



## [www.freertr.org](http://www.freertr.org) RIFT implementation

**Csaba MATE**

*GÉANT/KIFU – RARE/freeRtr Lead core developer*

**Frederic LOUI**

*GÉANT/RENATER – RARE/Technical leader*

IETF#116 –RIFT-WG

March 25<sup>th</sup> 2023

Public

[www.geant.org](http://www.geant.org)

## Agenda

- RARE/freeRtr in a nutshell
- RIFT draft implementation
- freeRtr interops with juniper crpd
- freeRtr interops with [github.com/brunorijsman/rift-python](https://github.com/brunorijsman/rift-python)
- Conclusion

## RARE project : Initial group focus

- GEANT project sub-task: RARE
  - Control plane software
  - Multiple data planes
  - Interface them and the result is ...
- Fully functional router
  - Running at hardware line rate
  - DIY “hackable/extensible” router
  - Control plane independence

One familiar platform



Multiple solutions



Each solution addresses



R&E

use case

## RARE project : Month 48

**RARE** is an open source routing platform, used to create a network operating system (NOS) on commodity hardware (a white box switch).



**RARE** uses FreeRtr as a control plane software  
and is thus often referred to as  
**RARE/FreeRtr**



# RARE latest news (beyond 2023 till 20xx ?)

- RARE p4 targets



bmv2 software switch



Intel/barefoot Tofino on WEDGE-BF100-32X, APS-BF2556X-T1, others



**under study**

- RARE “p4” emulation targets



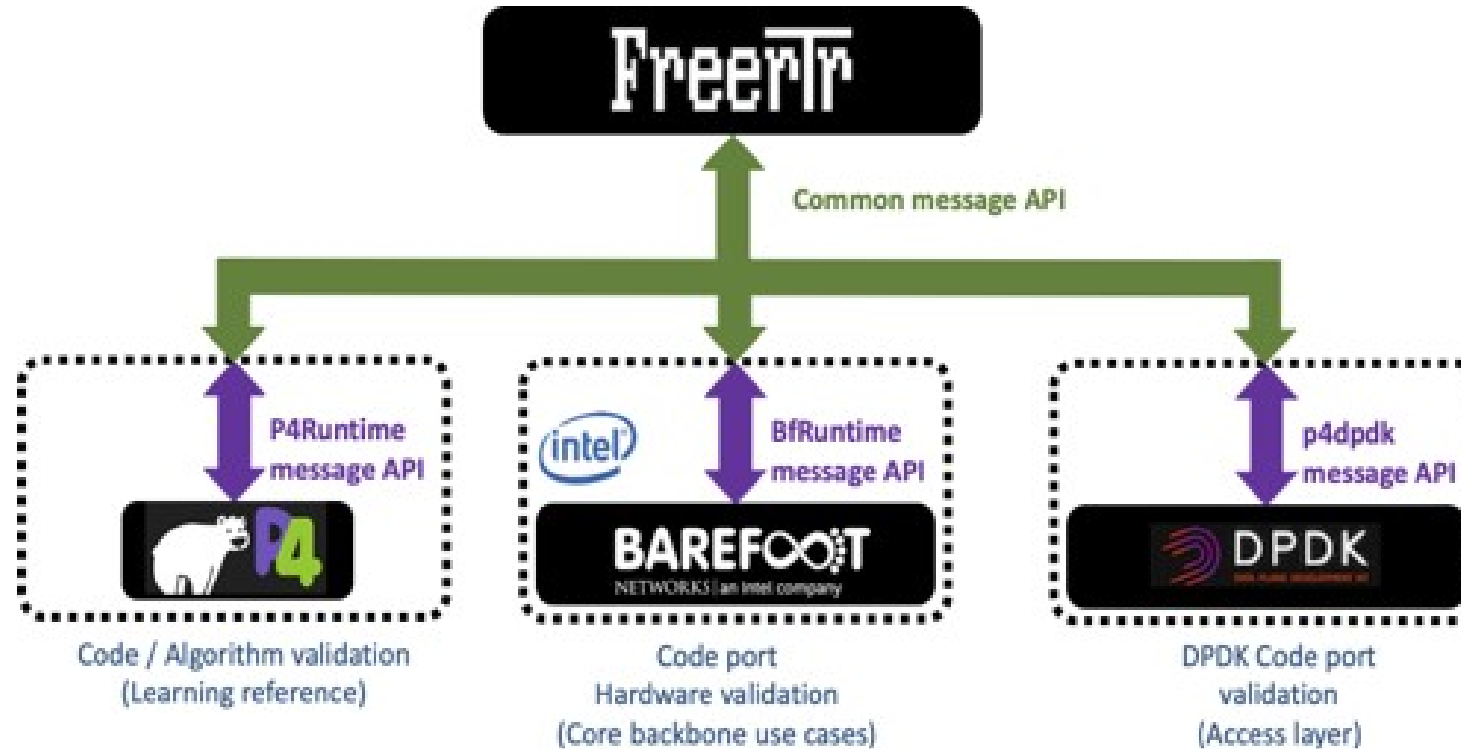
**XDP**

- RARE Network Programmable targets

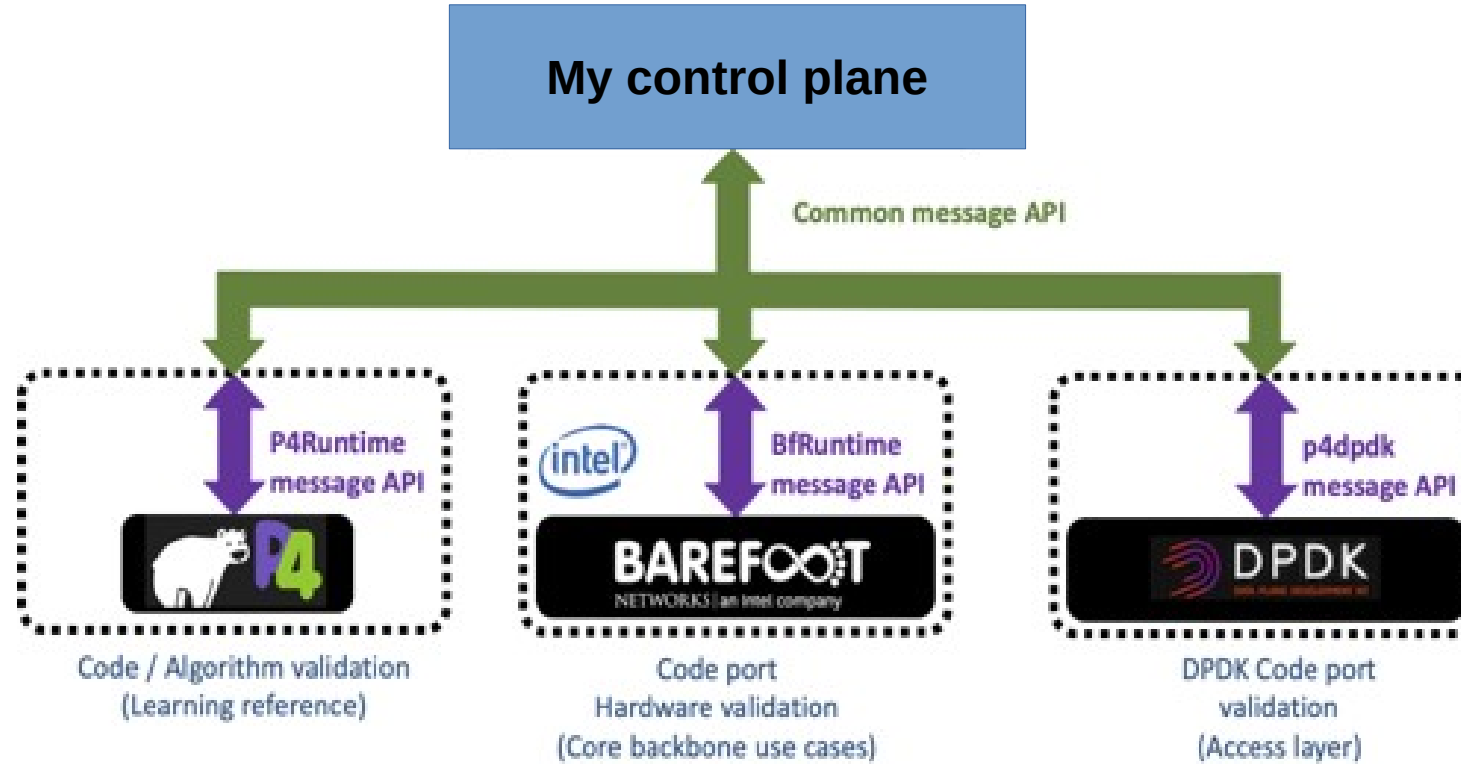


**under study**

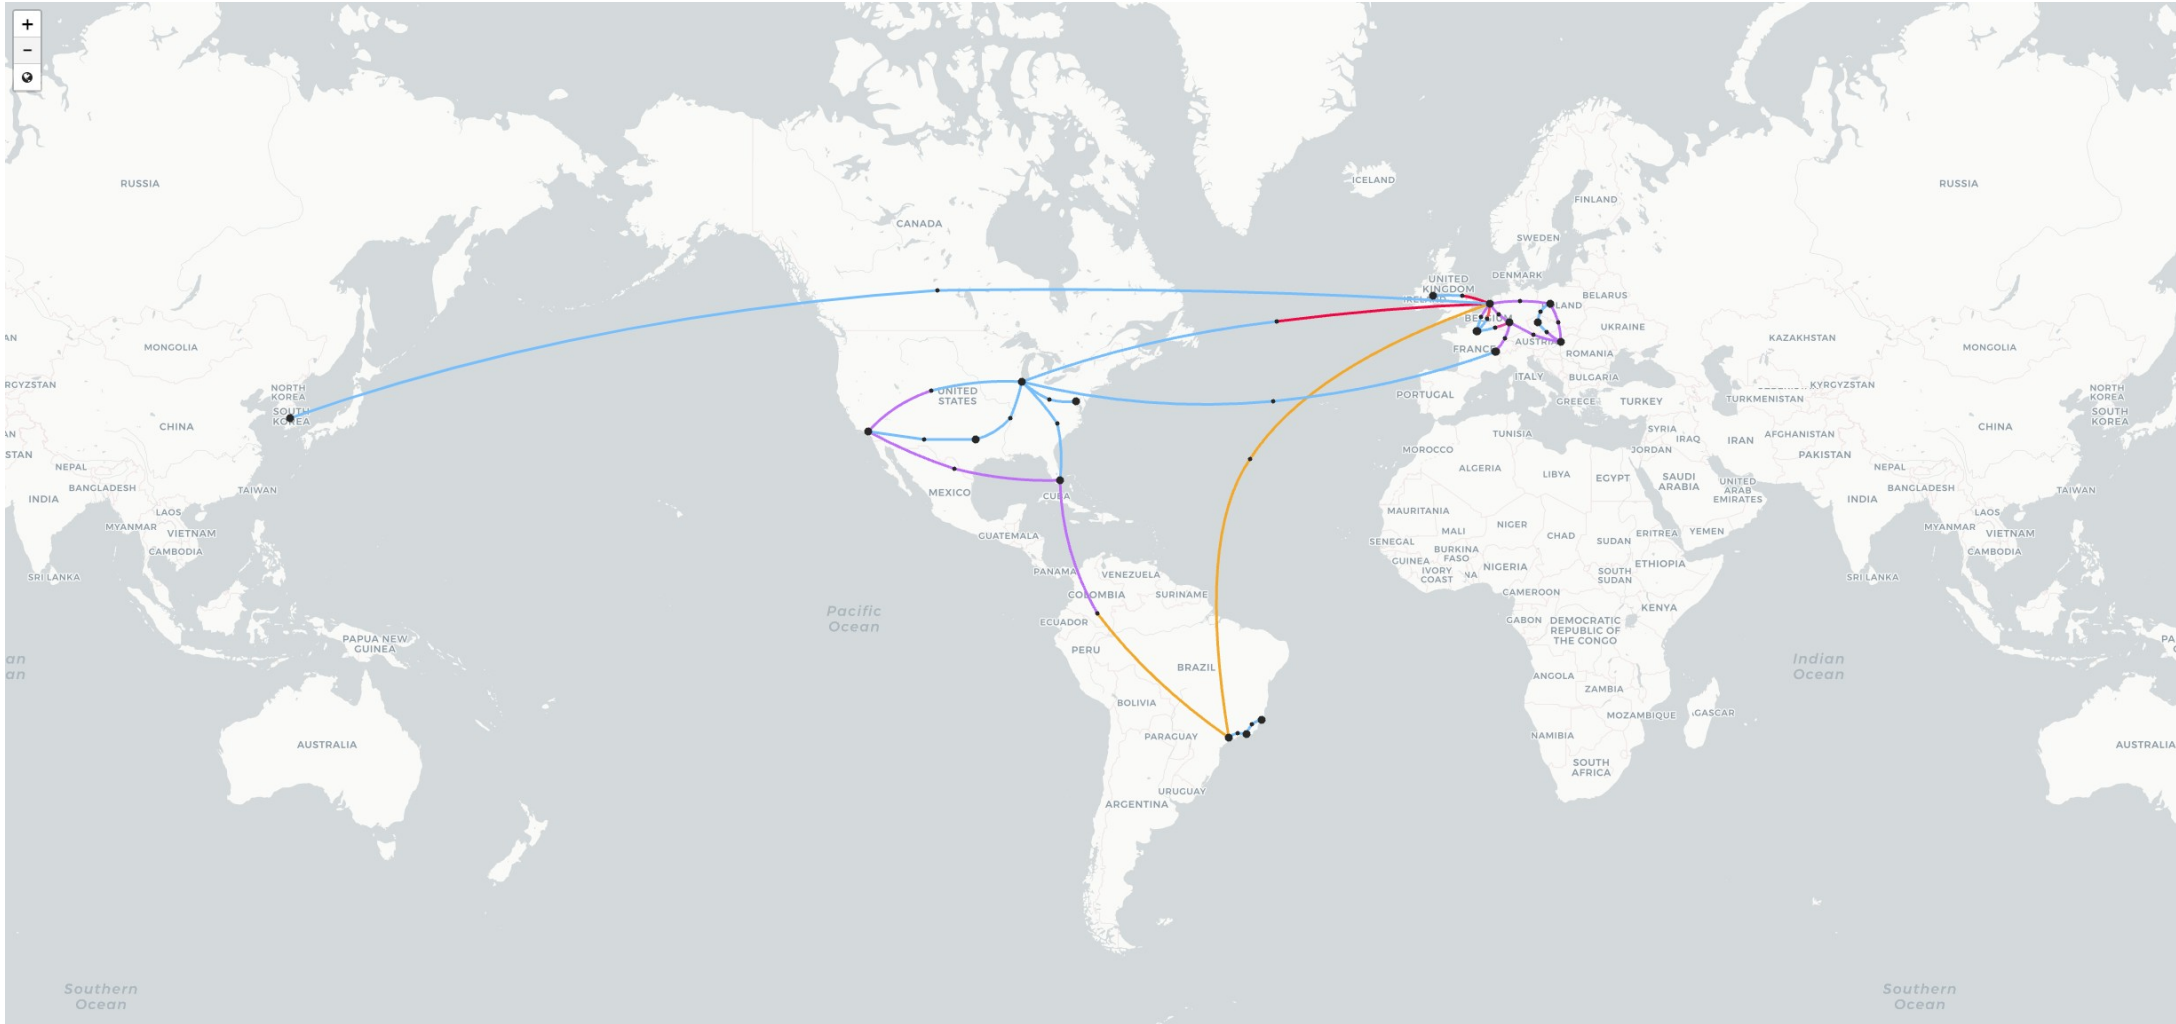
# RARE/freeRtr: One control plane to rule them all



# RARE/freeRtr: or not ...



# RARE Global P4 Lab





## What we have

- draft-rift-rift
  - Flooding procedures – interops
  - Data structures encoder/decoder - interops
  - Leaf procedures - interops
  - Some of the spine procedures – not too much tbh
  - Plans for draft-zzhang-rift-sr - once the TBDs are allocated
  - Plans for draft-zwx-rift-leaf-ring – once it's settled
- All the above for v4 and v6, covered by automated testing

```
mc36@noti: ~ x noti x mc [mc36@noti]:~ x mc [mc36@safe]:/d... x mc36@nass: /nfs/o... x sid x root@crpd01: / x
root@crpd01>
root@crpd01>
root@crpd01>
root@crpd01> show rift node status
System Name: crpd01, System ID: 00000000deadbabe
Level: 24, RIFT Encoding Major: 6, Minor: 1, EVPN ZTP Major: 1, Minor: 0
Flags: overload=False, acting_auto_evpn_dci_when_tof=None
Capabilities: flood-reduction: True, Hierarchy Indications: top_of_fabric
LIE v4 RX: 224.0.0.120, LIE v6 RX: ff02::a1f7, LIE RX Port: 914
Re-Connections: 0
Peers: 1, 3-way: 1, South: 0, North: 0, East-West: 1

root@crpd01> show rift interface statistics
Link ID: 257, Interface: ens4, Started: 2022/12/19 10:08:19.439
TIES RX: 1, TX: 3, REQ: 0, Neighbor REQ: 0
TIDES TX: 2, RX: 0, TIRES TX: 1, TIRES RX: 3
Pkt Rate/100msecs Highest: 61, Current: 60, Packet Sequence Losses: 1
Last TIE RX: 00000000badbabe/S/Node_____/00000001, Last TIE RX on: 2022/12/19 10:14:56.546
Last TIE TX: 00000000deadbabe/N/Node_____/10000001, Last TIE TX on: 2022/12/19 10:15:00.798
TIE TX Queue Len: 0
Largest TX'ed: TIE: 307, (00000000deadbabe/N/Node_____/10000001), TIDE: 354, TIRE: 120
Three-Way UP 1, DOWN 0, Last UP 2022/12/19 10:14:55.545
Last Reason DOWN: None
LIE TX 1429, RX 398
Last LIE TX 2022/12/19 10:16:12.301, RX 2022/12/19 10:16:12.572, Reject Reason: None
Current Level Self 24, Neighbor 24, Level Changes Self 1, Neighbor 1
Flood Leader: True, Changes: 1, Last Change: 2022/12/19 10:11:09.142

root@crpd01> █
```



```

mc36@noti: ~ x noti x mc [mc36@noti]:~ x mc [mc36@safe]:/d... x mc36@nass: /nfs/o... x sid x root@crpd01: / x
root@crpd01>
root@crpd01>
root@crpd01> show configuration | display set
set version 20221123.183731_builder.r1297844
set groups rift-defaults protocols rift traceoptions file size 1000000
set groups rift-defaults protocols rift traceoptions file files 4
set groups rift-defaults protocols rift traceoptions level notice
set groups rift-defaults protocols rift node-id 1108037632
set groups rift-defaults protocols rift level auto
set groups rift-defaults protocols rift lie-receive-address family inet 224.0.0.120
set groups rift-defaults protocols rift lie-receive-address family inet6 ff02::a1f7
set groups rift-defaults protocols rift interface <*> lie-transmit-address family inet 224.0.0.120
set groups rift-defaults protocols rift interface <*> lie-transmit-address family inet6 ff02::a1f7
set protocols rift apply-groups rift-defaults
set protocols rift node-id 3735927486
set protocols rift level top-of-fabric
set protocols rift interface ens4 mode advertise-subnets

root@crpd01> show rift database content
Dir Originator Type ID SeqNr Lifetime Origin Creation Time Origin Content Key ID
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
S 000000000badbabe Node 00000001 2 604664 604800 190 0
S 000000000deadbabe Node 10000001 63a039a27a00 604670 2022/12/19 10:15:00.798 604800 None
S 000000000deadbabe Prefix 20000017 63a039a0f456 604665 2022/12/19 10:14:55.548 604800 None
N 000000000deadbabe Node 10000001 63a039a2c5cb 604670 2022/12/19 10:15:00.798 604800 299 0
N 000000000deadbabe Prefix 20000009 63a039a02c6e 604665 2022/12/19 10:14:55.548 604800 174 0

root@crpd01> █

```

```
mc36@noti: ~ x noti x mc [mc36@noti]:~ x mc [mc36@safe]:/d... x mc36@nass: /nfs/o... x sid x root@crpd01: / x
sid#
sid#
sid#
sid#
sid#
sid#
sid#
sid#show config-differences
router rift4 2
vrf v2
router-id 195934910
exit
interface pwether1
mtu 1500
macaddr 0072.2175.5867
vrf forwarding v2
ipv4 address 10.123.123.122 255.255.255.0
router rift4 2 enable
pseudowire v1 loopback0 pckoudp 10.5.255.1 30042
no shutdown
no log-link-change
exit

sid#show ipv4 rift 2 database
dir origin num typ seq left
s 195934910 1 2 2 6d23h
s 195934910 2 3 1 6d23h
n 195934910 1 2 2 6d23h
n 195934910 2 3 1 6d23h
n 3735927486 268435457 2 109539812885963 6d23h
n 3735927486 536870921 3 109539812715630 6d23h

sid#
```

```
mc36@noti: ~ x noti x mc [mc36@noti]:~ x mc [mc36@safe]:/d... x mc36@nass: /nfs/o... x sid x root@crpd01: / x
```

```
sid#  
sid#  
sid#show config-differences  
router rift4 2  
vrf v2  
router-id 195934910  
exit  
interface pwether1  
mtu 1500  
macaddr 0072.2175.5867  
vrf forwarding v2  
ipv4 address 10.123.123.122 255.255.255.0  
router rift4 2 enable  
pseudowire v1 loopback0 pckoudp 10.5.255.1 30042  
no shutdown  
no log-link-change  
exit  
  
sid#show ipv4 rift 2 database  
dir  origin      num      typ  seq      left  
s   195934910     1        2    2        6d23h  
s   195934910     2        3    1        6d23h  
n   195934910     1        2    2        6d23h  
n   195934910     2        3    1        6d23h  
n   3735927486   268435457 2    109539812885963 6d23h  
n   3735927486   536870921 3    109539812715630 6d23h  
  
sid#show ipv4 rift 2 tree n  
`--sid  
  |--crpd01  
  
sid#
```

```
mc36@noti: ~ x noti x mc [mc36@noti]:~ x mc [mc36@safe]:/d... x mc36@nass: /nfs/o... x sid x root@crpd01: / x
```

```
mtu 1500
macaddr 0072.2175.5867
vrf forwarding v2
ipv4 address 10.123.123.122 255.255.255.0
router rift4 2 enable
pseudowire v1 loopback0 pckoudp 10.5.255.1 30042
no shutdown
no log-link-change
exit

sid#show ipv4 rift 2 database
dir  origin      num      typ  seq      left
s    195934910    1        2    2        6d23h
s    195934910    2        3    1        6d23h
n    195934910    1        2    2        6d23h
n    195934910    2        3    1        6d23h
n    3735927486  268435457 2    109539812885963 6d23h
n    3735927486  536870921 3    109539812715630 6d23h

sid#show ipv4 rift 2 tree n
`--sid
  |--crpd01

sid#show ipv4 rift 2 route
typ  prefix      metric  iface      hop      time
F    10.123.123.0/24 100/11  pwether1  10.123.123.123 00:02:37

sid#show ipv4 rift 2 othertree n 3735927486
`--crpd01
  |--sid

sid#
```

```
real_my> sho tie
```

Direction	Originator	Type	TIE Nr	Seq Nr	Lifetime	Contents
South	12345678	Node	1	2	604706	Name: real_my Level: 24 Capabilities: Flood reduction: True Neighbor: 00000033bd7a038e Level: 24 Cost: 1 Bandwidth: -1 Mbps Link: 1-693902388
South	12345678	Prefix	2	1	604706	Prefix: 0.0.0.0/0 Metric: 1 Prefix: ::/0 Metric: 1
South	222222222222	Node	1	2	604705	Name: sid Level: 24 Capabilities: Neighbor: 0000000000bc614e Level: 24 Cost: 10 Bandwidth: 4294967295 Mbps Link: 693902388-1
South	222222222222	Prefix	2	4	604704	Prefix: 1.1.1.0/30
South	222222222222	Prefix	3	3	604704	Prefix: 3.3.3.2/32
North	12345678	Node	1	2	604706	Name: real_my

```
mc36@noti: ~ x noti x mc [mc36@noti]:/s... x mc [mc36@safe]:/d... x mc36@nass: /nfs/o... x sid x mc36@rift: ~ x
sid#
sid#
sid#
sid#show config-differences
router rift4 2
  vrf v2
  router-id 222222222222
  redistribute connected
  exit
interface loopback2
  vrf forwarding v2
  ipv4 address 3.3.3.2 255.255.255.255
  ipv6 address 3333::2 ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
  no shutdown
  no log-link-change
  exit
interface pwether1
  mtu 1500
  macaddr 0055.3721.5f16
  vrf forwarding v2
  ipv4 address 1.1.1.2 255.255.255.252
  router rift4 2 enable
  pseudowire v1 loopback0 pckoudp 10.5.255.1 30042
  no shutdown
  no log-link-change
  exit

sid#
sid#
sid#
sid#
sid#
```



```
mc36@noti: ~ x noti x mc [mc36@noti]:/s... x mc [mc36@safe]:/d... x mc36@nass: /nfs/o... x sid x mc36@rift: ~ x
sid#
sid#
sid#
sid#
sid#
sid#
sid#
sid#
sid#show ipv4 rift 2 neighbor
iface      nodeid    name      peer      uptime
pwether1   12345678  real_my:ens4  1.1.1.1  00:00:59

sid#show ipv4 rift 2 database
dir  origin      num  typ  seq  left
s    222222222222  1    2    2    6d23h
s    222222222222  2    3    4    6d23h
s    222222222222  3    3    3    6d23h
n    12345678      1    2    2    6d23h
n    222222222222  1    2    2    6d23h
n    222222222222  2    3    4    6d23h
n    222222222222  3    3    3    6d23h

sid#show ipv4 rift 2 tree n
`--sid
   |--real_my

sid#show ipv4 rift 2 othertree n 12345678
`--real_my
   |--sid

sid#
sid#
sid#
sid#show config differences
```

```
mc36@noti: ~
x
noti
x
mc [mc36@noti]:/safe/img/rift
x
sid
x
mc36@nass: /nfs/own/web/fre...
x
sid#
sid#
sid#show ipv4 rift 2 database n      222222222222  2
direction=2
originator=222222222222
number=2
type=3
sequence=4
lifetime=6d23h
header= 0c 00 00 0c 00 01 0c 00 02 08 00 01 00 00 00 02 0a 00 02 00 00 00 33 bd 7a 03 8e 08 00 03 00 00 00 03 08 00 04
00 00 00 02 00 0a 00 03 00 00 00 00 00 00 00 04 00 08 00 02 00 09 3a 67 00
body= 0c 00 00 0c 00 02 0d 00 01 0c 0c 00 00 00 01 0c 00 01 08 00 01 01 01 01 00 03 00 02 1e 00 00 08 00 02 00 00 00 00
08 00 03 00 00 00 00 00 00 00
num=2 typ=12 val=(
  num=1 typ=13 val=(
    num=0 typ=12 val=(
      num=1 typ=12 val=(
        num=1 typ=8 val=16843008
        num=2 typ=3 val=30
      )
    )
  )
  num=0 typ=12 val=(
    num=2 typ=8 val=0
    num=3 typ=8 val=0
  )
)
)

sid#
sid#
sid#
sid#
```



## Key take-away – We are ready to roll into production

- Automated testing: <http://www.freertr.org/tests.html>
- 3rd party testing via Spirent usage
  - (thanks PSNC@WB team)
- P4 profile calibration
- DPDK is in operation
- Production instance



- Someone else? :)

## Useful links

- Project

<http://rare.freertr.org/>

<http://blog.freertr.org/>

<http://docs.freertr.org/>

- Contact


[rare-users@lists.geant.org](mailto:rare-users@lists.geant.org)

[rare-dev@lists.geant.org](mailto:rare-dev@lists.geant.org)

[freertr@groups.io](mailto:freertr@groups.io)

[https://twitter.com/rare\\_freerouter](https://twitter.com/rare_freerouter)

## Special thanks ...



And others ...  
Who make this possible !

# Thank you

Any questions?

[www.geant.org](http://www.geant.org)

