

Router for Academia Research Education

RARE/freeRtr in a nutshell



LOUI Frédéric

RENATER – RARE technical leader

ESNOG 28

May 12th 2022

Public

www.geant.org



GÉANT-4 programme


- 39 participants country
 - NREN participants
 - **But also NREN staff members !**
- GN4 phase 3 cooperation programme
 - 44 months
 - Budget 78 M€
- GN4-3 divided in “Work Package”
 - 9 WP divided in “Tasks”
 - Task 1 encompasses 6 Sub-tasks
 - RARE is one of these sub-task



**WHY are we doing
all of this ?**

Problems statements

- Needs of **network aware** applications
- Traditional procurement life cycle workload related to network equipment acquisition
- Reduce vendor lock-in
- Ability to implement ADHOC features
- Reduce digital divide affordable network solution without functionality trade off
- Software sovereignty ?



RARE/freeRtr mission statement

- One unified & Opensource routing platform
 - Multiple solutions (=dataplanes)
 - That fit R&E use case

At TREMENDOUS line rate !



Why RARE now?

- Starting from early 2010:
 - Several valuable Open Source control plane usage besides well know commercial vendor




- Starting from 2020:
 - Dataplane solution reached maturity ready to implement production grade use case



- NOS emergence

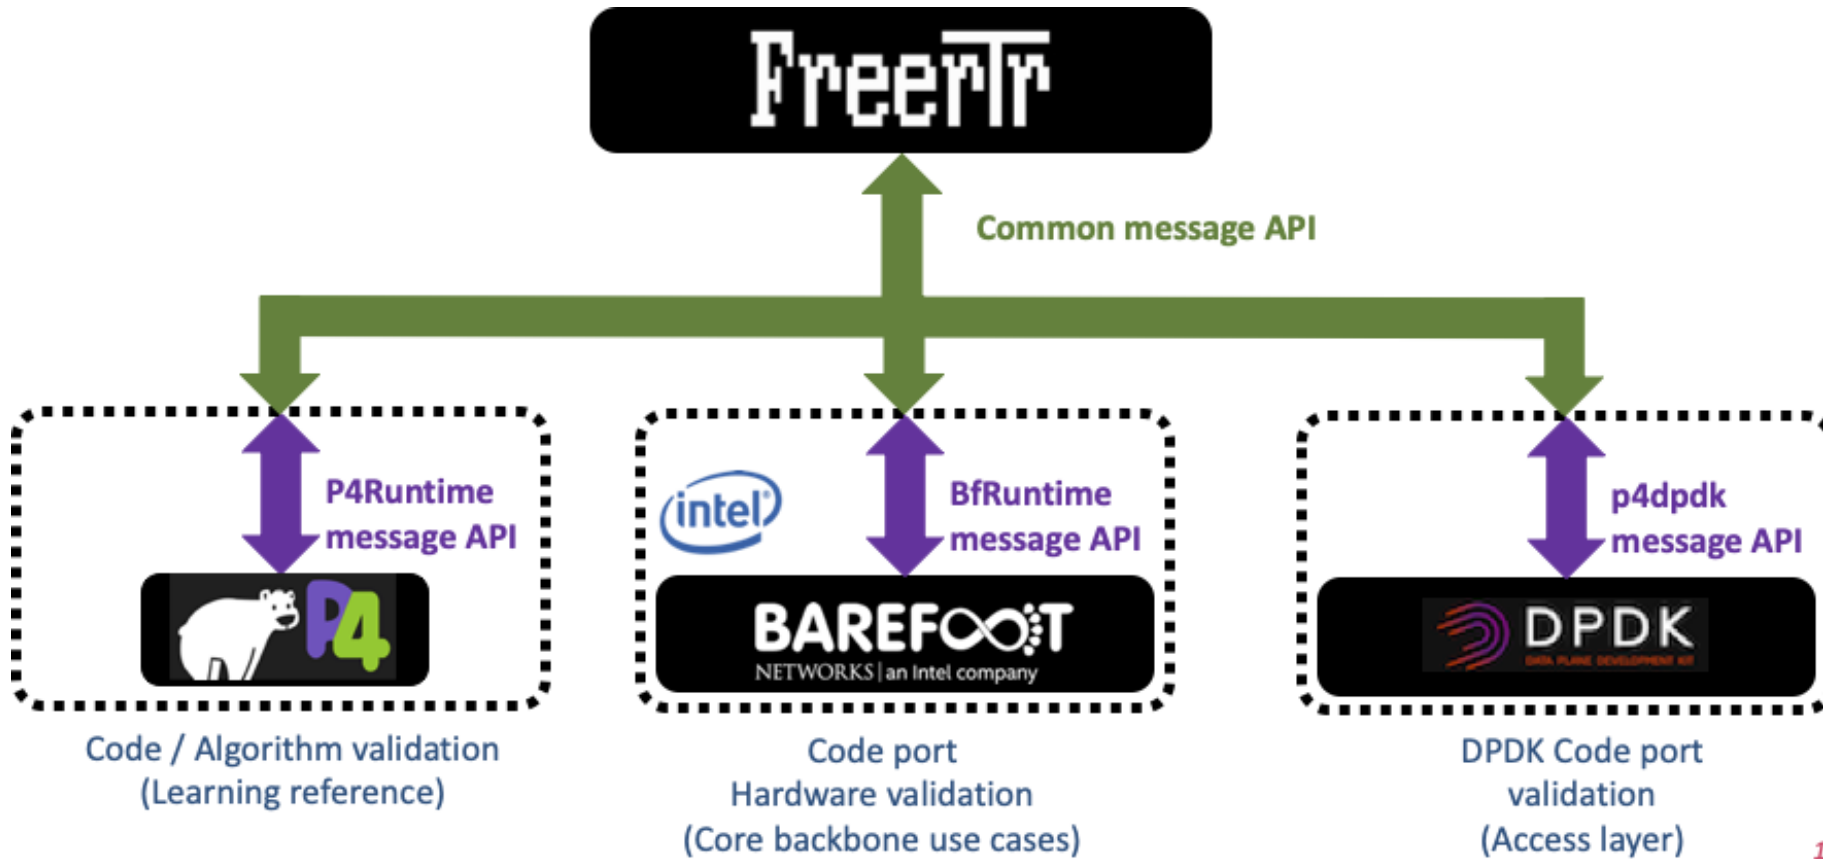


- Hypervisor Technology convergence
 - openstack.
 - kubernetes



It's a good time to tie Control Plane and Dataplane!

One control to rules all dataplanes





What about ISP and commercial use case ?

Practical use case #001 SOHO router


- DPDK flavor ideal for CPE
- nx1GE
- nx10GE small MAN ideal for small campus
- Couple of 100GE (Depending on server generation)





Practical use case #002 BRAS-BNG/LNS router

- DPDK and P4 dataplane
 - ➔ suitable for CAMPUS / EDGE BACKBONE router
- nx1GE, nx10GE, nx100GE



 **APS Networks**
BF-2556X-1T



Practical use case #003 LSR router

- P4 dataplane fits perfectly pure **LSR** core router
- NNI: 4 directions with (8x100GE) bundle



Edge-core
NETWORKS

WEDGE-100BF-32X



Practical use case #004 LER router

- P4 dataplane fits perfectly pure **LER** use case
- NNI: EST/WEST direction @ (8x100GE) bundle
- UNI: 16x100GE left for end user connection!



Edge-core
NETWORKS

WEDGE-100BF-32X





Practical use case #005 high performance BGP RR

- Recycling new server?
- Ideal for **K8s** cluster using **BGP** as **CNI** network plugin
- Taking advantage of server « huge » amount of RAM
- No need specific high performance dataplane



Practical use case #006 « small PE » Practical

Ideal for aggregation

- 2x10GE or 2x100GE NIC server side
- 2x10g+48x1g or 1x100g+48x1/10g switch





Practical use case #007 100GE Private Peering node

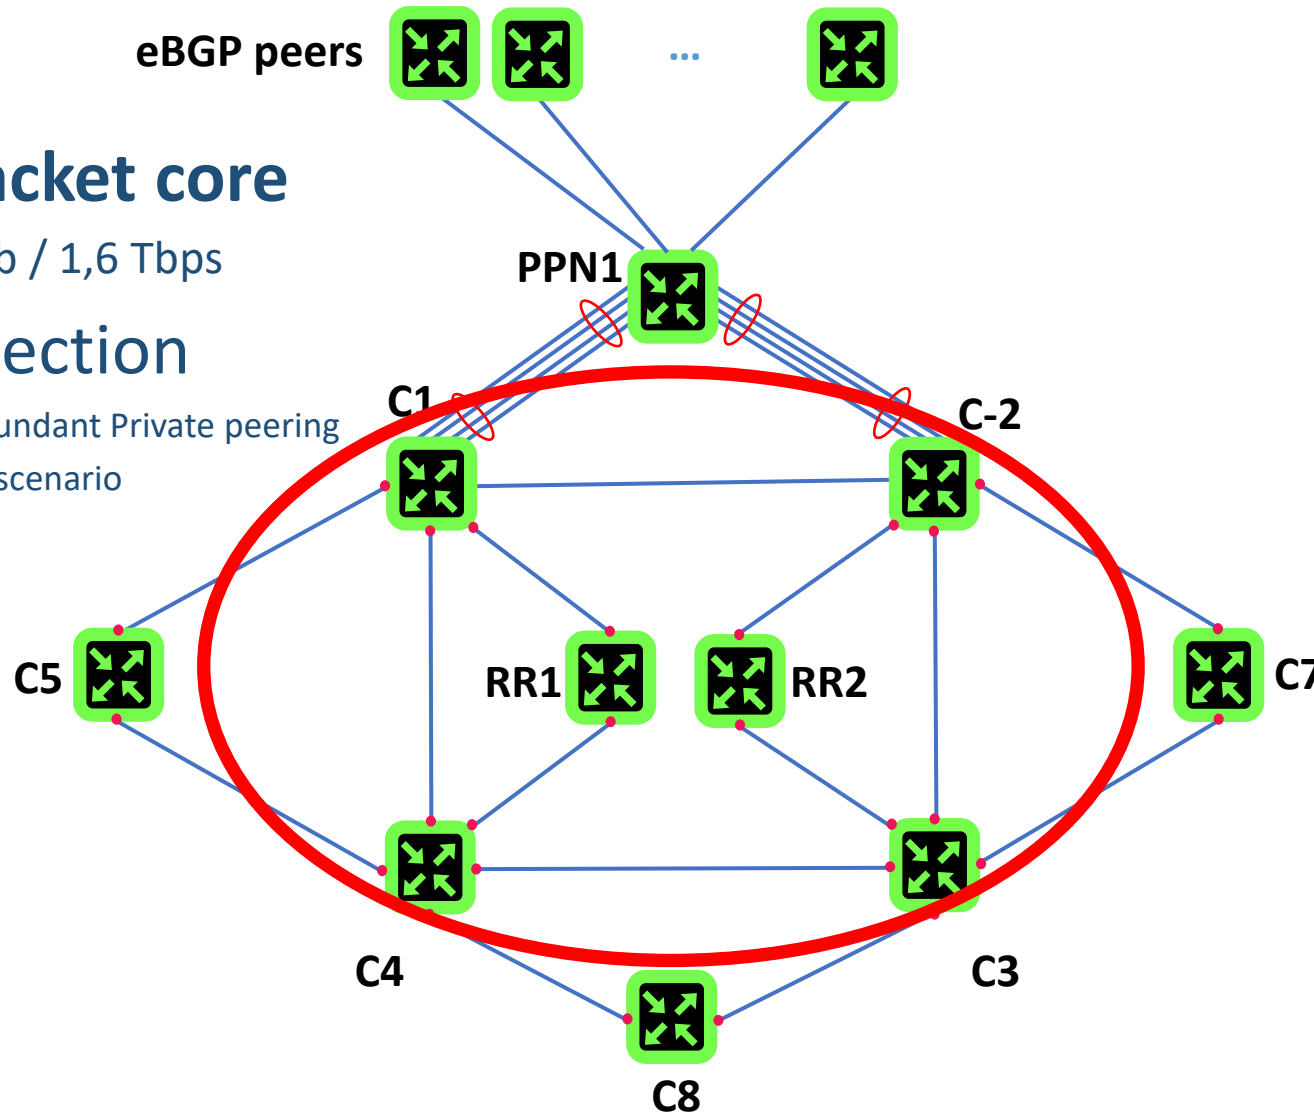


- High resilient **Packet core**

- 2 direction @ 400Gb / 1,6 Tbps

- User ports connection

- 24 ports left for 2x12 redundant Private peering
- 1:3 ratio with redundant scenario





Practical use case #xxx **The sky is the limit**

- Automation integration
- IXP with MPLS core
- ToR router combined to BGP aware network plugin
- Spine/Leaf DC router
- Global BGP monitoring for your entire BGP domain
- Global IGP guard for your entire IGP domain
- BGP flowspec aware anti DDOS
- AAA servers (TACACS, RADIUS)
- ...

We need YOUR creativity!



Acknowledgements ...



APS Networks





Useful links

- Project

freeRtr control plane's home: freertr.org

more information on dataplanes: rare.freertr.org

Project members' journey: blog.freertr.org

FreeRtr configuration guide: docs.freertr.org

- Contact

For daring RARE/freeRtr users: rare-users@lists.geant.org

For RARE/freeRtr JEDI developer wanabee: rare-dev@lists.geant.org

For RARE/freeRtr supporters  [@rare_freerouter](https://twitter.com/rare_freerouter)



IRC@DN42 #freertr



RARE/freeRtr guild



RARE/freeRtr 1st release is out !

- CI/CD deployment: ~ 2500 tests automated weekly
- ONIE image : one liner installation
- TOFINO based switches
 - Wedge 100BF-32X
 - Wedge 100BF-32QS
 - Wedge 100BF-64X
 - APS6064X
 - APS2556X1T
- An extensive list of specific tailored profile. RARE/freeRtr can act as P/PE,BNG,BRAS, CGNAT, traffic cleaner and more !
- An innovative software life cycle management leveraging **Nix package manager paradigm**.
 - it is possible to upgrade and rollback without any hitch at the software distribution level (think of an "apt upgrade" with the possibility to "apt downgrade #n » at the router level)
- And more !

For those who wants to start their journey:

Useful links: Source code!!!!



freeRtr core: sources.nop.hu/src/



TOFINO ASIC: sources.nop.hu/misc/p4bf/



P4Lang bmv2: sources.nop.hu/misc/p4lang/



p4emu: sources.nop.hu/misc/native/p4*



p4dpk: sources.nop.hu/misc/native/p4*



GitHub organisation: <https://github.org/rare-freertr>

Looking ahead: Finalize transition to production



Join the RARE project !

Extend HCL:

new TOFINO based hardware support
new DPDK release

New target:

TOFINO2
NVIDIA DPU
P4 SmartNIC
FPGA

New idea:

Polka / mPolka
Leverage Nix paradigm

And more ...

RARE/freeRrt GÉANT


P4 lab use cases




GP4L – GÉANT P4 LAB



GP4L – GÉANT P4 LAB



GP4L – GÉANT P4 LAB



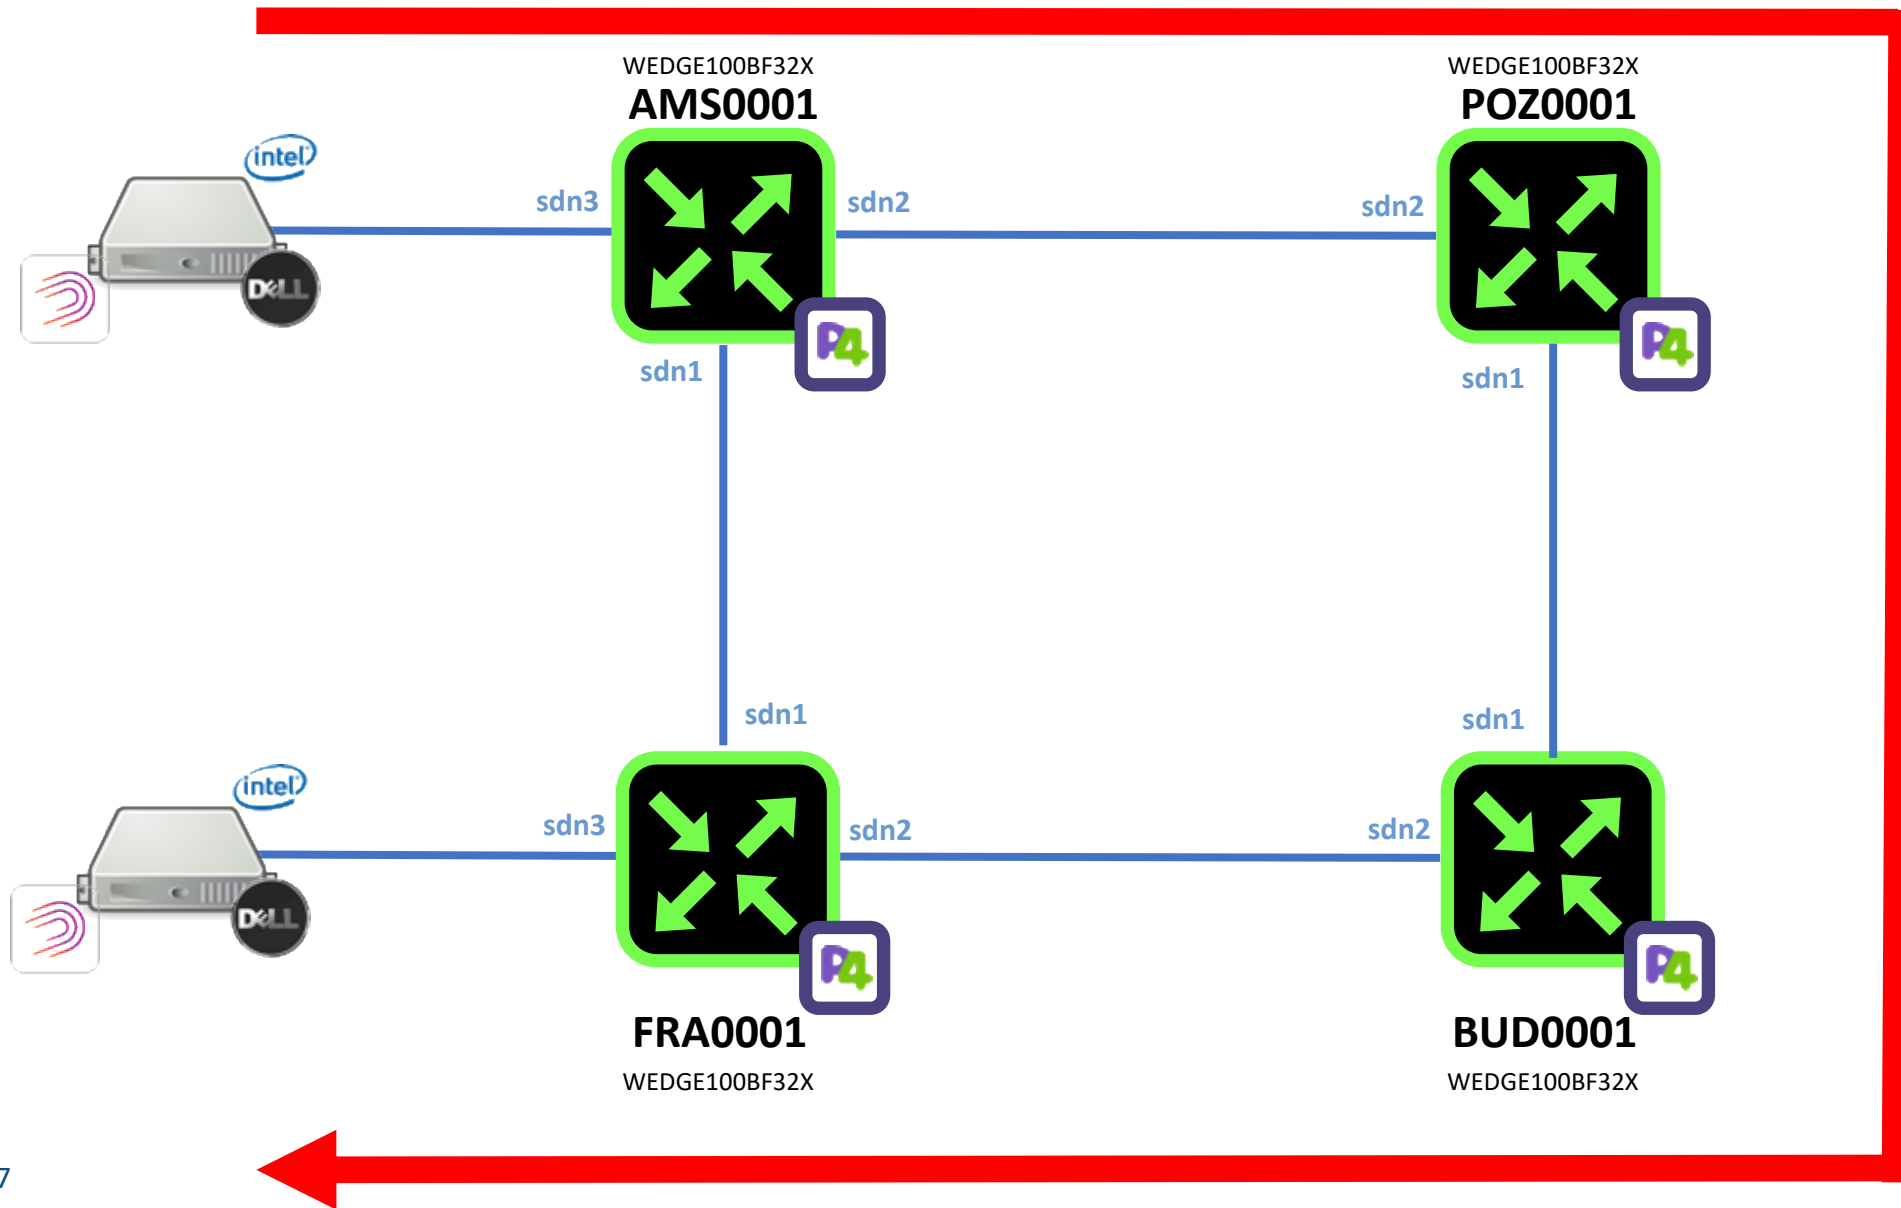
PolKA: [1]


Polynomial Key-based Architecture for Source Routing in Network Fabrics

How does PolKA work?


- Three polynomials:
 - **routeID**: a route identifier calculated using the CRT.
 - **nodeID**: to identify each core node.
 - Irreducible polynomial
 - **portID**: to identify the ports of each core node.
- The forwarding uses a **mod** operation (remainder of division):

$$\text{portID} = \langle \text{routeID} \rangle_{\text{nodeID}}$$

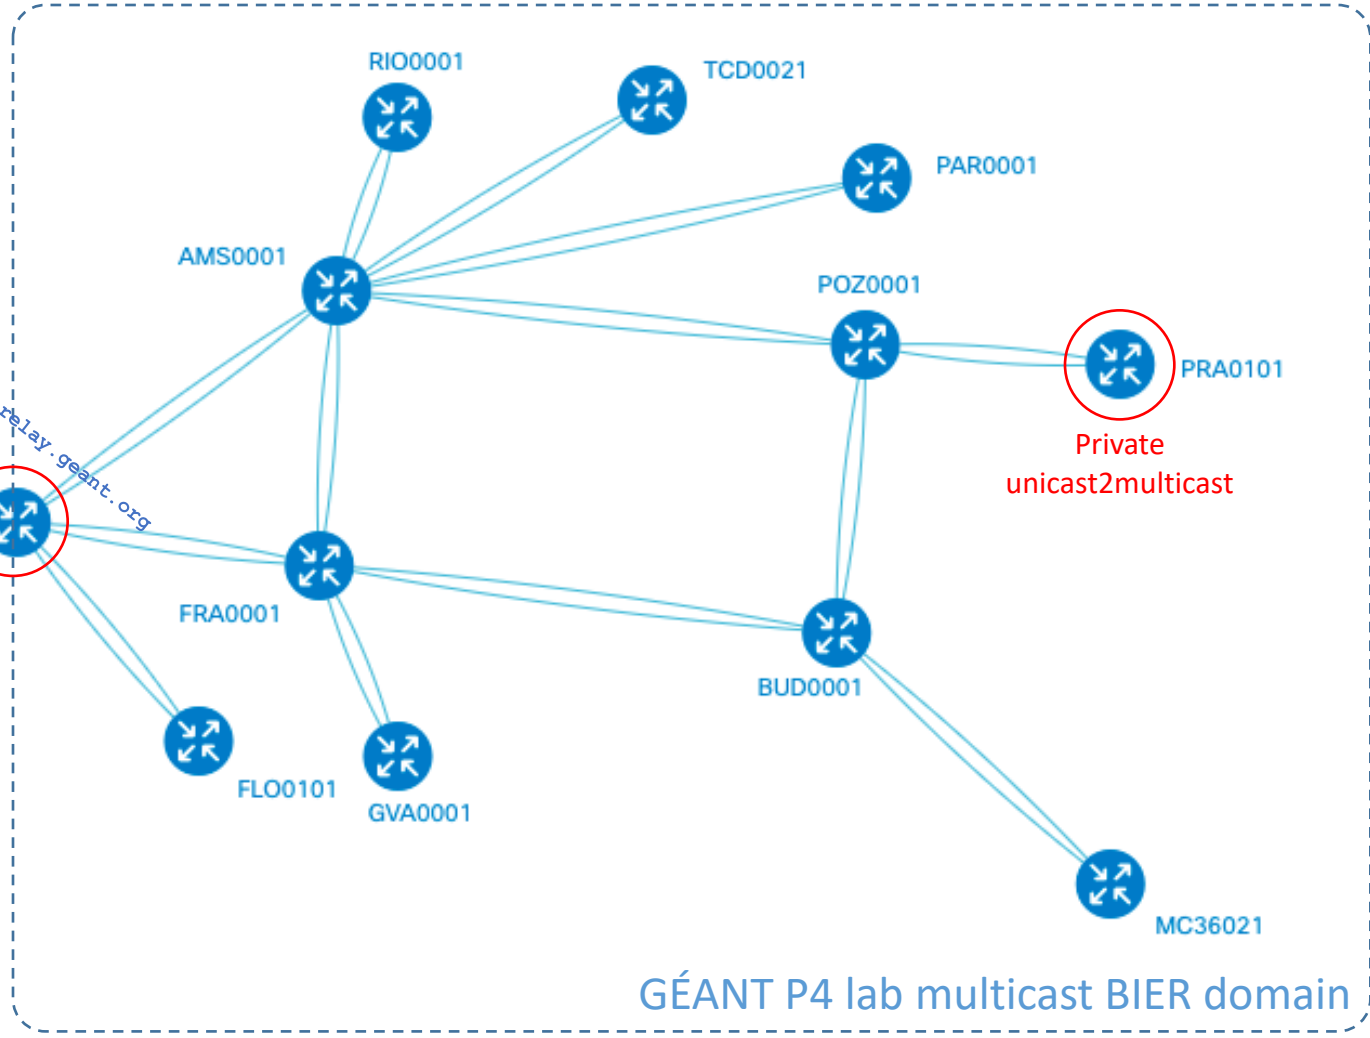





GP4L Topology rendering via BGP-LS [1]



GP4L AMT relay / AMT gateway / Unicast → Multicast



[1] <http://gp4l.geant.org/>
[2] <https://nightlies.videolan.org/>

AN INNOVATIVE MULTICAST STREAMING SERVICE !

AMT relay / AMT gateway / Unicast → Multicast

<http://mcast-menu.par.geant.org/>

<http://mcast-menu.pra.geant.org/>

sho ipv4 nat CLEARNET trans

	original		translated							
proto	source	target	source	target	age	last	timeout	pack	byte	
17	62.40.109.125 38337	62.40.109.31 2262	62.40.109.31 59487	232.123.86.28 1234	23:51:15	00:00:00	00:01:00	22711663	26527222384	
17	62.40.109.125 45392	62.40.109.31 2262	62.40.109.31 61222	232.123.45.36 1234	23:51:15	00:00:00	00:01:00	1850481	2161361808	
17	62.40.109.125 59210	62.40.109.31 2262	62.40.109.31 59051	232.123.33.114 1234	23:51:15	00:00:00	00:01:00	48333709	56453772112	

sho ipv6 nat CLEARNET trans

	original		translated							
proto	source	target	source	target	age	last	timeout	pack	byte	

sho ipv4 mrou CLEARNET

source	group	interface	upstream	targets	bytes
--------	-------	-----------	----------	---------	-------

sho ipv6 mrou CLEARNET

source	group	interface	upstream	targets	bytes
--------	-------	-----------	----------	---------	-------

sho ipv4 mrou CORE

source	group	interface	upstream	targets	bytes
--------	-------	-----------	----------	---------	-------

sho ipv6 mrou CORE

source	group	interface	upstream	targets	bytes
--------	-------	-----------	----------	---------	-------

- **INSTALL VLC 4**

<https://nightlies.videolan.org/>

- **Launch VLC AMT client**

`vlc amt://62.40.109.31@232.123.76.124:1234 --amt-relay amt-relay.geant.org`

Thank you

Any questions?

www.geant.org

