23-09-2019

Deliverable D6.13
White Box Evaluation

Deliverable D6.13

Contractual Date:	10-31-2019
Actual Date:	10-31-2019
Grant Agreement No.:	856726
Work Package	WP6
Task Item:	Task 1
Nature of Deliverable:	R (Report)
Dissemination Level:	CO (Confidential)
Lead Partner:	RENATER
Document ID:	GN4-3-19-1495F9
Authors:	Xavier Jeannin (RENATER), Mauro Campanella (GARR), Frederic Loui (RENATER), Edin Salguero (RENATER), Maxime Wisslé (RENATER), Christos Argyropoulos (GRNET), Jani Myyry (FUNET), Ivana Golub (PSNC), Tomasz Szewczyk (PSNC), Damian Parniewicz (PSNC) , Bojan Jakovljevic (AMRES), Pavel Benacek (CESnet), Marco Savi (GARR), Susanne Naegele Jackson (FAU/DFN)

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 presents if the new type of routers called white box could be used by research and education community and for which use cases. The Router for Academia, Research and Education project (RARE) investigates also if it is possible to use open source Network Operating System (NOS). The ability to program the data plane thanks to a high abstract level language (P4) opens the door to new applications for research and education. Two use cases are presented here: Telemetry and Distributed Denial of Service attack mitigation.

QASPER Review Sheet

<This section to be deleted before publication of deliverable>

REVIEW	**Main reviewer** (N. Surname)
General comments and summary of suggested changes	
Does the deliverable meet its contractual obligations as outlined in the Technical Annex? If not, why and is this acceptable?	
Do the overall approach, structure and content provide what is required under the EC contract (Technical Annex)?	
Is all content suitable for the EC or should certain parts of the content be kept confidential (for project participants only)?	
Does the deliverable adhere to the overall strategy and scope of the project as recorded in the Technical Annex?	
Does the deliverable have any content with potential PR value for project participants/NRENs and/or external audiences? If yes, please elaborate:	
Which information or parts of the document (or the whole content) would it be beneficial to the project to communicate/publicise further?	
For each part, please suggest in the table below which audiences should be addressed.	

GÉANT project participants	
GÉANT NREN partners (includes project participant and non-participant staff)	
End users (individuals, projects and organisations)	
Global R&E community end users	
Global projects (e.g. TEIN3, ALICE, etc.)	

EC, politicians, policy makers	
Industry, general public	

Recommendation	1) Major revision[1]	Y/N	2) Minor revision[2]	Y/N
Re-submitted for review – if 1)	DD/MM/YY			
Final comments				
Approved[3]:	DD/MM/YY			

[1] Deliverable must be changed and reviewed again before submission to the EC can be considered.
[2] Deliverable may be submitted to the EC after the author has made changes to take into account reviewers' comments as appropriate.
[3] For submission to EC.

Table of Contents

Table of Figures

Table of Tables

Executive Summary

A white box is a switch/router that allows running different Network Operating Systems (NOS). (The optical white box is out of scope of this report). The white boxes, born in data centre, offer an impressive forwarding capacity for a very low price. Although the currently available NOS are not providing all the features required by NRENs, the white box usage improves the NREN level of independence from the router vendor, and could thus change the way NRENs manage their network deployment. A white box is not a chassis and does not provide a second routing engine. The white box chipset forwarding characteristics (forwarding capacity, internal memory, size of buffer) determine the usage (IX switch, data centre, CPE, P/LSR, etc.).

Based on several use cases, WP6T1 has already been able to demonstrate that white box can be used for CPE and Internet eXchange point switch and is confident for the DC fabric use case, even if the technical analysis is not finished in this last case. From the present experience, the management decision to go into production is not only based on technical considerations and total cost ownership but also on internal organisational constraints (team workload, capacity to hire staff, strategic plan, etc). For use cases that require more routing features like Label Edge Router / Provider Edge (LER/PE), the currently available NOS could have limitations.

Router for Academia, Research and Education project aims to demonstrate that an open source control plane on a white box can be used as a router. RARE has already demonstrated that the development of open source data plane routing features and the integration of open source NOS (for instance FreeRtr) on the P4 data plane are feasible. RARE is currently working on CPE and P implementations, but there is no theoretical limitation, one can envision having a PE open source.

Thanks to data plane programmability (DPP), advanced network features can be programmed for NREN needs. DDoS mitigation algorithms were implemented on virtual P4 environment and the implementation on a P4-capable hardware is ongoing. In-band Network Telemetry (INT) allows very accurate network monitoring, debugging in a novel way and can significantly improve network management, using just a few nodes supporting INT.

1 Introduction

The networking industry landscape is evolving fast as the market trend is now directed toward data centre and cloud based services. The strategy of new players who want to enter this market is to propose not only lower prices and a higher ratio of port density, but also to decouple the network operating system (NOS) from the hardware in order to unlock the traditional monolith vendor router/switch market. Are we in the same situation at the network level as when Linux appeared in the UNIX world? Is white-box a real opportunity for NRENs and research and education?

The second significant evolution is the white box programmability thanks to recent advancements in data plane programmability and new chip implementation (e.g. Barefoot Tofino). A high-level language P4 for data plane programming has been developed in order to make the data plane programmable capitalizing from the OpenFlow experience.

The ability to integrate different pieces of software (control plane, data plane and intercommunication between these two components) is an opportunity to run an NOS (open source or commercial) over white box hardware. The Router for Academic, Research and Education (RARE) project will investigate, as a first stage, the feasibility to integrate an open-source network control plane that provides a complete feature set compliant to research and education ecosystem requirements, and to connect this control plane to a P4 data plane.

The data plane programmability (DPP) ensures line-rate packet processing. The DPP allows powerful algorithms to be compiled and executed directly in the data plane. This opens the door to the design and development of a lot of new features or improvements. Among them, WP6T1 selected new network monitoring solutions, In-band Network Telemetry (INT) and a new security solution for DDoS mitigation in order to demonstrate how DPP might be of interest to NRENs.

This document reports the evaluation of white box and data plane programming use in the NREN context. Section 2 details the investigation and the results regarding white box usage (White box for research and education). Section 3 presents the data plane programmability (DPP) work and section 4 reports the work of the RARE team. This work is then followed by a general conclusion in section 5.

2 White box

2.1 Definition and scope of the work

As there are several definitions of white box, WP6T1 gives here its definition that will define its work scope. A white box refers to a product that is sold without any trademark in opposition to a branded product. In the networking context, a white box is a switch/router on which different Network Operating System (NOS) could be installed. Note that WP6T1 will study the white box in the NREN context and not only in the context of data centres like the white box is often presented on the Internet. Moreover, the optical white box is out of scope of this task.

The business model for the proprietary hardware forces a router buyer to acquire a package made-up of a certain hardware, a proprietary NOS, the associated hardware maintenance, and NOS maintenance. In the case of a white box, the business model leaves the choice to the customer to buy the hardware with its maintenance from a hardware supplier and buy a commercial NOS with its maintenance from a software supplier. This gives two levels of independence: independence from the hardware (one can change the hardware vendor and keep the software) and independence from the NOS (one can change the NOS and keep the hardware). In order to evaluate the white box interest, "WB for research and education" will analyze white-boxes available on the market and their applicability and usability in the NREN context.

The Open Compute Project specifies an open source initiative called Open Network Install Environment (ONIE) that defines an open "install environment" for the installation of different NOSs on bare metal switches. Some white boxes can also be provided with a Linux system that allows to install the NOS as presented in Figure 1: White box architecture.

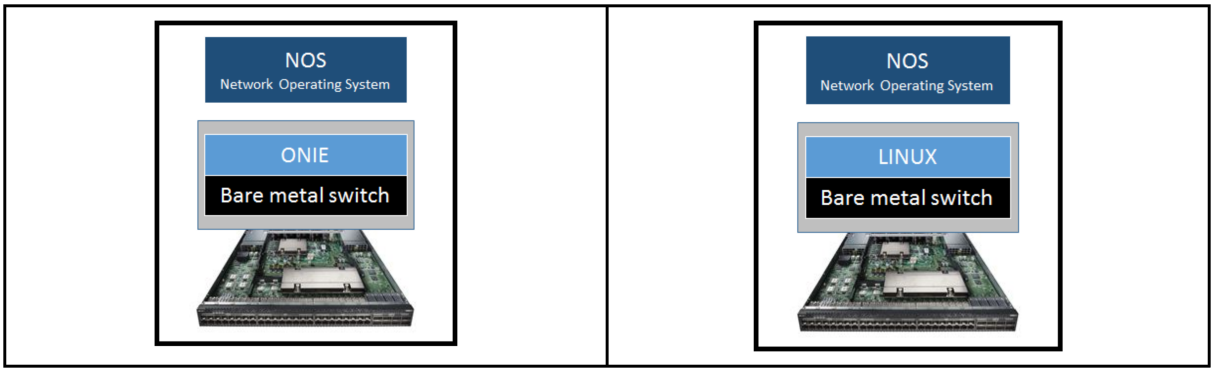

Figure 1: White box architecture

Since the white boxes were born in the context of data centres, the white box designers focus on a huge (several 100Gbps) forwarding capacity, but with features that aim to address the data centre market (a small number of routes, lots of layer 2 features, …). A strong difference with regards to traditional network provider chassis routers is that white box does not exist as a network chassis and does not provide several "route engine cards" (several CPUs). Some white boxes are equipped with the exact same chipsets that are used by traditional vendors [Merchant_Chips]. The price for this forwarding capacity is very competitive for this type of machine. There are different switch designs in relation to their usage: data centre, LAN, campus network or network backbone. The first white boxes were designed for data centre (DC) which implies a very short Round-Trip delay Time (RTT). The machine was designed to handle microbursts that could happen in DC (for instance TCP Incast traffic). This leads to a design with a relatively short buffer. As now the white boxes are deployed more commonly, the white box designers target new markets and a white box equipped with a large buffer forwarding chip is emerging now (Jericho 4GBytes) [Packet_buffers]. In section 2.2.4, there will be a discussion regarding the importance of the switch buffer size.

Recently, the server suppliers have put on the market a hardened X86 server specially designed to become a small router (switch form factor, no graphic card, hardware hardened, designed to be used without cooling, etc.). As different NOSs can be installed on this machine, it can also be considered as a white box. NRENs who express their interest to try white box want to test them with a minimal risk i.e. at the edge of their network, for instance for a site router use case, Customer Premises Equipment (CPE). As most white boxes previously available on the market are very powerful in terms of forwarding (several 100Gbps ports), they are not really adapted to fit in use cases that do not require such capacity. In this context, this new type of machine (the X86 server) can be appropriate for these types of use cases such as a CPE. Error: Reference source not found presents an example of a CPE design and its architecture.

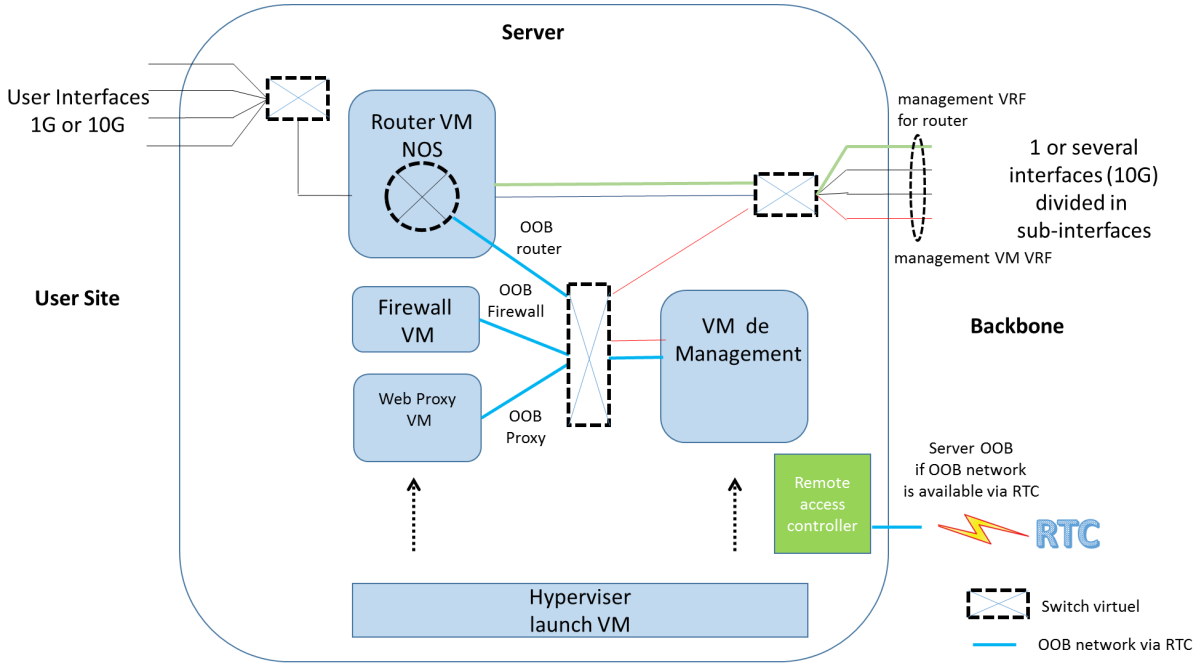

Figure 2: Example of CPE design over an X86 server with a NOS

2.2 White box for research and education

2.2.1 Objectives

The first step of the WP6T1 white box sub task is to evaluate if the current white boxes that are available on the market now or will become available during the project could be used. The rationale is to lead this evaluation by following NREN use cases, i.e. to verify if white box could be put in production in an NREN context. The benefit of this approach is that it takes into consideration all aspects required of these white boxes by the NRENs: routing, management (monitoring, authentication, maintenance model, etc.), security, license model and Total Cost of Ownership (TCO).
Another benefit of the WP6T1 approach is that WP6T1 will capture how NRENs will manage these white boxes in production, since the management has to be different from the management of a traditional router. The maintenance model is also not the same. There might be two different companies that will maintain the white box, one for the hardware and another one for the NOS. The NRENs' NOCs have to become familiar with the new NOS and also with the new way to deal with the maintainer.

2.2.2 NREN requirements and concerns

During the "White Boxing[4]" workshop[5] organised in Stockholm, April 4th 2019 (15 NRENs, 40 person registered), WP6T1 took the opportunity to conduct a survey to gather NREN interest, their potential use cases and their potential concerns. As can be seen from Error: Reference source not foundError: Reference source not found, they indicated three use cases they started with: CPE, cloud fabric and big science. Their concerns were related to the support, the quality of software and reliability. During its work on various use cases, WP6T1 will address these NRENs' concerns.

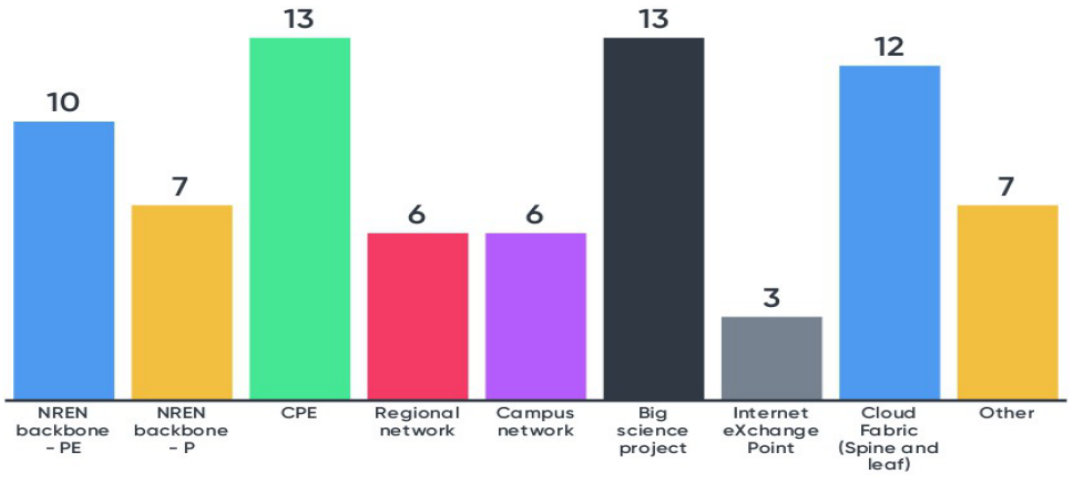

Figure 3: Use cases selected by European NRENs for white box usage

The NRENs identified the following points as critical, by the order of importance: the support, the software quality, the features availability, the stability and reliability.

[4] https://wiki.geant.org/display/SIGNGN/2nd+SIG-NGN+Meeting
[5] https://eventr.geant.org/events/3050

These points are addressed by the WP6T1 study. In order to be as complete as possible, WP6T1 will follow the test up to the implementation in production if the NRENs will decide to do so. This will be especially useful for stability and reliability assessment.
The NREN engineers had also expressed their concerns about the buffer size used in the white box forwarding chipset that could impact traffic in case of congestion and/or QoS usage. A specific study was started and the first results and analysis can be found later on in this section.

2.2.3 Methodology and use cases

At the beginning of the project, the WP6T1 team selected the use cases to work on, based on use cases that the participating NREN partners would consider realistic to implement in production. Initially the big science use case of the GÉANT team seemed to be a promising candidate for white box, but then the usual GÉANT network supplier proposed a new traditional solution with a very competitive price in order to avoid the risk of potentially losing a customer which led GÉANT to abandon exploring this white box avenue. The traditional vendor prefers to make an effort on the price than take the risk to lost users that would not come back to them.

Therefore WP6T1 is working on the following use cases:

- Customer-premises equipment (CPE) - [FUnet, RENATER],
- Provider Router (P) / Label Switch Router (LSR) - [PSNC] ,
- Data centre (or cloud) fabric - [GRNET, RENATER],
- Internet eXchange point (IX) - [RENATER].

Each of these use cases will stick to the following assessment process before going into production:

- Use case specification
- Technical validation - switch and routing features, management feature (monitoring, …), security features (ACL, …)
- Business model (License model and TCO)
- Qualification for production by NREN management - taking into account the previous analysis, NREN management will take a decision based also on the general context (manpower availability, strategic plan, etc.)
- Production - Deployment plan

2.2.4 Buffer Size and Performance Test

2.2.4.1 *Buffer size*

NREN engineers expressed their concerns regarding the buffer size available in white box in comparison with “traditional” routers. This point led WP6T1 to study the white box behaviour in case of congestion. The first white boxes targeted data centres where large buffers are not necessary but new white boxes are now available that are equipped with a new chipset (Jericho) and a large buffer.

The buffers function as microburst absorbers. One has to be cautious when speaking about the importance of buffer size in router and switch design. In the literature, what was true yesterday is now questioned, different models appear. Moreover, it is difficult to manage elephant and mice flows at the same time. Large buffers are needed to manage oversubscription with QoS mechanisms. Buffers delay the traffic a little bit for the microburst to be able to be absorbed by the overloaded

interface. Unfortunately, there is not a single definition of microburst and it is very difficult to obtain information from the router manufacturers on this topic.

The question is when, where the microbursts happen and which application is sensitive to delay variation. In [Beheshti_et_al] researchers from Stanford and the University of Toronto were able to conduct an experiment on the Level 3 commercial backbone (OC-48 links, buffer size = 60 MB / 190 msec. or 125,000 500B packets with no active queue management). The links were set to the experimental values: 1, 2.5, 5, or 10 msec. buffer. No drop was seen with the 5, 10, and 190 msec. buffers for the entire duration. Packet loss in the range of 0.02% to 0.09% was seen with 2.5 msec of buffering and correlated to the link utilization. There was a relatively large increase in packet loss with 1 msec. of buffering, but link utilization was still maintained. Most of the loss occurred when the link utilization was above 90% for a 30-second average. The packet drop level for the 1 msec. buffer was still below 0.2%.

In the data centre, TCP Incast traffic is generated by application request (Hadoop, Map Reduce, HDFS for instance) to several nodes that answer in general with very short lived flows but simultaneously generating microbursts. Researchers at the University of California at San Diego recently performed an in-depth analysis of traffic at Facebook [Roy_et_al]. Servers were 10Gbps attached, their utilization was under 10% (1% most of the time) and the data on buffer utilization was collected at 10 usec intervals for links to web servers and cache nodes. The buffers were constantly in use on the ToR switches (Facebook Wedge with Broadcom's Trident II ASIC, which has 12 MB of shared buffers). "Even though link utilization is on the order of 1% most of the time, over two-thirds of the available shared buffer is utilized during each 10-us interval."

From the study reviewed by WP6T1, the following table summarizes the applications that could be impacted by packet loss and delay variation:

Application	
High-Frequency Trading	Device latency must be eliminated and buffering minimized
Gaming	Usage of buffer could be beneficial if the latency is low but not if the RTT is close to 100 to 200 msec.
Non-live Streaming Video	Predicted to reach 80% of Internet traffic in 2019. Normally capable of sufficient host-side buffering to retransmit lost packets and tolerate moderate increases in latency. It is mainly the available bandwidth the major factor.
Live Streaming Video	Inherently bursty due to video compression algorithms. Applications will suffer similar issues with packet loss, latency, and delay variation.
Voice over IP	VoIP is sensitive to loss, jitter, and latency similar to video
DNS	Do not require special treatment, could be impacted if latency is very high
Web browsing	HTTP/1.1 use lot of parallel sessions and use buffer. HTTP/2 will limit the number of sessions and use larger initial congestion windows; this will lead buffer

	requirement reduction.
Peer to Peer	Distribute scientific data and Linux. no special consideration for buffering
Data Centre - Distributed Compute and Storage – MapReduce, HDFS	Such applications generate TCP Incast traffic, very short oversubscription due to the synchronized answer to request [Chen, Grifit, Zats & Katz]. A short buffer is efficient as seen in Facebook study. The buffer can also be set-up on the server and seems more efficient.
Data transfer	It was demonstrated [Jim_Warner] that large data transfers over long distances with high RTT using large pipes, when a 10 Gbps source sends to a 1 Gbps destination, benefit from large router buffers. In this case, few lost packets dramatically affect the transfer performance if RTT is high. This is a typical use case for NRENs in international projects.

Table 1: Application impacted by packet loss and delay variation

Last technical point, the buffer memory could be inside the NPU / Forwarding ASIC or in an external memory. The first case saves space and power consumption but does not allow for very large buffers. In the second case, additional memory needs to have a large bandwidth and therefore the technical solution is expensive.

In conclusion, the data centre and backbone scenarios differ a lot. The main point regarding buffer in the data centre are the applications used, as the RTT is very low. In the data centre case, a buffer usage appears often even with an almost empty network at 1% or 10% utilization, but a small buffer is enough to manage this. In telecom backbone, packet loss occurred only when utilization was above 90% for a 30-second average. A large buffer of five msec. seemed enough and significantly efficient.

Adding delay and variation of delay impact also some applications (VoIP, Live Streaming Video). On NREN backbone, it is especially the long distance data transfer in case of high speed transfer source sending to a slower speed transfer destination that large buffers are required. But this type of long distance data transfer is a use case that is widely served by NRENs.

Large buffer has to be considered if you expect to implement QoS or must deal with oversubscribed links. As we have already seen, white boxes are available with small or large buffers (Jericho 4 GByte per ASIC), and buffer size is one of the architectural parameters the network architect must optimize.

2.2.4.2 *Performance Tests*

In order to address the NRENs concerns regarding congestion and large buffers, PSNC built a testbed to be able to demonstrate buffering capabilities of a single white box platform. This test was led by PSNC in the context of the LSR/P router use case.

Performance Tests

The main goal of the test was to verify whether the head of line blocking and back pressure (according to RFC2889) appears on the tested white box platform.

For the test four 100GE interfaces were used. The white box platform was configured as MPLS LSR in order to switch MPLS packets. On the Spirent Test Center intermediate MPLS routers were emulated. On top of this setup, the RFC2889 Congestion Control script was started on traffic injector (Spirent).

The test indicated (see table X1 in Annex) that for a range of frame sizes starting from 64B to 1518B, load levels from 60 to 100% showed no head of line blocking and back pressure effects appearing on the tested platform.

Traffic bursts

The main goal of the test was to evaluate the burst handling capabilities of the MPLS LSR router built with white box platform and independent NOS. In the given case the Edgecore and IPinfusion were tested. The testbed shown in Figure 4 emulated the MPLS network with intermediate LSRs on the Spirent STC tester. From two 100GE interfaces traffic was sent to a single egress interface in order to emulate congestion conditions.

The traffic was sent for 10 seconds and its characteristic was changed in steps in order to measure burst handing performance. The source interface load was changed from 25% to 55% with a step of 5%. For each load value the number of burst packets was changed from 50000 to 1000000 with a step of 50000 packets per second.

Although the actual Internet traffic mix has changed over time, the standardized IMIX profiles used for testing have not been updated accordingly because the IMIX test results need to be comparable. The IMIX packet size distribution is shown in Table 2.

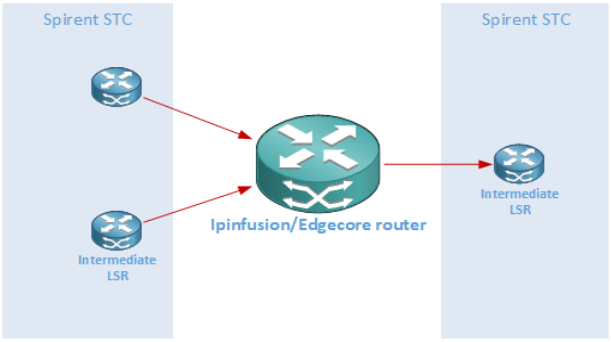

Figure 4: Traffic burst P/LSR tesbed

iMIX Distribution	Frame Length Mode	IP Total Length	Default Ethernet	POS Length	Weight	Percentage (%)
Default	FIXED	40	64	64	7	58,33
Default	FIXED	576	594	594	4	33,33
Default	FIXED	1500	1518	1518	1	8,33

Table 2: IMIX packet distribution

The Edgecore platform was controlled by IPinfusion OcNOS system. Although the actual Internet traffic mix has changed over time, the standardized IMIX profiles used for testing have not been updated accordingly because the IMIX test results need to be comparable.

The IMIX traffic was sent from two 100GE interfaces for 10 seconds to a single 100GE interface in order to generate a temporal congestion state. The tested platform was able to handle bursty traffic up to 350k packets per second (pps) without packet loss when the average load on the single source interface did not exceed 45% link utilization. At the same time, for properly switched packets, the average delay was lower than 20 us as it is show in Figure 5. For larger burst sizes, the tested platform was able to handle the traffic with packet loss lower than 1% keeping delay below 35us. As a result of the test, we can say, that the platform offers line-rate switching for time sensitive applications which do not require large buffers.

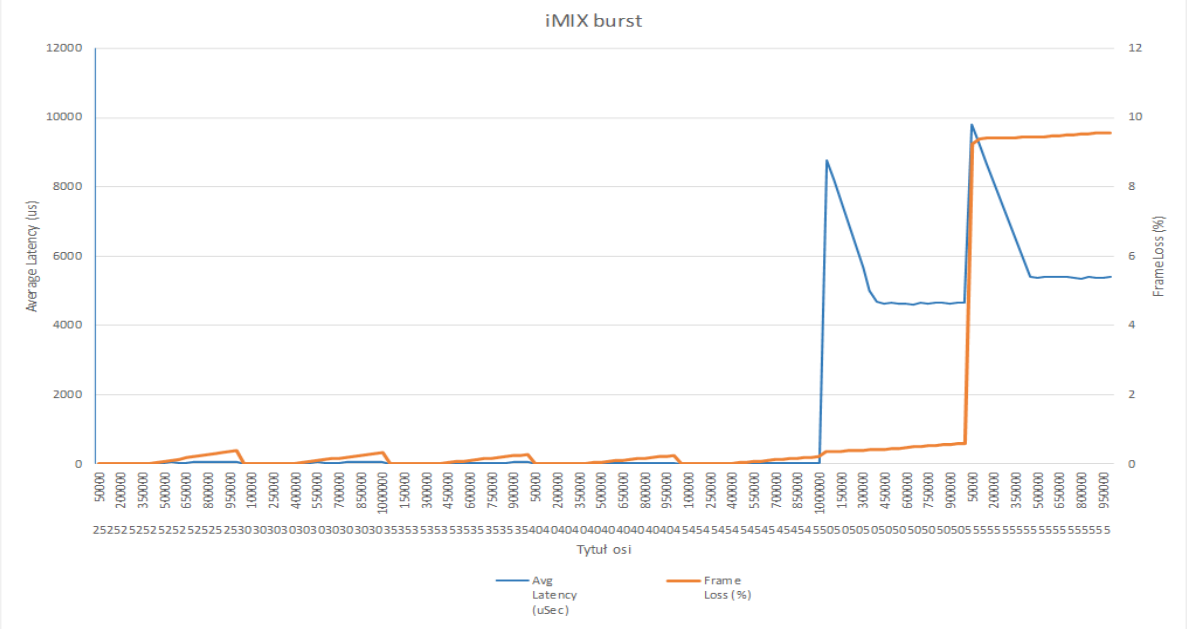

Figure 5: Bursty test results

2.3 Use cases

2.3.1 CPE Normandy

Currently in the region of Normandy, approximately 140 high-schools are connected through a network using an old version of CPE routers, limited in capacity. For this reason, the CPE has become a bottleneck, especially in cases where dark fiber is now available and needs to be renewed. The CPE specification requires to increase the bandwidth up to 1Gbps or more and a list of routing, management and security features was specified in the following document [Normandy_CPE_1]. Automation is also required. The cost cannot exceed the cost of the existing solution. Taking in consideration that physical white box is designed for data centre, even if they are cheaper, there are still not cost effective in comparison with a very small router. In this use case the connection throughput requirement is not very high, therefore a solution based on x86 servers with the switch form factor is suitable. Moreover, it is possible to add additional network functions like a firewall in the future. Figure 6Error: Reference source not found shows the basic architecture of Normandy

CPE. The chosen NOSs use Linux as platform, which gives the feasibility of implementing automation solutions based in software. This solution is very flexible at low cost.

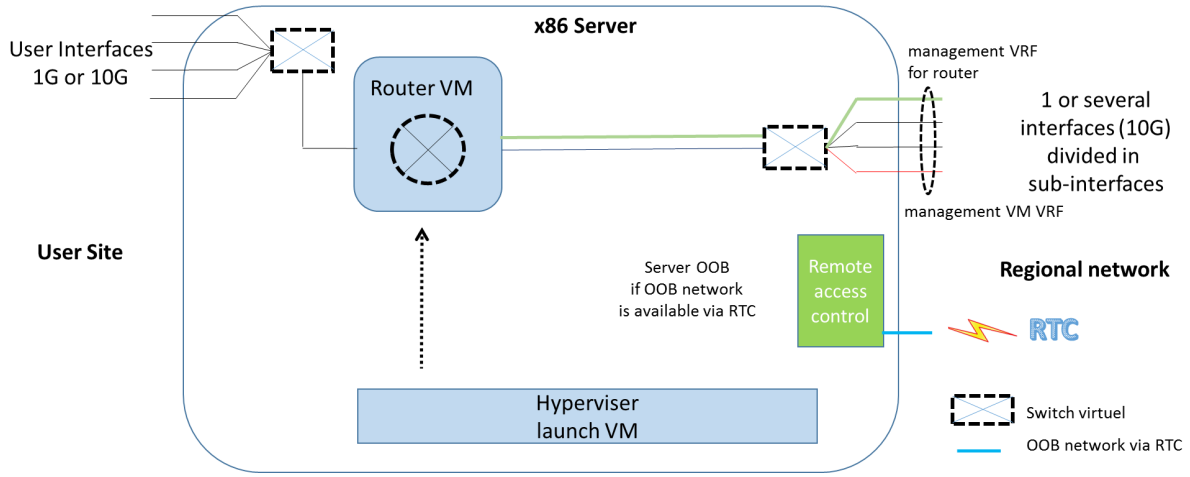

Figure 6: Normandy CPE architecture

The two different NOSs Free Range Routing (FRR) and Cumulus Networks were tested on a testbed and documented here [Normandy_CPE2]. For the proposed scenario, both of them worked perfectly. Cumulus Networks even uses FRR as part of their NOS to handle the IP Routing and Forwarding part. All the tests were successfully completed by both as far as the routing features were concerned. It was only in the management plane where Cumulus had advantages over FRR, providing user accounts with protocols like TACACS+ and Radius. Also in security, Cumulus supported ACL implementation with Linux IP list, whereas for FRR, IP table had to be used. However, since the white box was capable of running NFVs, if stronger security features were required, a firewall NFV could be configured.
For the automation part, both NOSs integrated successfully with Ansible. In fact, Cumulus offered a module inside Ansible to make provisioning in an easier way, facilitating the commands manipulation in Ansible Scripts. But even without a dedicated Ansible module, FRR could use the Linux shell module to be provisioned in an efficient way. Basically, the difference was that shell commands required to be more specific for any command. But as mentioned, it worked successfully in integration with FRR.
Based on the results of these tests (system and equipment management, security, routing protocols, automation deployment), and after considering advantages and disadvantages of both solutions, the client chose FRRouting as the solution suitable for the project deployment.

For the coming production phase, the idea is to start over with 2 high-schools connected through their own white box running FRR NOS. The original plan was to start with the first deployment (real scenario pilot test) in October 2019 and the customer already bought the 2 Dell x86 servers. Ultimate lessons will be learned when the machine will be in production.

2.3.2 FUNET CPE (F-CPE)

In the light of new nationwide network upgrade changes, FUNET is currently working on replacing existing CPE devices with those based on a white box. The motivation for this is found in the price of a white box combined with the set of functionalities that can be achieved and adjusted to the particular user needs.

The typical dual router setup is illustrated in the following Figure 7. The "FUNET White Box CPE" (F-CPE) would provide uplink connectivity with BGP and routed access for different subnets in the existing L2 networks. The routed access is typically protected against single-point-of-failures with Virtual Router Redundancy Protocol (VRRP). There might also be a customer owned firewall which usually is installed between the edge routers and L2 networks. If there are any special user requirements like the need to bypass the firewall, those users can be connected directly to the F-CPE routers.

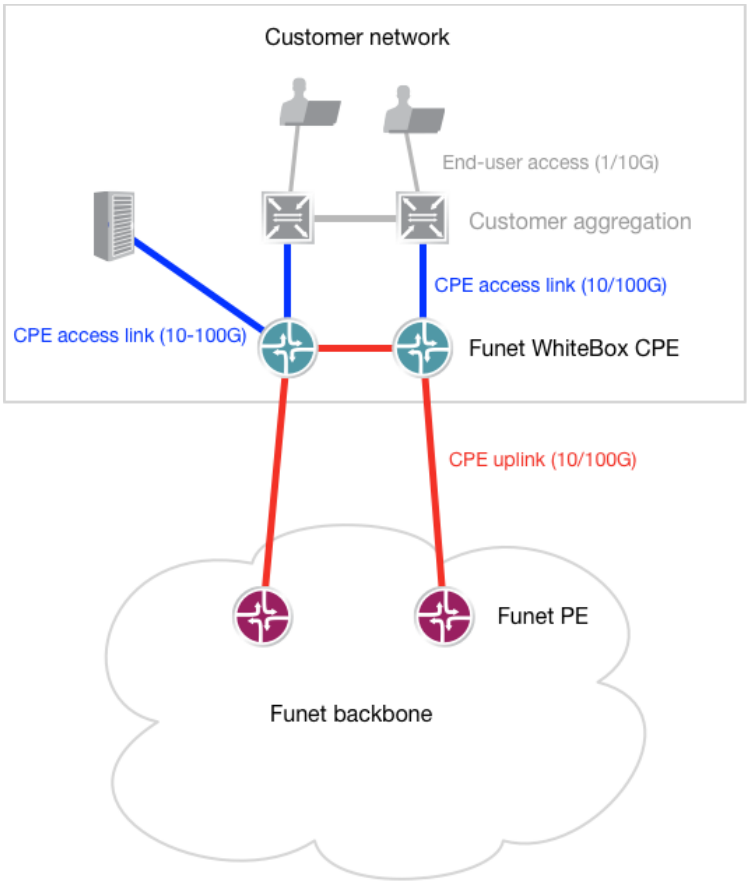

Figure 7: FUNET CPE project

The F-CPE use case was tested in a VMware environment with Cumulus VX NOS appliance to evaluate control plane support for required features. Cumulus provides commercial software support.

Feature tests were performed with two NOS instances connected directly to the FUNET backbone with BGP. Client connectivity was evaluated with a separate Ubuntu Linux virtual machine. The existing infrastructure was used for management, syslog, monitoring and DHCP/ DHCPv6 tests.

The following control plane features were successfully implemented: eBGP peering towards Funet backbone (both IPv4 and IPv6 unicast), iBGP peering between the Cumulus NOS instances (both IPv4 and IPv6 unicast), BGP route filtering, OSPFv2 (IPv4 unicast) and OSPFv3 (IPv6 unicast) as an IGP, VRRPv3 (IPv4 and IPv6) towards the client network, Loopback and interface access control lists (IPv4 and IPv6), Management VRF support, IPv4 DHCP relay support with a client host, IPv6 stateless address auto configuration (SLAAC), IPv6 DHCPv6 relay support with a client host, Jumbo frame support (IP MTU up to 9000 bytes).

Various management and monitoring features were also evaluated successfully: SSH management access with key-based authentication, remote syslog (via management VRF), DNS and NTP (via management VRF), SNMPv2 polling and feeding interface counters to InfluxDB/Grafana (via management VRF), configuration backup and restore, NOS software upgrade.

The control plane and BGP dynamic routing was stable for months and control plane responsive in general. In the next phase Cumulus NOS should be tested in the real hardware environment to evaluate forwarding plane performance and control plane and forwarding plane interoperability. The tests performed show that the main control plane functionality needed in the CPE use case is supported by Cumulus and thus the NOS can be considered for the purpose. However, as the system was only evaluated in a virtual environment, more tests would be needed with real hardware before decisions to proceed towards production can be done. The next step will be to prove the results

achieved in the virtual environment again on the selected hardware platform, and to complete the cost analysis for this solution. In CPE use case implemented with hardware white box, the biggest challenge would be finding a hardware which provides port density needed in CPE environment and is cost-efficient.

2.3.3 GRNET Data Centre

White box solution is evaluated in cloud fabric context with the expectation to benefit of cost-reduction, following large-scale cloud providers paradigm like Facebook, H/W-NOS decoupling, less vendor lock-in, no proprietary solutions, shorter life cycles, fully automated management and service provisioning.

For the DC use case speeds beyond 10 Gbps for the server ports are desirable, keeping speed of 10 Gpbs as the minimum requirement for server-facing interfaces.

The most typical setup for the server ports is a dual TOR switch link connectivity which enhances reliability requirements and gives more flexibility on the support side (e.g. TOR switch upgrades without traffic interruption), also doubling the link capacity. In this setup each server is connected with 1x10 Gpbs to the same rack TOR switch and using cross-rack cabling also to the adjacent rack TOR switch with another 1x10Gbps. However, basic hardware redundancy options like dual power feeds and hot-swappable power supplies and fan units might be beneficial. The basic node redundancy includes: non-redundant control plane, hot-swappable power supplies, hot-swappable fan units, Next Bussiness Day support~~NBD support~~ service.

The "DC Fabric" use case follows closely GRNET's current DC architecture, as shown in Figure 8. The basic service is the VLAN provisioning to customers/servers ports on TOR switches and the customer/server multi-homing on two adjacent TORs using EVPN/VXLAN Ethernet segment mechanism and LACP protocol for bonding. The mandatory features tested are EVPN/VXLAN Ethernet segment mechanism and LACP protocol, Ethernet interface setup, BGP/EVPN and VXLAN, SSH, DNS, NTP, Ansible, NETCONF but not VRFs in the first phase and multicast. Another issue under investigation is the potential use of the spine switches as the DC routers, running BGP protocol for inter-DC and IP network interconnections, as shown in the following Figure 8. Therefore, access-control lists (ACLs) are going to be tested for the L3 interfaces, evaluating their overall size and number that can be supported from the white box devices. The entire configuration should be manipulated by an automation mechanism such as Ansible for configuration creation and NETCONF for configuration deployment.

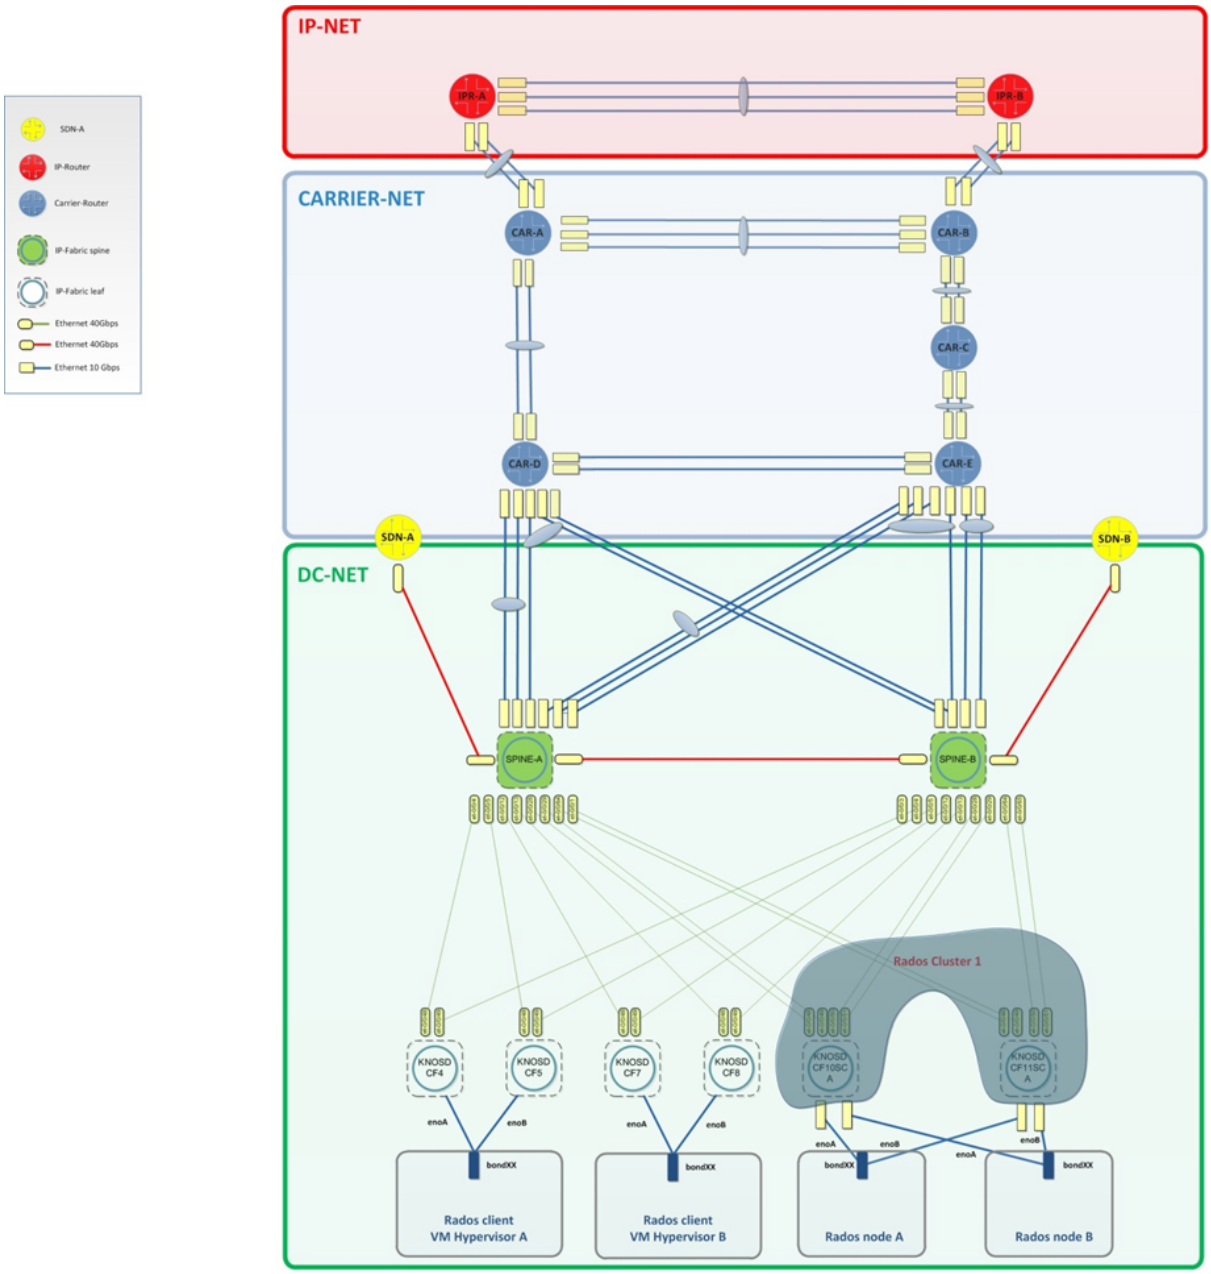

Figure 8: GRNET data centre project

Currently the fully fledged white box based DC fabric network laboratory is under procurement process. Hence, GRNET used a virtual environment, based on EVE-NG platform, in order to test EVPN/VXLAN based control plane on top of the expected CLOS topology. The virtual spines and leaves switches running Cumulus OS were deployed and the result is a fully functional environment. Moreover, the management plane was successfully handled thanks to NETCONF. Next round of tests (undergoing) is related to the EVPN control-plane and the routing models, including asymmetric, symmetric or centralized routing. An evaluation of the troubleshooting methods/tools will be done at this stage.

2.3.4 Normandy Data Centre

The Normandy regional network want to renew and increase the size of its data centre (50 racks up to 150 racks). It was decided to implement a first slice of 10 racks as prototype to assess the solution in a real operational environment. It was expected that white boxes would avoid vendor lock and save Capex. A tender was published and the result showed that traditional vendors answered with a very competitive proposal. The interesting point here is the rationale leading to the regional network's decision. Taking in consideration that the traditional vendor offer was at the same level as the white box vendor's, and the fact that the white box would require extra manpower for a team used to implement traditional vendor NOSs, the white box total cost ownership was judged higher than the traditional vendor's TCO. Therefore, it was decided to choose the traditional vendor solution.

This example proves, once again, that traditional vendors see white box solution as a real competition especially in the data centre domain. The traditional vendors do not expect that the white box users would be disappointed and would come back to them. Instead, they prefer to make a strong effort to keep this market.

2.3.5 Provider Router (P) / Label Switch Router

The objective is to deploy efficient and cost effective LSR routers in the GÉANT or NREN MPLS backbone. The LSR router allows PE router traffic aggregation and transport PE traffic between PoPs. LSR routers in most cases do not terminate particular services. Today there is a strong need for 100G and 400G interfaces in core networks. Scalability of MPLS network can be increased by adding another level in the network hierarchy. Moreover ability to install optical modules supporting selectable wavelength transmission increases potential capacity of fibre links. The number of required transponders increases with growing capacity of the network. Each time a pair of transponders is required, it increases Capex and Opex. In addition, the LSR devices are more efficient when it comes to power consumption, thus reducing Opex.

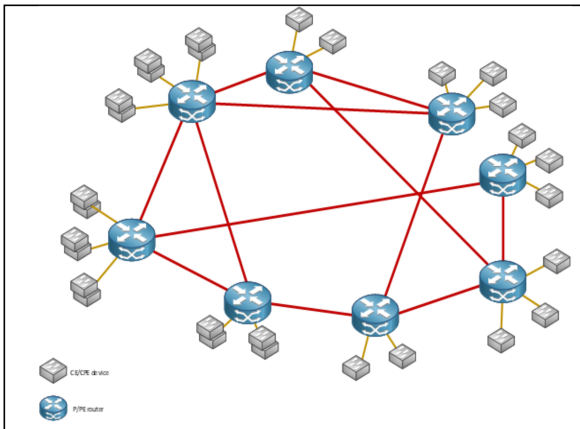

Figure 9: Collapse core architecture

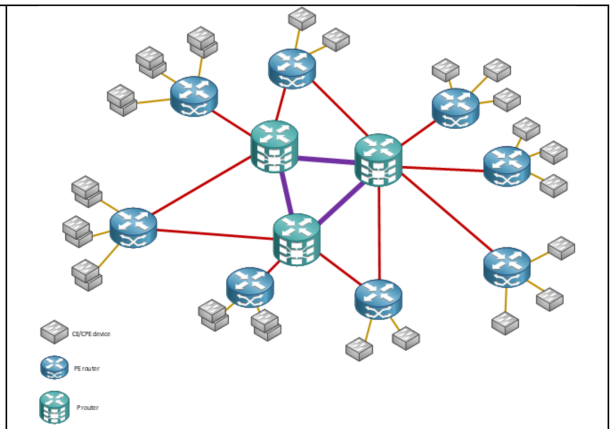

Figure 10: P/LSR core architecture

The main purpose of the test was to verify the basic LSR functionalities. Spirent TestCenter application and N11U chassis with 100G and 100G interfaces were used. The IPinfusion NOS was implemented on the white box device (EdgeCore) and configured as MPLS LSR. On the Spirent TestCentrer application, the PE and intermediate LSR routers were emulated. The LSPs were established on the testbed (Figure 11) between pairs of PE routers and passing through the device under test.

Figure 11: LSR/P testbed

Once configured on the tester the 10800 LSP paths were signalled. Next the traffic stream was generated in order to verify MPLS forwarding plane operation. During the test, all LSPs were up and no traffic loss was observed as shown in Figure 12.

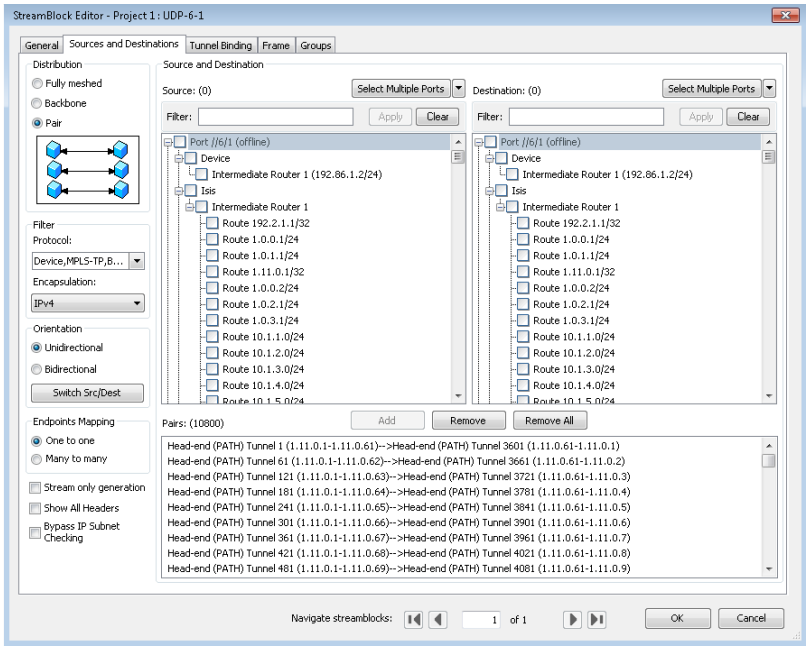

Figure 12: LSR/P testbed

The testing phase continued with the buffer size and performance validation of the white box device that is reported section 2.2.4 “Buffer Size and Performance Test”.

2.3.6 Internet eXchange point (IX)

The main purpose of a Global Internet eXchange point (GIX) is to connect multiple entities in a dedicated architecture and enable them to create direct peering among themselves. The most common architecture is the layer 2 architecture with switches that provide connectivity between all the customer routers in the same network, or with a route server service that eases the customer BGP peering establishment.

This has been selected as a use case for white box in order to test if this new concept can fulfill all the requirements of the RENATER's Internet Exchange point (SFINX) and at the same time provide a cheaper solution. OcNOS by IPInfusion was chosen as the NOS, as it is used in production by LYNX, a London GIX.

The testbed included 2 white boxes, 1 Juniper MX104 with a logical system to simulate a route reflector and the clients that send routes and perform BGP peering. In the first place, Open Network Install Environment (ONIE) was tested with all the required features listed here [SFINX].

Plane	Example of tested features
ONIE	Remote and automatic installation (DHCP, Web server), manual installation.
Management plane	SNMP, TACACS, RADIUS, LOGS, NETFLOW, SSH, …
Control plane	OSPF, VLAN, RSTP VXLAN,, …
Security	MAC/IP ACL, BPDU filters, …
Data plane	MAC address table limitation, ARP table limitation

Table 3: Except of GIX features tested

All tests were completed successfully, the results were conclusive and validated all required features. Some non-blocking limitations, such as the maximum numbers of characters of an interface description that is limited to 32 characters, were also identified. The NOS cost and the hardware were significantly lower compared to the existing costs regarding the current set-up. Therefore, it has been decided that the white box solution will replace the existing solution in the production. The transition will require detailed preparation to ensure minimal downtime. It will be performed by RENATER, and RENATER plans to finish the transition before the end of this year.

3 Data Plane Programmability

3.1 Definition

There has been a continuous effort to innovate in networking, starting from the control plane. The SDN paradigm has allowed to program the behaviour of the control plane in switches. Separating control from switching and centralizing the control function, allowed the definition of new behaviours based on standard packet headers. Modifications are implemented in the controller and sent, even frequently, in the form of switching rules to selected switching nodes.

However, the network chips operating in the switch data plane are still proprietary and mostly bound to fixed switching operations implemented in the silicon during the design phase.

A programmable data plane transform the network ASICs, allowing to program new forwarding behaviours in the packet processor itself. A full programming control on processor memory and functions permits almost complete freedom on packet header processing, including information insertion, change and removal. Data Plane programming requires adequate hardware like FPGA or Barefoot Tofino, [Barefoot], however the cost is not significantly higher than "traditional" Ethernet switching hardware.

The advantages are silicon programming using languages like P4, extreme flexibility in packet handling (telemetry), new function/protocol inclusion without hardware replacement, while maintaining wire speed performance. The effort on the Tofino chip is being standardized in the Protocol Independent Switch Architecture (PISA), by Barefoot, now Intel.

3.2 Use cases

3.2.1 Telemetry with Data Plane Programming

In order to extend the set of supported actions, protocols and monitoring approaches, software tools should be used which can be extended very easily. Unfortunately, the software performance solution cannot be sufficient, especially in a high-speed network environment like NREN backbones. Due to these reasons, the capability of programmable monitoring in P4 white box is very interesting because it allows easy deployment of new monitoring approaches at the speed of the network line.

WP6T1 decided to demonstrate the feasibility of this approach with In-Band Network Telemetry (INT) [INT] which is an emerging standard for real-time network monitoring. The INT technology is based on the insertion of telemetric data into each passing frame as it is presented in Figure 13. The telemetric data can carry information like the occupancy of switch queues, the time required for the crossing of autonomous system or security-related information. For example, it is possible to indicate malicious traffic with invalid protocols/protocol field values or traffic which is suspected to be DDoS traffic. This information can be added and processed directly at the data-plane layer and it can be used later in software which can perform the deep inspection.

This work aims to provide a precise measurement on network flows, without sampling, based on the initial parameters: delay, delay variation, packet loss, reordering. It will be performed with standard x86-64 based server with PCIe and Linux using FPGA cards (two 100G ports and powerful Virtex UltraScale+ FPGA [liberouter]. The Linux drivers support DPDK and 200Gbps transfers into RAM memory. The measurements can be used by a low latency application (such as the LoLa project) in a multi-domain environment. The telemetry use case is based on CESNET work related to the P4 INT which was previously presented during the P4 Workshop at Stanford, CA, USA [CESNET]. CESNET also developed a compiler from $P4_{14}$ language to VHDL which is one of HDL (Hardware Description Language) languages used for the description of digital circuits. This allows the implementation of a processing pipeline described in the P4 program. The generated hardware is capable of processing network data at the speed of 100 Gbps. More details about the generated architecture are available at [Benycze].

Figure 13: INT testbed plan

So far, we achieved the following results:

- Our P4 INT design was ported to the newest FPGA based SmartNIC which was developed at CESNET. The design is capable of managing the throughput of 100 Gbps. The P4 INT pipeline is used on both network interfaces. Therefore, the network card is capable of processing network traffic at the speed of 200 Gbps.

- The DPP research group has decided to use the newest $P4_{16}$ language revision. CESNET has started working on the new compiler implementation which generates the VHDL code from a provided $P4_{16}$ program, while the older $P4_{14}$ version will be also supported.

Regarding the use-case, the plan is to implement the tagging of malicious traffic. The information about malicious traffic will be inserted into the packet in the form of a P4 INT header. This information will then be moved into the software which decides the next step (mitigation or export of the information to the next connected systems).

3.2.2 DDoS detection

DDoS is a constant threat to NRENs' users and network services. The use case has the objective to validate data plane programming with P4 as a tool for service improvement in responsiveness and precision. In the first phase, the goals of the DDoS detection activity are, firstly, very fast detection of DDoS attacks on boundaries of NRENs/GÉANT network, secondly, provision of detailed information about DDoS traffic characteristics as it is shown in Figure 14. In phase two, the goal is to achieve almost immediate mitigation of the attack. In the activity, the WP6T1 team has been implementing the DDoS traffic detection and monitoring directly in the P4 programmable switch with use of big data streaming sketch memory structures.

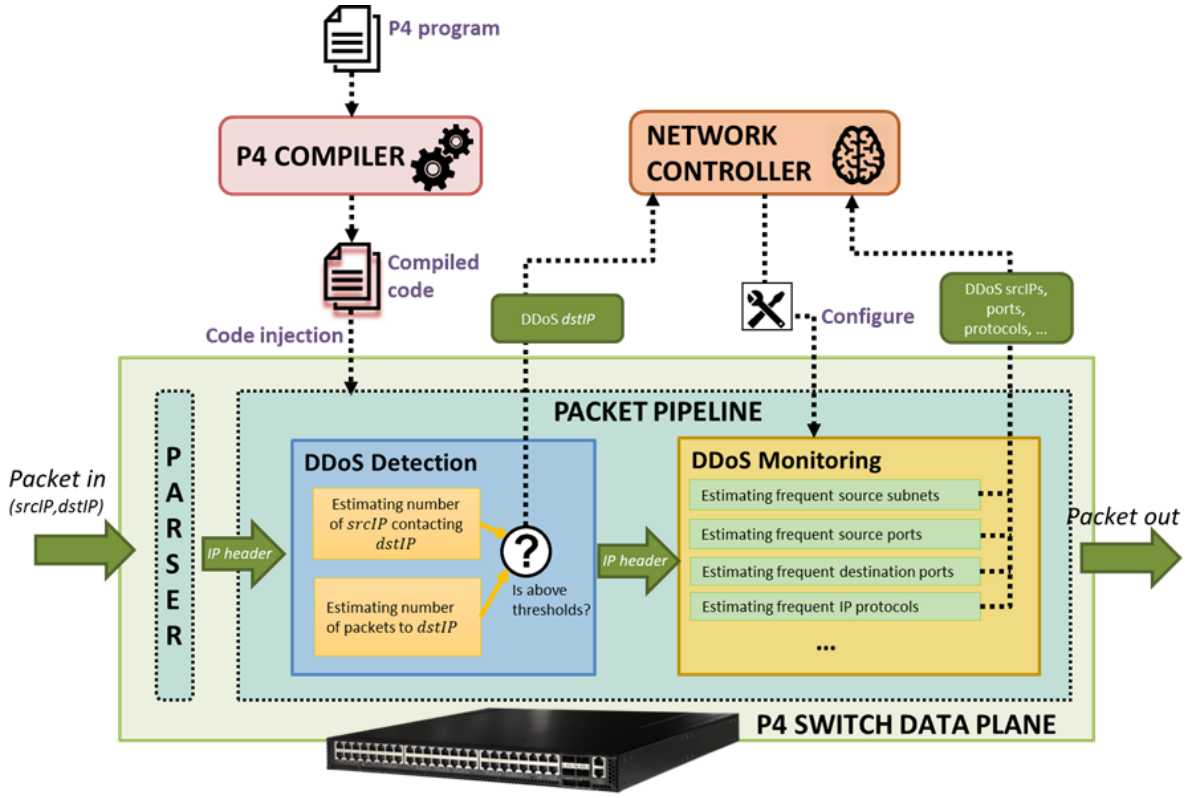

Figure 14: Overview of DDoS detection and monitoring prototype

The DDoS detection and monitoring prototype provides important information about DDoS active incoming attacks and can greatly enhance countermeasures against this type of attacks. The important benefit of P4 switch usage for handling DDoS traffic is that detection algorithms can be adapted for new types of DDoS attacks and added to the switch in quite a short amount of time.

3.2.2.1 *Sketch structures for DDoS traffic detection and monitoring*

DDoS detection and DDoS traffic monitoring can be performed directly in the data plane level of the white box thanks to the usage of the big data streaming sketch memory structures. The sketch structures allow for memory effective collection of summarized traffic statistics and have some interesting benefits in comparison to currently used monitoring techniques. They are processing packets with full wire speed and perform quite a simple set of actions per every packet. Moreover, all processed packets, without performance penalty, can contribute to traffic statistics. This is not true for the most popular approaches (sFlow and NetFlow) that require packet sampling (for example, for 10 Gbps the sampling is at a rate of about 1 packet every 2000 packets) and only based on the packets it reviewed. This is not good for the fast DDoS detection. Another benefit is that the sketch structures provide just aggregate information and require low volume communication between the network node and the network controller. This can be contrasted with Netflow, which is another monitoring technique. NetFlow sends information about every detected 5-tuple flow which means that in certain situations a lot of information can be passed to the network controller. It can be really problematic when a lot of short living flows are detected as in the case of DDoS attacks. The sketches do not require any additional CPU-intensive analytical processing which mostly is performed on powerful server machines in comparison with NetFlow and sFlow.

In general, the sketch structures (Figure 15) can provide a statistical estimation of items frequency in the big data stream (provided by Count-min sketch), items cardinality (by HyperLogLog algorithm) or membership (when Bloom filter is used). The simplicity of sketch structures, which are mostly based on hash algorithm operations, allows the implementation of sketch algorithms logic in the P4 programs which can be deployed on programmable white box switches.

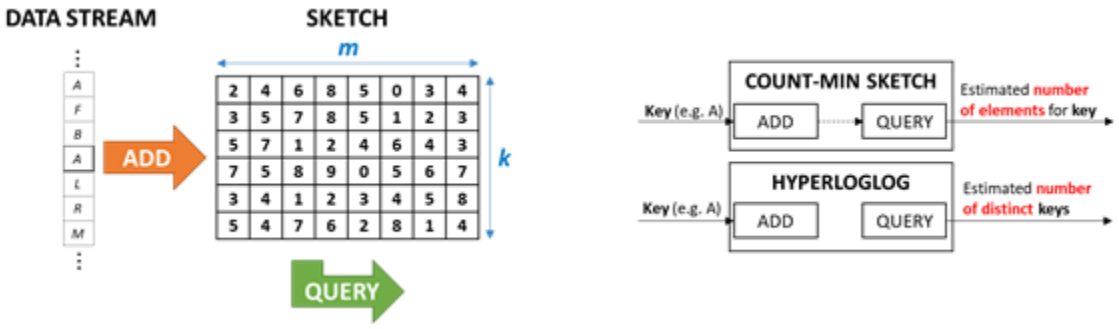

Figure 15: Sketch structure

When sketch structures are applied to network traffic monitoring they can generate aggregate information how many source IP addresses contacted a given destination IP address, how many packets were processed for each destination IP address, how many packets for each source UDP or TCP port were transmitted, and many others.

The current WP6T1 implementation includes two sketch structures for the DDoS detection. It detects all destination IP addresses that might be subject to a DDoS attack and reports it to the network controller. A destination IP address is considered to be under attack if it is contacted by a high number of source IP addresses and has received a high number of packets within a short time interval. The sketch-based DDoS detection algorithm is using a novel data structure that combines Count-min and HyperLogLog sketches (see Figure 16), with the aim of estimating the number of distinct source IP addresses that send at least one packet to a specific destination address. The P4 implementation of this complex, multi-dimensional sketch is a challenging task.

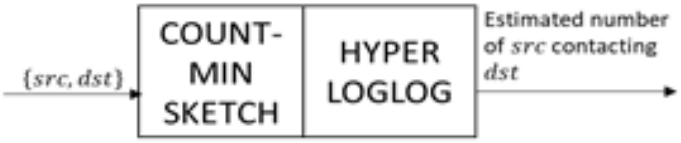

Figure 16: New sketch structure for the detection of DDoS attack targets

When the network controller is notified about the DDoS attack target at the specific destination IP address, the controller can activate one or more DDoS monitoring algorithms which provide more detailed DDoS traffic statistics:

- total traffic (packets and bytes) towards DDoS target
- most frequent source IP subnets (i.e. with prefix /16) originating the attack,
- most frequent source TCP/UDP ports - specific port numbers can suggest the usage of the DDoS amplification techniques based on vulnerable public network services (e.g.: DNS, NTP, SNMP),
- most frequent destination TCP/UDP ports which tell what service is under attack (e.g.: web portal),

- IP protocols used (is it UDP- or TCP-based DDoS attack),
- most frequent packet lengths (amplification attacks are using big packets).

The Figure 17 presents the workflow of DDoS detection and DDoS monitoring algorithms deployed on the programmable data plane device
More DDoS monitoring characteristics can be added in the future like the level of packet fragmentation, TCP flags used, etc.

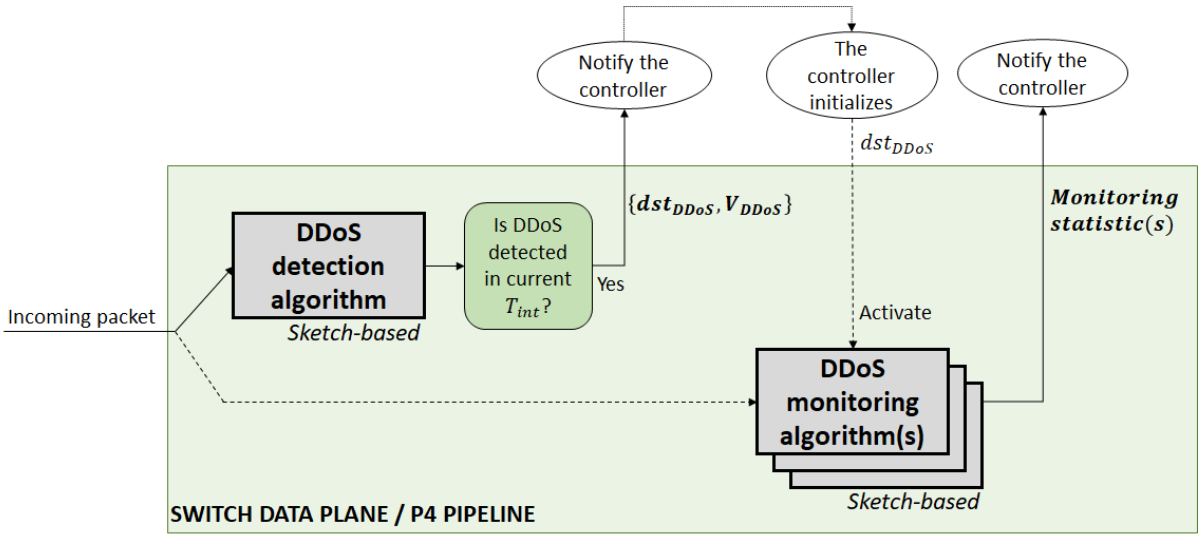

Figure 17: DDoS detection and DDoS monitoring workflow in a programmable data plane device

DDoS detection and monitoring service offered by this approach can provide immediate notifications when DDoS attacks start or disappear, together with frequently updated information about the characteristics of the DDoS traffic which can give insight about the phase of the attack.

3.2.2.2 *P4 implementation*

The DDoS detection and monitoring functionality is implemented with usage of the P4 switches. In the first phase, P4 behavioral model (BMv2) is used instead of white-box switches. BMv2 is a P4 software switch emulator which is broadly used for developing and testing P4 programs. Currently, DPP team is using a simple virtual network topology composed of three interconnected P4 switches. Each P4 switch is connected to a host. This topology (see Figure 18) is deployed on any developer machine with the usage of a Docker container called p4app which contains Mininet [Mininet] environment extended with BMv2 instead of default Open vSwitch [Open vSwitch]. When the Docker container is starting, the P4 program code is loaded into the P4 switches, sketch structures are allocated within each P4 switch and the network controller process starts. The DDoS traffic can be generated by the Python script activated after logging to a Mininet host.

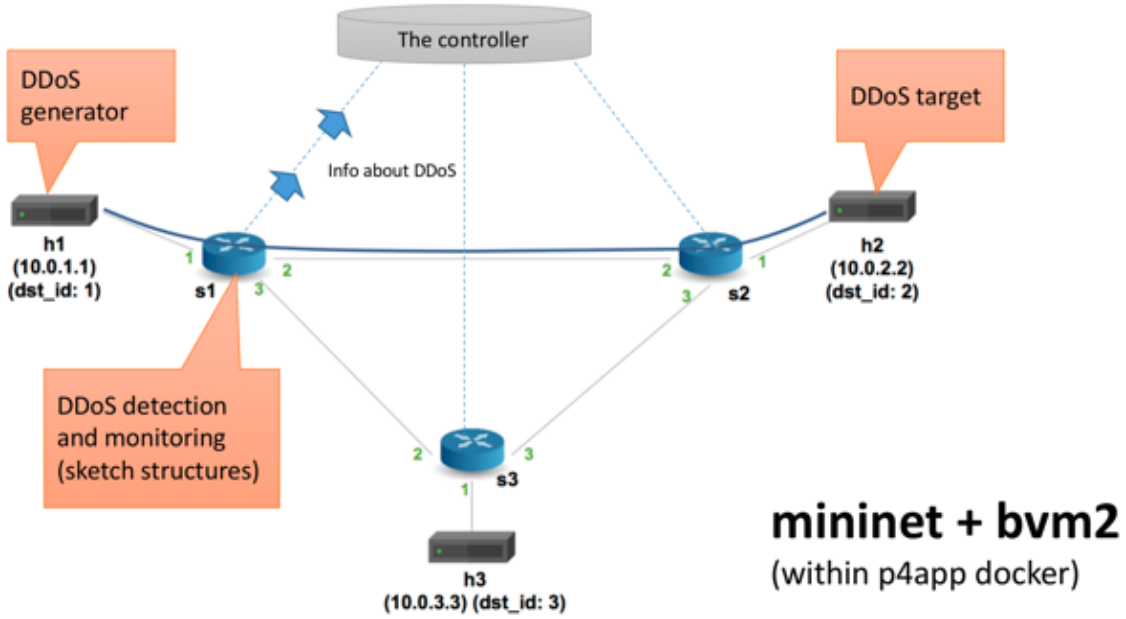

Figure 18: Virtual environment used for DDoS use case development

DPP team started with sketch structures implemented in $P4_{14}$ upgrading the prototype to $P4_{16}$. Apache Thrift is initially used and `simple_switch_CLI` as a way to communicate with the network controller and P4 switches. In the virtual testbed, the controller was very slow: ~0.1 sec for reading a few values from a register, ~0.3 sec for adding a table entry and ~0.1 sec for resetting a single register or table. However, it allowed to develop and test the code. Finally, the network traffic was observed for 3 seconds and then 5 seconds were required for performing all actions between the controller and switches. The prototype development also showed that for each new instance of the sketch structure a repetition of the same P4 code is needed. If P4 could support passing references or pointers to P4 register structures between the control functions then it would allow for writing sketch logic code only once. Another P4 limitation is the lack of logarithm operation and this is why the implementation of our HyperLogLog-based sketch is cumbersome in P4.
The virtual testbed is now being moved to two physical testbeds, one in PSNC Poland (based on two Arista 7170-32c) and the second in FBK Italy (based on three Edgecore Wedge100BF-32X connected in a triangle). The WP6T1 team tries to adapt the existing P4 code for the Tofino chipset which introduces a new set of P4 language constraints. P4 Runtime communication performance on Tofino should be much faster and it allows the WP6T1 team to set an observation time interval below 1 second, which in practice corresponds to how fast DDoS detection can be triggered. After successful validation of this DDoS prototype, P4 switches with the prototype code could be implemented at selected locations at the borders of PSNC and GARR production network domains. The switches based on Barefoot Tofino are equipped with 100Gbps interfaces and can easily handle traffic from different network operators or NRENs. However, P4 switch with uploaded P4 DDoS detection code could not replace currently existing switches because WP6T1 DDoS detection and monitoring prototype is not being integrated with any network protocol stack which can be found in most of the switches and routers (i.e.: VLANs, VXLANs, MPLS, IPv4/IPv6 routing, OSPF, ISIS, VPNs, etc). An alternative approach is that the P4 switch will receive traffic from one or many inter-domain links mirrored by already existing switches (it will be enough to copy only the first 64-bytes of each frame because only IP and UDP/TCP headers are read).

4 Router for Academia Research and Education (RARE)

4.1 Introduction

Router for Academic R&E (RARE) is a project that will assess the validity of control plane running (might be open source) on top of white label hardware in NREN context. The project proposes to find the optimal control plane coupled to the most suitable white labelled equipment. Beyond the technical aspect, the project aims to identify NRENs' use cases in order to verify the white box usage feasibility for the NRENs. The project aims to integrate different pieces of software related to these building blocks: control plane, data plane and communication between the control plane and data plane.

The NRENs cannot ignore this technological and economical breakthrough. More specifically, RARE expectations are:

- Being able to address specificities of education and research network features and services;
- Reduce Request For Improvement lead time development (for instance to test a new paradigm in network domain);
- Build a worldwide R&E community around "RARE project".

The Figure 19 shows a potential interest from the NREN community (Stockholm survey).

As can be seen, RARE is perceived as a bleeding edge project that still needs to have proven record of reliable hours of operation. The good news is that considering the survey results, the audience demonstrated a positive feedback.

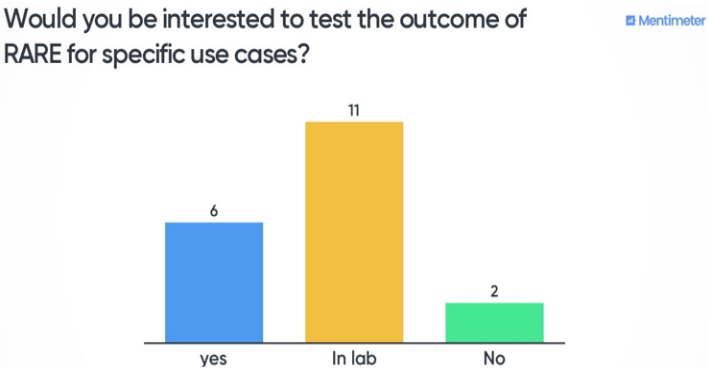

Figure 19: NREN survey results

RARE is an example that is combining the aforementioned work related to white box with the data plane programmability work, to create a router with all the functionalities needed for the academic, research and education community.

A key part of the work consists of enabling a control plane software to pilot a data plane via a programmatic interface. P4 is such a language proposing an interface that allows data plane programmability. P4 core language tries to be as much as possible independent from the target or NPU processor architecture, although architecture dependence is still prominent. For now, WP6T1 chooses to use Tofino Barefoot chipset that is available on different switches.

FreeRtr [http://freerouter.nop.hu/] is a first good control plane candidate, which has been used for years by KIFU, the Hungarian NREN. It is used as an operational route reflector but it can be used to

implement an LSR router and even LER. Several points will be investigated for the production stage: monitoring and security.

In parallel, different use cases for NRENs will be identified, especially the ones that could be used for starting the usage of these white boxes in production.

Use cases list has been elaborated following various community's feedback: from WP6-T1-WB cross-activity, from a specific WP6-T1 White Box event survey in Stockholm, from STF community.

RARE will consider the following use cases by order of implementation simplicity:

- Baseline feature set common to all use cases below: SSH transport for management, TACACS/RADIUS for management, infrastructure ACL in order to protect router interface also known as CoPP, LPTS, CP protection, monitoring capability providing link utilization and CPU counters if relevant.
- Service Provider grade P router (P function relates to the capacity of a router to only switch traffic at a high line rate): IPv4/IPv6 addressing, ISIS (or OSPF) IGP routing, MPLS/LDP, Segment Routing over [MPLS|IPv6]
- Telecom Service Provider grade IX switch: VLAN, BUM Storm control, etc.
- CLOUD Service Provider grade Spine/Leaf and Virtual Gateway router: VXLAN, EVPN, etc.
- Telecom Service Provider Edge grade PE router with a minimum feature set: L3VPN (IPv4 might be sufficient as a start?), L2VPN (evpn might be sufficient as a start?) Point to point, point to multipoint, multipoint to multipoint
- Performance validation: Full line rate performance, table size scalability performance.

The RARE team can validate the code using BAREFOOT bf_switchd virtual switch. So technically everything can be developed in this virtual environment. However, high scale data plane throughput can be tested only with real hardware. The BAREFOOT WEDGE-100BF-32X p4 switch is a hardware equipped with 32x100GE interfaces and powered by TOFINO NPU that can reach 6.4 Tbps. There will be two sites connected by GÉANT equipped with 2xEdgecore Wedge100BF-32X (6 x 40G QSFP adapter, 6 x 10G SFP). The location of the testbed is not known for now and the testbed is under procurement by WP7.

Additionally, Jisc and SWITCH will connect this testbed (see Figure 20) with one switch and RENATER will do the same with two switches, increasing the European testbed to eight machines spread over European NRENs.

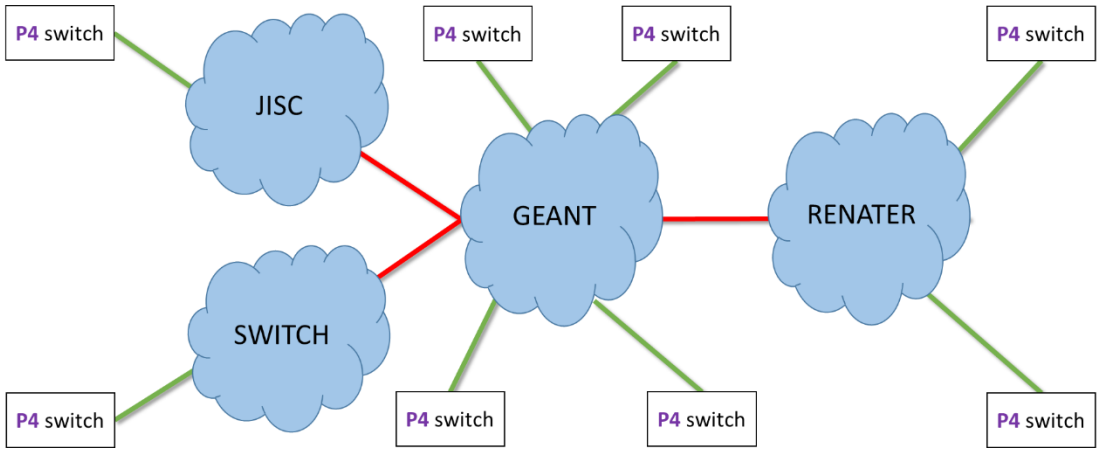

Figure 20: GN4-3 WP6T1 RARE European testbed

4.2 Work progress

Liaison has been established with P4lang P4.org organization with BAREFOOT by subscribing to the FASTER program [FASTER] dedicated to education and research. The project participants attended an official P4 training in Amsterdam. It helped the group to get a strong understanding of the TOFINO architecture and insight on how to move forward in order to reach the project objective.

Self-paced training was also important, as P4 is still a bleeding edge technology. For dissemination and learning purposes, RARE team uses P4lang public BMv2 software switch. In that respect a set of packages have been created and are daily build upon GitHub source code modification on Ubuntu 16.04 and Ubuntu 18.04.

The BAREFOOT switch uses ONL which is inherently a Debian stable distribution. Similar to the previous section a set of P4lang packages have been elaborated for Debian 10 (stable).

P4lang software has also been ported on NixOS which is an OS with the particularity to be minimal. Only the software needed to run is installed on the OS. This is a good OS for system appliance and a good candidate for a P4 generic software CPE, for example.

A RARE repository on GitHub is a public repository that has the particularity to use Open Source BMv2 P4 software switch. Essentially, this repository helped to team organize the work in a geographically distributed way and validate P4 algorithm, this public repository is also a valuable resource for P4 technology dissemination/training that the team might conduct within the R&E community.

RARE GÉANT BitBucket Git is a private git repository, with the particularity to contain RARE code for a proprietary switch image from BAREFOOT Company. This switch uses a silicon NPU called TOFINO. Essentially all the code from RARE GitHub and RARE GÉANT git will get the same structure but in this private git instead the target architecture is TOFINO bf_switchd.

Several unit labs can be created in this development environment in order to address all features and use cases for the NREN context. Figure 21 shows the topology that could be created in order to test an MPLS P router use case that requires a combination of IPv4 forwarding, IS-IS, MPLS.

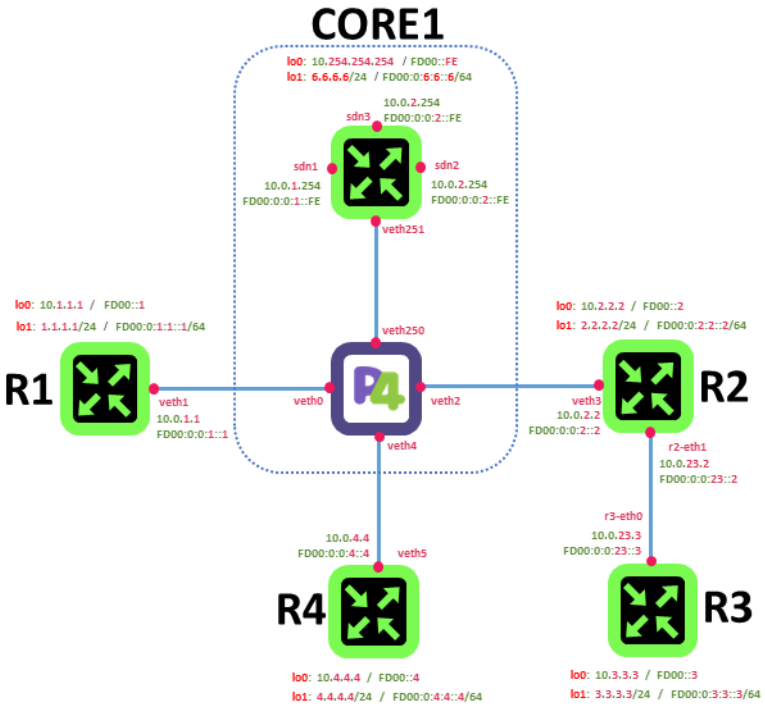

Figure 21: RARE lab topology

The integration made very good progress and now FreeRtR is able to pilot a P4 data plane. This breakthrough allows the RARE team to implement a lot of use cases like IPv4, IPv6, MPLS, SR-MPLS. The next step will be to implement it on Barefoot Switchd (virtual environment) and after that on real P4 hardware.

5 Conclusion

The NRENs cannot ignore the technological and economical breakthrough that white box and DPP represent. There is now a possibility to run different network operating systems (open or commercial) over commodity hardware; moreover it is now possible to implement a data plane allowing to create new features or to have an open source data plane.

White box for research and education demonstrate that white box can be used for CPE and Internet eXchange point switch. Even if the technical analysis is not finished, it is likely that white box can be used for DC fabric use case. The management decision is not only based on technical considerations and TCO but also on the workload of their team, their capacity to hire staff, the flexibility solution, the independence from a vendor, the strategic plan, etc. For use cases that require more routing features like LER/PE, the NOS available now could be limited. WP6T1 will continue working on all these use cases (Data centre, GIX, CPE, P/LSR) and will provide a technical analysis.

RARE project already demonstrated thanks to the development of unit labs (IPv4, MPLS, Segment Routing, …) that data plane routing features are feasible. RARE also demonstrated that the integration of a control plane (FreeRtr) on P4 data plane is feasible, even if there is a lot of work to finish the full integration. The next step is to integrate this on a virtual environment specific to a chipset (TOFINO) and then to implement it on P4 switch.

DDoS algorithms were implemented on a virtual P4 environment and the implementation on P4 hardware is ongoing. It will require adaptations as some functions used in the virtual environment are not available on the P4 hardware for its implementation on the P4 switch.

In-band Network Telemetry allows network monitoring and debugging in novel ways and can significantly improve network management, using just a few nodes supporting INT. A new implementation on $P4_{16}$ will be provided.

References

[Barefoot] https://www.barefootnetworks.com/products/brief-tofino/

[Benycze] https://github.com/benycze/PhD-Thesis

[Beheshti_et_al] Neda Beheshti, Emily Burmeister, Yashar Ganjali, John E. Bowers, Daniel J. Blumenthal, and Nick McKeown – Optical Packet Buffers for Backbone Internet Routers

[CESNET] 100G In-Band Network Telemetry with P4 and FPGA
Pavel Benáček, Viktor Puš (CESNET, a.l.e.); Michal Kekely, Lukáš Richter (Netcope Technologies a.s.); Pavel Minařík, Jan Pazdera (Flowmon Networks a.s.). https://p4.org/events/2017-05-09-p4-workshop/

[FASTER] https://forum.barefootnetworks.com/hc/en-us

[INT] https://p4.org/assets/INT-current-spec.pdf

[Jim_Warner] https://fr.slideserve.com/kasper-wolfe/speed-match-buffers

[liberouter] https://www.liberouter.org/combo-200g2ql/

[Merchant_Chips] A DEEP DIVE INTO CISCO'S USE OF MERCHANT SWITCH CHIPS, https://www.nextplatform.com/2018/06/20/a-deep-dive-into-ciscos-use-of-merchant-switch-chips/

[Mininet] http://mininet.org/

[Normandy_CPE_1] GN4-3 WP6T1 Wiki, https://wiki.geant.org/download/attachments/120497358/GN4Ph3_CPE_RENATER_20190314-V0.01.docx

[Normandy_CPE_2] Topology and test result, https://wiki.geant.org/display/gn43wp6/RENATER+CPE+use+case

[OCP] https://en.wikipedia.org/wiki/Open_Compute_Project

[Open_vSwitch] https://www.openvswitch.org/

[Packet_buffers] https://people.ucsc.edu/~warner/buffer.html

[Roy_et_al] Inside the Social Network's (Datacenter) Network – Arjun Roy, Hongyi Zeng, Jasmeet Bagga, George Porter, and Alex C. Snoeren

[SFINX] https://wiki.geant.org/pages/viewpage.action?pageId=126976348

[Wiki] GN4-3 WP6T1 Wiki https://wiki.geant.org/download/attachments/120497358/GN4Ph3_CPE_RENATER_20190314-V0.01.docx

Glossary

DPP	Data Plane Programming
FPGA	Field Programmable Gate Array
IX	Internet eXchange point
NOS	Network Operating System
P4	Programming Protocol-Independent Packet Processors - programming language
TCO	Total Cost Ownership
ToR	Top of rack switch

Appendices

White box for research and education

RFC2889 Congestion control

Frame Size	Burst Size	Intended Load (%)	Head of Line Blocking	Back-Pressure	Frame Size	Burst Size	Intended Load (%)	Head of Line Blocking	Back-Pressure
64	1	60	No	No	512	5	80	No	No
64	1	65	No	No	512	5	85	No	No
64	1	70	No	No	512	5	90	No	No
64	1	75	No	No	512	5	95	No	No
64	1	80	No	No	512	5	100	No	No
64	1	85	No	No	512	7	60	No	No
64	1	90	No	No	512	7	65	No	No
64	1	95	No	No	512	7	70	No	No
64	1	100	No	No	512	7	75	No	No
64	3	60	No	No	512	7	80	No	No
64	3	65	No	No	512	7	85	No	No
64	3	70	No	No	512	7	90	No	No
64	3	75	No	No	512	7	95	No	No
64	3	80	No	No	512	7	100	No	No
64	3	85	No	No	512	9	60	No	No
64	3	90	No	No	512	9	65	No	No
64	3	95	No	No	512	9	70	No	No
64	3	100	No	No	512	9	75	No	No
64	5	60	No	No	512	9	80	No	No
64	5	65	No	No	512	9	85	No	No

64	5	70	No	No	512	9	90	No	No
64	5	75	No	No	512	9	95	No	No
64	5	80	No	No	512	9	100	No	No
64	5	85	No	No	1024	1	60	No	No
64	5	90	No	No	1024	1	65	No	No
64	5	95	No	No	1024	1	70	No	No
64	5	100	No	No	1024	1	75	No	No
64	7	60	No	No	1024	1	80	No	No
64	7	65	No	No	1024	1	85	No	No
64	7	70	No	No	1024	1	90	No	No
64	7	75	No	No	1024	1	95	No	No
64	7	80	No	No	1024	1	100	No	No
64	7	85	No	No	1024	3	60	No	No
64	7	90	No	No	1024	3	65	No	No
64	7	95	No	No	1024	3	70	No	No
64	7	100	No	No	1024	3	75	No	No
64	9	60	No	No	1024	3	80	No	No
64	9	65	No	No	1024	3	85	No	No
64	9	70	No	No	1024	3	90	No	No
64	9	75	No	No	1024	3	95	No	No
64	9	80	No	No	1024	3	100	No	No
64	9	85	No	No	1024	5	60	No	No
64	9	90	No	No	1024	5	65	No	No
64	9	95	No	No	1024	5	70	No	No
64	9	100	No	No	1024	5	75	No	No
128	1	60	No	No	1024	5	80	No	No
128	1	65	No	No	1024	5	85	No	No
128	1	70	No	No	1024	5	90	No	No
128	1	75	No	No	1024	5	95	No	No
128	1	80	No	No	1024	5	100	No	No
128	1	85	No	No	1024	7	60	No	No
128	1	90	No	No	1024	7	65	No	No
128	1	95	No	No	1024	7	70	No	No
128	1	100	No	No	1024	7	75	No	No
128	3	60	No	No	1024	7	80	No	No
128	3	65	No	No	1024	7	85	No	No
128	3	70	No	No	1024	7	90	No	No
128	3	75	No	No	1024	7	95	No	No
128	3	80	No	No	1024	7	100	No	No
128	3	85	No	No	1024	9	60	No	No
128	3	90	No	No	1024	9	65	No	No
128	3	95	No	No	1024	9	70	No	No
128	3	100	No	No	1024	9	75	No	No
128	5	60	No	No	1024	9	80	No	No
128	5	65	No	No	1024	9	85	No	No
128	5	70	No	No	1024	9	90	No	No
128	5	75	No	No	1024	9	95	No	No
128	5	80	No	No	1024	9	100	No	No
128	5	85	No	No	1280	1	60	No	No
128	5	90	No	No	1280	1	65	No	No
128	5	95	No	No	1280	1	70	No	No
128	5	100	No	No	1280	1	75	No	No
128	7	60	No	No	1280	1	80	No	No
128	7	65	No	No	1280	1	85	No	No
128	7	70	No	No	1280	1	90	No	No
128	7	75	No	No	1280	1	95	No	No
128	7	80	No	No	1280	1	100	No	No
128	7	85	No	No	1280	3	60	No	No
128	7	90	No	No	1280	3	65	No	No
128	7	95	No	No	1280	3	70	No	No
128	7	100	No	No	1280	3	75	No	No
128	9	60	No	No	1280	3	80	No	No
128	9	65	No	No	1280	3	85	No	No
128	9	70	No	No	1280	3	90	No	No
128	9	75	No	No	1280	3	95	No	No
128	9	80	No	No	1280	3	100	No	No
128	9	85	No	No	1280	5	60	No	No
128	9	90	No	No	1280	5	65	No	No
128	9	95	No	No	1280	5	70	No	No
128	9	100	No	No	1280	5	75	No	No
256	1	60	No	No	1280	5	80	No	No
256	1	65	No	No	1280	5	85	No	No
256	1	70	No	No	1280	5	90	No	No
256	1	75	No	No	1280	5	95	No	No
256	1	80	No	No	1280	5	100	No	No
256	1	85	No	No	1280	7	60	No	No
256	1	90	No	No	1280	7	65	No	No
256	1	95	No	No	1280	7	70	No	No
256	1	100	No	No	1280	7	75	No	No
256	3	60	No	No	1280	7	80	No	No
256	3	65	No	No	1280	7	85	No	No

256	3	70	No	No	1280	7	90	No	No
256	3	75	No	No	1280	7	95	No	No
256	3	80	No	No	1280	7	100	No	No
256	3	85	No	No	1280	9	60	No	No
256	3	90	No	No	1280	9	65	No	No
256	3	95	No	No	1280	9	70	No	No
256	3	100	No	No	1280	9	75	No	No
256	5	60	No	No	1280	9	80	No	No
256	5	65	No	No	1280	9	85	No	No
256	5	70	No	No	1280	9	90	No	No
256	5	75	No	No	1280	9	95	No	No
256	5	80	No	No	1280	9	100	No	No
256	5	85	No	No	1518	1	60	No	No
256	5	90	No	No	1518	1	65	No	No
256	5	95	No	No	1518	1	70	No	No
256	5	100	No	No	1518	1	75	No	No
256	7	60	No	No	1518	1	80	No	No
256	7	65	No	No	1518	1	85	No	No
256	7	70	No	No	1518	1	90	No	No
256	7	75	No	No	1518	1	95	No	No
256	7	80	No	No	1518	1	100	No	No
256	7	85	No	No	1518	3	60	No	No
256	7	90	No	No	1518	3	65	No	No
256	7	95	No	No	1518	3	70	No	No
256	7	100	No	No	1518	3	75	No	No
256	9	60	No	No	1518	3	80	No	No
256	9	65	No	No	1518	3	85	No	No
256	9	70	No	No	1518	3	90	No	No
256	9	75	No	No	1518	3	95	No	No
256	9	80	No	No	1518	3	100	No	No
256	9	85	No	No	1518	5	60	No	No
256	9	90	No	No	1518	5	65	No	No
256	9	95	No	No	1518	5	70	No	No
256	9	100	No	No	1518	5	75	No	No
512	1	60	No	No	1518	5	80	No	No
512	1	65	No	No	1518	5	85	No	No
512	1	70	No	No	1518	5	90	No	No
512	1	75	No	No	1518	5	95	No	No
512	1	80	No	No	1518	5	100	No	No
512	1	85	No	No	1518	7	60	No	No
512	1	90	No	No	1518	7	65	No	No
512	1	95	No	No	1518	7	70	No	No
512	1	100	No	No	1518	7	75	No	No
512	3	60	No	No	1518	7	80	No	No
512	3	65	No	No	1518	7	85	No	No
512	3	70	No	No	1518	7	90	No	No
512	3	75	No	No	1518	7	95	No	No
512	3	80	No	No	1518	7	100	No	No
512	3	85	No	No	1518	9	60	No	No
512	3	90	No	No	1518	9	65	No	No
512	3	95	No	No	1518	9	70	No	No
512	3	100	No	No	1518	9	75	No	No
512	5	60	No	No	1518	9	80	No	No
512	5	65	No	No	1518	9	85	No	No
512	5	70	No	No	1518	9	90	No	No
512	5	75	No	No	1518	9	95	No	No
					1518	9	100	No	No

Table 4: RFC2889 Congestion control test results

Traffic bursts

iMIX Distribution	Duration (Seconds)	Load (%)	Burst	Tx Frames	Rx Frames	Rx Expected Frames	Rx Out of Sequence Frames	Tx Frame Rate (fps)	Frame Loss	Frame Loss (%)	Avg Latency (uSec)
Default	10	25	50000	164101488	164101488	164101488	0	16368398	0	0,00000	2,14600
Default	10	25	100000	164101300	164101300	164101300	0	16368398	0	0,00000	4,37700
Default	10	25	150000	164101642	164101642	164101642	0	16368398	0	0,00000	16,44600
Default	10	25	200000	164101095	164101095	164101095	0	16368398	0	0,00000	16,86200
Default	10	25	250000	164194995	164194995	164194995	0	16368398	0	0,00000	14,08300
Default	10	25	300000	164199921	164199921	164199921	0	16368398	0	0,00000	21,58100
Default	10	25	350000	164100548	164100548	164100548	0	16368398	0	0,00000	21,17900

Default	10	25	400000	164150470	164121913	164150470	0	16368398	28557	0,01740	38,68400
Default	10	25	450000	164197258	164117655	164197258	0	16368398	79603	0,04850	26,84300
Default	10	25	500000	164250366	164121390	164250366	0	16368398	128976	0,07850	30,28200
Default	10	25	550000	164197867	164018179	164197867	0	16368398	179688	0,10940	41,46500
Default	10	25	600000	164295526	164066935	164295526	0	16368398	228591	0,13910	39,50500
Default	10	25	650000	163807821	163529568	163807821	0	16368398	278253	0,16990	39,34100
Default	10	25	700000	164395509	164065910	164395509	0	16368398	329599	0,20050	42,03100
Default	10	25	750000	164394516	164014411	164394516	0	16368398	380105	0,23120	45,69500
Default	10	25	800000	164250220	163822137	164250220	0	16368398	428083	0,26060	51,15800
Default	10	25	850000	164000206	163520650	164000206	0	16368398	479556	0,29240	56,96000
Default	10	25	900000	164050194	163520624	164050194	0	16368398	529570	0,32280	55,57100
Default	10	25	950000	164396417	163815033	164396417	0	16368398	581384	0,35360	64,40900
Default	10	25	1000000	164350174	163721202	164350174	0	16368398	628972	0,38270	62,60500
Default	10	30	50000	196839065	196839065	196839065	0	19642076	0	0,00000	2,19300
Default	10	30	100000	196846787	196846787	196846787	0	19642076	0	0,00000	3,25100
Default	10	30	150000	196845980	196845980	196845980	0	19642076	0	0,00000	4,35100
Default	10	30	200000	196847934	196847934	196847934	0	19642076	0	0,00000	13,47900
Default	10	30	250000	196849498	196849498	196849498	0	19642076	0	0,00000	11,02100
Default	10	30	300000	196750788	196750788	196750788	0	19642076	0	0,00000	14,45000
Default	10	30	350000	196847121	196847121	196847121	0	19642076	0	0,00000	18,76600
Default	10	30	400000	196913736	196884850	196913736	0	19642076	28886	0,01470	20,91900
Default	10	30	450000	196849930	196771077	196849930	0	19642076	78853	0,04010	27,24200
Default	10	30	500000	196659603	196530621	196659603	0	19642076	128982	0,06560	35,00800
Default	10	30	550000	196724692	196546360	196724692	0	19642076	178332	0,09070	42,29500
Default	10	30	600000	196790273	196560460	196790273	0	19642076	229813	0,11680	35,09000
Default	10	30	650000	196849565	196571773	196849565	0	19642076	277792	0,14110	38,16400
Default	10	30	700000	196950304	196621174	196950304	0	19642076	329130	0,16710	35,53300
Default	10	30	750000	196700282	196320933	196700282	0	19642076	379349	0,19290	42,02200
Default	10	30	800000	197046648	196617131	197046648	0	19642076	429517	0,21800	40,40000
Default	10	30	850000	196849165	196368985	196849165	0	19642076	480180	0,24390	43,98500
Default	10	30	900000	196593094	196063186	196593094	0	19642076	529908	0,26950	45,46900
Default	10	30	950000	197100220	196520859	197100220	0	19642076	579361	0,29390	49,13600
Default	10	30	1000000	196650208	196018798	196650208	0	19642076	631410	0,32110	52,03400
Default	10	35	50000	229537298	229537298	229537298	0	22915756	0	0,00000	2,23900
Default	10	35	100000	229545347	229545347	229545347	0	22915756	0	0,00000	7,38200
Default	10	35	150000	229507689	229507689	229507689	0	22915756	0	0,00000	4,28700
Default	10	35	200000	229546270	229546270	229546270	0	22915756	0	0,00000	9,16500
Default	10	35	250000	229508145	229508145	229508145	0	22915756	0	0,00000	8,14900
Default	10	35	300000	229542986	229542986	229542986	0	22915756	0	0,00000	11,54300
Default	10	35	350000	229500766	229500766	229500766	0	22915756	0	0,00000	16,96300
Default	10	35	400000	229508339	229479266	229508339	0	22915756	29073	0,01270	17,80800
Default	10	35	450000	229508398	229430236	229508398	0	22915756	78162	0,03410	21,57100
Default	10	35	500000	229543770	229414266	229543770	0	22915756	129504	0,05640	22,17600
Default	10	35	550000	229500460	229321348	229500460	0	22915756	179112	0,07800	24,26700
Default	10	35	600000	229382692	229153920	229382692	0	22915756	228772	0,09970	35,40700
Default	10	35	650000	229668881	229390158	229668881	0	22915756	278723	0,12140	30,29400
Default	10	35	700000	229450354	229120475	229450354	0	22915756	329879	0,14380	30,70000
Default	10	35	750000	229508492	229129606	229508492	0	22915756	378886	0,16510	37,83200
Default	10	35	800000	229541074	229111009	229541074	0	22915756	430065	0,18740	35,95400
Default	10	35	850000	229600288	229121480	229600288	0	22915756	478808	0,20850	37,73500

Default	10	35	900000	229544475	229015535	229544475	0	22915756	528940	0,23040	40,68800
Default	10	35	950000	229500256	228920918	229500256	0	22915756	579338	0,25240	41,69900
Default	10	35	1000000	229382754	228751961	229382754	0	22915756	630793	0,27500	44,66000
Default	10	40	50000	262193010	262193010	262193010	0	26189436	0	0,00000	2,30900
Default	10	40	100000	262206199	262206199	262206199	0	26189436	0	0,00000	3,79600
Default	10	40	150000	262202623	262202623	262202623	0	26189436	0	0,00000	4,41300
Default	10	40	200000	262201897	262201897	262201897	0	26189436	0	0,00000	5,63100
Default	10	40	250000	262201312	262201312	262201312	0	26189436	0	0,00000	10,18600
Default	10	40	300000	262247487	262247487	262247487	0	26189436	0	0,00000	11,42400
Default	10	40	350000	262201476	262201476	262201476	0	26189436	0	0,00000	14,03800
Default	10	40	400000	262150750	262121575	262150750	0	26189436	29175	0,01110	15,94500
Default	10	40	450000	262161299	262082519	262161299	0	26189436	78780	0,03010	19,02900
Default	10	40	500000	262275876	262147033	262275876	0	26189436	128843	0,04910	20,93900
Default	10	40	550000	262252597	262073394	262252597	0	26189436	179203	0,06830	24,85600
Default	10	40	600000	262301326	262071727	262301326	0	26189436	229599	0,08750	23,48100
Default	10	40	650000	262200438	261921234	262200438	0	26189436	279204	0,10650	27,21200
Default	10	40	700000	262112931	261783359	262112931	0	26189436	329572	0,12570	29,17500
Default	10	40	750000	262299048	261920460	262299048	0	26189436	378588	0,14430	35,55600
Default	10	40	800000	262136796	261707224	262136796	0	26189436	429572	0,16390	32,22300
Default	10	40	850000	262161387	261681283	262161387	0	26189436	480104	0,18310	32,88000
Default	10	40	900000	262296644	261766301	262296644	0	26189436	530343	0,20220	34,73000
Default	10	40	950000	262064204	261484509	262064204	0	26189436	579695	0,22120	36,71300
Default	10	40	1000000	262202491	261571831	262202491	0	26189436	630660	0,24050	38,68900
Default	10	45	50000	294849125	294849125	294849125	0	29463116	0	0,00000	2,43800
Default	10	45	100000	294848328	294848328	294848328	0	29463116	0	0,00000	8,41200
Default	10	45	150000	294845844	294845844	294845844	0	29463116	0	0,00000	4,29700
Default	10	45	200000	294846800	294846800	294846800	0	29463116	0	0,00000	8,45400
Default	10	45	250000	294845662	294845662	294845662	0	29463116	0	0,00000	7,73200
Default	10	45	300000	294835538	294835538	294835538	0	29463116	0	0,00000	9,30300
Default	10	45	350000	294866891	294866891	294866891	0	29463116	0	0,00000	13,94700
Default	10	45	400000	294859688	294831361	294859688	0	29463116	28327	0,00960	22,10300
Default	10	45	450000	294814216	294734983	294814216	0	29463116	79233	0,02690	23,21900
Default	10	45	500000	294814255	294685207	294814255	0	29463116	129048	0,04380	18,22700
Default	10	45	550000	294835623	294656501	294835623	0	29463116	179122	0,06080	21,02300
Default	10	45	600000	294846473	294617583	294846473	0	29463116	228890	0,07760	24,42100
Default	10	45	650000	294814274	294535675	294814274	0	29463116	278599	0,09450	27,21800
Default	10	45	700000	294824699	294494190	294824699	0	29463116	330509	0,11210	25,53100
Default	10	45	750000	294793284	294414255	294793284	0	29463116	379029	0,12860	31,25000
Default	10	45	800000	294782893	294353420	294782893	0	29463116	429473	0,14570	29,42000
Default	10	45	850000	294887566	294408019	294887566	0	29463116	479547	0,16260	33,54100
Default	10	45	900000	294917568	294387412	294917568	0	29463116	530156	0,17980	32,79200
Default	10	45	950000	294814265	294233322	294814265	0	29463116	580943	0,19710	34,92400
Default	10	45	1000000	294876621	294247064	294876621	0	29463116	629557	0,21350	34,89800
Default	10	50	50000	327451017	326319690	327451017	0	32736796	1131327	0,34550	8775,83300
Default	10	50	100000	327465288	326286457	327465288	0	32736796	1178831	0,36000	8175,67500
Default	10	50	150000	327457928	326256249	327457928	0	32736796	1201679	0,36700	7559,59700
Default	10	50	200000	327466168	326238793	327466168	0	32736796	1227375	0,37480	6935,72200
Default	10	50	250000	327466144	326211390	327466144	0	32736796	1254754	0,38320	6305,14800
Default	10	50	300000	327465279	326187016	327465279	0	32736796	1278263	0,39040	5659,52300
Default	10	50	350000	327466160	326157994	327466160	0	32736796	1308166	0,39950	5002,06100

Default	10	50	400000	327466145	326126182	327466145	0	32736796	1339963	0,40920	4673,09700
Default	10	50	450000	327451389	326053729	327451389	0	32736796	1397660	0,42680	4613,80100
Default	10	50	500000	327466155	326026481	327466155	0	32736796	1439674	0,43960	4645,17100
Default	10	50	550000	327450815	325958786	327450815	0	32736796	1492029	0,45560	4626,77000
Default	10	50	600000	327466485	325925899	327466485	0	32736796	1540586	0,47050	4638,51600
Default	10	50	650000	327466151	325867605	327466151	0	32736796	1598546	0,48820	4608,51700
Default	10	50	700000	327466156	325822125	327466156	0	32736796	1644031	0,50200	4655,03500
Default	10	50	750000	327466161	325770815	327466161	0	32736796	1695346	0,51770	4637,43000
Default	10	50	800000	327454364	325715450	327454364	0	32736796	1738914	0,53100	4659,79200
Default	10	50	850000	327467163	325681087	327467163	0	32736796	1786076	0,54540	4668,19200
Default	10	50	900000	327467609	325629322	327467609	0	32736796	1838287	0,56140	4638,45200
Default	10	50	950000	327466155	325575600	327466155	0	32736796	1890555	0,57730	4640,73300
Default	10	50	1000000	327453358	325517622	327453358	0	32736796	1935736	0,59110	4646,89400
Default	10	55	50000	360044665	326752099	360044665	0	36010474	33292566	9,24680	9794,51100
Default	10	55	100000	360030787	326219494	360030787	0	36010474	33811293	9,39120	9243,45700
Default	10	55	150000	360028227	326193668	360028227	0	36010474	33834559	9,39780	8692,39800
Default	10	55	200000	360031340	326178605	360031340	0	36010474	33852735	9,40270	8141,44100
Default	10	55	250000	360031458	326171470	360031458	0	36010474	33859988	9,40470	7590,44100
Default	10	55	300000	360038964	326149178	360038964	0	36010474	33889786	9,41280	7039,34800
Default	10	55	350000	360034172	326130388	360034172	0	36010474	33903784	9,41680	6488,59600
Default	10	55	400000	360015371	326067088	360015371	0	36010474	33948283	9,42970	5933,12400
Default	10	55	450000	360033241	326074055	360033241	0	36010474	33959186	9,43220	5413,36300
Default	10	55	500000	360029663	326036569	360029663	0	36010474	33993094	9,44170	5381,21600
Default	10	55	550000	360036527	326011528	360036527	0	36010474	34024999	9,45040	5411,33400
Default	10	55	600000	360004803	325949206	360004803	0	36010474	34055597	9,45980	5397,20300
Default	10	55	650000	360033230	325915299	360033230	0	36010474	34117931	9,47630	5393,65000
Default	10	55	700000	360054469	325894947	360054469	0	36010474	34159522	9,48730	5396,52000
Default	10	55	750000	360011953	325839280	360011953	0	36010474	34172673	9,49210	5386,85800
Default	10	55	800000	360025787	325771655	360025787	0	36010474	34254132	9,51440	5338,41900
Default	10	55	850000	360034822	325755656	360034822	0	36010474	34279166	9,52110	5396,36900
Default	10	55	900000	360083160	325723176	360083160	0	36010474	34359984	9,54220	5386,35300
Default	10	55	950000	360030605	325661928	360030605	0	36010474	34368677	9,54600	5376,32400
Default	10	55	1000000	360038900	325679145	360038900	0	36010474	34359755	9,54330	5401,97200

Table 5: Traffic burst test results