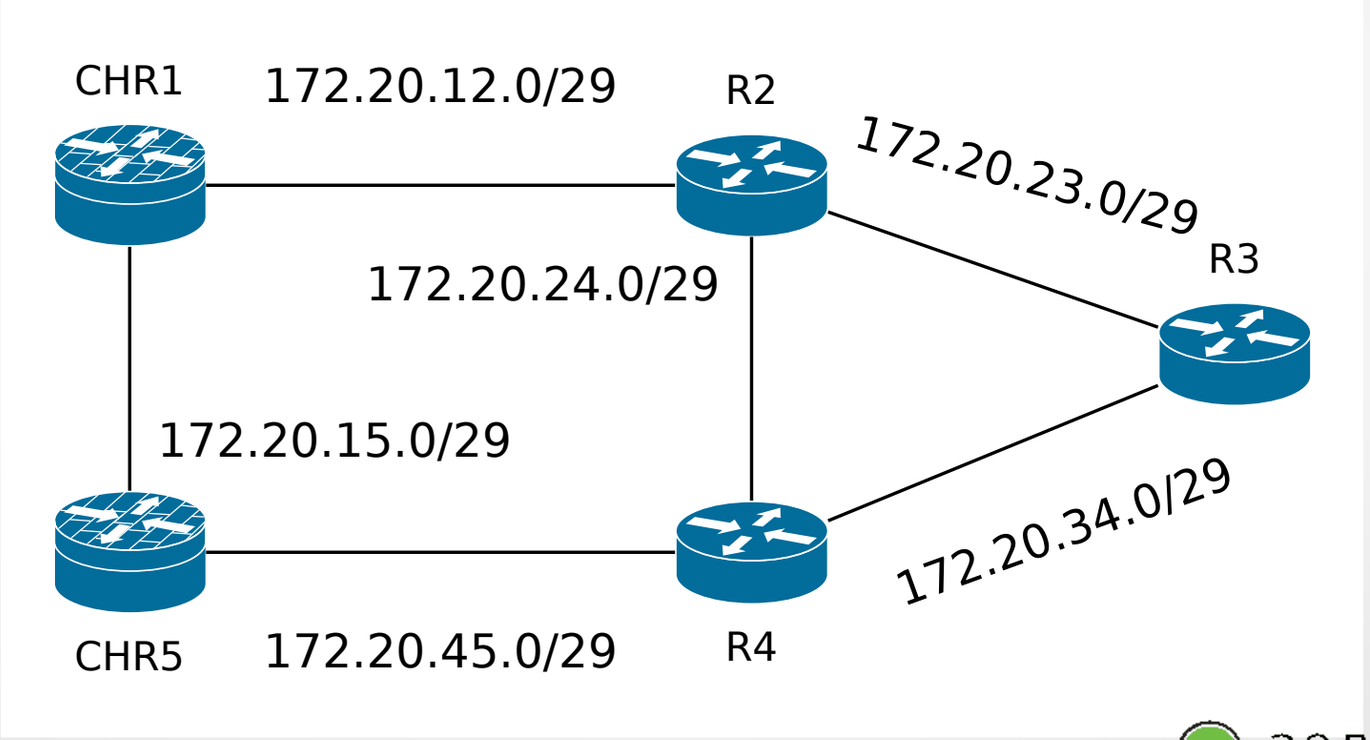

Сетевая связность.

Выделим адресацию.

- Линковочные сети:

CHR1-R2	172.20.12.0/29
R2-R3	172.20.23.0/29
R3-R4	172.20.34.0/29
R2-R4	172.20.24.0/29
R4-CHR5	172.20.45.0/29
CHR1-CHR5	172.20.15.0/29

Loopbacks

CHR1	172.20.0.1
R2	172.20.0.2
R3	172.20.0.3
R4	172.20.0.4
CHR5	172.20.0.5


OSPF

OSPF. Настройка

- Cisco R2

```
router ospf 220
 router-id 172.20.0.2
 log-adjacency-changes
 passive-interface default
 no passive-interface FastEthernet0/0
 no passive-interface FastEthernet1/0
 no passive-interface FastEthernet2/0
 network 172.20.0.0 0.0.63.255 area 0.0.0.0
```

- MikroTik CHR1

```
/routing ospf instance
set [ find default=yes ] router-id=172.20.0.1

/routing ospf network
add area=backbone network=172.20.0.0/18
```

OSPF neighbor

- Cisco R2

```
R2#sh ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.20.0.4        1   FULL/DR         00:00:35    172.20.24.4     FastEthernet2/0
172.20.0.3        1   FULL/DR         00:00:35    172.20.23.3     FastEthernet1/0
172.20.0.1        1   FULL/BDR        00:00:34    172.20.12.1     FastEthernet0/0
```

- MikroTik CHR1

```
[admin@CHR1] > /routing ospf neighbor print brief
 # ROUTER-ID       ADDRESS         STATE                    STATE-CHANGES
 1 172.20.0.2      172.20.12.2     Full                             5
 2 172.20.0.5      172.20.15.5     Full                             5
```

Loopback“s доступны

- Cisco R2

```
R2#ping 172.20.0.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.20.0.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

- MikroTik CHR1

```
[admin@CHR1] > ping 172.20.0.2 count=5
  SEQ HOST                                     SIZE TTL TIME
    0 172.20.0.2                                 56 255 1ms
    4 172.20.0.2                                 56 255 1ms
    sent=5 received=5 packet-loss=0% min-rtt=1ms avg-rtt=1ms max-rtt=2ms
```

MPLS. Настройка

- Cisco R2

```
mpls ip

mpls label protocol ldp

mpls ldp router-id Loopback0


interface FastEthernet0/0

 description #CHR1

 ip address 172.20.12.2 255.255.255.248

 duplex half

 mpls ip

end
```

- MikroTik CHR1

```
/mpls ldp

set enabled=yes lsr-id=172.20.0.1 transport-
address=172.20.0.1

/mpls ldp interface

add interface=ether2

add interface=ether3
```

LDP neighbor

- Cisco R2

```
R2#sh mpls ldp neighbor
    Peer LDP Ident: 172.20.0.1:0; Local LDP Ident 172.20.0.2:0
        TCP connection: 172.20.0.1.646 - 172.20.0.2.27299
        State: Oper; Msgs sent/rcvd: 13171/11522; Downstream
        Up time: 1w0d
        LDP discovery sources:
          Fa0/0, Src IP addr: 172.20.12.1
        Addresses bound to peer LDP Ident:
          172.20.0.1      172.20.12.1     172.20.15.1     192.168.1.240
```

- MikroTik CHR1

```
[admin@CHR1] > /mpls ldp neighbor print
Flags: X - disabled, D - dynamic, O - operational, T - sending-targeted-hello, V - vpls
 #      TRANSPORT       LOCAL-TRANSPORT PEER                      SEN ADDRESSES
 1 DO   172.20.0.2      172.20.0.1      172.20.0.2:0              no            172.20.0.2
                                                                                172.20.12.2
                                                                                172.20.23.2
                                                                                172.20.24.2
```

MPLS forwarding-table

- Cisco R2

```
R2#sh mpls forwarding-table
Local  Outgoing      Prefix            Bytes Label   Outgoing   Next Hop
Label  Label or VC   or Tunnel Id      Switched      interface
18     Pop tag     172.20.0.1/32     8193       Fa0/0      172.20.12.1
19     Pop tag     172.20.0.3/32     7579       Fa1/0      172.20.23.3
20     Pop tag     172.20.0.4/32     140        Fa2/0      172.20.24.4
```

- MikroTik CHR1

```
[smithy@CHR1] > /mpls forwarding-table print
Flags: H - hw-offload, L - ldp, V - vpls, T - traffic-eng
 #    IN-LABEL      OUT-LABELS   DESTINATION                    INTERFACE   NEXTHOP
 0    expl-null
 1  L 22                        172.20.23.0/29                 ether4     172.20.12.2
 2  L 17            16          172.20.34.0/29                 ether4     172.20.12.2
```

VPLS

- CHR1:

  /interface vpls add disabled=no name=vpls5 remote-peer=172.20.0.5 vpls-id=0:0

- CHR5:

  /interface vpls add disabled=no name=vpls5 remote-peer=172.20.0.1 vpls-id=0:0

```
[admin@VPLS-CHR] > /ip dhcp-server lease pri
Flags: X - disabled, R - radius, D - dynamic, B - blocked
 #   ADDRESS      MAC-ADDRESS       HOST-NAME
 0 D 10.50.50.49  00:50:79:66:68:0B VPC51
[admin@VPLS-CHR] > /ping 10.50.50.49 count=3
  SEQ HOST                                     SIZE TTL TIME  STATUS
    0 10.50.50.49                                56  64 19ms
    1 10.50.50.49                                56  64 40ms
    2 10.50.50.49                                56  64 6ms
    sent=3 received=3 packet-loss=0% min-rtt=6ms avg-rtt=21ms max-rtt=40ms
```

```
VPC5> ip dh
DDORA IP 10.50.50.49/24 GW 10.50.50.50

VPC5> ping 10.50.50.50

84 bytes from 10.50.50.50 icmp_seq=1 ttl=64 time=10.075 ms
84 bytes from 10.50.50.50 icmp_seq=2 ttl=64 time=6.318 ms
84 bytes from 10.50.50.50 icmp_seq=3 ttl=64 time=5.205 ms
84 bytes from 10.50.50.50 icmp_seq=4 ttl=64 time=5.566 ms
84 bytes from 10.50.50.50 icmp_seq=5 ttl=64 time=6.281 ms
```

L3VPN. VRF MikroTik

https://clck.ru/EScVY

```
CHR1:
/ip route vrf
add export-route-targets=20866:101 \
    import-route-targets=20866:101 \
    interfaces=ether6 \
    route-distinguisher=20866:101 \
    routing-mark=vrf1
add export-route-targets=20866:102 \
    import-route-targets=20866:102 \
    interfaces=ether1 \
    route-distinguisher=20866:102 \
    routing-mark=vrf2

/ip address
add address=192.168.2.1/30 interface=ether6
add address=192.168.2.5/30 interface=ether1
```

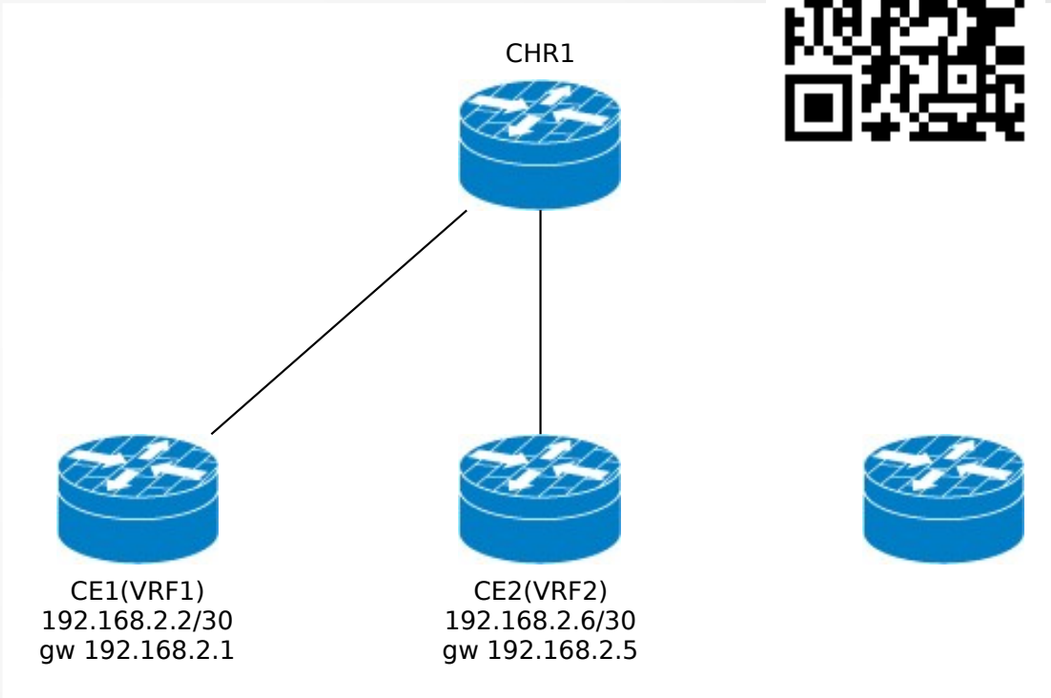

L3VPN. VRF MikroTik

Route List

Routes | Nexthops | Rules | VRF

Find vrf1

	Dst. Address	Gateway	Distance	Routing Mark	Pref. Source
DAC	192.168.2.0/30	ether6 reachable	0	vrf1	192.168.2.1
DAo	192.168.3.0/27	192.168.2.2 on vrf1 reachable ether6	110	vrf1	

Find vrf2

	Dst. Address	Gateway	Distance	Routing Mark	Pref. Source
DAC	192.168.2.4/30	ether1 reachable	0	vrf2	192.168.2.5

```
[smithy@CE1] > ping 192.168.2.1 count=2
  SEQ HOST                                     SIZE TTL TIME  STATUS
    0 192.168.2.1                                56  64 0ms
    1 192.168.2.1                                56  64 0ms
    sent=2 received=2 packet-loss=0% min-rtt=0ms avg-rtt=0ms max-rtt=0ms

[smithy@CE1] > ping 192.168.2.6 count=2
  SEQ HOST                                     SIZE TTL TIME  STATUS
    0 192.168.2.6                                             timeout
    1 192.168.2.6                                             timeout
    sent=2 received=0 packet-loss=100%

[smithy@CE1] > ping 192.168.2.5 count=2
  SEQ HOST                                     SIZE TTL TIME  STATUS
    0 192.168.2.5                                56  64 0ms
    1 192.168.2.5                                56  64 0ms
    sent=2 received=2 packet-loss=0% min-rtt=0ms avg-rtt=0ms max-rtt=0ms

[smithy@CE1] >
```

L3VPN. VRF MikroTik

```
[smithy@CE1] >
[smithy@CE1] > ping 192.168.2.1 count=2
  SEQ HOST                                     SIZE TTL TIME  STATUS
    0 192.168.2.1                                56  64 0ms
    1 192.168.2.1                                56  64 0ms
    sent=2 received=2 packet-loss=0% min-rtt=0ms avg-rtt=0ms max-rtt=0ms

[smithy@CE1] > ping 192.168.2.6 count=2
  SEQ HOST                                     SIZE TTL TIME  STATUS
    0 192.168.2.6                                             timeout
    1 192.168.2.6                                             timeout
    sent=2 received=0 packet-loss=100%

[smithy@CE1] > ping 192.168.2.5 count=2
  SEQ HOST                                     SIZE TTL TIME  STATUS
    0 192.168.2.5                                56  64 0ms
    1 192.168.2.5                                56  64 0ms
    sent=2 received=2 packet-loss=0% min-rtt=0ms avg-rtt=0ms max-rtt=0ms

[smithy@CE1] >
```

L3VPN. VRF MikroTik

https://clck.ru/EScVY

- Пинг до шлюза, всё ок.
- Пинг до CE2, который находится в другом VRF, связности нет, всё правильно.
- Пинг до шлюза, который находится в другом VRF, связность есть, но её не должно было быть.

VRF-то ненастоящий!

- Проблема реализации VRF
- Ждём фикса в ROS7.

L3VPN. VRF Cisco

- Вернёмся к нашей модели и подключим CE3(VRF1) к R3 (C7201)

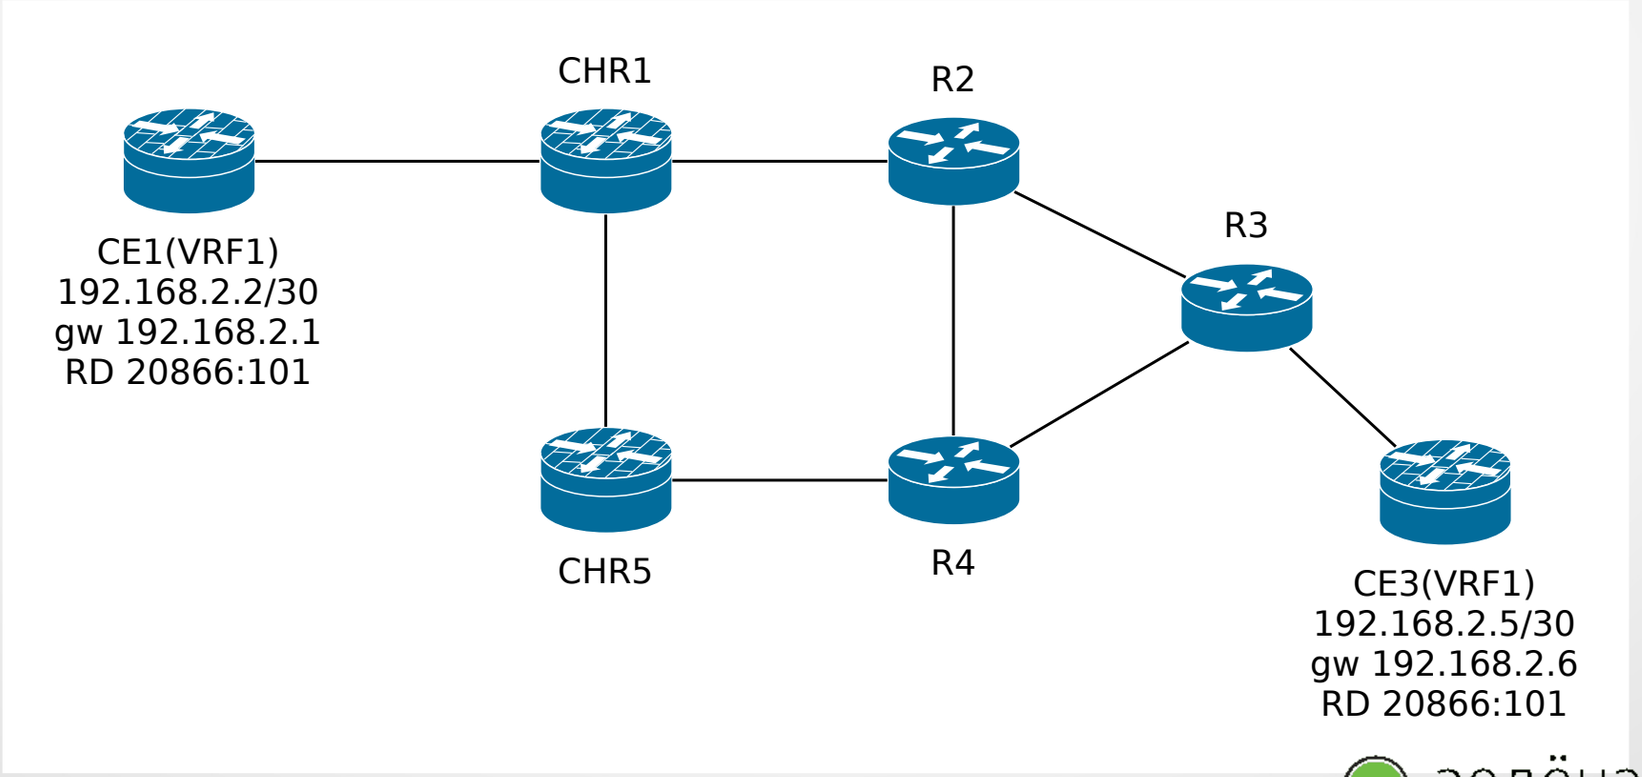

L3VPN. VRF Cisco

- R3:

```
ip vrf vrf1
 rd 20866:101
 route-target export 20866:101
 route-target import 20866:101
!
interface Fa2/0
 description #VRF1-CE3
 ip vrf forwarding vrf1
 ip address 192.168.2.6 255.255.255.252
```

L3VPN. VRF RD

- **Route Distinguisher**

Для того, чтобы различать маршруты различных VPN, обычный IPv4 префикс дополняется специальной приставкой длиной 8 байтов – RD – Route Distinguisher.

Например, маршрут одного vrf будет выглядеть так:

64500:**100**:10.10.10.10/32,

а другого vrf так:

64500:**200**:10.10.10.10/32.

Теперь процесс BGP сможет их друг от друга отличить.

L3VPN. VRF RT

- **Route Targets (RT)** - служит для определения в какой VRF PE-маршрутизатор поместит iBGP-маршрут.

  RT анонсируется в обновлениях BGP как атрибут BGP extended community.

  У конкретного префикса может быть несколько RT.

L3VPN. BGP

L3VPN. iBGP

- Для того, чтобы информация о маршрутах в VRF распространялась между PE необходимо запустить iBGP

L3VPN. BGP

R3:

```
router bgp 20866
 bgp router-id 172.20.0.3
 bgp log-neighbor-changes
 neighbor 172.20.0.1 remote-as 20866
 neighbor 172.20.0.1 update-source Loopback0
 !
 address-family ipv4
 neighbor 172.20.0.1 activate
 neighbor 172.20.0.1 send-community extended
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family vpnv4
 neighbor 172.20.0.1 activate
 neighbor 172.20.0.1 send-community extended
 exit-address-family
 !
```

CHR1:

```
/routing bgp instance
set default as=20866 client-to-client-reflection=no router-id=172.20.0.1


/routing bgp peer
add address-families=ip,vpnv4 name=peer-R3 remote-address=172.20.0.3 \
    remote-as=20866 ttl=default update-source=br-lo0
```

L3VPN. BGP Cisco MikroTik

```
R3:
router bgp 20866
 address-family ipv4 vrf vrf1
 redistribute connected
 redistribute ospf 110 vrf vrf1
 no auto-summary
 no synchronization
 exit-address-family
```

```
CHR1:
/routing bgp instance vrf
add redistribute-connected=yes \
    redistribute-ospf=yes routing-mark=vrf1
```

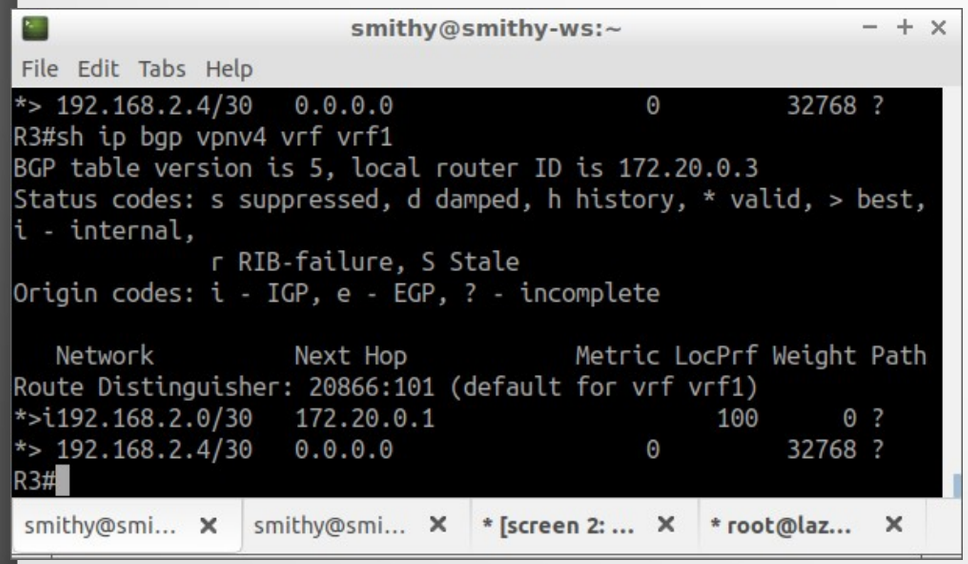

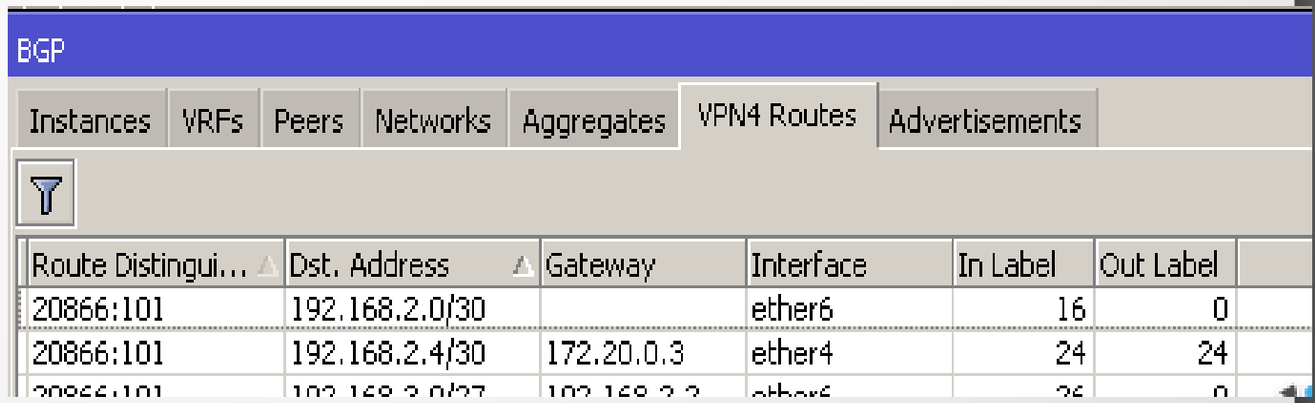

CE1 CE3 VRF1

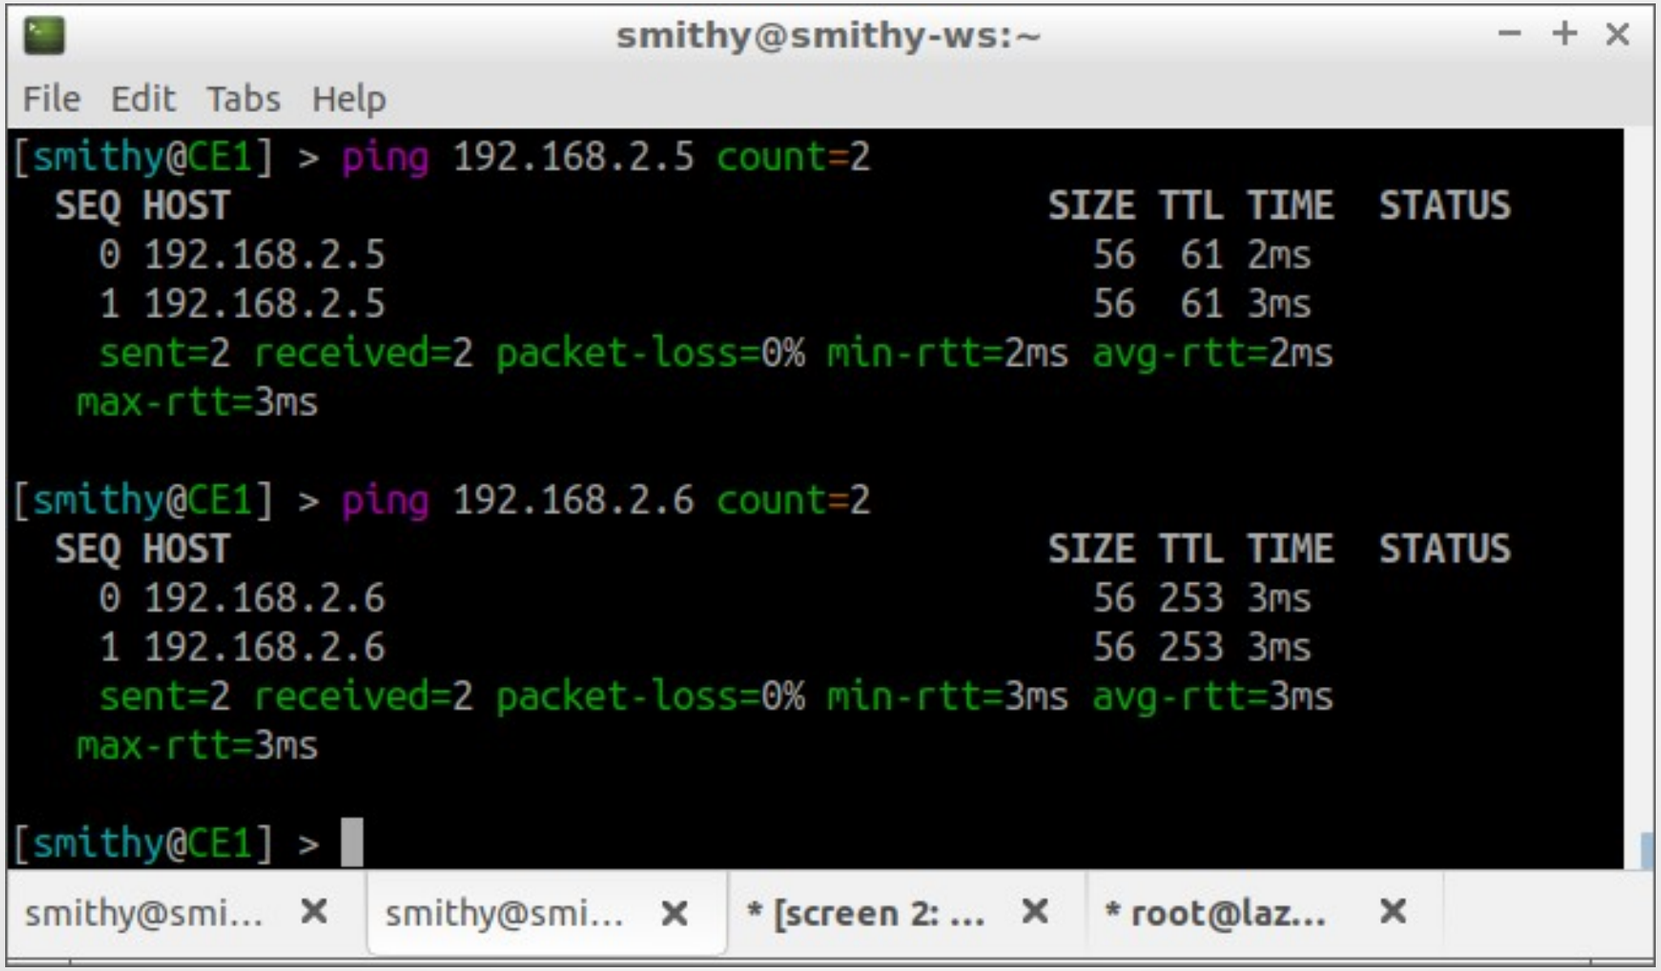
smithy@smithy-ws:~
File Edit Tabs Help
[smithy@CE1] > ping 192.168.2.5 count=2
SEQ HOST SIZE TTL TIME STATUS
0 192.168.2.5 56 61 2ms
1 192.168.2.5 56 61 3ms
sent=2 received=2 packet-loss=0% min-rtt=2ms avg-rtt=2ms
max-rtt=3ms
[smithy@CE1] > ping 192.168.2.6 count=2
SEQ HOST SIZE TTL TIME STATUS
0 192.168.2.6 56 253 3ms
1 192.168.2.6 56 253 3ms
sent=2 received=2 packet-loss=0% min-rtt=3ms avg-rtt=3ms
max-rtt=3ms
[smithy@CE1] >
smithy@smi...
smithy@smi...
* [screen 2: ...
* root@laz...

Тематические ссылки

- Сети для самых маленьких. MPLS L3VPN
- https://habr.com/post/273679/

- VRF (L3VPN) в Mikrotik RouterOS: defective by design
- https://clck.ru/ESkKw

OSPF, MPLS, LDP, BGP, VRF

- Спасибо за внимание!

- Ссылка на презентацию

  https://yadi.sk/i/HxreyVwB8w-CFQ
- Конфиги

  https://yadi.sk/d/Slm1tvStEPccUw