

30-03-2020

Deliverable D6.5

Network Technology Evolution Report

Deliverable D6.5

Contractual Date: 31-03-2020
Actual Date: 30-03-2020
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: GN4-3-20-325EED
Authors: Claudio Allochio (GARR), Mauro Campanella (GARR), Tim Chown (JISC), Ivana Golub (PSNC), Xavier Jeannin (RENATER), Frédéric Loui (RENATER), Jani Myyry (FUNET), Damian Parniewicz (PSNC), Lefteris Poulakakis (GRNET), Edin Salguero (RENATER), Nicolas Quintin (RENATER), Piotr Rydlichowski (PSNC) Marco Savi (GARR), Tomasz Szewczyk (PSNC), Maxime Wisslé (RENATER)

© GÉANT Association on behalf of the GN4-3 project.

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

The deliverable reports on the work on network technologies performed as a part of the Network Technology Evolution task in the Network Technologies and Services Development Work Package of the GÉANT 4-3 project. Technologies in the scope of the work included White box (WB), Data Plane Programming (DPP), Router for Academia, Research and Education (RARE), Data Transfer Node infrastructure (DTN), Low Latency transfer (LoLa), Optical Time and Frequency Networks (OTFN) and Quantum Key Distribution (QKD).

Table of Contents

Executive Summary	1
1 Introduction	3
2 White Box for Research and Education	5
2.1 CPE Normandy	5
2.2 FUNET CPE	6
2.3 GRNET Data Centre	6
2.4 Provider Router (P) / Label Switch Router (LSR)	6
2.5 Internet eXchange Point (IX)	7
2.6 White Box Performance Test	8
2.7 Future Work	8
3 Data Plane Programmability	9
3.1 Telemetry with Data Plane Programming	9
3.2 DDoS Detection	11
3.3 Future Work	12
4 Router for Academia Research and Education (RARE)	13
4.1 Achievements	13
4.2 Future Work	15
5 Data Transfer Node Infrastructures (DTNs)	17
5.1 Achievements	17
5.2 Future Work	18
6 Low Latency/Low Jitter Infrastructure	19
6.1 Achievements	20
6.2 Future Work	22
7 Optical Time and Frequency Networking (OTFN)	23
7.1 Achievements	23
7.2 Future Work	27
8 Quantum Key Distribution	30
8.1 Achievements	30
8.2 Future Work	31
9 Conclusions	32

References	34
Glossary	35

Table of Figures

Figure 2.1: Universal CPE running several network functions	5
Figure 2.2: PSNC interoperability testbed	7
Figure 3.1: INT demo, Jitter data per packet showed at 5 sec. time scales	10
Figure 3.2: INT demo, Jitter data per packet showed at 1 sec. time scales	10
Figure 3.3: Detection results of the testing with real traffic traces including five DDoS attacks	11
Figure 4.1: European P4 testbed	14
Figure 4.2: European P4 testbed management	15
Figure 6.1: Incomplete map of locations where low latency applications are in common use	20
Figure 6.2: LoLa monitoring analysis software architecture	22
Figure 7.1: Time/Frequency setup	25
Figure 7.2: Suggestion of T/F setup in in-line-amplifier site	26
Figure 7.3: Schematic of Optical Add and Drop Multiplexer used to insert or extract T/F signals	27
Figure 7.4: Fibre footprint for this interconnection	28

Table of Tables

Table 7.1: T/F existing techniques/services, performances and TRL	24
---	----

Executive Summary

This document reports on the evaluations of the applicability of a number of emerging technologies conducted by the *Network Technology Evolution* task (WP6 T1) within the *Network Technologies and Services Development* work package (WP6).

White boxes (WB) can be useful for some R&E network scenarios. Their properties of disaggregation and independence allow NRENs to manage their switch/router estate in a totally different manner. The first commercial deployments have started in cloud scenarios; WP6 T1 is considering many use cases and has implemented its first WB-based CPEs.

Data Plane Programming (DPP) focuses on new types of programmable switch/router devices. The work highlights the advantages and limitations of the P4 language and P4 chipset for DDoS and network monitoring use cases.

Router for Academia, Research and Education (RARE) was very successful in coupling an open source control plane with a P4 data plane on top of a P4 hardware switch. As a result, an almost complete LER/PE router (IPv4, IPv6, ISIS, MPLS, SR-MPLS, L3VPN, XConnect, VPLS, EVPN, 6VPE, etc.) was implemented. RARE is now focusing on implementing control plane policy and monitoring/telemetry support. This project continues the network innovation within R&E as it is relatively easy to integrate a new network protocol with a RARE solution.

The Data Transfer Node infrastructure (DTN) work intends to support end users in transferring large scale data across the network with good performance. Effective end-to-end data transfer services include technical aspects (i.e. choosing hardware and software solutions), and also non-technical considerations, which WP6 will address in a focus group.

Low Latency infrastructure (LoLa) aims to optimise performance of real-time applications avoiding high latency and high jitter routes. The LoLa project is working on designing a segment-by-segment low latency and low jitter weather map solution for the GÉANT network backbone. This solution would improve the operation of real-time applications by supporting latency and jitter metrics and increasing network transparency, paving the way for enabling more optimal routing of such traffic.

The Optical Time and Frequency Networking (OTFN) team is working towards interconnecting national T/F services. Apart from contributing to metrology science with the interconnection of the European National Metrology Institutes (NMIs), scientific experiments, civil and strategic applications might also benefit from T/F services and might see value in an official European time service.

The work on Quantum Key Distribution (QKD) in WP6 T1 started with a review of the quantum landscape and its importance to the GÉANT and NREN communities. QKD is an important step in the evolution of quantum communication systems that will be essential to interconnect distributed quantum computing systems.

1 Introduction

The *Network Technology Evolution* task (WP6 T1) within the *Network Technologies and Services Development* work package (WP6) evaluates the pertinence of a variety of emerging technologies and their usefulness and applicability for European National Research and Education Networks (NRENs). During the first project year the following technologies were evaluated:

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

Section 2 of this report evaluates White Box, which disrupts the economic model of the router market by disaggregating the hardware and the network operating system. The work reviews whether the GÉANT community is in the same situation now as when Linux appeared in the UNIX world, and whether white boxes provide a real opportunity for NRENs and research and education (R&E) networks.

Section 3 reports on DPP. The data plane is now readily programmable thanks to the P4 language. DPP investigates applications, like DDoS detection and In-band Telemetry (INT), that could be implemented in new ways and with enhanced capabilities on P4-compliant hardware (the work covers P4 switches and FPGA cards).

Section 4 evaluates RARE, which uses the same opportunity P4 provides to demonstrate that an open source control plane coupled with a P4 data plane, running on white box, can be used as a router for R&E use cases.

Section 5 looks at DTN infrastructures. The objective of the data transfer node (DTN) stimulation activity is to explore large scale data transfer infrastructures and services for the GÉANT community, analysing the missing pieces that slow down deployment of such infrastructures or reduce their effectiveness.

Section 6 reports on LoLa. The network path used by real-time applications might not be an optimal path in cases where the routes are determined based on bandwidth capacity rather than latency and jitter. The first missing piece is a monitoring tool that could provide information about the segment-by-segment latency and jitter. This measure would be the foundation for identifying, debugging and possibly improving the network path used by real-time applications. This is particularly important for

applications supporting real-time collaborative performances such as those run by the LoLa project [[LoLa](#)].

Section 7 explores OTFN. Time and Frequency (T/F) measurement is a foundation for science as many sciences require a very accurate time source, pushing technologies to their ultimate precision. T/F services are already in production in several NRENs. Since the types of signals differ from the usual data signals, T/F services and also Quantum Key Distribution are examples of emerging new usage of fibre infrastructure beyond traditional data services. The European National Metrology Institutes (NMIs) require the comparison of high precision time and frequency signals for their work. NRENs and GÉANT are looking at the possibility to transport such high precision signals over existing interconnected European infrastructure fibre links. Other European experiments would also like to access these very high accurate T/F services because of their very demanding requirements (such as for example Smart grid, 5G, Global Navigation Satellite System, and Hyper-trading). OTFN aims to deploy T/F services between European NRENs, and in doing so provide European science and society a more accurate time service.

Section 8 evaluates Quantum Key Distribution (QKD), which allows network-based communications to be secured through a step change in cryptographic capability, using quantum mechanics theory. However, comparing theory and reality, some early implementations in the real world have introduced flaws that currently make the system imperfect. Still, QKD is the most mature technology the quantum revolution has produced so far. From the state of the art in NRENs, WP6 T1 will explore current work and existing solutions, world-wide projects and possibly test technology in order to design QT solutions for GÉANT. WP6 T1 has joined the team that will work on quantum standardisation.

Section 9 offers conclusions and next steps for the WP6 T1 work on the explored technologies, based on the evaluations so far.

2 White Box for Research and Education

A white box (WB) is a switch or router manufactured from commodity components that allows different Network Operating Systems (NOSs) to be run on the same hardware, decoupling the NOS software from the hardware. White box technology has only very recently come to the fore commercially, and experience of its use in NREN environments is relatively limited.

This deliverable is an update to Deliverable D6.3 *White Box Evaluation* in which the objectives and first results were reported [D6.3]. The WP6 T1 approach is to investigate the feasibility of the use of WB for NRENs through use case investigation and performance analysis. The following sections provide a brief summary of the achievements for each use case.

2.1 CPE Normandy

Based on the FRRouting NOS and an x86 server (Dell VEP 4600) running a Linux instance in a virtualised environment, WP6 T1 has completed a solution that was successfully put into production in two Normandy high schools from October to November 2019. The solution provides an appropriate reliability and performance level. One of the most important benefits of this solution is the ability to implement additional Virtual Network Functions (VNFs) on the white box (see Figure 2.1). The first potential VNF could be a virtual firewall. However, the timeline for any new features will need to be agreed with the users, who at the moment are not considering such improvements.




Figure 2.1: Universal CPE running several network functions

2.2 FUNET CPE

The FUNET CPE was successfully tested in a VMware environment with a Cumulus VX NOS connected directly to the FUNET backbone. Client connectivity was simulated with a separate Ubuntu Linux virtual machine. The CPE features required by FUNET related to management (SSH, syslog, monitoring, DNS, NTP, SNMPv2, DHCP/DHCPv6), routing (Jumbo frame support, IPv4, IPv6, BGP IPv4 and IPv6 unicast, BGP route filtering, OSPFv2-v3, VRRPv3, VRF support) and security (interface access control lists IPv4 and v6) were validated.

The next step will be to validate the same results on a hardware machine to assess the performance and the hardware and software compatibility. Automation will also be part of the year 2 work plan (Ansible, NETCONF).

2.3 GRNET Data Centre

As a first step, the target configuration of the data centre was successfully validated with Cumulus NOS on a virtual lab (EVE-NG platform). The IP fabric based on EVPN/VXLAN includes BGP, using the eBGP auto-configuration feature, EVPN type 2 (IP/MAC) and type 5 (IP prefix) route advertisements, server multi-homing (configured using MLAG), inter-VLAN routing capability on the tunnel terminating devices, active-active links and fast failover.

GRNET published a request for a proposal for the procurement of six white box switches to be deployed in a small data centre. As result of this tender, a solution based on white box technology was selected; during Q2 of 2020, GRNET will receive four AS5812-54T (leaf) and two AS7712-32X (spine) Edgework switches that will run Cumulus.

The next step is to deploy the six switches and set up a production-grade data centre network infrastructure. The real-life deployment will allow re-evaluating the capabilities detailed above in production, especially any potential constraints due to the hardware/chipset not observed in the virtual lab evaluation. GRNET will explore Zero Touch Provisioning (ZTP) and the basic capabilities of the management plane (remote configuration modification, sanity/syntax checking, rollback).

2.4 Provider Router (P) / Label Switch Router (LSR)

For this use case, WP6 T1 chose OcNOS software, installing it on an Edgework white box. During the test a number of transit LSPs were successfully signalled with the RSVP protocol and thus established on the white box platform. The IS-IS with traffic engineering extension was tested as an IGP protocol. The tests were performed in the lab with network topology emulated by the Spirent TestCenter system. The tests demonstrated the ability of the white box platform to efficiently handle MPLS traffic. This proves that the control plane processes and data plane switching works properly on the tested platform.

The next important scope of the WB architecture validation process is testing interoperability with other vendors. For that purpose, either hardware or virtual platforms can be used, as the tests are focused more on the control plane functions than switching performance. PSNC built a lab, as shown in Figure 2.2, consisting of OcNOS WBs and some legacy vendor devices (for example, Juniper, Arista, etc.). In the virtual testbed based on GNS3 software [GNS3] a basic MPLS network was implemented and the testing of some basic and advanced MPLS functionalities has been started: RSVP signalled LSPs, LSP with predefined path (Traffic Engineering), LSP with primary and secondary paths, LSP with primary and secondary paths.

The control plane interoperability tests will continue during year 2 to verify if OcNOS can provide advanced MPLS functionalities. The interoperability test will focus on features currently implemented by NRENs such as traffic engineering, protection mechanisms, troubleshooting (MPLS ping/traceroute), and inter-domain signalling.




Figure 2.2: PSNC interoperability testbed

2.5 Internet eXchange Point (IX)

RENATER is working on replacing the current traditional vendor switches in their Internet eXchange point with two white box switches (Dell EMC S4048-ON) providing the same level of layer 2 services. After completing the investigation and mastering the NOS installation based on Open Network Install Environment (ONIE), W6 T1 selected OcNOS [OcNOS]. This NOS successfully completed a set of tests for the device management plane (NTP, SNMP, SSH, etc.) and for the control and data plane (STP, VLAN, ACL, etc).

The two white box switches were installed in two RENATER PoPs. A full migration plan [SFINX] was written and the two white boxes were fully configured and interconnected thanks to a dedicated 10 Gbps lambda.

In order to continue with the migration plan, some preparation work will need to be done in one of the racks in RENATER which will host the client interconnection using 100 Mbps copper Ethernet interfaces (for initial test purposes). Electrical installations will need to be improved in the data centre to provide sufficient electrical power in the rack. The white box solution must be able to hand over all the IX traffic at the very beginning of the migration, as the first step after interconnecting the white box to the IX; the STP algorithm has to be changed and this operation might be risky for the old hardware.

2.6 White Box Performance Test

Continuing the work reported in Deliverable D6.3 [D6.3], the WP6 T1 team ran performance tests of a white box switch and compared the results with the legacy vendor devices. The performance was tested with regard to the packet buffering capabilities. Spirent TestCenter was used to compare the buffering capabilities of the OcNOS/Edgework platform with PTX from Juniper and an Arista 7280 series switch. The result of the test was exactly as expected, demonstrating the influence of buffer size on switching performance and usage of the buffer: the machines equipped with deep buffers have the same behaviour regardless of whether the forwarding chipset is a commodity or traditional vendor one. The general conclusion of the tests is that, as for traditional routers, it is very important to analyse the white box platform hardware architecture in order to select the best equipment for a given network requirement. Two white boxes built on the same chipset may provide slightly different switching performance results.

2.7 Future Work

The 'white box for research and education' task has made good progress in its investigation and completed almost all of its initial objectives. The focus of the remaining period until the end of Year 2 (December 2020) will be to gather and analyse feedback from the production use cases. The lessons learned could be very useful for NREN network operations, and to assess if this new economic model fits well in the research and education context. In addition, the total cost of ownership (TCO) will be analysed. WP6 T1 will provide a guide to help NRENs identify the elements that need to be considered when estimating TCO, according to their use case(s) and their context. This analysis will also be driven by the use cases investigated by WP6 T1. Finally, additional performance tests will be conducted in order to complete the performance analysis.

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 both improve existing services and facilitate the design of new services. In both the use cases explored in this work, i.e. network telemetry and DDoS detection, the novel programming language P4 has been used on dedicated hardware: second generation FPGA, and switches with the Tofino chip. Such use cases require validation before being candidates for use in NREN and R&E scenarios.

The objectives were therefore to validate the new functionalities offered by data plane programming, to evaluate the complexity, challenges and the boundaries in programmability in data plane hardware and to explore the balance between the data computations conducted by the switch/router CPU in the data plane versus computation done on external nodes.

3.1 Telemetry with Data Plane Programming

One of the DPP team's first achievements was to provide a new version of the compiler that translates the P4 language to VHDL. The previous version of the compiler developed by CESNET supported only P4₁₄ and the lack of support for the latest P4₁₆ release was a severe limitation for the project. Therefore, CESNET has prepared a new P4₁₆ to VHDL language translator. In year 1, WP6 T1 finished the initial implementation, including real deployment on an FPGA card, with a demo.

The effort focused first on In-band Network Telemetry (INT) as a tool to gather information on traffic at the level of individual packets. The use case entailed the definition of telemetric data [INT], exchanged with the collector, and a specification of the basic architecture of INT-ready network probes. WP6 T1 uses the P4 language for its INT sink implementation, described in the In-band Network Telemetry standard [INT].

Telemetric data such as outgoing/destination timestamp and sequence number are inserted in the first INT nodes. The sink node then removes, collects and sends the telemetric data to the collector. From the collector, different applications can be used to plot graphs of observed metrics allowing the users to zoom in on an individual packet. Another big advantage of INT is the possibility to insert user-defined telemetric data in each network frame. For example, one-way delay, jitter and frame reordering were chosen for a first prototype that was demoed during the GÉANT Symposium 2020 in Ljubljana. The external analysis and data presentation limitations are an indication of the need to find a balance between the measurement precision, the breadth of information collected at the data plane level, and the external analysis engine. The demo showed that such a choice is now possible; it can be programmed and dynamically adjusted to the monitoring needs, almost in real time.

This prototype was based on a FPGA SmartNIC, a standard x86-64 server board, a P4₁₆ to VHDL compiler and a visualisation layer based on Grafana and InfluxDB. The demo presented the delay, jitter and reordering of locally generated traffic, with artificial delay inserted.




Figure 3.1: INT demo, Jitter data per packet showed at 5 sec. time scales




Figure 3.2: INT demo, Jitter data per packet showed at 1 sec. time scales

Figure 3.1 and Figure 3.2 show jitter measurements thanks to the P4 INT implementation over a FPGA card at very high precision (the X axis is time and the Y axis is the value of jitter). Figure 3.2 shows that an operator is able to zoom in with a sub-second level of detail.

3.2 DDoS Detection

In Year 1, WP6 T1 worked on volumetric DDoS attack detection and DDoS traffic monitoring directly in the data plane using the P4 language. WP6 T1 investigated the possibility of using sketches [Chaignon], probabilistic data structures that allow summarising packets and flow statistics. This approach was successfully implemented in an emulated switch using BMv2.

WP6 T1 then proceeded to port the sketch-based code to a carrier-grade P4-programmable switch (EdgecoreWedge100BF-32X) hosted in Fondazione Bruno Kessler (FBK). The porting process posed various technological challenges, and during the process the team discovered some implementation differences between the virtualised BMv2 environment and the actual physical setup of the switch.

The main current boundaries in hardware are the total amount of available memory (used to store information for the sketches) and the number of P4 pipeline stages in the chip (i.e. atomic actions on information within a packet) that are 12 for ingress and egress. The hard boundaries were overcome with the implementation of a more simplified strategy, which allows detecting those destination IP addresses that are ‘outliers’ statistically with respect to the number of flows with specific source IP address(es). The results show that with a sketch memory of only 400 KB, it is possible to immediately detect those outliers when they are targeted by a significantly higher number of IP sources than for average IP destinations. Figure 3.3 shows the detection results of the testing with real traffic traces when five DDoS attacks are injected into the regular traffic. All IP destinations targeted by more than 400 sources (and assumed to be under DDoS attack) are detected, with no false positives or false negatives.




Figure 3.3: Detection results of the testing with real traffic traces including five DDoS attacks

Additionally, preliminary tests on switch performance, in terms of packet processing time and throughput, showed that the execution of the P4 DDoS detection program in the chip does not impact packet processing time and is not a function of the packet size. This is ensured by the strictly imposed hardware boundaries that have forced the team to implement a simpler (yet effective) strategy, while safeguarding the device’s line rate packet processing performance.

The following major achievements have been completed:

- Proficiency level was reached with the P4 language, also via the participation in dedicated P4 training (provided by the GÉANT GLAD team), with the design and implementation in BMv2, and the publication of some network monitoring strategies [Ding_1][Ding_2] (not strictly related to DDoS detection).
- Successful implementation and testing of a strategy for volumetric DDoS attack detection in P4 using BMv2, which proved the feasibility of the approach (no hard memory boundaries).
- Successful implementation of DDoS attack traffic monitoring in P4 using BMv2 which provides more detailed information about a DDoS attack, such as total traffic volume, subnets that are the main sources of the attack, top source and destination UDP/TCP ports and the distribution of IP protocol numbers in the IP headers of DDoS traffic.
- Successful implementation and testing of a simplified DDoS detection strategy in a carrier-grade P4-programmable switch that takes into account realistic memory and hardware processing stages per packet limitations.

3.3 Future Work

Concerning telemetry, the focus of effort in Y2 (to December 2020) will be on improving the scalability of the existing INT prototype, and validating INT measurements in a mixed FPGA-Tofino environment. It will be tested in an environment that is closer to a production one using the testbed deployed by the RARE team.

All work shall be performed with frequent consultations and dialogue with other relevant activities in the project, with the GÉANT community, as well as with interested external NRENs (e.g. ESnet) and researchers (e.g. the National Technical University of Athens).

WP6 T1 also plans to investigate how to improve the telemetric collector so that it scales to be able to monitor specific traffic flows of up to 10 Gbps or more.

The current version of the P4₁₆ compiler supports the initial set of functionalities, which allows deployment of a basic INT functionality. The compiler and the metadata specification will be improved according to experimental needs.

Regarding the DDoS work, the Y2 plan is to complete testing on the limitation and performance of the algorithmic detection strategy. A sensitivity analysis will be performed related to the sketch size, accuracy and memory/stages ('switch resources') occupied. A feasibility study of a P4 mitigation program is planned, with external activation (i.e. triggered by the control plane or the detection strategy), ideally with fine-grained resolution of flows to minimise the drop of legitimate traffic.

The effort in Data Plane Programming demonstrated the capabilities of embedding P4 programs in the data plane. It has also demonstrated its use in the area of In-band Network Telemetry which is seen to be of interest to some other NRENs (for example, see the ESnet presentation [ESnet]). The results of this year's work, as well as the NREN community interest in use cases of data plane P4 programmability, will be used as input for possible future work.

4 Router for Academia Research and Education (RARE)

RARE is an effort that focuses on determining if a novel software routing platform based on coupling an open source control plane with a P4 data plane can fit with R&E use cases. The project aims to integrate different pieces of software related to these building blocks:

- **A control plane:** RARE uses FreerTr as the control plane component.
- **A data plane:** P4 is used to describe the RARE data plane.
- **A communication interface** between the control plane and data plane. This interface has to be compliant with the P4Runtime specification.

The key part of the work consists of enabling a control plane software to drive a data plane via a programmatic interface. To achieve this, the P4 language is used as an interface that facilitates data plane programmability.

The P4 core language has been designed to be as independent as possible from the target or Network Processing Unit (NPU) architecture. However, in practice architecture dependence is still prominent. Code adjustments followed by a target-specific compilation are necessary if a P4 program needs to be run on a specific architecture and achieve high performance forwarding.

4.1 Achievements

In Year 1, activities were focused on ‘technology scouting’ and building the structure for the RARE project. The following major achievements have been completed:

- The RARE group received training (thanks to the GÉANT GLAD team in WP1 T5) in the P4 language which is dedicated to driving the data plane.
- To ease future dissemination activities and training, P4 packages have been developed to quickly set up an operational environment:
 - On Ubuntu:
 - Nightly build: <https://launchpad.net/~frederic-loui/+archive/ubuntu/p4lang-stable-bionic-nightly>
 - Frozen working build: <https://launchpad.net/~frederic-loui/+archive/ubuntu/p4lang-master-bionic-snapshot-20191127>
 - On Debian buster:
 - <https://build.opensuse.org/project/show/home:frederic-loui:p4lang:p4c:master>

- Creation of a RARE work repository:
 - Public git repository using BMv2: <https://github.com/frederic-loui/RARE>
 - Private git repository using BAREFOOT bf_switchd on GÉANT git: <https://bitbucket.software.geant.org/projects/RARE/repos/rare/browse>
- A working LSR(P) and LER(PE) router platform that is almost ready (only control plane protection and telemetry are missing):
 - A communication interface has been implemented between the control plane and the data plane using a GRPC P4Runtime/BfRuntime specification.
 - Using BMv2 target (simple_switch_grpc).
 - Using BAREFOOT bf_switchd target powered by TOFINO NPU (WEDGE-BF100-32X).
 - Automated P4 tests have been implemented.
- A community documentation and learning web site has been published: <https://rare-freertr.mp.ls/>

At the beginning of Y2, the RARE team, with support from GÉANT Network Operations, started the deployment of a European P4 testbed. Four core P4 switches were deployed at GÉANT PoPs in Poznan, Amsterdam, Frankfurt and Budapest, as shown in Figure 4.1. The figure also shows NRENs that are planning to connect to the core testbed, i.e. currently Jisc, RedIris (University of Murcia), RENATER and SWITCH. These switches are currently being deployed in the NREN labs and are expected to be connected to the GÉANT P4 testbed during Year 2.




Figure 4.1: European P4 testbed

The management of the P4 testbed is currently provided by NMaaS (WP6 T3) (see Figure 4.2). For now, a bastion server controlling P4 switch access is deployed and an Elastic Logstash Kibana is under deployment to support telemetry reporting.

- RARE adoption support:
 - By implementing and demonstrating specific use cases relevant for the NREN community
 - Apart from the technical development, and depending on available manpower and resources, the main focus will be on adoption and use case-focused deployment, workshops and training.

5 Data Transfer Node Infrastructures (DTNs)

The objective of the data transfer node (DTN) stimulation activity under WP6 Task 1 is to explore best practices for deploying and testing large scale data transfer infrastructures and services for the GÉANT community. While there are some infrastructures in place in the NREN context to support large scale data transfers (such as the LHCONe and LHCOPN networks used by the Large Hadron Collider research experiments at CERN), widespread use of performant data transfer infrastructures, including DTNs, remains somewhat in its infancy, particularly in emerging research disciplines that require a data-intensive approach. This work aims to explore the issues in stimulating wider adoption of good principles in deployment and use of DTNs within easy-to-use data transfer infrastructures.

5.1 Achievements

In year 1, the DTN infrastructure work began by determining the current status of DTN deployments in the research and education community - which projects in this area exist, which hardware and/or software is used, how and for what. A survey among European NRENs was conducted to determine their involvement with supporting large scale data transfers for their users, and their use of DTNs. Responses were received from 29 NRENs [Survey].

One of the first things discovered was that there is no unique hardware or software solution that is dominantly used, or that fits all scenarios. There are variations between organisations in how DTNs are built, and which architecture, hardware or software is used for specific use cases. In some cases, special devices were used, while in other cases specific software was installed on certain endpoints.

While 'Science DMZ' principles are commonly followed, the details of the DTN systems, their software and configurations, varies. One common element of the Science DMZ model, network performance measurement, is most commonly done with perfSONAR, which is developed within WP6 T3 as part of an international perfSONAR consortium.

The results of the survey, as well as information about transfer tools and services, the specifications and the examples collected, were published in a wiki page [Wiki] (milestone M6.4) and presented and discussed at the 18th STF meeting. The wiki is intended to be a reference resource for the GÉANT community; there is currently no such single resource otherwise available. During the discussions, one conclusion was that the DTN subject, and the broader issue of effective end-to-end data transfer services for scientists and researchers, includes both technical and non-technical considerations (including researcher engagement with NRENs). A focus group was created to gather this information and follow up on next steps.

Based on the survey results on the different technical solutions, the team in the DTN focus group is exploring the creation of a DTN testbed on the GÉANT Testbed Service (GTS) where the tools used for DTNs and/or to support demanding data transfers will be tested. This will provide an infrastructure on which different scenarios and different solutions can be tested, the results achieved compared with a live environment, and collaboration between different NREs, research organisations or projects established. Early phases of this work were presented at the Performance Management Workshop in Zagreb, Croatia in March 2020.

The Performance Management Workshop was organised by WP6 and hosted by CARNET. The workshop had a session dedicated to DTN use cases and solutions, where NREs who reported some work on DTNs were invited to present. The session goals were to provide an opportunity to NREs to promote their solutions and use cases related to DTN, and to pave the way for discussions, knowledge sharing and collaboration in the community related to DTN.

All WP6 teams, including the initial DTN team in Task 1 and the current DTN focus group team endeavour to keep their work transparent and to disseminate information about their work objectives, methodologies and results. So far, apart from the information on the public wiki, the work of the team has been presented at several events, including SIG-MSP (March 2019), STF18 (October 2019), the DeIC Conference (October 2019) and the Performance Management Workshop (March 2020).

5.2 Future Work

The survey conducted among European NREs revealed that 11 of them indicated that they were interested in some form of further work on DTN infrastructures, although few specific suggestions or details were noted. The focus group will continue to gather information and work with the members of the community to specify and select use cases of interest to the NREs.

In addition, the team will continue to maintain and further develop the wiki pages, extending them with more information and more specific topics of interest to the community, such as a public space and a reference point for the GÉANT community for topics related to DTN infrastructures. The wiki will be open for contributions and collaboration from all interested parties and will contain the results and evolution of the DTN on GTS testbed.

The team will also continue its dissemination work, not just of the results achieved by the DTN focus group, but also supporting the dissemination of other organisations in this area. Building upon the success of the Performance Management Workshop, the WP6 team will consider organising similar workshops in the future where DTN infrastructure, use cases and tools will be presented, demonstrating solutions and points of view from different organisations, searching for commonalities in the underlying principles and methodologies and thus potential areas of collaboration.

6 Low Latency/Low Jitter Infrastructure

Real-time applications like LoLa, Ultragrid, MVTP-4K and Nimbra, which are used for activities such as online teaching, collaborative performing arts, and virtual meetings, depend for their best performance not only on available bandwidth, but also on having a low latency and low jitter path along the network. At present, the network performance monitoring tools deployed in NREN and R&E networks focus on metrics associated with bandwidth and throughput, rather than reporting on combinations of latency and jitter.

There are currently at least 23 NRENs where the above applications are in common use, and the geographical spread covers most of the GÉANT backbone footprint (see Figure 6.1). The LoLa team in GN4-3 WP6 T1 is exploring possible solutions to set up a real-time monitoring system which will present latency and jitter measurements gathered from the GÉANT backbone, to complement the existing measurements that report the bandwidth used. Such monitoring systems would be made available for users - real-time application users, NREN and GÉANT operations - to provide better insight to the network status and thus help them optimise network services for these and similar use cases.

Currently, NRENs that are using the GÉANT backbone for time-sensitive applications do not have information about latency and jitter over all network segments. A path that is optimised for bandwidth or capacity thanks to the routing metrics currently used can be different from one that ensures the lowest latency and minimal, stable jitter.

An example is the bandwidth-oriented routing from southern European NRENs towards Nordic NRENs which go via the London PoP, and can add nearly 20ms (RTT) to the compared latency via the Hamburg PoP or, as was observed some time ago, the east-west NRENs' bandwidth routing which was not using the Milan-Vienna link, adding in that case about 15ms RTT.

While latency is generally stable unless there is a routing change, jitter can vary rapidly, therefore the current (active manual) monitoring approach does not ensure enough reactivity with respect to time-sensitive users' requirements. Even if today it is possible to fine-tune manually the network path and the configuration of routing for time-sensitive applications, this approach is no longer scalable.

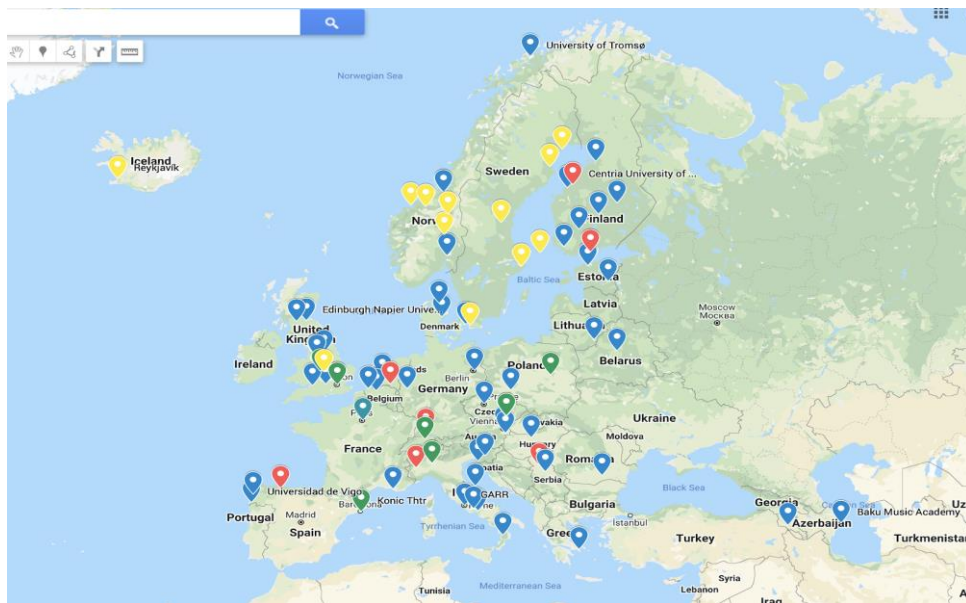


Figure 6.1: Incomplete map of locations where low latency applications are in common use

The latency and jitter monitoring service would provide a user interface that shows latency and jitter segment by segment on a weather map, built on the metrics provided by JunOS TWAMP/RPM/JTI features and implemented with modern time-series and dashboard stacks (e.g., Influx+Grafana, ElasticSearch+Kibana).

6.1 Achievements

From the beginning of the GN4-3 project, WP6 T1 has tested different measurement technologies, which are usually used for capacity assessment, and has applied them to the specific LoLa traffic use case (but the other real-time applications are very similar). After an extensive measurement campaign focusing in particular on perfSONAR, WP6 T1 has identified the potential and the limitations of current active measurement tools for this specific use case.

The existing perfSONAR tools are good enough for generating a LoLa-like traffic and measuring its delay and jitter statistics relevant for evaluation of quality LoLa sessions. At the same time the tests showed that the hardware for a system running such tests must be powerful enough for the generation of LoLa-like UDP traffic (hundreds Mbps to several Gbps of UDP traffic). Besides this, the tests showed that the monitoring will need to combine these tools into a holistic monitoring system that will use both host-based and router-based latency measurements.

The use of the existing Performance Measurement Platform (PMP) was explored to perform tests and visualise jitter and latency measurement results. However, since current PMP nodes have been designed for a different use and certain traffic types (such as TCP iPerf tests), they are not as capable in their current specification for more demanding application traffic (such as heavier UDP-based tests). Therefore, different types of (future) nodes should be considered for real-time network applications like LoLa, which use UDP-based protocols.

In detail, the measurement activity considered a test plan with seven LoLa setups at increasing image resolution and refresh rate, with streams ranging from 640x480 @30fps up to 1920x1080 @60fps. The test plan has been executed over two LoLa deployments: a short-distance setup located in Milan, where LoLa and perfSONAR measurements were run concurrently to compare the observations on latency and jitter, and a long-distance deployment between Vilnius and Trieste. On the Vilnius-Trieste setup, additional measurement tools including the TWAMP implementation from Juniper and native latency-jitter from Mikrotik have been assessed.

The test results led to the following conclusions:

- The generation of application-like traffic is desirable when latency and jitter measurements are being performed. The probes must be able to generate traffic that mimics the packet size and their packets per second (pps) rate of the real-time application in order to provide an environment as similar to the actual use case as possible.
- Tools available in the perfSONAR suite and based on OWAMP/TWAMP protocols (twping, owping and powstram) proved useful for performing delay and jitter measurements during periods of LoLa traffic generation. These perfSONAR tools provide statistics such as median, max, min and mean delays, and STD (Standard Deviation) and percentiles for delay and jitter. However, the current perfSONAR GUI shows only the average delay.
- The precision of twping/owping measurements depends on the quality of the NTP sync. Between Vilnius and Trieste, stable delay values of 27 +/- 2-3 ms with 1-2 ms jitter have been observed.
- Twping/owping itself could be used for LoLa-like traffic generation, using a single tool for both load generation and latency/jitter sampling. However, more powerful hardware nodes would be needed in such cases.
- Further tests with TWAMP on Juniper routers showed that it interoperates with the perfSONAR TWAMP client thanks to compliance with RFC 5357. In addition, TWAMP servers on the routers along a path can be used for fault detection and localisation by tw-pinging from some end hosts (without the need for TWAMP clients on Juniper routers). Finally, TWAMP sessions between Juniper TWAMP clients and servers can be used for per-segment fault detection and localisation or per-segment weather maps without using end hosts.

For these reasons, WP6 T1 together with the GÉANT Operation Centre are exploring what is currently available in the GÉANT network, their current monitoring systems, and how real-time applications fit into the current operational framework, towards creating a new service that embraces delay and jitter measurements relevant for those applications.

The benefit that such a monitoring service would provide is a better, simpler, real-time view of the latency and jitter across the entire GÉANT backbone network during the configuration and debugging phases for LoLa-like real-time collaborative applications.

Figure 6.2 shows the current high-level design for the monitoring tool.

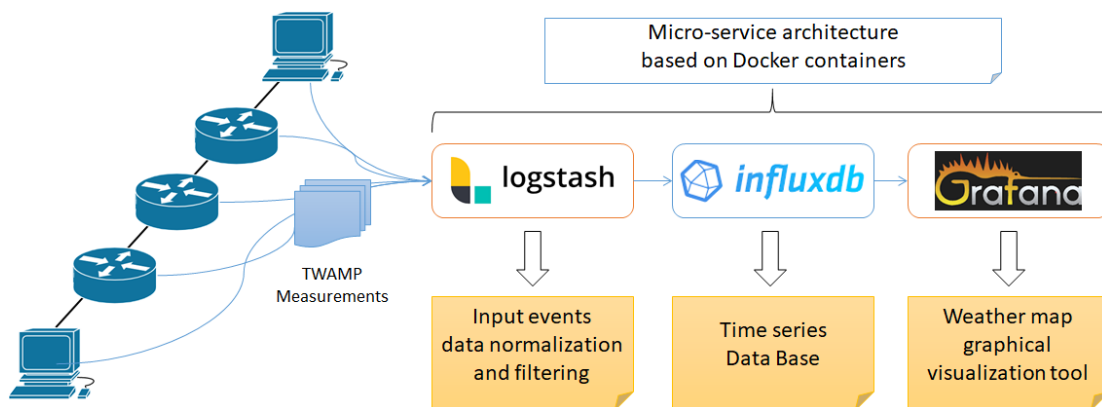


Figure 6.2: LoLa monitoring analysis software architecture

6.2 Future Work

To identify and potentially in the future provide even more automated procedures to optimise the network path for latency and jitter, there is a need to monitor as many GÉANT backbone segments as possible, ideally all of them.

WP6 T1 is discussing with the GÉANT Operation Centre procedures to get access to the monitoring features available on the current Juniper devices deployed on the network backbone, such as the Real-time Performance Monitoring (RPM) feature already active on GÉANT routers, and the Juniper Telemetry Interface (JTI) that is currently under investigation by the GÉANT NOC. Through RPM, and JTI at a later stage, WP6 T1 might be able to collect the latency and jitter measurements that will be processed and exposed through a dedicated monitoring service.

Future work might include, but should not be limited to, extending the latency and jitter monitoring inside NREN backbones to achieve a multi-domain solution, integration of new probes (similar to nodes used for Performance Management Platform (PMP) service) with jitter and delay measurements, using perfSONAR and/or telemetry data from different routers/switch vendors solutions, etc. Exploration of these options will be done in parallel with the current work and might be used as an input for a proposal for potential future work after Year 2.

7 Optical Time and Frequency Networking (OTFN)

Beyond the strong metrology interest, Time and Frequency (T&F) services are critical to many industrial sectors, both directly and indirectly, including telecommunications, positioning, energy, finance and advanced science use cases. Access to precise time and frequency signals is of major interest for European industry, research and the economy, and thus the technology is likely to be of interest to many NRENs.

During recent years, European National Metrology Institutes (NMIs) have successfully developed and tested new T&F techniques (using optical fibres) which are showing performance stability at least three times better than the current best commercial services, and which can respond effectively to the challenges of tomorrow. These outstanding performances are required by NMIs, which are willing to compare their new generation of optical clocks (frequency stability measurements) to pave the way to the redefinition of the International Unit for time - Second (SI), which is currently based on RF clocks. Unfortunately, these comparisons can only be done through optical fibres, which is why there is an urgent need to interconnect European NMIs.

However, these new metrological services are not compatible with standard telecommunication networks because a bi-directional propagation within the fibre is mandatory. This means NRENs need to explore how to connect distant end-users to their NMIs over their current or future infrastructures. The NRENs are in a position to undertake applied research in a way that commercial operators may find difficult, to optimise the delivery of T&F services on their fibre infrastructures. This has led to the development of national and international testbeds and pilot services.

OTFN members believe it is the role of NRENs to stimulate innovation in the research and education community and bring forward technological advances. The availability of cost effective and cutting edge NREN network services enables and encourages new technologies to spill over into the commercial sector, which ultimately benefits society as a whole.

With that goal in mind, the OTFN subtask members within the GN4-3 project began working on two key aspects to catalyse the emergence of metrological networks in Europe. The WP6 T1 subtask's main objective is to study the challenges of deploying new international T/F services, drawing on the CLONETS (H2020 project) early conclusions and best practice guide [\[CLONETS\]](#). WP6 T1 has participated in several dissemination events (and plans to continue doing so) so that more and more European NRENs become aware of the potential of these new services for their users and communities.

7.1 Achievements

Building upon the conclusions of the CLONETS project, WP6 T1 has recognised the urgent need to interconnect the NMIs, which will then subsequently disseminate their signals to their own communities (sub networks) through new international interconnections. It is also of major concern that the current NMI interconnections must be converted into permanent links to guarantee further

progress and allow the exploitation of the developed infrastructure. The WP6 T1 OTFN team has focused on giving an international T/F service proof of concept, expanding the metrological network and making the infrastructure perennial. In addition, the WP6 T1 team is also looking closely at other promising services that could be provided by NRENs within the GÉANT community such as Quantum Key Distribution (QKD), requiring equivalent access to fibres to interconnect its users.

The work of the OTFN team is interlaced with the work conducted in WP7 and the GN4-3N project. The inputs and requirements from the GÉANT network and its users were considered, and possible solutions were also evaluated from the perspective of its implementation in GÉANT and the NRENs' networks.

The OTFN team has provided technical inputs to GÉANT for their equipment procurement process. The team has also explored existing solutions for time and frequency distribution, including types of technique, performances and the current TRL level of such solutions. The summary of this analysis is provided in Table 7.1. Public quotes (T/F hardware) were updated and a high-level costing estimate for implementing T/F was studied.

Existing advanced techniques			Performances Frequency criteria based on instability (the lowest = the best) Time criteria based on precision (the lowest = the best)	TRL
FREQUENCY	Optical Carrier (Carrier Wavelength)	Active cancellation	10^{-15} @1s ; 10^{-20} @1d	8
		Two-way comparison	10^{-17} @1s ; 10^{-21} @1d	5
	RF Carrier (Modulated Wavelength)	Active cancellation with electronic delays (ELSTAB)	10^{-16} @1d (unidirectionnal)	8
		White Rabbit PTP	10^{-15} @1d (unidirectionnal)	8-9
TIME	Two way comparison		TDEV ~ 2ps	5-6
			TDEV ~ 30ps calibration through GPS (unidirectionnal)	6
	Protocol based (White Rabbit PTP)		Verified with GPS disagreement ± 2 ns	8-9

Table 7.1: T/F existing techniques/services, performances and TRL

Different realistic scenarios (spectral occupation in C-Band, L-Band, etc.) were elaborated and are shown in Figure 7.1 and Figure 7.2. The pros and cons of the setup can be found below the two figures.

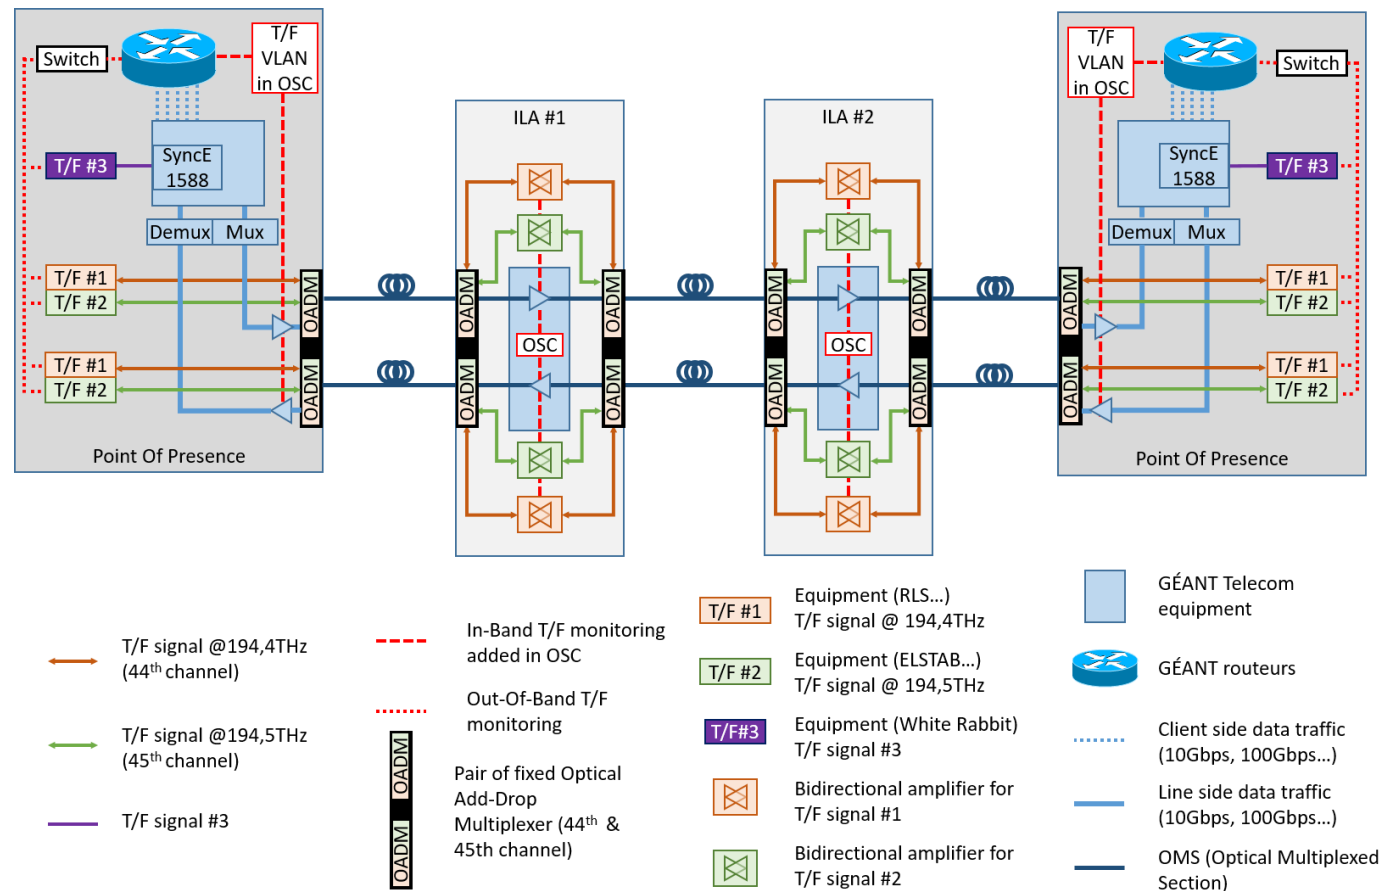


Figure 7.1: Time/Frequency setup

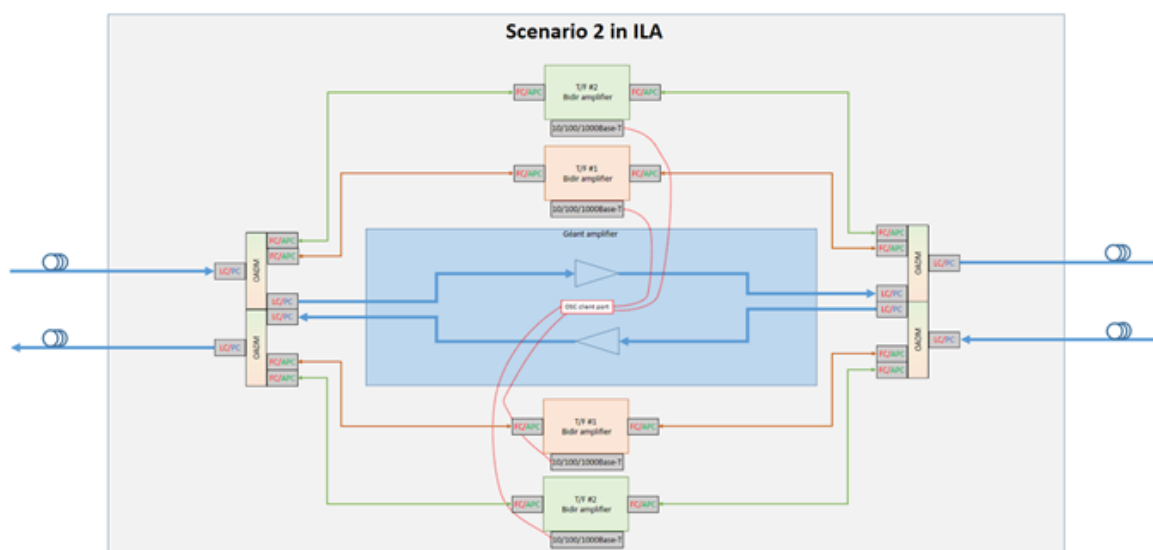


Figure 7.2: Suggestion of T/F setup in in-line-amplifier site

Pros:

- High TRL (8) + highest performance requirements are matched.
- Three TF services at the same time (wide range end-users), as stated in the CLONETS deliverables.
- Only one stage of additional OADM (additional losses for data traffic minored down to <1,4dBs per couple of OADM).
- There is a guarantee that both T/F signals will be inserted through APC connectors (great advantage for lowering oscillations threshold).
- Two amplifiers are depicted but a common amplifier design will be more cost efficient.

Cons:

- Two channels in C-Band are required (could be two times 50GHz bandwidth).
- There could be some coupling between bidirectional amplifiers.
 - intra sites coupling if each T/F service uses its own amplifier.
- One single amplifier for both T/F signals.
 - gain won't be optimised for each signal.
- Need to have SyncE+1588 modules to map (AMP) unidirectional data T/F services to higher ODU with other data traffic.
- No flexibility to reduce C-Band occupation afterwards (drawback for GÉANT).

WP6 T1 organised a meeting with representatives of different T/F research projects including GÉANT, TiFOON, REFIMEVE and CLONETS at GIP-RENATER in Paris on 11th June 2019. A variety of technical solutions were presented, and the technical and organisational aspects of a potential bidirectional service implementation were discussed. It has been recognised that there are very few international connections that are easily accessible so far, and there are different issues related to the co-propagation of data and T/F services in the same network, including risks and limitations of

implementing such services while protecting the integrity of data traffic (which remains the main objective of telecommunication networks).

The reviewed options include the use of C band, L band, dedicated fibre to T/F signal or shared fibre (T/F + data signal). The option to use the channel in the middle of the C band (44th channel of ITU-T DWDM grid @194.4 THz) is seen as optimal for one team (because of the lowest attenuation) but at the same time as unacceptable for another (as it falls in the middle of the frequency band used for data transfer).

Implementing the necessary T/F hardware equipment adds losses due to the additional Optical Add and Drop Multiplexers (OADMs) (see Figure 7.3) used to insert or extract the metrological signal which consequently degrades OSNR. It should be noted that OADMs are designed so that data signals have minimum attenuation.

It was concluded that there is no unanimous consensus between metrologists about the technique that should be used. However, several NRENs (RENATER, CESNET, PSNC) at the moment have production implementations of OTFN based on a dark channel setup in the middle of the C-Band (44th channel of ITU-T DWDM grid @194.4 THz): a dedicated ITU-T channel within a shared fibre and using only its own optical equipment such as an amplifier.

HLD setup in OADM, thin film filters

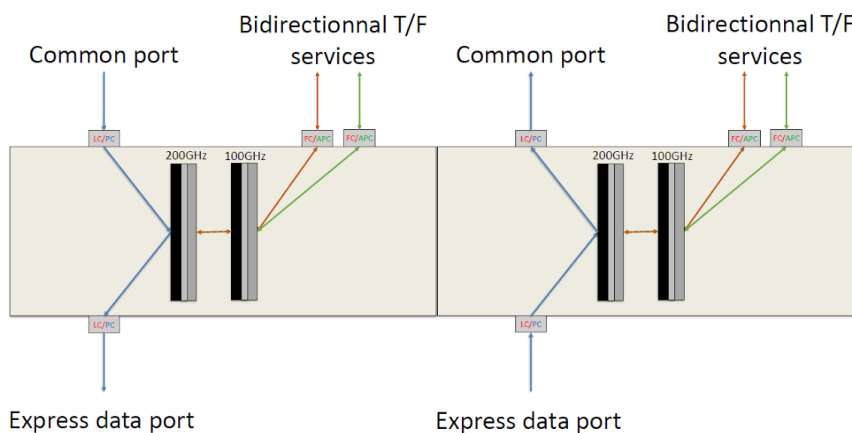


Figure 7.3: Schematic of Optical Add and Drop Multiplexer used to insert or extract T/F signals

7.2 Future Work

The OTFN team is exploring the opportunity to set up an international T&F connection in Geneva. The city is at the crossroads of several countries in Europe and several NRENs either already have some fibre infrastructure there or have plans to install it (PSNC, GÉANT, RENATER, SWITCH, NORDUnet). This could open the door to new interconnections between NMIs. There are, moreover, advanced science use-cases (anti-matter experiments) at CERN that might directly benefit from performant metrological signals from multiple NMI.

The major aspects of this interconnection are based on the following statements:

- The REFIMEVE+ (French T/F production service provided as a bidirectional signal based on optical carrier wavelength 44th/194.4 THz @1542.14 nm) ends currently in Grenoble and will be extended to Geneva.
- OPTIME (Polish T/F production service) would use a 3 118 km fibre to connect Geneva. This is provided as a mono-directional signal sent thanks to Repeater Laser Station (RLS) (@194.4 THz). The needed RLS (estimated to four but maybe two would be enough) could be lent by REFIMEVE+ project for a 12 months duration. This still needs to be confirmed.
- CESNET and PSNC will provide the possibility of comparing signals between CMI (Czech NMI) and GUM (Polish NMI) using unidirectional transmission. Gateway in joint point will be necessary. The direct connection of CESNET in Geneva is mainly subject of availability of necessary GÉANT Spectrum Connection Service between Praha and Geneva that would add another international connection.
- SWITCH plans to provide a high frequency signal to CERN. This activity is expected to start by the end of 2020 and the signal is expected to be available at CERN in 2021. It could be used by the CERN research community and for comparing it with other NMIs as well. The latter requires a kind of gateway to match the different frequencies used by SWITCH/METAS and other NMIs - e.g. an automated frequency comb. Having a difference in the wavelength would open the challenge of setting up the first operating gateway between two NMIs. It is based on bidirectional @1572.06 nm. Gateway (automated frequency comb) would be required at CERN to match the two signals.

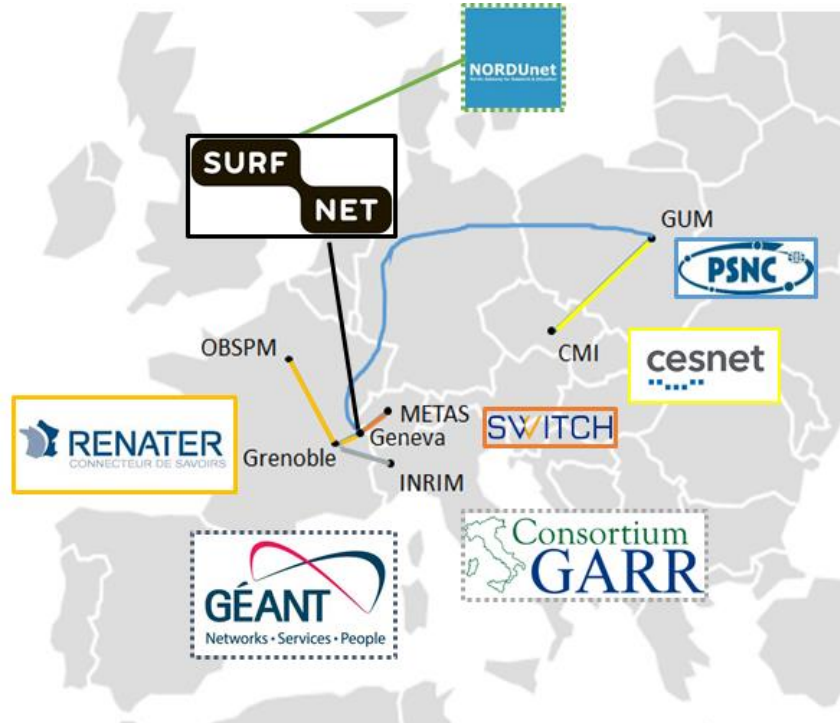


Figure 7.4: Fibre footprint for this interconnection

Figure 7.4 presents the current presence of European NRENs in Geneva. The colours in the picture represent the NRENs present in Geneva and their fibre infrastructure, which could be used for the interconnection.

In addition, the OTFN team is actively participating in dissemination activities with the goal to share knowledge about the OTFN technology, prerequisites and possibilities, as well as about existing solutions and planned future work. A white paper is being prepared describing the experience and lessons learnt during the first phase of the WP6 T1 OTFN work.

The OTFN team in WP6 T1 will continue to work on exploring possibilities to establish interconnectivity between NRENs, which would then be used to provide an OTFN service to NMIs and other users interested in high precision time and frequency services. Apart from the technical perspective, evaluation of such an interconnection should include the financial perspective, including warranty, shipping, installation costs, racks in Geneva, etc. It should also include an assessment of the interest from potential users in NREN member institutions and projects including, but not limited to, CERN. If there is a positive evaluation, future plans would include finalising the interconnection and implementation of a T/F service in Geneva, as well as numerous dissemination activities and workshops, as appropriate.

8 Quantum Key Distribution

The work on Quantum Key Distribution (QKD) started within WP6 T1 officially in Y2, in January 2020. QKD is an implementation of a quantum cryptography concept, a process of encrypting the transmitted data using the principles of quantum mechanics that guarantees the security to new higher levels of assurance than have been possible before. A QKD system uses a dedicated quantum communication channel to generate and transmit secure private keys that can be used in various cryptographic algorithms and devices to secure user data transmission and services.

QKD technology is in its relative infancy. NRENs are interested in QKD technology development and implementation towards potential new services for their users or user communities. The NREN community is already involved in discussions on QKD technology, on the preparation of documents related to quantum networks, and the testing of QKD systems through their participation in different projects.

A QKD system uses keys to encrypt the data. The keys are sent over the quantum communication channel, so the security of key transmission is guaranteed by the laws of quantum mechanics. For the transmission system with QKD technology to be fully secured, the encryption algorithm used in the regular communication channel for user data needs to be a post quantum algorithm, which is resistant to potential hacking attacks using quantum computing infrastructure. Over the years, a number of different variations of QKD systems have been developed and targeted at different scenarios and use cases.

There are two solutions for the implementation of QKD systems in current optical data transmission systems. The first approach uses separate dark fibre as a quantum communication channel for keys and complements the existing optical data transmission infrastructure. The second approach assumes a co-propagation technique, and the quantum communication channel is transmitted as part of the spectrum of the existing optical data transmission system. This approach requires advanced spectrum filtering and management techniques. Quantum communication (using single photons and its quantum states) cannot be regenerated and extended in a traditional way as the optical amplifiers destroy the quantum states of the signal. Currently the only solution to achieve significantly longer QKD transmission, i.e. several hundreds of kilometres, is to deploy a number of trusted node sites along the path that retransmit the signals.

8.1 Achievements

Since the work on QKD is new to the GÉANT project, the first activities initiated by the WP6 T1 QKD team were to determine the current position of activities within and outside of the GÉANT community. This work included determining state-of-the-art in GÉANT NRENs, Europe (and in particular EU projects) as well as world-wide projects, and the availability and maturity of vendor solutions.

To gather inputs from the NRENs about their involvement and plans for the deployment of QKD solutions in their networks, the WP6 T1 QKD team constructed a survey that, with the help of WP3, was distributed to all NRENs. At the time of writing, the survey is still ongoing, however, the preliminary results have shown that while 50% of NRENs consider QKD strategic or important, only a small percentage are currently actually involved in quantum projects.

In 2018 the European Commission launched its Quantum Flagship initiative, which is specifically aimed at supporting quantum technologies projects. This is focused on quantum sensing, computing, simulation and communication. QKD technology is one of the Quantum Flagship priorities. In 2019, the OpenQKD project was launched with a focus on establishing QKD testbeds in Europe and implementing various use cases in existing networks, services and infrastructure. The project involves 38 partners including several European NRENs.

Apart from the European Commission flagship projects, countries like Austria, France, Germany, Italy, Netherlands, Switzerland and the United Kingdom have launched their own national projects to support quantum technologies development.

With respect to QKD systems, there are currently three important big projects in Europe: CiVIQ, OpenQKD and QUAPITAL, all including NREN partners. The projects establish important international collaboration and standardisation activities in terms of QKD, quantum networks and quantum Internet.

The second area of WP6 T1 QKD work includes dissemination activities with the aim to increase the engagement of NRENs in discussion, design and later in implementation of QKD solutions. Current work and plans related to QKD vary among NRENs, so dissemination and engagement activities also vary, from extending the initial team with representatives from NRENs that are already looking for QKD solutions and preparing for testing and deployment, to informative presentations for NRENs that do not (yet) have QKD in their short-term plans.

A WP6 T1 team member has accepted an invitation to join the standardisation group formed by the European Commission. They will participate in the group, represent the view of the GÉANT project and its members, and validate and use the work as input for WP6 T1 activities.

The team also plans to engage the NREN community and identify relevant applications, and scenarios for quantum technologies within NRENs and GÉANT environment, beginning with QKD.

8.2 Future Work

In the forthcoming period, the QKD team will continue the scoping work and use any opportunity to bring benefits from other projects to the work within its task, and to coordinate NREN QKD/quantum activities within existing and potential national and/or European projects. It will also liaise with vendors in the search for optimal architectural solutions.

The team will also continue dissemination activities to increase the level of knowledge and awareness of QKD and quantum technologies in European NRENs, particularly those that do not yet have firm plans to deploy any QKD solutions. Last but not least, the team will continue to explore available QKD technical solutions for early implementations in relevant NREN environments.

9 Conclusions

The *Network Technology Evolution* task of the Network Technologies and Services Development Work package in the GÉANT 4-3 project is exploring several technologies and their relevance and applicability in the European NREN community: White box, DPP, RARE, DTN, LoLa, OTFN and QKD. Based on the evaluations so far, the conclusions and next steps for the WP6 T1 work on these technologies are as follows:

White box will certainly be deployed first in the cloud world, which is where the most important market growth can be observed. White box can be used in R&E but not for all use cases. Disaggregation and independence allows NRENs to manage their switch/router estate in a very different manner. Beyond the technical and economic aspects, the strategic criteria are very important, i.e. how important it is to become independent from the hardware router provider and whether sufficient manpower is available to handle this new task.

DPP explored the use of the P4 language for telemetry and DDoS detection use cases. The work demonstrated that the limitation of the P4 language and P4 chipset memory means the DDoS algorithms need to be relatively simple, but that such algorithms can still also be relatively effective. Monitoring using In-band Network Telemetry (INT) in the data plane is very promising as it allows a very detailed view of all events that happen in the network. This paves the way for a new monitoring approach that could help NRENs to better understand their traffic and diagnose performance issues.

RARE was very successful in implementing the coupling of an open source control plane with a P4-driven data plane, on top of a P4 hardware switch. As a result of this achievement, an almost complete LER/PE router (IPv4, IPv6, ISIS, MPLS, SR-MPLS, L3VPN, XConnect, VPLS, EVPN, 6VPE, etc.) was implemented. The two missing pieces to test in production are the control plane policy and monitoring/telemetry which are under study. Beyond the availability of a full open source router for R&E networks, this project returns network innovation to R&E as it is very easy to integrate a new network protocol with the RARE solution.

DTN centralises all information regarding DTN infrastructures in a wiki that is intended to be a reference resource for the GÉANT community. As the DTN development includes both technical and non-technical considerations, a focus group was created to gather information, and work with the members of the community to specify and select use cases of interest to the NRENs. The ongoing implementation of DTN on a GTS testbed will allow European NRENs to test different DTN implementations more easily.

The **LoLa** project is designing a segment-by-segment latency and jitter weather map solution for the GÉANT backbone to allow network operations teams to monitor the performance of real-time applications.

After looking at existing **OTFN** solutions in European NRENs, the OTFN team has started work to explore interconnecting national T/F services. Beyond the contribution to metrology science with the interconnection of the European National Metrology Institutes, some European experiments (anti-matter experiments for instance) could also benefit from T/F services and one would expect that these T/F services will help civil and strategic applications by providing an official time service.

The **QKD** activity started in January 2020 by scoping the existing and planned QKD work. A survey to determine the status in NRENs showed that while many NRENs view QKD as important, only a few are involved in QKD projects so far. These preliminary results indicate the need for knowledge sharing and dissemination that should precede any broader technological implementation. In parallel, the team is exploring the possibilities of a specific use case and technology availability for potential testing, which will depend on further development of vendor solutions, results of European and world-wide projects, as well as standardisation efforts.

References

- [Chaignon] P. Chaignon, K. Lazri, J. Francois and O. Festor, *Understanding disruptive monitoring capabilities of programmable networks*, IEEE Conference on Network Softwarization (NetSoft), 2017
- [CLONETS] <https://www.clonets.eu/>
- [D6.3] https://www.geant.org/Projects/GEANT_Project_GN4-3/GN43_deliverables/D6-3_White-Box-Evaluation.pdf
- [Ding_1] D. Ding, M. Savi, D. Siracusa, *Estimating Logarithmic and Exponential Functions to Track Network Traffic Entropy in P4*, to appear in *IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2020
- [Ding_2] D. Ding, M. Savi, G. Antichi, D. Siracusa, *An Incrementally-deployable P4-enabled Architecture for Network-wide Heavy-hitter Detection*, to appear in *IEEE Transactions on Network and Service Management*, 2020
- [ESnet] <https://wiki.geant.org/display/SIGNGN/4th+SIG-NGN+Meeting>
- [GNS3] <https://www.gns3.com/>
- [INT] <https://github.com/p4lang/p4-applications/blob/master/docs/INT.pdf>
- [LoLa] <https://lola.conds.it/>
- [OcNOS] <https://www.ipinfusion.com/products/ocnos/>
- [SFINX] <https://wiki.geant.org/download/attachments/120497358/SFINX%20Migration%20Guide%20v3.00.docx?api=v2>
- [Survey] <https://wiki.geant.org/pages/viewpage.action?pageId=133765748>
- [Wiki] <https://wiki.geant.org/display/dtn>

Glossary

ACL	Access Control List
APC	Angled Physical Contact
BGP	Border Gateway Protocol
DPP	Data Plane Programming
DTN	Data Transfer Node
eBGP	external Border Gateway Protocol
ECMP	Equal-Cost Multi-Path
EVPN	Ethernet Virtual Private Network
FBK	Fondazione Bruno Kessler
FPGA	Field Programmable Gate Array
GUI	Graphical User Interface
INT	In-band Telemetry
IGP	Interior Gateway Protocol
IP	Internet Protocol
IS-IS	Intermediate System to Intermediate System
IX	Internet eXchange point
JTI	Juniper Telemetry Interface
LAN	Local Area Network
LoLa	Low Latency
LSP	Label Switched Path
LSR	Label Switch Router
MPLS	Multi-Protocol Label Switching
NMI	National Metrology Institute
NOS	Network Operating System
NPU	Network Processing Unit
NREN	National Research and Education Networks
NTP	Network Time Protocol
OADM	Optical Add-Drop Multiplexer
ONIE	Open Network Install Environments
OSNR	Optical Signal-to-Noise Ratio
OTFN	Optical Time Frequency Network
P4	Programming Protocol-Independent Packet Processors - programming language
PMP	Performance Measurement Platform
PoP	Point of Presence
QKD	Quantum Key Distribution
R&E	Research and Education
RARE	Router for Academia, Research and Education

RIOT	Routing In and Out of Tunnels
RLS	Repeater Laser Station
RPM	Real-time Performance Monitoring
RSVP	Resource Reservation Protocol
SNMP	Simple Network Management Protocol
SSH	Secure Shell
STD	Standard Deviation
STP	Spanning Tree Protocol
T&F	Time and Frequency
TCAM	Ternary Content-Addressable Memory
TCO	Total Cost Ownership
VHDL	(VHSIC-HDL) Very High Speed Integrated Circuit Hardware Description Language
VLAN	Virtual Local Area Network
VNF	Virtual Network Function
VXLAN	Virtual Extensible LAN
WB	White Box
WP	Work Package
ZTP	Zero Touch Provisioning