25-11-2021

Deliverable D6.7
Network Technology Evolution Update

Deliverable D6.7

Contractual Date:	30-11-2021
Actual Date:	25-11-2021
Grant Agreement No.:	856726
Work Package	WP 6
Task Item:	Task 1
Nature of Deliverable:	R (Report)
Dissemination Level:	PU (Public)
Lead Partner:	RENATER
Document ID:	GN4-3-21-41D9AE
Authors:	Xavier Jeannin (RENATER), Piotr Rydlichowski (PSNC), Krzysztof Turza (PSNC), Wojbor Bogacki (PSNC), Frédéric Loui (RENATER), Claudio Allochio (GARR), Mauro Campanella (GARR), Tim Chown (Jisc), Ivana Golub (PSNC), Tomasz Szewczyk (PSNC), Maxime Wissle (RENATER), Jani Myyry (FUNET), Theodore Vasilopoulos (GRNET), Jordi Ortiz (RedIRIS, University of Murcia)

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 856726 (GN4-3).

Abstract

This document reports on the evaluations of the applicability of a number of emerging technologies in the GÉANT community conducted by the *Network Technology Evolution* task (WP6 T1) within the *Network Technologies and Services Development* work package (WP6) of the GÉANT4-3 project. It is an update of the deliverable D6.5 *Network Technology Evolution Report*, March 2020, and related to the milestone document M6.10 *Evaluation of Planned Network Technology Experiments and Potential for Use by the GÉANT Community*.

Table of Contents

Table of Figures

Executive Summary

This document reports on the evaluations of the applicability of a number of emerging technologies in the GÉANT community conducted by the *Network Technology Evolution* task (WP6 T1) within the *Network Technologies and Services Development* work package (WP6) of the GÉANT 4-3 project. It is an update of the deliverable D6.5 *Network Technology Evolution Report* [D6.5], and related to the milestone M6.10 document, *Evaluation of Planned Network Technology Experiments and Potential for Use by the GÉANT Community* [M6.10]. Technologies in the scope of the work included white box (WB), Data Plane Programming (DPP), Router for Academia, Research and Education (RARE), low latency networking, Optical Time and Frequency Networking (OTFN) and Quantum Key Distribution (QKD).

The assessment of the use of white boxes in a production context for several R&E use cases (CPE, GIX, data centre) demonstrated that hardware and NOS disaggregation can allow greater flexibility in the way NREN networks are managed. The use case evaluations were complemented by performance and Total Cost of Ownership (TCO) analysis. The validation of white boxes spans over the chosen hardware, software, use case and into performance testing. It should be noted that the use and adoption depend not only on technical but also strategic and economic aspects.

The DPP team implemented a monitoring tool using In-band Network Telemetry (INT) in the data plane in a testbed that includes three European countries, spanning over a production network. Three types of INT nodes are developed: source, transit and sink, and measurements of inter-arrival time, timestamps, packet reordering etc. are provided for UDP traffic. Considering this INT tool provides finer granularity; it can be used as an excellent addition to traditional monitoring systems. The team is sharing its experiences, open issues, and challenges via infoshares and documents which can be found on the DPP wiki pages [DPP_Wiki].

RARE is an open-source routing platform well suited to a number of R&E networking use cases. It has successfully implemented an open-source control plane (FreeRtr) over multiple data planes, including P4 and DPDK. RARE is deployed in the GÉANT P4 Lab which has grown from an initial deployment at four GÉANT PoPs to an infrastructure that spans worldwide, including North and South America, Asia and Europe. Its adaptability was demonstrated by the PolKA team that won a Google Research Scholar Prize.

The low latency networking subtask designed and successfully implemented a segment-by-segment latency and jitter measurement and visualisation solution called TimeMap, which allows network operations teams to support troubleshooting of real time applications such as LoLa. The tool is being deployed by the GÉANT Operation team for monitoring the GÉANT backbone.

The OTFN team began by reviewing Time & Frequency (T&F) deployments in the GÉANT and NREN community. It became evident that there are multiple methods by which T&F services can be implemented, over different frequency bands. The focus of the OTFN task, therefore, was to consolidate and share knowledge about these current different approaches to implementing time and frequency services, while also assisting GÉANT in its own L-band testing for GN4-3N circuits. The team has also worked to support interconnections for the T&F services between various NRENs.

The QKD team has assessed the state of QKD projects within the GÉANT and NREN community, including the available QKD hardware and software, and complemented this with an NREN survey. The team's focus is now directed towards the dissemination of quantum and QKD technology among the GÉANT community through infoshares and white papers. Quantum simulators were used to explore NREN QKD use cases, broadening the NREN interest and engagement for quantum technologies. Work has been done to prepare a long distance QKD proof of concept between two GÉANT PoPs, which is planned to be run in 2022 in collaboration with Toshiba, and the OpenQKD project.

1 Introduction

The *Network Technology Evolution* task (WP6 T1) within the *Network Technologies and Services Development* work package (WP6) of the GÉANT 4-3 project evaluates the applicability of a variety of emerging technologies for adoption in European National Research and Education Networks (NRENs).

Network technologies evolve by nature very quickly and require WP6 to be agile and responsive. At the beginning of the project, the evaluation work started on the following areas with an initial assessment period of two years:

- White Box for research and education (WB)
- Data Plane Programming (DPP)
- Router for Academia, Research and Education (RARE)
- Data Transfer Node Infrastructures (DTNs)
- Low Latency networking
- Optical Time Frequency Network (OTFN)
- Quantum Key Distribution (QKD)

Some of the work was completed (DTN, WB), and some continued (RARE, DPP, low latency networking, OTFN and QKD). For each of the work topics, this document briefly summarises individual achievements and provides references to further information.

This document is an update of the deliverable D6.5 *Network Technology Evolution Report* [D6.5], delivered in March 2020 and reports in more detail on the work areas presented in the M6.10 milestone document *Evaluation of Planned Network Technology Experiments and Potential for Use by the GÉANT Community* [M6.10].

Although presented in D6.5, the Data Transfer Node infrastructure (DTN) work was moved to the *Network Services Evolution and Development* task of WP6 (Task 2), where it has since concluded with the work results published in a white paper [DTN], an infoshare and a wiki page [DTN_Wiki]. Therefore, it is no longer described in this document.

The structure of this document is as follows:

Section 2 presents the work done on white box for research and education, with summary of all considered use cases. Section 3 covers the DPP work on In-band Network Telemetry (INT). RARE is presented in Section 4, low latency networking and TimeMap in Section 5, OTFN in Section 6 and QKD work in Section 7, followed by the Conclusions section.

2 White Box for Research and Education

A white box is a switch/router manufactured from commodity components that allows different Network Operating Systems (NOSs) to be run on the same piece of commodity hardware, decoupling the NOS software from the hardware. This subtask aimed to validate whether the white box solutions available on the market can be useful in an NREN context, including features, functionalities, as well as performance [WB_Perf] and Total Cost of Ownership (TCO) [TCO], [TCOxls]. The approach, defined in detail in deliverable D6.5 *Network Technology Evolution Report* [D6.3], is based on potential NREN use cases for which a white box solution is validated, and then put into production if the outcome is acceptable to the NREN.

The organisations that participated in the WB evaluations had similar objectives: explore if a selected white box solution can replace the existing solution, increase their independence from hardware and software vendors, and determine whether it would be possible to reduce TCO while providing an acceptable solution.

The white box evaluations planned in WP6 have all been completed, with each use case published as a white paper and/or in a deliverable.

There are four main conclusions of the work:

- White boxes can be used in production for different use cases, many of them relevant in the NREN context.
- Implementing a white box solution does not require significantly different knowledge and expertise to that already existing in most NRENs.
- The decision to use white boxes in production depends not only on technical and TCO aspects, but also on prior experience and knowledge as well as organisational strategy.
- To validate a white box for production, a thorough (ideally automated) performance testing of the hardware, software, control and data plane in the specific context of a given use case is needed.

Another WP6 team extended the white box paradigm further and developed an open-source RARE Operating System (ROS) [RARE] running over a programmable white box. ROS is now sufficiently mature and available for production use.

2.1 Normandy CPE

The regional network of Normandy [RN] was looking for a solution for their network SYVIK [SYVIK] to replace existing customer premises equipment (CPE) (Cisco routers) at their high schools in order to improve the capacity and flexibility of these devices, as well as to increase independence from its hardware provider and be able to bring additional services to the high schools. The equipment in use at the time of this evaluation was not capable of sustaining the increasing network bandwidth demand, and also is not the cheapest solution. The routing requirements for the CPE were limited: BGP peering, IGP, VLAN, logical interface, VRF light [Normandy_CPE].

In 2019, the available bare-metal white boxes were generally equipped with high forwarding capacity interfaces, due to their usual data centre use case, which increased the price drastically and were not necessary for use as a CPE. For a lower cost solution, an x86 white box server running a virtualised environment with a virtual router on top of the hypervisor was considered to be more attractive. To be specific, the virtual NOS, implemented as a VM running on top of a Linux VM. This architecture also provides a way to support deployment at scale via Linux automation tools. The chosen solution was a small machine from the Dell VEP (Virtual Edge Platform) family [VEP], suitable for running a virtualised environment, and providing 2*10 Gbps and several 1Gbps ports at a competitive price.

Two different NOSs were tested and validated: Cumulus Networks and Free Range Routing (FRR). Both of them run on an Ubuntu server Linux VM and have been proven to be suitable for a router CPE implementation. Based on the evaluation, the SYVIK administrators decided to use FRR [FRR]. FRR is an open-source solution which provides the necessary set of features for a router CPE device, and is free of charge, easy to understand, and very familiar for those who have worked with most of the popular vendors' networking operating systems in the market.

A great advantage of a white box running in a virtualised environment is that it can support Virtual Network Functions (VNFs). Depending on the hardware components (RAM, CPU, etc.), it can support one or several VNFs on the same hardware, delivering a single box solution for many purposes at a customer's premises (example: router + firewall + web proxy). By doing so, the customer can have a lower TCO with a single white box device.

An important consideration during the design phase of a white box CPE solution is the choice of the hypervisor and the orchestrator. This choice determines the structure of the whole management solution. In this particular case, the white box CPE solution was tested over free Proxmox/KVM [PROX] and commercial VMWare ESXi [VMW] hypervisors. Both are fully compatible with a white box x86 server, and Proxmox/KVM was chosen for economic reasons.

The solution was deployed in production in twelve high schools in Normandy. For a future CPE deployment, SYVIK is now considering evaluating RARE/FreeRtr, which was developed during this white box project.

2.2 Funet CPE

The Funet CPE use case was evaluated to find alternative platforms for a campus edge router to be used within the Finnish University and Research Network (Funet). The existing solution at the time used fully fledged traditional routers with L3-capable switches but it was not seen as a suitable solution due to the overall cost as well as the available features. It was, therefore, considered that a white box solution might do better regarding both price and performance/features.

Two solutions were analysed: a network operating system (NOS) which typically supports data centre-specific equipment, and a NOS running on hardware primarily designed for mobile networks but sharing similar features to those required in a campus environment.

The first solution was based on the Cumulus Linux virtual NOS and virtual forwarding plane where the focus was to test the control and management plane features. The hardware supported by Cumulus was designed for data centre environments, and, even though the solution worked in a stable manner for several months, it did not have all the features required for a CPE, such as deep buffers, and was, therefore, considered as not suitable for a CPE use case.

The second solution was based on Edgecore AS7315-27X [E_AS7315-27X] devices, which are 100Gbps-capable Disaggregated Cell Site Gateways (DCSGs) specified by the Telecom Infra Project (TIP) [TIP_DCSG], [E_AS7315-27X_DCSG]. As Cumulus NOS was not supported, ADVA Ensemble Activator NOS, provided by the supplier of the equipment, was used. It was clear that the ADVA NOS was more focused on supporting customers in mobile networks, but it also seemed to be able to provide the features needed in CPE environments. NOSs designed for mobile aggregation seem to lack some features, such as Layer 2 bridging and storm protection, commonly used in campus networks. On the other hand, they implement many unnecessary protocols that were not seen as needed in the Funet use case.

The targeted topology for Funet's CPE use case is presented in Figure 2.1. More detail about the use case can be found in the use case white paper [CPE_Funet].

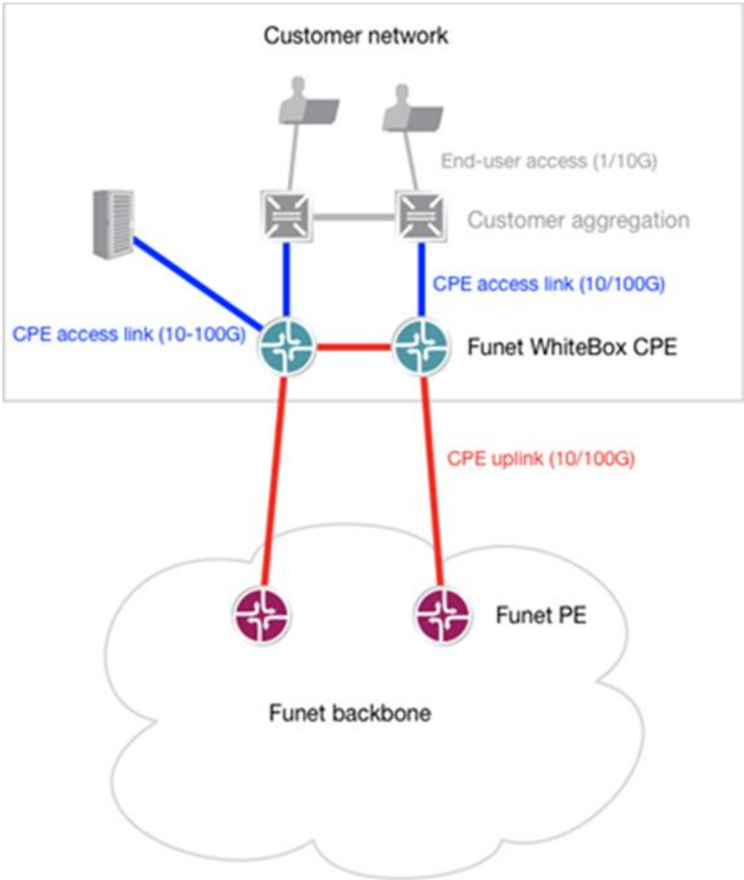

Figure 2.1: FUNET CPE project

Considering the end benefits and shortcomings of the explored solutions for the given scenario, it is unlikely that either of the two solutions would currently be the primary choice for the CPE scenario in Funet. However, the evaluation has shown the potential for white boxes, and so a white box solution will continue to be considered, for example, for less demanding users looking for cheaper solutions.

2.3 CPE Murcia University - RARE/FreeRtR

This use case was designed and deployed as a part of the RARE project (see Section 4). The RARE project provides a platform that is able to interconnect a NOS over several data planes, which typically would be FreeRtr as the control plane with a P4 or DPDK data plane.

The GAIA laboratory targets the research of Future Internet architectures, and houses several facilities for the University of Murcia and the Spanish government. For this use case, the goal was to renew the CPE router (a Cisco 7200) used in the GAIA laboratory at University of Murcia.

The technical requirements were Ethernet, MSTP and VLAN at L2, as well as IPv4, IPv6, LACP, NAT, DHCP, HRSP, L3 ACLs, tunneling, Wireguard, and OpenVPN. In terms of throughput, the minimum requirement was 1Gbps. At the management level, the cost of the solution is a very strong constraint, as the maintenance cost of the traditional vendor CPE is not insignificant.

Vendor independence is important. On one hand, an open-source approach provides release management independence from vendor lock-in and ensures continuity that is only limited by the interests of the GAIA laboratory. On the other hand, however, an open-source approach permits the GAIA research laboratory to perform proof of concept implementations beyond state of the art in a real (production) environment.

The University of Murcia has participated in the RARE project from its start, and, as soon as the routing operating system based on RARE/FreeRtr software achieved a production-grade maturity, a process was started to explore if this solution could be suitable for the laboratory. It was determined that the solution should have all the features needed for this use case, and be able to support a sufficient number of NAT sessions, i.e., to be able to store a large table within the data plane.

A RARE/FreeRtr solution based on a server with DPDK support was chosen as it was possible to forward several tens of Gbps of traffic with such a platform and not suffer from a memory shortage. To be sure that a backup solution was available, an HSRP session was set up between the RARE/FreeRtr system and the Cisco 7200. The production deployment of a DPDK DELL server running RARE/FreeRtr was completed in October 2020, and has been running without any problem since then. This deployment supports all the network services of the laboratory. A second white box was deployed in May 2021 - a WEDGE 100BF32X equipped with an Intel Tofino 1 chipset (with a chipset forwarding capacity of 6.4 Tbps) [Intel_Tofino]. This white box supports all services of the laboratory except the NAT service (due to hardware limitations for storing a full NAT table); the performance of this 100G-capable device with RARE/FreeRtr was also tested, providing good results.

The introduction of this new network device required configuring the appropriate HSRP sessions between the Cisco 7200, the WEDGE 100BF32X and the DPDK DELL server to provide a backup for all the network services, including the NAT service. Figure 2.2 shows the architecture of the network deployed in production at the GAIA laboratory.

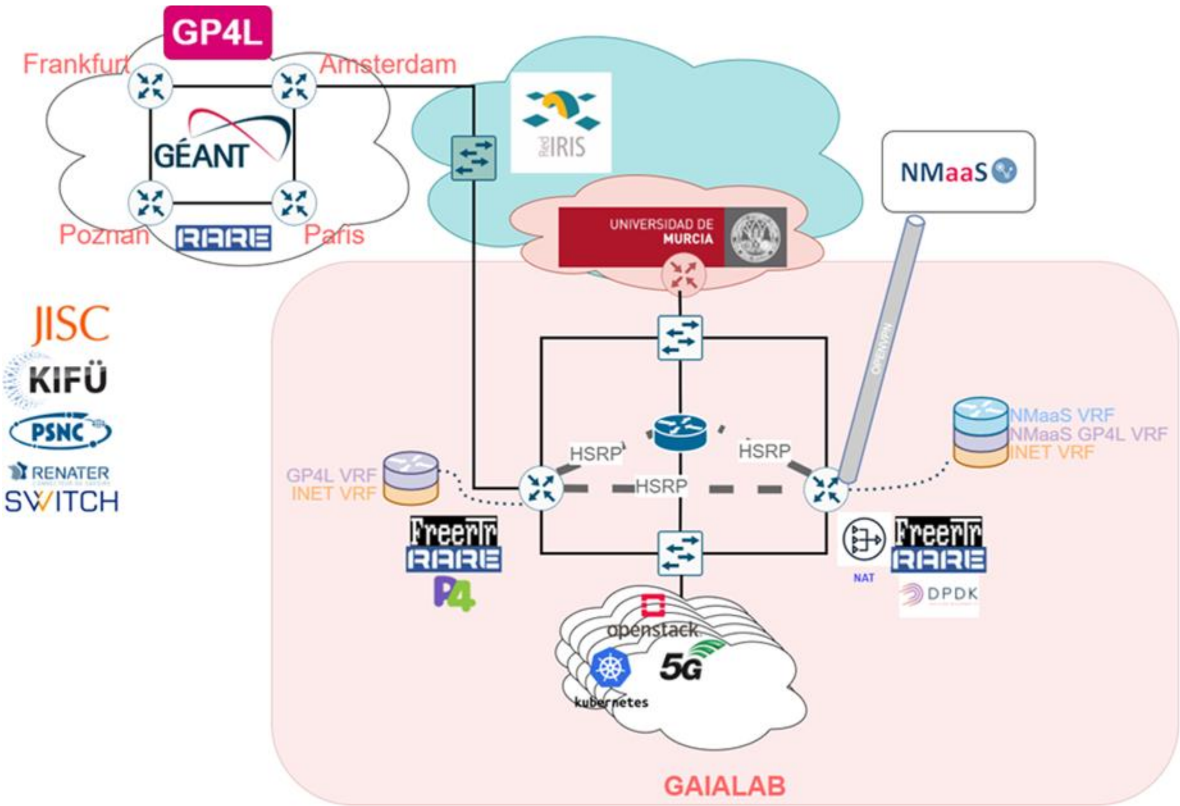

Figure 2.2: Current network architecture of Gaia Lab

This solution demonstrates the stability of RARE/FreeRtr in production, and that it provides an affordable CPE solution for the laboratory. It is a very good example of how the combination of RARE/FreeRtr on a white box can be successfully used in R&E organisations to both provide good production connectivity, and the ability to rapidly test new protocols and services.

2.4 GRNET Data Centre

For GRNET, the Greek NREN, the motivation for a solution based on a white box was to have better independence from vendors, more flexibility in choosing their preferred NOS and hardware, cost savings, and to be able to use open-source solutions.

The data centre design was based on an IP Clos EVPN/VxLAN fabric, and on a spine-leaf folded Clos topology. The network devices selected to implement this design are the Edgecore AS7712-32X (equipped with the chipset forwarding Tomahawk 3.2Tbps) for the spines, and the Edgecore AS5812-54T (equipped with the chipset forwarding Broadcom Trident II+ 720Gbps) for the leaves. The chosen NOS was Cumulus Linux, recently acquired by NVIDIA. The topology of the data centre is shown in Figure 2.3.

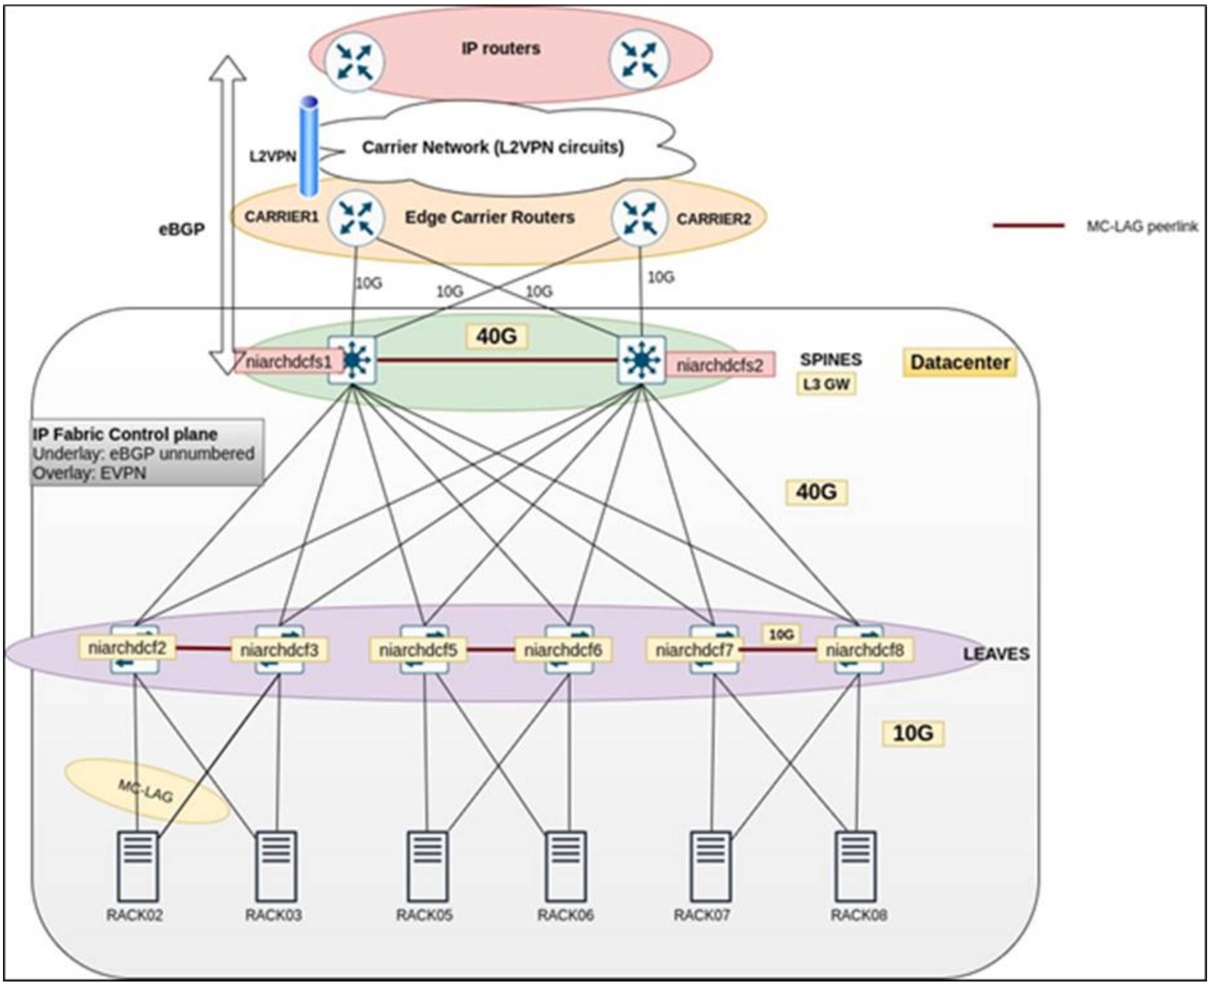

Figure 2.3: Data centre topology - spine-leaf folded Clos topology

The implementation of this solution shows that a white box switch can be used in a data centre fabric in a production network. All the fundamental functionalities and features (such as BGP, BFD, ECMP, EVPN, VXLAN, LACP, LLDP, and ACLs (L3 and L2)) were achieved. Nevertheless, some advanced features were not available for the chosen switches equipped with Broadcom forwarding chipset. As NVIDIA also owns the chipset supplier, Mellanox, the NVIDIA strategy is to provide some Cumulus features only on the Mellanox chipset (for instance the EVPN Multihoming feature). The fallback was to use MC-LAG to implement active-active server connectivity and some control plane redundancy features. Moreover, from release 4.4 (July 2021), Cumulus Linux will support only NVIDIA Spectrum-based Application-Specific Integrated Circuit (ASIC) platforms [CMS], putting the white box based on Broadcom ASICs at its end of life within the next few years. However, since the software and hardware are decoupled in a white box, GRNET can either keep white box switch hardware and try another NOS (e.g., Pluribus), or try out white box switches with cheap NVIDIA chipsets and keep the switches based on Broadcom for another use case.

The white box switch performance results are sufficient for regular operation and can handle several failure scenarios. Further details can be found in the white paper *White Box: GRNET Data Centre Use Case* [DataCentre].

The network devices are now implemented in production. At the time of writing, this solution supports the beginning of this small data centre that might gradually gain momentum as new computing requests emerge.

2.5 Provider Router (P) / Label Switch Router (LSR) in PSNC

PSNC, the Polish NREN, was looking for a scalable high-capacity MPLS platform for their network PIONIER, that would be capable of meeting future demands as traffic volumes increase.

One architecture used in NREN MPLS core networks uses the functionality provided by a Label Edge Router (PE/LER) and a Label Switching Router (P/LSR) on the same platform, a PE/LER (see Figure 2.4). This is a very cost effective and flexible solution. Each node in the network can be used as a service endpoint and transit for other traffic at the same time.

In this collapsed core architecture, the PE/LER routers are hosted in the NREN PoPs that are interconnected by optical fiber. Often a DWDM transmission system is implemented on optical links. In this case, the number of required transponders increases with the growing capacity of the network, and each time a pair of transponders is required, the CapEx and OpEx are increased.

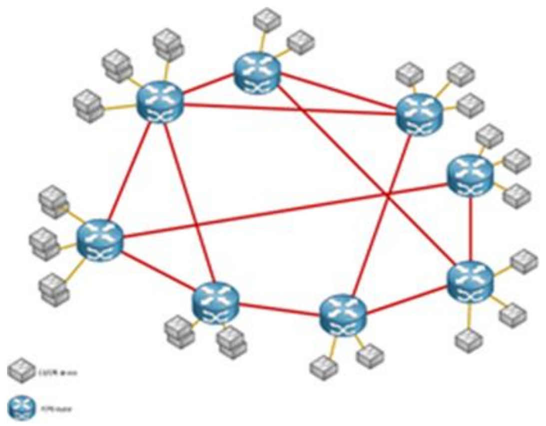

Figure 2.4: Collapsed core architecture

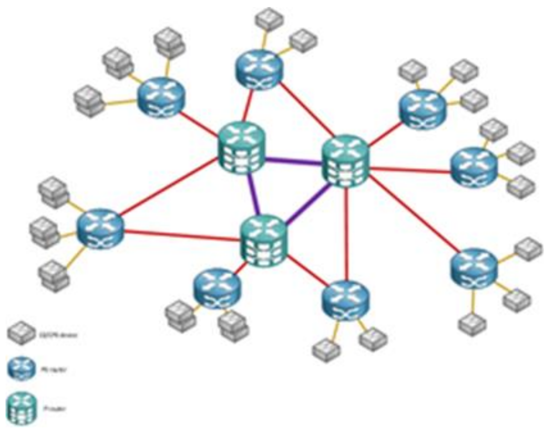

Figure 2.5: P/LSR core architecture

The scalability of MPLS networks can be increased by adding another level to the network hierarchy. P/LSR routers are used mainly as aggregation and transit nodes for traffic generated by PE/LER routers (see Figure 2.5). In most cases, the LSR routers do not terminate the network services.

One potential solution was a white box platform with NOS software running on it. There are two key factors which need to be considered for the P/LSR architecture. The first is the software and hardware compatibility of the white box platform. The second is the interoperability of the device with the existing network elements. All these aspects were analysed in the PSNC laboratory.

During the first step, the NOS supporting the MPLS control plane protocols was selected. At the moment the most mature solution in this area is delivered by IPinfusion's OcNOS system. Next, a white box platform supporting OcNOS was installed in the lab. The traffic switching performance tests were conducted with hardware boxes, and the functional tests were conducted in the virtual environment.

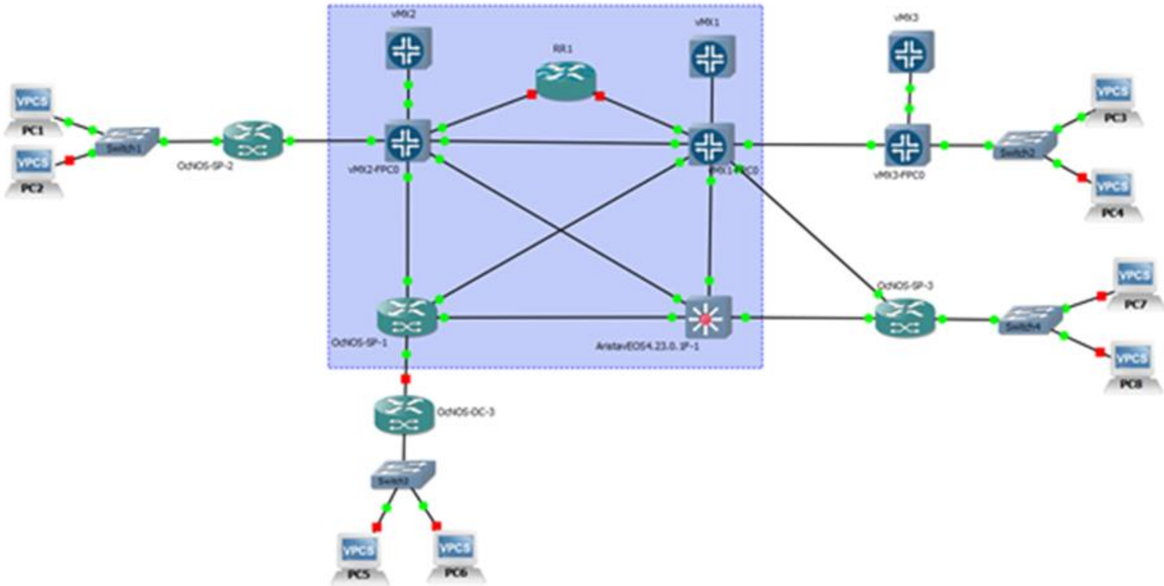

Figure 2.6: Functional and interoperability tests network diagram

The results obtained during the tests showed that MPLS technology is actively developed on the white box platforms. Due to the origin of white box vendors (like Dell or Edgecore), a strong relation to data centre networks is noticeable. Therefore, a number of features are available that are heavily used in these networks.

In the case of existing MPLS core networks, interoperability is a key topic to consider. During the testbed tests, RSVP signalled LSPs and E-Line services were successfully established between OcNOS and the Juniper vMX. Some incompatibilities and some advanced TE functionalities were detected for the MPLS traceroute functionality.

As a result of the conducted tests, PSNC decided not to consider the OcNOS-based white box platform as suitable for country-wide deployment. This was mainly because:

- A set of specific features already used in the NREN core network are currently not available on the white box platforms or are causing some interoperability problems. Examples include LSP auto-bandwidth, fast reroute in link-node protection mode, and streaming of detailed telemetry data directly from the network node.
- The currently used vendor platform has proven over time to support new features that are added to it, without the need for hardware replacement).

However, the white box solution might be interesting for research projects that are building their own lab infrastructures as it offers a good CapEx level with a set of flexible functionalities.

2.6 Internet eXchange Point in RENATER

RENATER, the French NREN, wanted to renew one of its Global Internet Exchange points (GIX), SFINX, with a similar set of features and for a significantly lower cost. The objectives of this work were to evaluate the white box solution in production, increase independence from traditional vendors, provide new solutions for future deployments and reduce the Total Cost of Ownership.

The current solution is based on Brocade MLX switches that are providing a Layer 2 connection service to SFINX users. For a new white-box-based solution, a Dell (Dell EMC S4048-ON) machine and a NOS called OcNOS from IP Infusion were chosen. The solution was evaluated at the technical level to consider all needed functionalities, at the operational level to explore the level of maintenance and required knowledge needed to support the solution in production, and at the strategic level to consider further development directions. Figure 2.7 shows the white box implementation in SFINX.

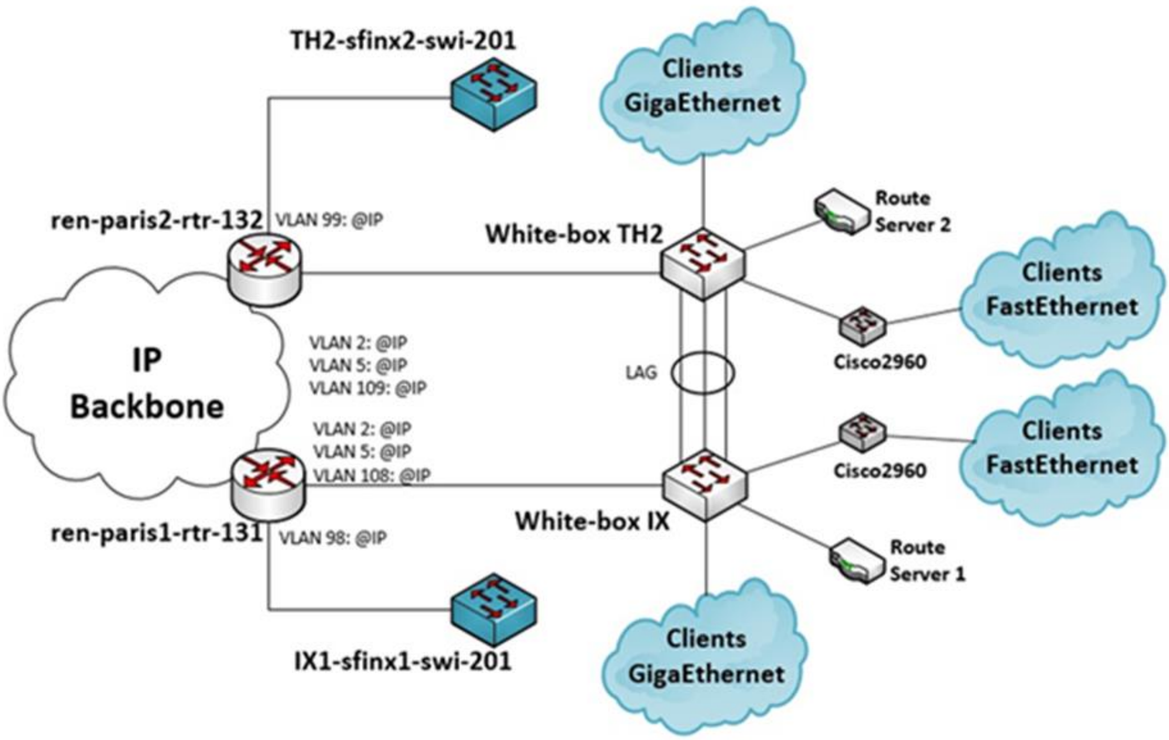

Figure 2.7: SFINX GIX topology with white box switches

At the technical level, it was noticed that OcNOS does not provide the same set of features as traditional vendor equipment, especially for MPLS. However, in this use case, it was concluded that the feature set provided by this solution satisfied all requirements, and that it provides space for a possible evolution of SFINX. During tests, when a configuration was changed extensively, the white box solution showed some signs of instability, but, after the configuration was implemented, the switch ran in production for several months without any problems. In addition, it was not possible to integrate this solution with the TACACS system, but this was resolved through a workaround.

At the operational level, it was concluded that the same level of maintenance is needed for both - hardware and software. A question was raised about how quickly and easily the NOC engineers could become familiar with a new NOS. However, the NOS is very similar to the Cisco IOS and it was very easy to learn how to work with the new switch. Except for the TACACS limitation, the integration with the NOC tools was straightforward (OOB access, monitoring, SSH connection). The hardware maintenance was provided by the supplier, and the NOS maintenance by the software supplier. If it is not obvious where to direct a required maintenance request, the RENATER NOC will first ask the software supplier who is the best party to identify the origin of the problem. For example, a BIOS problem was solved correctly thanks to this process. More information about this use case is available in a white paper [GIX].

At the strategic level, by implementing the white box solution, the RENATER management achieved greater flexibility in addressing its deployment requirements. The white box solution was already considered for network edge (like in the CPE Normandy use case), or specific network functions (like for SFINX GIX), or for other future network areas.

Finally, a white box-based implementation of the GIX use case was considered successful: all the needed features are provided, and there is space for future growth. The TCO was reduced, however, given the market's responsiveness to competition, the cost comparison should be revisited as appropriate.

Similar to the CPE Normandy use case, after this GIX deployment project was finished, the RARE Operating System [RARE] was developed to the point where it was sufficiently mature for production use, and it is considered to be a candidate for the next generation of the SFINX.

2.7 White Box Performance Testing

The performance of a device depends heavily on its hardware and software, as well as the specific usage and placement of the device in the topology of a network. It is worth looking at hardware and software from a control plane and data plane perspective separately. The approach to evaluating the performance of a Device Under Test (DUT) should take different aspects into consideration: - hardware, software, control and data plane, and the use case. The final results should clearly state the observed environment.

This is particularly important in the white box context, as hardware can originate from one vendor and software from another. Multiple combinations of NOSs and hardware are possible, and the performance assessment can depend heavily on the selected use case, which may require a specific set of features. It is possible that one feature available in a NOS can be used only on a specific hardware, or that some parts of the software platform could require optimisation depending on the selected hardware.

When testing hardware, different factors can impact the performance. For example, for forwarding chipsets, the following elements should be considered:

- The number of match/action stages in the pipeline
- Static Random Access Memory (SRAM)

- Ternary content addressable memory (TCAM)
- Packet buffer size
- Cryptography capabilities

These factors are important, as, for example:

- Some chipsets support more labels that can be processed in a single lookup than in the case with commodity forwarding chipsets.
- Ternary content-addressable memory (TCAM) can limit the number of routes the device can manage which also limits its applications.
- Some forwarding chipsets could embed cryptography capabilities allowing for instance MACSec and IPSec support at line rate (BCM88850 Jericho2c+).
- With more stages, the chipset can manage more protocols and features, so the number of stages could limit the number of features that the NOS can implement.

For control plane testing, it is important to define a set of protocols, technologies, and functionalities that are relevant to the selected use case. For example, in this particular P/LSP device use case, performance should be evaluated for a specific number of prefixes in memory, a list of established protocol sessions or adjacencies, etc. The control plane performance would then include not only verification of how many prefixes or LSP can be handled by the platform, but also whether the device switches the traffic correctly in a given setup.

For data plane testing, a set of benchmarking methodologies is defined in two RFC documents. RFC 2544 defines a number of tests that may be used to verify the performance characteristics of a network device and the setup of a DUT (number of routes, filters, protocols configured). RFC 2889 provides the methodology for the benchmarking of a local area network (LAN) switching devices and focuses on switching performance tests. Most of the tests are specific to Layer 2 switching, but some can be used in other switching environments such as an MPLS backbone network. For example, a congestion control test can be used to validate MPLS P/PE devices. The Spirent Test Centre [STC] provides some pre-prepared tests aligned with RFC 2889.

One of the tests performed by the WP6 team was related to switch buffer sizes, for which four platforms were tested: Edgecore AS5912-54X-O-AC-F, Dell EMC S4248FBL-ON, Arista 7280SR, and Juniper PTX. The egress interface was overloaded by generating two packet streams from two ingress interfaces to a single egress interface. Each stream kept the same mean transmission speed in bits per second. However, each stream also contained a number of burst packets in order to reproduce a 'real network traffic' profile (Internet MIX [CAIDA]) according to RFC 2544. The results have shown that even with line-rate switching performance and without head of line blocking, the results between the tested hardware platforms were different.

In another test, the WB performance test team chose a P/LSP use case as an example to elaborate on a more general approach to network device validation. This use case was implemented on two platforms: Edgecore AS5912-54X-O-AC-F and Dell EMC S4248FBL-ON tested under the OcNOS operating system, using the Spirent traffic injector. To check how the DUT would operate in a mid-size MPLS core network, a network consisting of 11 neighbouring routers was defined on interfaces on the network tester. One of the routers was a DUT neighbour intermediate router and 10 were last-hop routers (destination routers). Next, RSVP was used to signal a number of transit LSPs (2x900 LSPs) passing through the DUT. Finally, traffic was mapped to emulated LSPs and the accuracy of the forwarding mechanism verified. In this case, the test results were positive. The tests are explained in more detail in the *White box performance testing and evaluation* white paper [WB_Perf].

The performance tests emphasised the importance of the dedicated approach to the hardware-software selection in the specific use case environment. In addition, this work provided a standards-based model for network device testing that can be used in or outside of an NREN environment. For further details, see [WB_Perf].

2.8 White Box Total Cost of Ownership Analysis

Beyond an organisation's strategic interest in a white box approach, there is an expectation of a reduction in the Total Cost of Ownership (TCO) of a white box router or network device. The TCO is an estimate of all the direct and indirect costs involved in acquiring and operating a product or system over its lifetime. It is usually calculated as the sum of capital expenditures (CAPEX) and operating expenses (OPEX).

The WB team undertook work to consider the different components of that cost and give guidance on how to calculate the TCO for an R&E use case. The principles used for that calculation were:

- Identify and describe the components and parameters that influence the cost of white box technology implementation and are required to make the calculation.
- Make a fair comparison that includes equipment with similar capabilities, capacities, and performance.
- Software comparisons must accurately consider the features each product provides. If feature gaps exist, the costs of supplementary software from the vendors or third parties must be included for a true comparison.
- In situations where there are no other reliable reference prices for the required equipment and services related to a particular institution, community, or country, use vendors' global price lists (GPLs).

The team's white box TCO analysis showed that choosing a hardware or software solution always involves taking several dimensions and aspects into account. Technology, technical features, prior knowledge and experience, price range, existing solutions, and systems are only some of them. For any of the solutions in scope, the TCO is extremely important – primarily to business stakeholders, but it should also be recognised and considered by the technical teams.

Also, TCO is not the only criterion for the hardware choice. Being locked in by a vendor is likely to be a drawback for flexibility. The analysis has resulted in the documentation of a method to calculate the TCO that will allow NRENs and R&E institutions to compare several solutions easily. It is only after looking at the TCO of a particular use case that the NRENs and R&E institutions can recognise clear benefits of the technology and decide whether to proceed with a production implementation

Further details can be found in the white paper [TCO]. The team also produced an Excel spreadsheet [TCO-xls] that helps organisations to compare different solutions, and included an example of such a comparison in the white paper.

3 Data Plane Programmability

The work of the Data Plane Programmability (DPP) team in WP6 T1 aims to investigate how the innovation of data plane programmability can help to improve existing services and facilitate the design of new services. After assessing the use of P4 data plane programmability on switches equipped with a Tofino chipset to detect DDoS attacks [DPP_DDOS], the DPP sub-task focussed on high precision network monitoring using In-Band Network Telemetry (INT).

In-band Network Telemetry is a powerful magnifying glass on network behaviour, and its specification is defined by the P4.org Application Working Group [INTspec]. Figure 3.1 shows an overview of the three types of INT nodes, each with a dedicated function:

- The INT 'source' node, which adds a small INT header to every chosen packet containing for example Switch IDs, Interfaces IDs, Timestamps, Link and queue utilisation.
- An INT 'transit' node, placed on the path between the source and the sink node which may add specific local information.
- The INT 'sink' node, which exports all INT data to the collector.

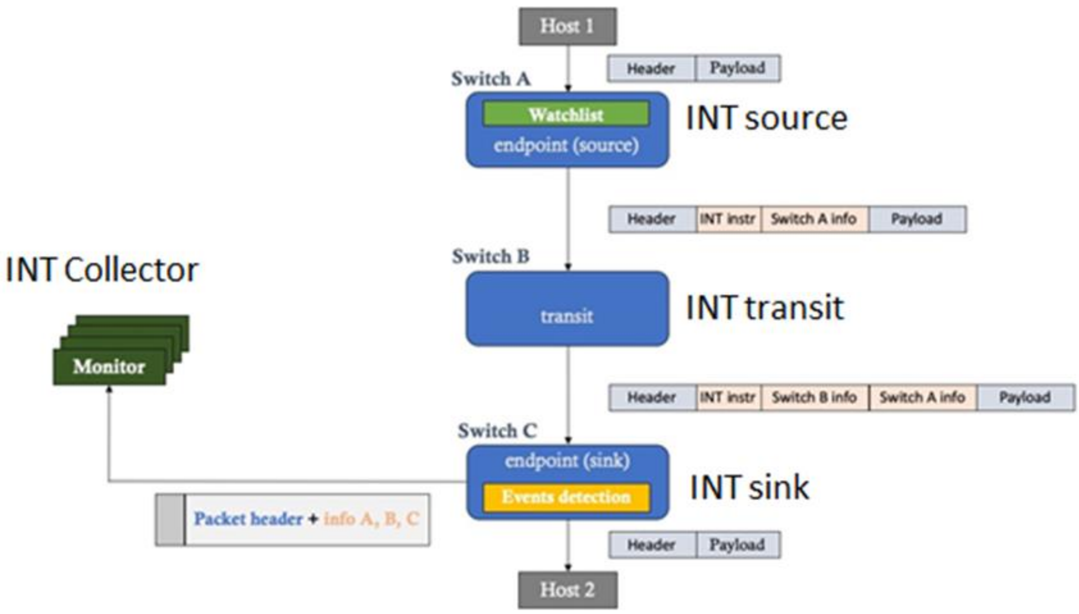

Figure 3.1: In-Band Network Telemetry overview

To implement and test INT, the WP6 DPP team deployed a testbed over the production European R&E network (see Figure 3.2).

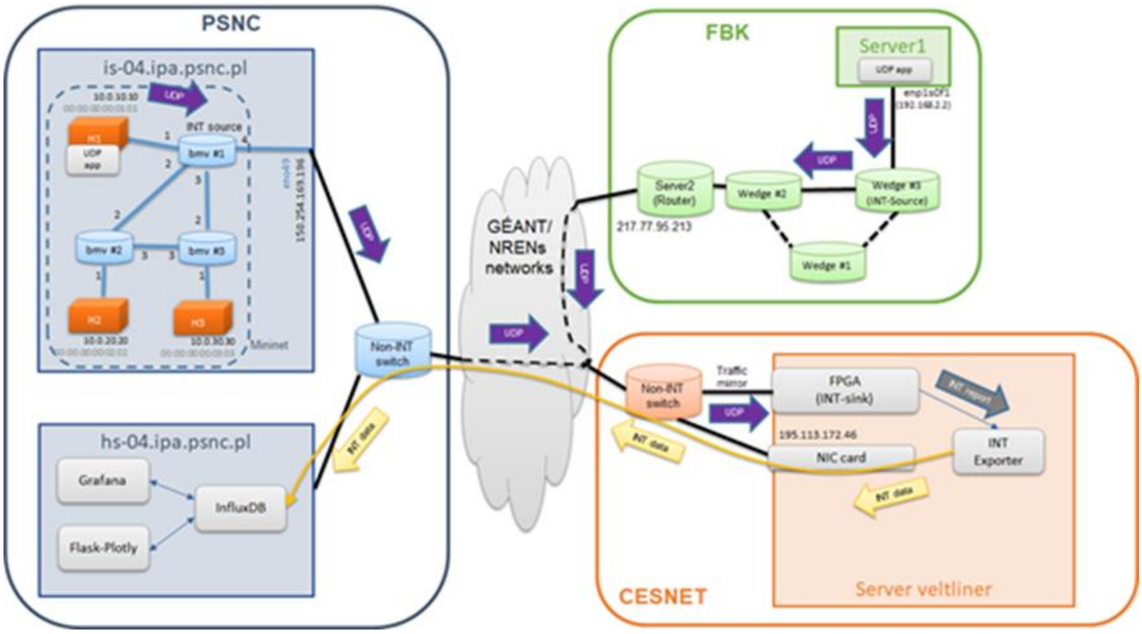

Figure 3.2: DPP In-band Network Telemetry testbed

All three types of nodes are being developed for three types of network devices: a BMv2/Mininet virtual environment implemented on servers, a server equipped with a COMBO-200G2QL FPGA, and several Wedge100BF-32X Tofino programmable switches. CESNET is testing a fourth platform based on a commodity PC, and the Data Plane Development Kit (DPDK) software [DPDK], as it can provide an off-the-shelf hardware solution with adequate performances for link speeds of up to 10 Gbps.

The testbed uses public IP addresses for all its nodes to ensure that the monitored packets flow over the NRENs' production networks. The UDP packets are transported from the INT source to the sink node over the testbed, the INT data are sent to a collector located in PSNC, with packet rates of up to 260K packets per second currently supported. The addition and removal of INT headers is performed using an EdgeCore Wedge Tofino switch and FPGA-based network card, respectively. The experiments which generated 260K INT reports per second demonstrated that each packet in a fully loaded 1Gbps UDP flow can be monitored using INT.

Due to the potentially large amount of gathered data where INT headers are added to every packet in a flow, INT data storage and post-processing is a significant challenge. In this testbed, the data are stored in an InfluxDB database, and visualised using Grafana and Plotly. Figure 3.3 shows an example of Grafana visualisation for an UDP flow of about 2Mbps.

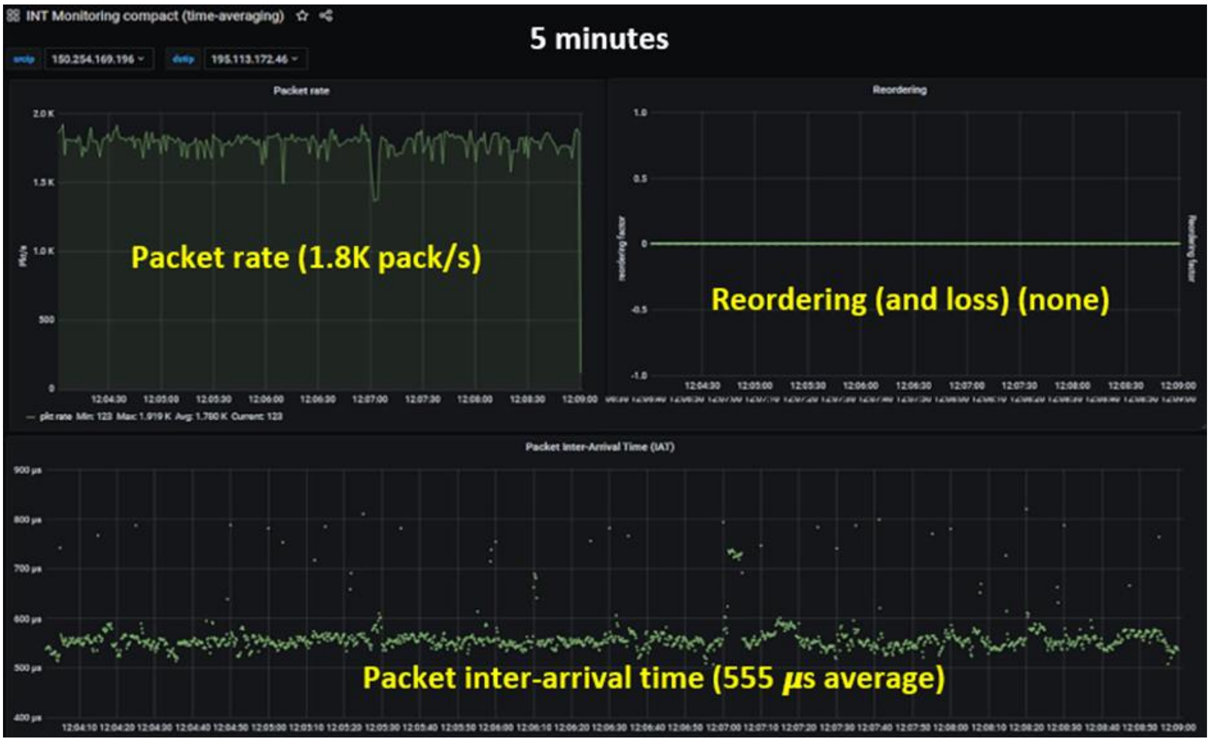

Figure 3.3: Plotly visualisation of INT sink and INT source timestamps difference during 100ms for a UDP flow from PSNC to CESNET (1.8K packets/s)

Plotly, an in-house developed tool [PSNCPlotly], was used to present the INT data with finer granularity than provided by Grafana. Plotly can present data gathered at the packet-level resolution down to the nanosecond level, which is particularly important for high-speed data transfers. Four types of graphs can be presented using Plotly: timestamps difference, timestamps differences histogram, inter-arrival time and inter-arrival time histogram (an example of time difference between packets at arrival is shown in Figure 3.4).

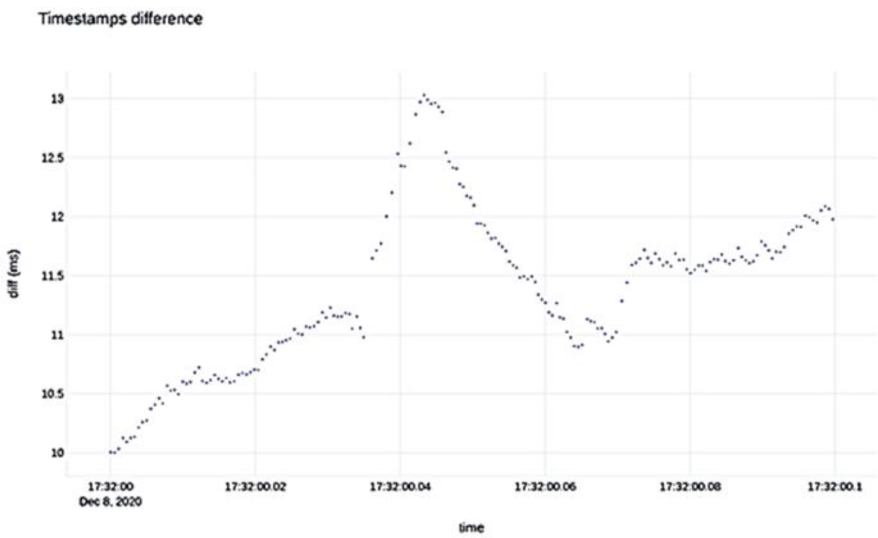

Figure 3.4: Packet inter-arrival time graph using Plotly tool based

One lesson learned from this implementation is that the time synchronisation needs to be tuned to a sufficient precision (of at least a few microseconds) to provide consistent analysis and delay measurements. Some platforms require specific effort to ensure that timestamps are reliable at the microsecond or lower precision for specific use cases. The challenges also stem from a node's internal clock stability, the time synchronisation protocol, and precise time source. To address this, the team is exploring time synchronisation aspects for INT use cases further, which might include the potential to integrate time synchronisation techniques already developed for the Tofino chipset on other platforms. A white paper on this topic is in development.

There are several other areas of work in focus for the team. One is to continue the development of the code in-line with the standard for selected platforms. The team is developing a solution to support (two-way) TCP flows, having initially focused on (one way) UDP flows. The development of an INT version over DPDK is in progress, and the group is also considering packaging the code in containers or virtual machines to more readily allow people that are interested in the technology to learn more about data plane programming and INT in practice, including data collection and representation. In addition, the team may explore the usage of other databases (e.g., NoSQL) to store the data and be used to handle INT metadata or analytics.

Future work might also include assessing the need and defining good practice for INT for high-speed links (i.e., above 10 Gbps). It can be expected that for many use cases it would not be necessary to capture INT headers for every packet, so determining the selection criteria and frequency could be useful. Therefore, several techniques may be considered to apply INT to 10Gbps or even 100Gbps flows, including:

- In-data plane reduction of the number of INT reports by skipping INT data insertion in a packet if the INT data is very similar to data inserted in the previous packet.
- In-data plane generation of INT data aggregates, histograms, and events when such final monitoring output is required by users/applications of the monitoring system.

- Intermediate processing of INT data using message broker systems (e.g., Kafka and Kafka Streams) allowing the reduction of the amount of data to be stored in the INT database.
- Using InfluxDB only to store INT data that, for example, represents "large" variations with respect to the expected value and data received from low-level INT data reduction/aggregation.

The direction for future work will be determined based on the available resources, but, most importantly, users' needs and further technology developments.

More information about this work is available on the INT wiki [INT_WIKI] which also contains links to the white paper *In-Band Network Telemetry Tests in NREN Networks* [INT_WP], the *Deploying In-band Network Telemetry (INT) on R&E networks* infoshare [INT_Info] given in March 2021, and the INT source code [DPP_Github].

4 Router for Academia Research and Education (RARE)

The WP6 team working on the Router for Academia Research and Education (RARE) has created an open-source routing platform for multiple data planes with the goal to support research and education community use cases. Its use is being validated in many scenarios, including through the deployment of the GÉANT P4 Lab which features a core of four RARE nodes located at GÉANT PoPs.

The RARE platform includes FreeRtr software as the control plane component and supports five data planes: UNIX sockets, Libpcap, DPDK, BMv2 (P4), and the Intel Tofino chipset. The software can run on a variety of platforms, such as the Edgecore Wedge 100BF-32X, which features 32 100G interfaces at a very competitive price.

RARE can be used for different use cases that include:

- Service Provider: CPE (SOHO, Small schools, University Campus, small MAN, etc.), BGP Route Reflector (RR), MPLS core, Provider (P) or Provider Edge (PE) routers
- Data Centre: spine, leaf, ToR
- Internet Exchange Point (IXP): switch or route server

The list of currently available features is listed on the RARE project portal [RARE] together with the documentation, step-by-step guides, and a blog [RARE_Blog]. The documentation also includes control plane documentation [FreeRtr_Doc] and code [RARE_Code]. RARE has been presented at many events such as at IETF [IETF110], RIPE [RIPE82], UKNOF [UKNOF47], 2021 P4 Workshop [P4_WS], RNP's national conference, and a proposal for presentations was accepted for JRES RENATER's national conference.

To support the testing and development of RARE, and to provide a platform for others to gain experience with the platform, the GÉANT P4 Lab (GP4L) [Lab_Manual], which comprises a core of four RARE P4 switches (Edgecore Wedge) powered by Intel Tofino ASICs, has been deployed at four GÉANT PoPs. The Lab has proven to be successful, with RARE and GP4L now attracting strong interest from the GÉANT community (such as HEAnet and Jisc) and, increasingly, the worldwide R&E community as well. As a result, non-European sites (STARLIGHT, RNP) are now connected and additional sites CALTECH, KREONET are considering connecting (see Figure 4.1).

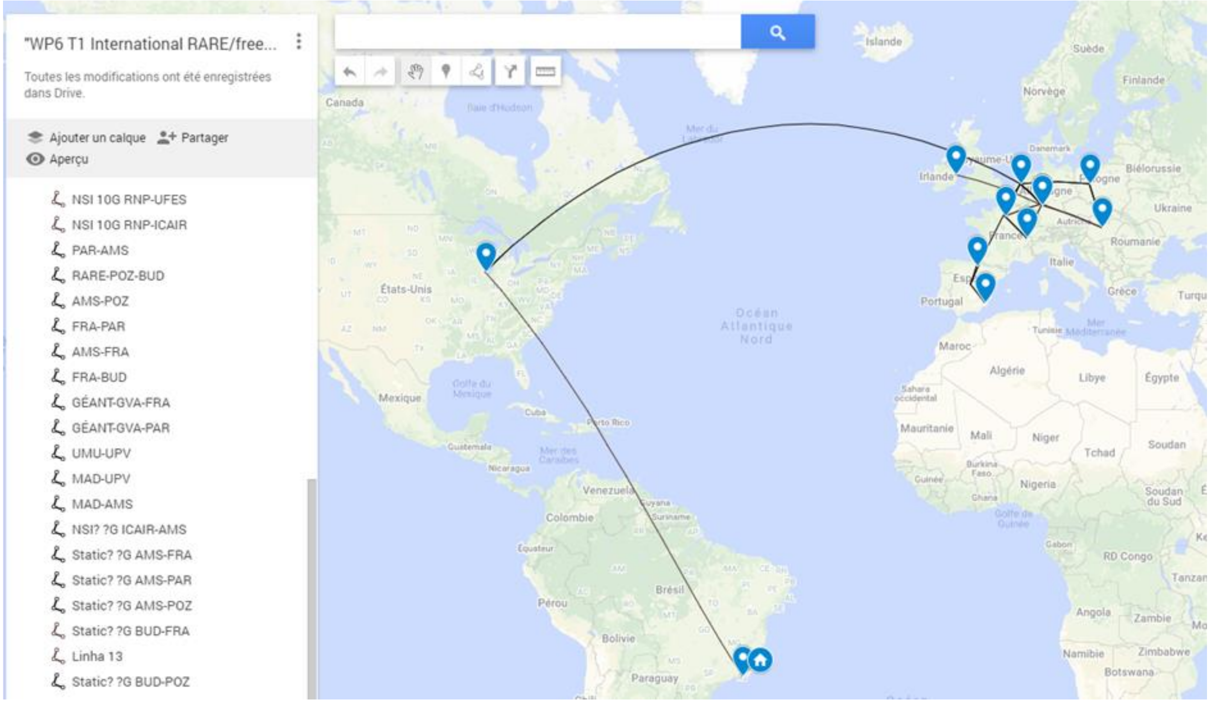

Figure 4.1: The GÉANT P4 lab with its international connections

To manage the GÉANT P4 Lab, the team is using NMaaS, the Network Management as a Service platform provided by the GÉANT project [NMaaS-info], [NMaaS-portal]. NMaaS provides a portfolio of network management and monitoring applications on a per-user secured network monitoring infrastructure. Prometheus is used as a monitoring system, Booked is used to schedule the use of the lab, and the collaboration on documentation is done using CodiMD, all through the NMaaS instance created for RARE.

The GÉANT/RARE project team actively contributed to the PolKA project's success. The project is led by Brazilian researchers from the Federal University of Espirito Santo (Ufes), who used the GÉANT P4 Lab to implement and verify the PolKA packet forwarding paradigm [IEEE_Polka]. This work received a Google Research Scholar Program award (Networking section at the Google Award Winning Recipient page [Google_Award], Ufes GitHub page[Ufes_Github]). Since the award, support for PolKA packet forwarding has been added to RARE/FreeRtr.

After a presentation at the IETF110 meeting [IETF110], the RARE team has collaborated with engineers from Juniper, Cisco, and Akamai on testing the FreeRtr implementation of the BIER multicast protocol, and implemented Automated Multicast Tunnelling (AMT, defined in RFC 7450) on FreeRtr to support relaying multicast content to unicast-only receivers. It was tested by transmitting video content over the GÉANT P4 lab, and development and testing are still in progress.

RARE/FreeRtr has reached a state of maturity that allows it to be considered for use in production environments. One example is that RARE/FreeRtr is in production for the GAIA networking laboratory at the University of Murcia, as explained in Section 2.3. With production use cases in mind, the RARE team is working on several capabilities that would facilitate a simpler distribution of RARE/FreeRtr to the world, including an ONIE-based image release development process for the WEDGE100BF-32X, enhancing the Continuous Integration/Continuous Delivery (CI/CD) release with automated testing suites, documentation, and dissemination, as well as continuing and strengthening collaboration with standardisation bodies, industry, and research initiatives such as those with the GNA-G Data Intensive Science (DIS) WG [GNA_G_DIS], Trinity College Dublin (via HEAnet), and UFES (PolKA).

Both RARE and its associated GÉANT P4 Lab have been very successful in attracting interest from both the GÉANT and worldwide R&E communities, as well as industry representatives at events like the IETF. A number of future areas of work are foreseen, including supporting more types of hardware and NOSs, as well as new features or elements. It would be desirable to add additional automated testing suites for the community to use, to support CI/CD based on FreeRtr's existing fully automated test suite (which ensures that no release induces code regression), and to consider the types of performance tests described in Section 2.3, e.g., to include control and data plane tests for specific NOS/hardware combinations.

Another direction for RARE/FreeRtr development might be on the Peering Policy Tools suite to leverage FreeRtr BGP Monitoring Protocol (RFC 7854) capability and create a global BGP watch guard, exploit active data sources (IPFIX, Netflow, INT etc) to elaborate a list of significant reports useful on a daily basis for peering management activity, implement simple peering management primitives automation, etc.

The additional features to be implemented for RARE will be selected based on inputs from the existing and potential RARE user communities. This was done, for example, with the addition of AMT support in response to the collaboration with the IETF mboned WG [MBONED], or the addition of PolKA packet forwarding support in collaboration with the Ufes researchers.

5 Low Latency Networking

The Low Latency networking subtask aims to provide a measurement infrastructure to support the monitoring and troubleshooting of real-time networked applications, in particular the LoLa audio-visual streaming application used by performing arts communities [LoLa].

The initial work explored various options for measuring and monitoring latency and jitter, which led to more specific tests using both Juniper Real-time Performance Monitoring (RPM) probes and TWAMP measurements (as defined in RFC5357, and supported by many vendors and open-source packages such as perfSONAR).

This effort led to the design, development and implementation of an open-source platform called TimeMap which provides per-segment latency/jitter measurements on the GÉANT backbone routers using RPM measurements. An interactive visualisation of the gathered measurement data is available through the TimeMap dashboard. The system is implemented using Telegraf, InfluxDB, and Grafana.

TimeMap was created with two specific objectives: to use measurement tools included in the existing equipment to remove the need to deploy extra hardware (i.e., a zero-footprint monitoring solution), and to make the service as easily deployable inside an NREN's backbone as possible. To achieve these objectives, TimeMap is based on micro-services that can readily be adapted for deployment in backbone networks of any NREN interested in having access to detailed latency and jitter measurements and visualisations.

TimeMap is available online [TimeMap], where it is accessible via eduGAIN federated login [eduGAIN]. Figure 5.1 presents the TimeMap landing page, where the topology of the GÉANT backbone can be seen and explored, and Figure 5.2 shows an example of the available measurements presented through Grafana.

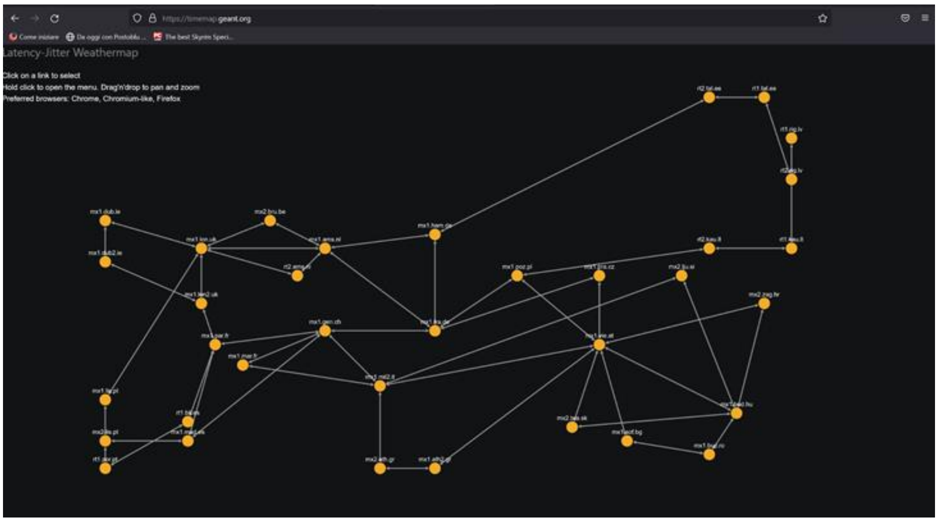

Figure 5.1: TimeMap landing page

Figure 5.2: TimeMap Dashboard with RPM measurements between routers

The tool is designed to help debug issues for real-time services which need a low latency and low jitter network transport, as it can identify activities on the backbone which affect such parameters.

TimeMap helped to discover some latency/jitter spikes in the measurements performed on the GÉANT backbone, and the TimeMap team is working together with GÉANT Operations, in collaboration with Juniper, to try to fix this issue.

Following an assessment of the tool, the GÉANT Operations team decided to adopt it for production use and is working with WP6 towards this goal. To prepare the deployment, the documented code was published with a first version of a user guide and an admin guide on the GÉANT GitLab [TimeMap_GitLab]. The team is also following up with several NRENs interested in deploying the TimeMap inside their own backbones.

Further work might include the investigation and prototyping of other active monitoring tools that are also available in the network, and/or in the existing equipment, such as TWAMP. It could be explored if TimeMap can be an addition or a replacement of an existing solution, either to a hop-by-hop and/or for end-to-end measurements. In addition, the use of an automated alert system/tool could be explored as it could help NOC operators to react faster in cases where latency and/or jitter issues arise. Other potential further work might include exploring inter-domain views of latency and jitter measurements across multiple NREN (and GÉANT) backbones, allowing the service to facilitate wider area debugging and troubleshooting in support of optimising the performance of real-time applications for which latency and jitter are key considerations.

6 Optical Time and Frequency Networking (OTFN)

Time and frequency services are important in a wide range of application areas, such as precision spectroscopy, remote clock comparisons in fundamental metrology, relativistic geodesy, or synchronisation in large-scale facilities

OTFN performance indicators are defined by accuracy (closeness of agreement between a measured quantity value and a true quantity value of a measurand), stability (a property of a measuring instrument, whereby its metrological properties remain constant in time), and uncertainty (a non-negative parameter characterising the dispersion of the quantity values being attributed to a measurand, based on the information used).

The application areas mentioned above require an ultra-precise time based on ultrastable dissemination of frequency signals (10-18 relative frequency uncertainty).

The two-way time and frequency transfer (TWTFT) or global navigation satellite systems (GNSS) achieve fractional frequency stabilities at the 10-16 level with measurement times of a few days. To overcome this limitation, several optical fibre networks have been implemented for the dissemination of ultrastable and accurate optical frequencies, spanning thousands of kilometres and providing transfer stabilities of a few 10-15 at 1 second with ultimate accuracies beyond 10-19.

To share knowledge around time and frequency technologies within the research and education community, the OTFN group published an initial white paper [TF_NREN] that introduces the basic concepts (accuracy, stability, uncertainty), the quality indicator (Allan variance and Allan deviation), and the basic elements of time and frequency transmission systems. A comparison of the existing solutions according to the different spectral regions (S-band, C-band, L-band) is presented, as well as the pros and cons of the dark channel and dark fibre solutions. The white paper also described four T&F national networks:

- the REFIMEVE+ from RENATER [REFIMEVE]
- the Swiss metrology network from SWITCH [TF_1572] [TF_LBand]
- the Czech metrology network from CESNET [CITAF]
- the NLPQT and PIONIER-LAB networks (derived from the OPTIME network) from PSNC [NLPQT], [PIONIER_LAB], [OPTIME]

The WP6 OTFN group also organised and delivered the GÉANT infoshare *European Time and Frequency Services - Principles, Challenges and Use Cases* [OTFN_Info].

The OTFN team cooperates with European projects like TIFOON [TIFOON] and CLONETS-DS [CLONETS_DS] that deal with time and frequency issues. It has also been seeking opportunities to assist NRENs in establishing T&F service interconnects, for example, RENATER and GARR have connected French and Italian NMIs in Grenoble. In addition, the WP6 OTFN group is aiming to establish an international test link between partners from France, Poland, and the Czech Republic. The connection between Poland and France would have its connecting point at CERN in Geneva, and between the Czech Republic and Poland in Cieszyn. However, due to the ongoing pandemic, this work has been delayed.

The OTFN group has also supported the GÉANT team in performing tests on providing time and frequency services in the L-band, using their newly acquired Infinera equipment, particularly by assisting with analysing and solving test issues. The use of L-band in this case results from the choice of equipment that is used in the new GN4-3N network rollout.

NRENs will need to determine how to deploy non-IP services such as T&F alongside existing data services. A spectrum-sharing solution is more cost-effective than using separate dark fibres for data and T&F services. Such an approach is implemented in France around Channel 44 of the C Band [REFIMEVE]. However, network operators and engineers may have concerns about placing a fixed alien channel inside the C-band, given the perceived potential for interference between the signals. Therefore, it is important for NRENs to gain experience with spectrum sharing for such scenarios, to understand the practical deployment considerations.

SWITCH, in collaboration with the National Metrology Institute (NMI) METAS, INRIM, and CESNET, has designed such a spectrum sharing solution between the frequency signal and traditional data network by exploiting dense wavelength division multiplexing (DWDM). A stabilised frequency-metrology network in Switzerland, spanning over 456 km of optical fibres, and operating in the L-band (1565nm to 1625nm), in ITU-T CH07 (λ=1572.06nm, f=190.7THz) was deployed between the METAS and two dedicated research institutes at the University of Basel and ETH Zurich. This work demonstrated that a metrological signal can be implemented on L-band, sharing the same fibre with data traffic in the C-band (being used by SWITCH NREN users), without any measurable disturbance on the data transmission. The metrological signal reached the same performance as systems operating in the C-band. It should be noted that this solution currently has the drawback of limited availability of L-band components, introducing longer lead times and higher manufacturing costs for early implementers. More detail about this solution from the NREN point of view has been provided in a white paper [TF_LBand].

The OTFN team has started work on the management and monitoring of T&F services in a production environment, to be summarised in a white paper. Some additional future work might include a feasibility study on T&F services on L-band using Raman pumps on long distance links, which might be useful for GÉANT who use this type of pump and are investigating a T&F services implementation on L-band on their existing Infinera equipment.

Considering the interest in the community, as well as the relatively slowly progressing service footprint, dissemination and knowledge sharing should continue to be an important focus for the team, as could collaboration with other projects and organisations regarding the implementation of T&F services.

7 Quantum Key Distribution

The quantum technologies and the second quantum revolution carry many expectations but are still at a very early stage. The future quantum technologies programs, framework and standardisation initiatives are still being heavily discussed at national and international levels, with numerous partners and institutions involved, which creates a very dynamic environment and community.

Within quantum communication, Quantum Key Distribution (QKD) is seen as the first application that could be deployed in practice, although judging by the first studies undertaken, it appears that the current technology still has strong limitations, for instance in terms of distance between QKD nodes.

Taking into account both the need and importance of getting involved with quantum technologies as soon as possible, and the still relatively small number of NRENs that are already engaged in this area, the WP6 QKD team has initiated a number of activities to enhance collaboration within the community towards faster onboarding and early deployment of quantum technologies. These activities have included some experimental technical work, knowledge sharing, and community engagement.

Although practical quantum solutions are still being developed, with very few vendors offering production-grade solutions, the QKD team was looking for opportunities to demonstrate the technology in use cases that are relevant to the R&E community. To this end, collaboration was established with several vendors as well as different research projects. One such initiative resulted in an agreed collaboration between the WP6 QKD team, GÉANT, Toshiba, and the OpenQKD project, in order to demonstrate the use of next-generation QKD Twin Field technology [TwinField] between two relatively distant PoPs on the GÉANT network, and to demonstrate using operational infrastructure and services. At the time of writing, this proof of concept is in its exploration phase, where fibre paths and equipment placement options are determined and ensured, with an initial project timeline agreed among partners.

The second thread of the technical work is related to quantum simulators. Considering that quantum solutions are not yet widely available (due to their price, limited number of vendors and maturity, and availability of their solutions, the applicability of use cases and overall readiness of the existing networks and operational teams), the WP6 QKD team determined that testing simulators and sharing knowledge about their potential use would help to bring quantum technologies and their usage closer to the community. Several simulators were considered (some of them were mentioned in the *Quantum Technologies Status Overview* white paper [QUANTUM WP]), and software QuISP (Quantum Internet Simulation Package) [QUISP] and QKDNetSim (Quantum Key Distribution Network Simulation Module) [QKD2020] were chosen for initial testing by the QKD team.

PSNC and CESNET (both NRENs are members of the WP6 QKD team) have progressed further with their own deployments of physical testbeds for quantum technologies. These physical deployments were used to establish a first physical quantum connection between Cieszyn (Poland) and Ostrava (Czech Republic). Figure 7.1 presents the testbed in PSNC. Gaining practical experience of physical QKD testbeds has been important for the GN4-3 WP6 team.

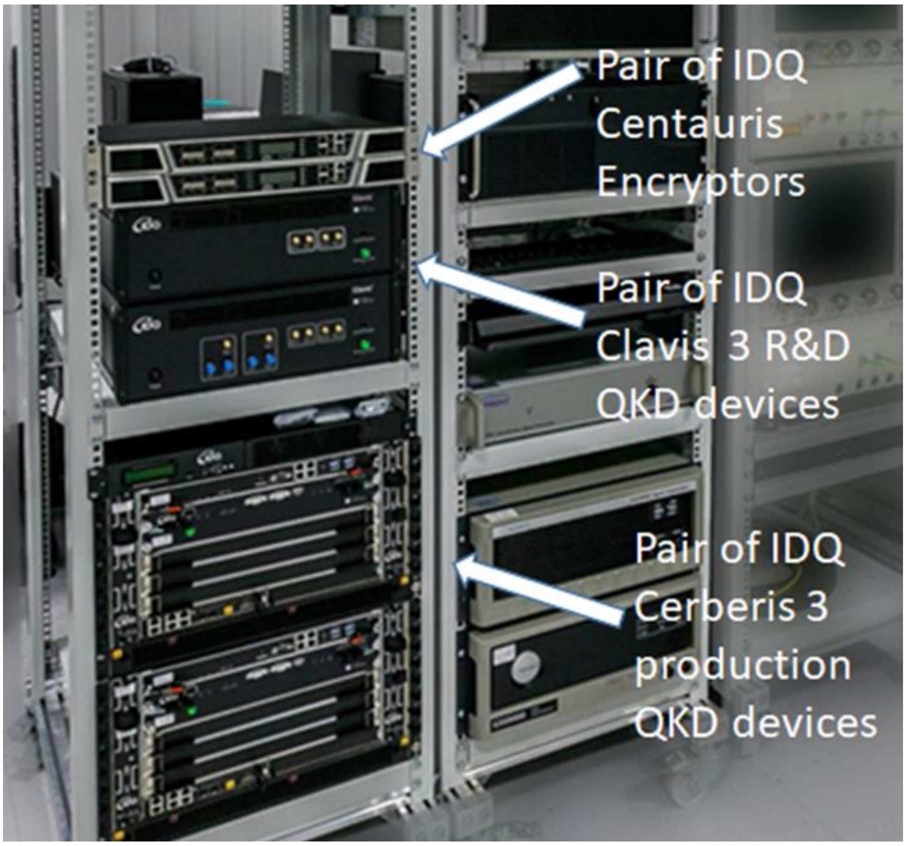

Figure 7.1: PSNC QKD testbed

Considering that NREN network engineers are not trained in these technologies, the QKD group looked for ways the GÉANT and NREN communities could be involved in the development and implementation of quantum technologies, especially quantum computing and quantum communication.

A white paper [QUANTUM_WP] was published that explained quantum key distribution, related technologies, and projects, and gave a general state-of-the-art overview. This provided a base for two dedicated QKD Technology infoshares *Quantum Technologies - Principles, Challenges and Applications* [QInfo1] and *Quantum Key Distribution - Practical Implementations, Challenges, R&E Use Cases and Standardisation outlook* [QInfo2].

Two further infoshares then focused on Quantum simulators [Qinfo3] and physical testbeds [Qinfo4]. These infoshares provided on a wide view on QKD: both technical and practical approaches, and ongoing European Commission activities. In addition, the TNC21 conference included a special session devoted to QKD technology progress and WP6 activities.

To complement infoshares and conference sessions whose focus is on uni-directional information sharing, the QKD team has started to organise monthly coordination calls, open for anyone who is interested in quantum technologies, regardless of the status of their involvement or deployment. The meetings are attended by representatives of NRENs, universities, and research institutes. Participants are given an opportunity to present their work, questions, or topics of interest, all with the aim to further collaboration towards faster development and deployment of quantum technologies in the GÉANT community.

The team will also continue to progress the planned proof of concept to be run between two PoPs on the GÉANT backbone.

Potential for a joint activity is being explored between the QKD group and Optical Time and Frequency Networking group. Quantum communication, networking, and QKD technologies require and will benefit from precise time and frequency synchronisation signals, both in research and real-life scenarios, and it is also important to consider the possibilities for both services to co-exist within the same fibre infrastructure.

As the Quantum Technologies programs and activities are very dynamic, it is hard to predict accurate progress and areas where the GÉANT and NREN contribution will be crucial. Further development is needed (such as quantum repeaters) to be able to fully implement and test the quantum networking schemes. QKD technology also requires substantial work in the areas of standardisation and certification, to which, as this is connected with real network implementation and applications, the GÉANT and NREN community can make important contributions.

8 Conclusions

The *Network Technology Evolution* task of the *Network Technologies and Services Development* work package in the GN4-3 project is exploring several technologies, and their relevance and applicability in the European NREN community, specifically: white box, DPP, RARE, Lola, OTFN and QKD.

This document has presented a summary of the work in these areas completed to date in the project. The various strands of work are at different levels of maturity, with different focuses, from software development for a new latency and jitter measurement and visualisation platform that can be used now, to knowledge sharing on a key technology for the future - QKD.

The **White box for R&E** activity is now completed. It has demonstrated that white boxes can be used for a number of NREN use cases in a production context. The hardware and NOS disaggregation allows NRENs to manage their network device estate in a much more flexible way, thus achieving a greater level of independence from traditional equipment vendors. This activity has shown that a feature and performance validation platform for network devices would be very useful for the European NRENs, and that the final decision on WB adoption depends not only on technical but also economic and strategic criteria.

DPP explored use cases for data plane programmability in the R&E community, and implemented a network monitoring tool using In-band Network Telemetry (INT). When adding an INT header to each packet transmitted, this approach allows a very detailed, fine grained view of all events in the network. This paves the way for new monitoring capabilities that could help NRENs to better understand their traffic, and diagnose performance issues that were previously not possible or very challenging to detect. Future work may focus on scaling the technology to high-capacity networks, where a significant amount of data has to be stored and analysed and presented to end-users.

RARE was successful in implementing an open-source control plane over five open-source data planes. RARE/FreeRtr offers a very rich set of features and has demonstrated that it could be used for a wide range of specific use cases including CPE, route reflector, LSR P and PE routers, data centre spine and leave, or Internet exchange point route server. There is now a worldwide interest in RARE/FreeRtr (North and South America, Asia and Europe). The associated GÉANT P4 Lab, initially created with a core set of devices at four GÉANT PoPs now has global reach, further extending the possibilities for collaboration.

The **low latency networking** team has developed TimeMap, a segment-by-segment latency and jitter measurement and visualisation solution for the GÉANT backbone, which allows network operations teams to monitor network performance in support of real-time applications such as LoLa. This tool is currently under deployment by the GÉANT Operations team, and discussions are ongoing with a number of NRENs about deploying such a solution in their backbones.

After looking at existing **OTFN** solutions in European NRENs, the OTFN team has started to explore interconnecting national T/F services, as well as knowledge dissemination via white papers and infoshares. The interconnection of the NMIs is ongoing in several directions, including, PSNC's connections with CESNET and RENATER. The OTFN team has assisted GÉANT in tests on using L-band (as supported by the new GN4-3N optical equipment)- for T&F services on the GÉANT network. The work also includes collaboration with other T&F projects, such as CLONETS-DS [CLONETS_DS].

The QKD team performed an initial assessment of the available QKD hardware and software on the market, together with an NREN survey, and published a quantum technologies status white paper, which also highlighted existing projects. With QKD in its relative infancy, the focus has been on knowledge exchange and dissemination, through many infoshares and coordination meetings. While a QKD proof of concept is planned between two GÉANT PoPs in collaboration with Toshiba and OpenQKD, work also progresses in testing quantum simulators, and sharing knowledge on physical testbeds such as those in PSNC and CESNET.

Overall, a wide range of network technologies has been evaluated within GN4-3 WP6 in the Network technology Evolution Task (WP6T1). The evaluations have taken a variety of forms. The work to develop TimeMap began with an assessment of the requirements for network measurements in support of more demanding real-time applications such as LoLa, and led to the design and implementation of a system supporting per-segment monitoring. It demonstrated the value of strong collaboration between those working with the applications, those operating the networks, and those with expert knowledge of specific protocols such as RPM and TWAMP. Through this close collaboration, a valuable platform has been developed, which is being adopted in production by GÉANT Operations.

RARE has shown that the innovative combination of an open-source routing platform with a programmable P4 data plane provides a powerful, feature rich, and performant platform that can serve a wide range of use cases. Seeding this innovation in the GÉANT P4 Lab, where RARE devices are hosted at four GÉANT PoPs, has helped spread interest in the solution around the world. Again, collaboration between researchers (such as PolKA), network operators and open-source router developers has proven valuable in producing a platform with strong potential for wider adoption. Similarly, the work on INT within the DPP team has, through a testbed spanning three NRENs, shown the potential for innovative ways to gain new insights into network behaviour that were not previously possible, or certainly challenging to achieve.

Other areas of work have been more focused on knowledge sharing within the community. Both the OTFN and QKD teams have been very active in dissemination, writing multiple white papers showing current state of the art in the R&E community, holding infoshares, and promoting discussion at a variety of events. Such intense collaboration, even at an early stage of future service development, is invaluable for the GÉANT community. In a similar way, the white box evaluations, undertaken for a range of use cases, and supported by separate analyses of total cost of ownership and performance testing, are helping NRENs and the wider community to understand the applicability of the technology to their own R&E use cases.

Whether the collaborative activities of WP6T1 are developing specific tangible software or network solutions, or raising the bar in the level of knowledge shared by and within the community, such evaluations continue to show they are a key function for GÉANT, the NRENs and the communities they serve.

References

[CAIDA] *Trends in Wide Area IP Traffic Patterns*, Sean McCreary and KC Claffy, CAIDA, San Diego Supercomputer Center, University of California, San Diego

[CITAF] https://citaf.org/en/index

[CLONETS_DS] https://clonets-ds.eu/

[CMS] https://docs.nvidia.com/networking-ethernet-software/cumulus-linux-44/Whats-New/

[CPE] https://www.geant.org/Resources/Documents/GN4-3_White-Paper_CPE_Router.pdf

[CPE_Funet] https://www.geant.org/Resources/Documents/GN4-3_White-Paper_White_Box_Funet-CPE.pdf

[D6.3] https://www.geant.org/Projects/GEANT_Project_GN4-3/GN43_deliverables/D6-3_White-Box-Evaluation.pdf#search=d6%2E3

[D6.5] https://www.geant.org/Projects/GEANT_Project_GN4-3/GN43_deliverables/D6.5_Network_Technology_Evolution_Report.pdf#search=d6%2E5

[eduGAIN] https://edugain.org/

[DataCentre] https://www.geant.org/Resources/Documents/GN4-3_White-Paper_White-Box-GRNET-Data-Centre-Use-Case.pdf

[DPDK] https://www.dpdk.org

[DPP_Wiki] https://wiki.geant.org/display/NETDEV/DPP

[DPP_DDOS] D. Ding, M. Savi, F. Pederzolli, M. Campanella and D. Siracusa, "*In-Network Volumetric DDoS Victim Identification Using Programmable Commodity Switches*," in IEEE Transactions on Network and Service Management, vol. 18, no. 2, pp. 1191-1202, June 2021, doi: 10.1109/TNSM.2021.3073597

[DPP_Github] https://github.com/GEANT-DataPlaneProgramming

[DTN] https://www.geant.org/Resources/Documents/GN4-3_White-Paper_DTN.pdf

[DTN_Wiki] https://wiki.geant.org/display/NETDEV/DTN

[E_AS7315-27X] https://www.edge-core.com/_upload/images/CSR310_AS7315-27X-DS-R05-20210906.pdf

[E_AS7315-27X_DCSG] https://www.opencompute.org/documents/edgecoreas7315-27xocp-pdf

[FreeRtr_Doc] https://docs.freertr.net/

[FRR] https://frrouting.org/

[GIX] https://www.geant.org/Resources/Documents/GN4-3_White-Paper_White_Box_GIX.pdf

[GNA_G_DIS] https://www.gna-g.net/join-working-group/data-intensive-science/

[Google_Award] https://research.google/outreach/research-scholar-program/recipients/

[IEEE_Polka] https://ieeexplore.ieee.org/document/9165501

[IETF110] https://datatracker.ietf.org/meeting/110/materials/slides-110-bier-01-bier-implementation-01

[INT_WP] https://www.geant.org/Resources/Documents/GN4-3_White-Paper_In-Band-Network-Telemetry.pdf

[Intel_Tofino] https://www.intel.com/content/www/us/en/products/network-io/programmableethernet-switch/tofino-series/tofino.html

[INT_Info] https://events.geant.org/event/558/
https://www.youtube.com/watch?v=_1W_psrCXLM&t=3s

[INT_WIKI] https://wiki.geant.org/display/NETDEV/INT

[INTspec] https://github.com/p4lang/p4-applications/tree/master/docs

[Lab_Manual] https://wiki.geant.org/display/RARE/Laboratory+Manual

[LoLa] https://conts.it/art/lola-project/old-lola-project-web-site/lola-low-latency-audio-visual-streaming-system

[M6.10] https://www.geant.org/Resources/Documents/M6-10_Evaluation-of-Planned-Network-Technology.pdf

[MBONED] https://datatracker.ietf.org/wg/mboned/charter/

[NLPQT] http://nlpqt.fuw.edu.pl/en/the-national-system-for-generation-and-distribution-of-reference-optical-carrier/

[NMaaS-info] https://www.geant.org/Services/Connectivity_and_network/NMaaS/

[NMaaS-portal] https://nmaas.eu/welcome/login

[Normandy_CPE] https://www.geant.org/Resources/Documents/GN4-3_White-Paper_CPE_Router.pdf

[OPTIME] http://optime.org.pl/

[OTFN_Info] https://events.geant.org/event/451/

[P4_WS] https://opennetworking.org/wp-content/uploads/2021/05/2021-P4-WS-Frederic-Loui-Slides.pdf

[PIONIER_LAB] https://pionier-lab.pionier.net.pl/

[Plotly] https://plotly.com/python/

[PSNCPlotly] https://github.com/GEANT-DataPlaneProgramming/int-analytics/tree/master/utils

[PROX] https://www.proxmox.com/en/

[QInfo1] https://events.geant.org/event/353/

[Qinfo2] https://events.geant.org/event/453/

[Qinfo3] https://events.geant.org/event/991/

[QInfo4] https://events.geant.org/event/1006/

[QKD2020] https://www.qkdnetsim.info/

[QUANTUM_WP] https://www.geant.org/Resources/Documents/GN4-3_White-Paper_Quantum-Technologies-Status-Overview.pdf

[QUISP] https://aqua.sfc.wide.ad.jp/quisp_website/

[RARE] https://wiki.geant.org/display/RARE/Home

[RARE_Blog] https://wiki.geant.org/pages/viewrecentblogposts.action?key=RARE

[RARE_Code] https://bitbucket.software.geant.org/projects/RARE/repos/rare/browse

[REFIMEVE] http://www.refimeve.fr/index.php/en/

[RIPE82] https://ripe82.ripe.net/wp-content/uploads/presentations/29-rare-freertr-ripe82.pdf

[RN] https://www.normandie.fr/

[STC] https://spirent.com/

[SYVIK]	https://www.criann.fr/reseau-syvik/
[TCO]	https://www.geant.org/Resources/Documents/GN4-3_White-Paper_White-Box-TCO.pdf
[TCO-xls]	https://www.geant.org/Resources/Documents/TCO-Calculator.xlsx?web=1
[TF_1572]	https://www.osapublishing.org/oe/fulltext.cfm?uri=oe-29-16-24592&id=453418
[TF_Lband]	https://www.geant.org/Resources/Documents/GN4-3_White-Paper_Ultrastable-Frequency-Transfer-in%20L-Band.pdf
[TF_NREN]	https://www.geant.org/Resources/Documents/GN4-3_White-Paper_Time_and_Frequency.pdf
[TIFOON]	http://empir.npl.co.uk/tifoon/
[TimeMap]	https://timemap.geant.org/ (Accessible to eduGAIN users)
[TimeMap_GitLab]	https://gitlab.geant.org/gn4-3-wp6-t1-lola/timemap_public
[TIP_DCSG]	https://cdn.brandfolder.io/D8DI15S7/as/q6nbao-b8vfa0-fvfoyt/DCSG_Technical_Specification_-_Telecom_Infra_Project.pdf
[TwinField]	https://www.nature.com/articles/s41566-021-00811-0
[Ufes_Github]	https://github.com/nerds-ufes
[UKNOF47]	https://indico.uknof.org.uk/event/53/timetable/?view=standard_inline_minutes
[VEP]	https://www.dell.com/en-us/work/shop/povw/virtual-edge-platform-4600
[VMW]	https://www.vmware.com/cloud-solutions.html
[WB_Perf]	https://www.geant.org/Resources/Documents/GN4-3_White-Paper_White-Box-Testing-and-Evaluation.pdf

Glossary

ACL	Access Control List
AMT	Automated Multicast Tunnelling
ASIC	Application-Specific Integrated Circuit
BFD	Bidirectional Forwarding Detection
BGP	Border Gateway Protocol
CapEx	Capital Expenditure
CD	Continuous Delivery
CI	Continuous Integration
CPE	Customer Premises Equipment
CPU	Central processing Unit
DCSG	Disaggregated Cell Site Gateway
DHCP	Dynamic Host Configuration Protocol
DIS	Data Intensive Science
DPDK	Data Plane Development Kit
DPP	Data Plane Programming
DTN	Data Transfer Node Infrastructure
DUT	Device Under Test
E-Line	Ethernet Line
ECMP	Equal-Cost Multi-Path
EVPN	Virtual Extensible LAN
FRR	Free Range Routing
Gbps	Gigabytes per second
GIX	Global Internet eXchange point
GNSS	Global Navigation Satellite System
GPL	Global Price List
HRSP	Hot Standby Router Protocol
IGP	Interior Gateway Protocol
INT	In-band Network Telemetry
IOS	Internetwork Operating System
IP	Internet Protocol
IPSec	Internet Protocol Security
IXP	Internet Exchange Point
L	Layer
LACP	Link Aggregation Control Protocol
LAN	Local Area Network
LER	Label Edge Router
LLDP	Link Layer Discovery Protocol
LoLa	LOw LAtency
LSP	Language Server Protoco

LSR	Label Switching Router
Mbps	Megabits per second
MC-LAG	Multi-Chassis Link Aggregation Group
MPLS	Multi Protocol Label Switching
MSTP	Multiple Spanning Tree Protocol
NAT	Network Address Translation
NMaaS	Network Management as a Service
NMI	National Metrology Institute
NOC	Network Operation Centre
NOS	Network Operating System
NREN	National Research and Education Networks
OcNOS	Open Compute Network Operating System
ONIE	Open Network Install Environment
OOB	Out-of-Band
OpEx	Opeartional Expenditure
OTFN	Optical Time Frequency Network
P	Provider
P4	Programming Protocol-independent Packet Processors
PE	Provider Edge
PolKA	Polynomial Key-based Architecture
PoP	Point of Presence
QKD	Quantum Key Distribution
R&E	Research and Education
RAM	Random-Access Memory
RARE	Router for Academia, Research and Education
RFC	Request for Comments
ROS	RARE Operating System
RPM	Real-time Performance Monitoring
RR	Route Reflector
RSVP	Resource Reservation Protocol
SRAM	Static Random Access Memory
SSH	Secure Shell
STC	Spirent Test Centre
T	Task
T&F	Time & Frequency
TCAM	Ternary Content Addressable Memory
TCO	Total Cost of Ownership
TE	Traffic Engineering
TIP	Telecom Infra Project
TWAMP	Two Wire Active Measurement Protocol
TWTFT	Two-Way Time and Frequency Transfer
UDP	User Datagram Protocol
VEP	Virtual Edge Platform
VLAN	Virtual LAN
VM	Virtual Machine
VNF	Virtual Network Function
VRF	Virtual Routing and Forwarding
VXLAN	Virtual Extensible LAN
WB	White Box
WP	Work Package