

Project Overview: Router for Academia Research & Education (RARE)

Frédéric Loui

RARE/freeRtr technical leader

RENATER/GÉANT

Paris, France

frederic.loui@renater.fr

Csaba Mate

RARE/freeRtr lead developer

KIFU/GÉANT

Budapest, Hungary

cs@mp.ls

Alexander Gall

RARE/freeRtr lead release manager

SWITCH/GÉANT

Zurich, Switzerland

alexander.gall@switch.ch

Maxime Wisslé

RARE/freeRtr developer

RENATER/GÉANT

Paris, France

maxime.wissle@renater.fr

Abstract—In this paper, we give a brief overview of the GÉANT project Router for Academia Research & Education (RARE). Its objective is to create a modern open routing platform. It leverages the FreeRtr control plane and programmable data planes that aims to support Research and Education (R&E) specific use cases. RARE/freeRtr provides different data plane implementations that allow for a high degree of deployment flexibility. Most router features are developed with P4 for the software target BMv2, and where possible also for the Intel Tofino hardware target. Additionally, a DPDK-based data plane called P4DPDK for x86 platforms is also available that addresses use cases requiring lower bandwidth.

I. INTRODUCTION

RARE [3] is an acronym with a rich history [5] in the European Research and Education (R&E) networking community. The acronym dates from 1987 and was the first name for the fledgling European R&E organization before becoming TERENA and later GÉANT [1]. In our context, RARE stands for Router for Academia Research & Education. The goal of the RARE project is to create a modern open router operating system that can quickly be adapted to meet the numerous challenging use cases from the R&E community.

This paper outlines some of the latest challenges faced by the R&E network community today, and describes the corresponding solutions utilizing the RARE platform. Additionally, while presenting the RARE vision, it will also provide the reader with a high-level understanding of the modular RARE architecture and how it contributes to the rapid development and automated testing of new protocols and features. While the RARE team is working on several use cases, this paper will focus on the use of RARE as a fully-featured router operating system for P4 hardware devices such as Intel Tofino, as well as for software targets such as DPDK.

II. THE RARE MISSION STATEMENT

The RARE project aims to provide an open, programmable router operating system which will meet the challenging demands of current and future R&E networking use cases. Examples of some of the R&E use cases we are working on include:

- Route reflector on virtual or bare metal server

This project is funded by the GÉANT 4th EU programme.

- High speed (up to 100 Gbit/s) university border router
- Data transfer nodes for research
- Science DMZ routers
- Small routers for schools CPE on Dell VEP platforms
- Source routing for R&E networks
- Telemetry for IP networks
- In-band Network Telemetry (INT) for P4 networks
- BIER (Bit Index Explicit Replication) multicast forwarding using MPLS encapsulation for MPLS [2]
- PolKa (Source routing in P4 networks using polynomials) [6]
- Multidomain networking architectures with the GNA-G and Caltech for future solutions for CERN data transfers

III. RARE/FREERTR OVERALL ARCHITECTURE




Fig. 1: FreeRtr overall architecture.

RARE utilizes freeRtr open source control plane [4]. The relationship between RARE and freeRtr is similar to the GNU/Linux relationship, where GNU is part of the free software foundation project and Linux is the kernel used. Similarly, RARE is the umbrella project that uses freeRtr as control plane. As shown in Figure 1, the RARE/freeRtr architecture consists of 3 components:

A. Open Source Control Plane

freeRtr open source control plane is used. It provides all features and protocols expected from an enterprise grade router operating system, with additional support for P4 hardware devices and custom-tailored, novel protocols.

B. Programmable Data Plane

RARE provides implementations for different programmable data plane targets that can interact with freeRtr control plane through a well-defined common message API. BMv2 data plane is supported for development purposes, as well as Intel Tofino for high-performance use cases. Additionally, a DPDK-based data plane called P4DPDK is provided, which supports the complete RARE router feature list [3] on x86 platforms.

C. Common interface between the Control Plane and the Data Plane

Different data plane targets use different native message APIs. For example, BMv2 and Intel Tofino are using the P4Runtime and BfRuntime API respectively. FreeRtr implements a simple common text message API to interact with these native data plane APIs, thereby enabling programmable access from the control plane while supporting interchangeable data planes.

IV. RARE/FREERTR VALUE PROPOSITION

In the following, the results of the RARE project are presented. Generally, special care has been taken to make the solutions simple and sustainable.

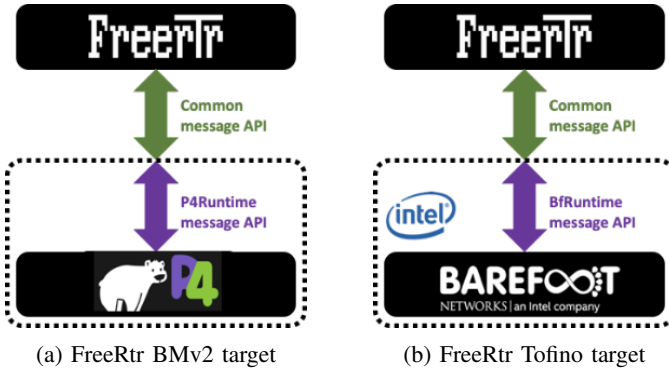


Fig. 2: FreeRtr P4 targets

The BMv2 target is mainly used in development for validating code logic, allowing automated nightly unit- and regression testing. BMv2 conveniently offers unlimited stages and resource memory, which allows focusing solely on the logic of packet-forwarding algorithms. While BMv2 is great for algorithm sketching and educational purposes, the Intel Tofino platform offers a hardware-based target with a total switching capacity of 6.4 Tbit/s for high-performance use cases. P4 programs written for BMv2 do not necessarily run on the Tofino platform due to resource limits and other constraints related to high-speed packet processing [7]. While porting the RARE code to Intel Tofino, some features such as cryptographic operations could not be implemented due to these hardware limitations. Since the BMv2 and Intel Tofino code base share the same skeleton and coding style, a RARE/freeRtr Tofino image using BMv2 has been developed that allows to verify the data plane implementation without

requiring actual hardware. The freeRtr architecture for the P4 targets is shown in Figure 2.

For use cases that require complex network features on low bandwidths, the Intel Tofino hardware platform therefore is not flexible enough, as it imposes too many resources limitations. In order to resolve this gap, the RARE team implemented a custom DPDK-based data plane which is suitable for x86 targets, called P4DPDK. As shown in Figure 3, the P4DPDK data plane offers the same common message API so it can be seamlessly used together with the freeRtr control plane. As DPDK performance is based on the available hardware resources such as CPU, it provides a great solution for use cases which require lower data rates beneath 100 Gbit/s.




Fig. 3: freeRtr DPDK Target

V. SUMMARY

P4 is the first domain-specific language that has unlocked the door to programmable data plane. User applications that were limited in terms of throughput in the past, now have the ability to run at hardware line rate of up to 400 Gbit/s with Tofino2. The RARE/freeRtr open source routing platform provides a networking swiss army knife solution based (among others) on the P4 language. New target devices that combine hybrid technologies such as FPGA and Intel Tofino together with technology like DPDK or RDMA will unlock new use cases. Suggestions for use cases or questions to the RARE team can be addressed by writing an email to rare-users@lists.geant.org

REFERENCES

- [1] GÉANT Project. [Online]. Available: <https://www.geant.org/>
- [2] IETF BIER (Bit Index Explicit Replication) Working Group. [Online]. Available: <https://datatracker.ietf.org/wg/bier/about/>
- [3] RARE (Router for Academia, Research Education). [Online]. Available: <https://wiki.geant.org/display/RARE/Home>
- [4] RARE/FreeRtr - freeRouter. [Online]. Available: <http://www.freertr.net/>
- [5] H. Davies and B. Bressan, *A History of International Research Networking: The People who Made it Happen*. John Wiley & Sons, 2010, ISBN: 978-3-527-32710-2. [Online]. Available: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/9783527629336>
- [6] C. Dominicini, D. Mafioletti, A. C. Locateli, R. Villaca, M. Martinello, M. Ribeiro, and A. Gorodnik, "PolKA: Polynomial Key-based Architecture for Source Routing in Network Fabrics," in *2020 6th IEEE Conference on Network Softwarization (NetSoft)*, 2020, pp. 326–334. [Online]. Available: <https://doi.org/10.1109/NetSoft48620.2020.9165501>
- [7] F. Hauser, M. Häberle, D. Merling, S. Lindner, V. Gurevich, F. Zeiger, R. Frank, and M. Menth, "A Survey on Data Plane Programming with P4: Fundamentals, Advances, and Applied Research," 2021. [Online]. Available: <https://arxiv.org/abs/2101.10632>