

## Programmable Networking with P4, GEANT RARE/freeRtr and SONiC/PINS

Marcos Schwarz, RNP, marcos.schwarz@rnp.br

### Abstract

We will demonstrate a L3 overlay network composed of P4 programmable switches from the GNA-G AutoGOLE / SENSE Persistent Multi-Resource Infrastructure and the GEANT RARE P4 testbed, stitched through L2 circuits aggregating capacity from multiple trans continental, international and regional capacity. This network will be composed of 100G and 400G P4 devices running SONiC and RARE/FreeRtr.

SONiC and RARE/FreeRtr are full featured open source network operating systems (NOSes) that at the same time implement traditional networking protocols and can be extended to integrate next generation features as an opt-In strategy. Which provides means to implement production oriented networks and enable experimentation features and protocols only to a given set of users, without impacting the rest of the network.

SONiC is a well established Open Source NOS with contributions and adoption by major merchant silicon manufacturers and switch platform vendors. Its initial focus was on datacenter requirements and protocols with support for both fixed-function and programmable switches, more recently it has been improved to support Telco requirements and protocols like MPLS and Segment Routing. Another important recent extension was PINS (P4 Integrated Networking Stack) by Google, ONF (Open Networking Foundation) and Intel, which implements a hybrid control plane with support to P4Runtime as an optional interface enabling SDN, with which ONF is adopting SONiC as a basis for it's next generation architectures as in SD-FABRIC.

RARE/FreeRtr, developed by GEANT, is a R&E oriented routing platform with a rich support for Telco protocols like MPLS, Segment Routing, Path Computation, BGP-LS and many more, implemented on P4 switches and CPU packet processing platforms (DPDK, eBPF/XDP). It has been developed in partnership with REN (Research and Education Network) and Academic communities, to support advanced/novel protocols and use cases, like PolKA (Polynomial Key-based Architecture) a stateless source routing protocol based on arithmetic operations over a polynomial encoded route label, and IPv6 Flow Label Packet Marking initiative from WLCG.

Exploring the advancements from SONiC and RARE/FreeRtr, using persistent P4 testbeds from GNA-G AutoGOLE and GEANT RARE, we are able to build state of the art networks and provide a pre-production testbed to integrate, validate and showcase: 1) emerging industry architectures/protocols (i.e. SRv6), 2) novel protocols from the NRE community (i.e. PolKA, Packet marking) and 3) a base platform for the evolution of intelligence and orchestration initiatives (i.e. SENSE)

### Goals

1. Demonstrate a persistent global L3 overlay network based on P4 switches
2. Demonstrate the capability to support multiple virtual networks, that implement on the same devices different choices of routing stacks, traditional and SDN based
3. Demonstrate an intercontinental high capacity transfers (100G and over) exploring multiple source routing solutions (Segment Routing and PolKA)
4. Demonstrate the management infrastructure and tools used to operate this global network

### Resources

Capacity from trans continental, international and regional networks through AutoGOLE/SENSE

P4 devices and DTNs from AutoGOLE/SENSE and GEANT/RARE P4 testbed sites: Caltech, CERN, RNP, KISTI/KREONET, Starlight, Tennessee Tech, Trinity College of Dublin, UCSD, UNESP

### Involved Parties

- **AmLight/FIU:** Julio Ibarra ([Julio@fiu.edu](mailto:Julio@fiu.edu)), Jeronimo Bezerra ([jbezerra@fiu.edu](mailto:jbezerra@fiu.edu)), Vasilka Chergarova ([vchergar@fiu.edu](mailto:vchergar@fiu.edu)), Italo Valcy ([idasilva@fiu.edu](mailto:idasilva@fiu.edu))
- **Caltech HEP:** Harvey Newman ([newman@hep.caltech.edu](mailto:newman@hep.caltech.edu)), Justas Balcas ([jbalcas@caltech.edu](mailto:jbalcas@caltech.edu)), Raimondas Sirvinskas ([raimis.sirvis@gmail.com](mailto:raimis.sirvis@gmail.com)), Catalin Iordache, Preeti Bhat, Andres Moya, Sravya Uppalapati
- **CERN/SWITCH:** Edoardo Martelli ([edoardo.martelli@cern.ch](mailto:edoardo.martelli@cern.ch)), Carmen Misa ([carmen.misa@cern.ch](mailto:carmen.misa@cern.ch))
- **GEANT/RENATER:** Frederic Loui ([frederic.loui@renater.fr](mailto:frederic.loui@renater.fr))

- **KISTI/KREONET:** Buseung Cho ([bscho@kisti.re.kr](mailto:bscho@kisti.re.kr)),  
Tergel Munkhbat ([tergelmunkbat@kisti.re.kr](mailto:tergelmunkbat@kisti.re.kr))
- **HEAnet:** Eoin Donal
- **SouthernLight/Rednesp:** Antonio J F Francisco,  
Ney Lemke (**UNESP**) ([ney.lemke@unesp.br](mailto:ney.lemke@unesp.br)),  
Carlos Antonio Ruggiero (**USP**)  
([toto@ifsc.usp.br](mailto:toto@ifsc.usp.br)), Bruno Baldim  
([bruno@ansp.br](mailto:bruno@ansp.br)), Rogério Motitsuki  
([rogerio@ansp.br](mailto:rogerio@ansp.br))
- **RNP:** Marcos Schwarz ([marcos.schwarz@rnp.br](mailto:marcos.schwarz@rnp.br)),  
Leandro Ciuffo ([leandro.ciuffo@rnp.br](mailto:leandro.ciuffo@rnp.br))
- **Starlight/MREN/iCAIR:** Joe Mambretti ([j-mambretti@northwestern.edu](mailto:j-mambretti@northwestern.edu)), Jim Chen ([jim-chen@northwestern.edu](mailto:jim-chen@northwestern.edu)), Fei Yeh  
([fyeh@northwestern.edu](mailto:fyeh@northwestern.edu))
- **TCD:** Frank Slyne ([fslyne@tcd.ie](mailto:fslyne@tcd.ie))
- **Tennessee Tech:** Susmit Shannigrahi  
([sshannigrahi@tntech.edu](mailto:sshannigrahi@tntech.edu))
- **UCSD:** John Graham
- **UFES/IFES:** Magnos Martinello  
([magnos.martinello@ufes.br](mailto:magnos.martinello@ufes.br)), Moises Ribeiro  
([moises@ele.ufes.br](mailto:moises@ele.ufes.br)), Christina Dominicini  
([crisrina.dominicini@ifes.edu.br](mailto:crisrina.dominicini@ifes.edu.br)), Everson Borges  
([everson@ifes.edu.br](mailto:everson@ifes.edu.br)), Rafael Guimaraes  
([rafael@rafaelguimaraes.net](mailto:rafael@rafaelguimaraes.net))
- **UMd/MAX/ESnet:** Xi Yang (**ESnet**) ([xiyang@es.net](mailto:xiyang@es.net)),