SDN-based IP and Layer 2 services with Open Networking Operating System in the GÉANT Service Provider Network

Pier Luigi Ventre - GARR / University of Rome Tor Vergata
Stefano Salsano - University of Rome Tor Vergata / CNIT
Matteo Gerola, Elio Salvadori - Create Net
Mian Usman, Sebastiano Buscaglione – GÉANT
Luca Prete, Jonathan Hart, William Snow – ON.LAB

Abstract

The migration of Service Providers' wide area IP networks towards SDN is a challenging task. In this paper, we consider the critical requirements of GÉANT, the 500Gbps pan-European provider interconnecting 38 National Research and Educational Networks (NRENs), for a total of 50 million users. A long term evolution path towards the *softwarization* of GÉANT is discussed, consisting of four steps to be realized in the future years, from providing SDN-based connectivity, to the so-called Software Defined Infrastructure (SDI). As a first step, the *softwarization* of some basic services currently offered by GÉANT is considered: *GÉANT IP*, *GÉANT plus* and *GÉANT open*. This paper reports the concrete experience in the SDN-based design and implementation of these services, which have been called L3-SDX and L2-SDX. Both use cases have been addressed with the use of the open source Open Network Operating System (ONOS®). The L3-SDX service has been developed on top of an existing ONOS application, called SDN-IP. SDN-IP allows interconnections between SDN and legacy networks through BGP. The L2-SDX service has been realized as a new ONOS application. Both services are currently deployed on the GÉANT Testbed Service (GTS), a continental facility offering geographical virtual testbeds to the research community. The article reports the experience gained from this experimental deployment and discusses the benefits for a Service Provider like GÉANT.

Introduction

Software Defined Networking (SDN) is a recent paradigm [1] potentially able to transform the design of both datacenter and wide area networks. The promise of SDN is to foster innovation and flexibility thanks to centralized network control and standard interfaces. The fundamentals of the SDN approach are: i) the separation of control and data planes, ii) the logical centralization of the former as a software layer called Controller or Network Operating System (NOS), iii) the introduction of a flexible forwarding paradigm (based on filtering matches and actions) and iv) the direct control of the hardware through common management interfaces (e.g. OpenFlow). SDN can be seen as a part of an even wider trend towards the *softwarization* of networks [2][3], which implies a complete rethinking of how Service Provider networks are now structured. It is expected that this process will greatly increase the flexibility and efficiency of networks, reducing equipment and operational costs.

In this paper, we start from the analysis of current services offered by GÉANT, the 500Gbps pan-European network interconnecting 38 National Research and Educational Networks (NRENs), for a total of 50 million users. We refer to the NRENs as the *customers* of GÉANT. We identify the GÉANT needs and requirements toward the upgrading of its infrastructure. A long term evolution path towards the *softwarization* of GÉANT is discussed, consisting of several steps to be realized in the future years, from providing SDN-based connectivity services to the so-called Software Defined

Infrastructure (SDI) [4][5], which is also able to dynamically offer a wide range of computing / storage / network resources.

The first step of the GÉANT migration process consists of the *SDNization* of some operational services. In particular, we consider *GÉANT IP*, that is the basic service providing Internet connectivity to the NRENs, and two Layer-2 connectivity services called *GÉANT plus* and *GÉANT open*. These services are currently delivered through 26 Point of Presences (PoPs) located in Europe and 2 Open eXchange Points (OXP) in London and Paris (see next section for further details). The OXPs are similar to the standard Internet eXchange Points (IXPs), allowing NRENs to exchange traffic with external (non-*GÉANT*) networks. The introduction of SDN technologies in an IXP is referred to as SDX (Software Defined internet eXchange) [6]. In this paper, we give a wider meaning to the SDX concept, extending its potential applicability not only to the exchange points (IXPs or OXPs) but also to the PoPs. We designed and developed two SDN-based services called L3-SDX and L2-SDX (L3 and L2 stands for Layer 3 and Layer 2). These services represent the fundamental building blocks of the augmented SDX mentioned above and have been developed using ONOS [7], a promising open source solution for SDN control plane.

GÉANT network and services

As one of the largest and most complex Research and Education Networks in the world, GÉANT needs to support different services, such as standard IP transit connectivity or ultra-high capacity data center interconnections. The GÉANT infrastructure offers extensive links to networks in other world regions. External peers (other NRENs and external Autonomous Systems) are interconnected through 26 POPs, located all over Europe, and two Open eXchange Points (OXPs) (Figure 1). GÉANT offers a wide range of *connectivity and network management services* as described in [8]. We focus on a subset of them, called *GÉANT IP*, *GÉANT plus* and *GÉANT Open*.

GÉANT IP provides IP transit services to interconnect participating NRENs together and with other approved research organization and providers. It provides a peering service for IP traffic, isolated from general-use Internet access. *GÉANT plus* offers the NRENs point-to-point Layer 2 (Ethernet) circuits among end points at GEANT PoPs [8]. The PoPs constitute the backbone of the dual (optical transmission and packet) layer network through which GÉANT supplies connectivity to its customers. From the GÉANT perspective, the PoPs are the end-points of *GÉANT IP* and *GÉANT plus*. Finally, with the *GÉANT Open* service, NRENs can connect with external (non-*GÉANT*) networks through the OXPs. Inside an OXP, the customers (NRENs or external participants) request the establishment of Layer 2 circuits between end-points, which are manually provisioned through VLAN tunnels. The customers can use the Layer 2 circuits for whatever reason, including private BGP peering. Therefore, OXPs are different from traditional IXPs, which provide a switched layer 2 infrastructure used by multiple participants to exchange traffic through public BGP peering.

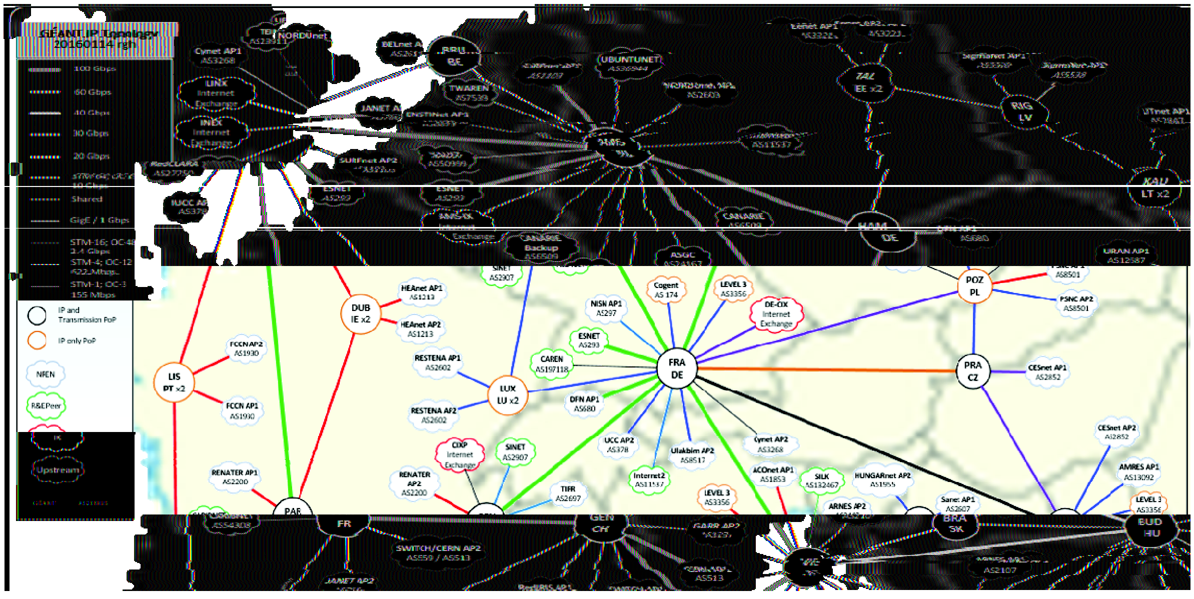
100 Gbps
60 Gbps
40 Gbps
30 Gbps
20 Gbps
Shared
GigE / 1 Gbps
IP and Transmission PoP
IP only PoP
NREN
R&EPeer
Upstream
Cynet AP1 AS3268
NORDUnet
UBUNTUNET
TWAREN AS7539
LINX Internet Exchange
INEX Internet Exchange
TAL EE x2
RIG LV
HAM DE
DFN AP1 AS680
URAN AP1 AS12687
HEAnet AP1 AS1213
HEAnet AP2 AS1213
DUB IE x2
FCCN AP2 AS1930
LIS PT x2
FCCN AP1 AS1930
RESTENA AP1 AS2602
LUX LU x2
RESTENA AP2 AS2602
CAREN AS197118
ESNET AS293
NISN AP1 AS297
Cogent AS 174
LEVEL 3 AS3356
DE-CIX Internet Exchange
FRA DE
POZ PL
PSNC AP2 AS8501
PRA CZ
CESnet AP1 AS2852
DFN AP1 AS680
UCC AP2 AS378
Ulakbim AP2 AS8517
Cynet AP2 AS3268
ACOnet AP1 AS1853
Internet2 AS11537
CIXP Internet Exchange
SINET AS2907
TIFR AS2697
RENATER AP1 AS2200
RENATER AP2 AS2200
LEVEL 3 AS3356
SILK AS132467
HUNGARnet AP2 AS1955
ARNES AP2
Sanet AP1 AS2607
CESnet AP2 AS2852
AMRES AP1 AS13092
GEN CH
BRA SK
SWITCH/CERN AP2 AS559 / AS513

GÉANT Softwarization path

The migration of GÉANT network and its services towards a Software Defined Infrastructure is a challenging task, which cannot happen overnight. We decomposed the transition of GÉANT to the SDN paradigm in incremental steps as shown in Figure 2. Each step enhances GÉANT infrastructure making it more sustainable, manageable, less expensive and introduces new services/functionalities to the portfolio. The transition path also takes into account the operational requirements of a production network. This migration strategy rationalizes and extends some ideas already presented by Monga in [5]. The idea behind the strategy is to initially introduce the concept of SDX, replacing the current architecture based on PoPs and OXPs, and then to progressively enhance its functionality. The first step of the path is the realization of the L3–SDX and L2–SDX, referred to as *pure connectivity* SDX [5][10]. L3-SDX supports *GÉANT IP*, while L2-SDX represents the *SDNization* of *GÉANT plus* and *GÉANT Open*. With this first step, there is no full migration to SDN technology, because NRENs can run legacy technologies without direct interaction with the GÉANT SDX operations.

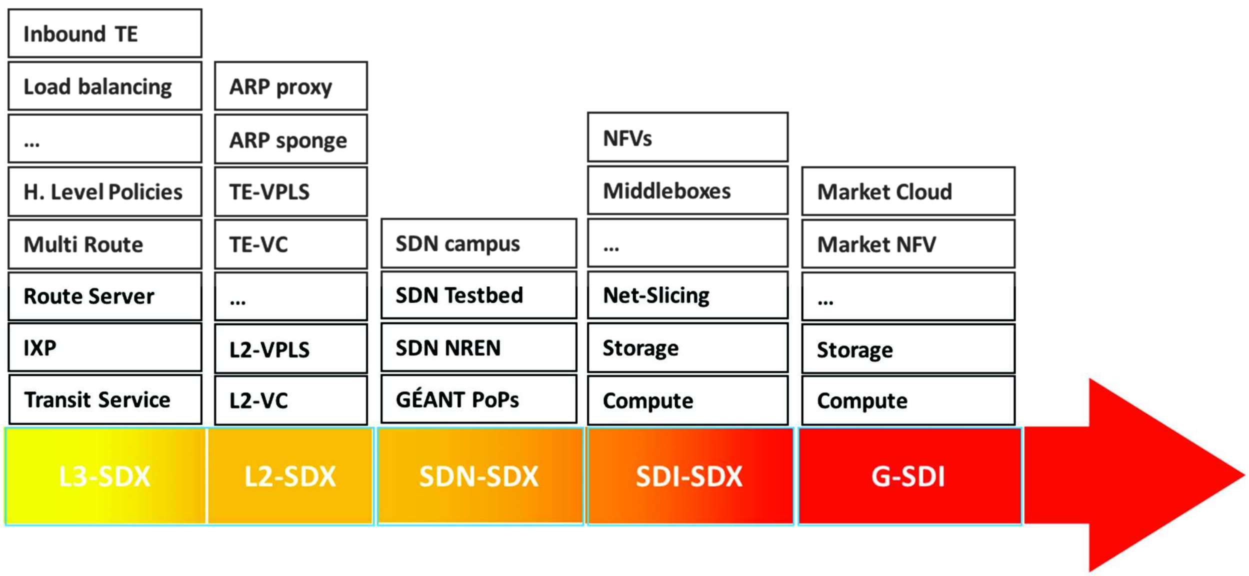

Figure 2 - Softwarization path

Assuming that the NRENs will setup their own SDN infrastructure, the next step, called SDN–SDX consists in the interconnection and harmonization of the SDN infrastructures between GÉANT and the NRENs. To understand the advantages of this step, consider that the NRENs have their customers (research organizations) which requires connectivity toward other research organization in the same NRENs, or in remote NRENs or outside the GÉANT's NRENs. Thanks to the SDN–SDX, the NRENs and the research organizations could fully exploit the advantages of SDN paradigm, leveraging end-to-end SDN based services spanning the GÉANT backbone and the NRENs.

The next step, denoted as SDI–SDX refers to a GÉANT SDN infrastructure augmented with cloud resources: a NREN can request not only networking services but also compute and storage resources, outsourcing part of its computations to the SDX cloud.

The final step is referred to as G-SDI, for Global SDI. It foresees a wider adoption of the full *softwarization* (SDN augmented with storage and computing resources) by all NRENs. Following this approach, an end-user (research organization) can obtain compute and storage resources from other NRENs or from GÉANT leveraging the resources offered through a logically global GÉANT SDX. At this step, considering GÉANT's position in the European scenario, it will be possible to exploit the Buyya vision of a market-oriented cloud computing [11]. GÉANT's role fits well as a "super party" that manages the market. This evolution of the GÉANT infrastructure encompasses an economic model, useful for the auto–sustainability of the GÉANT project and its participants (the NRENs themselves).

GÉANT Requirements for SDN-based IP and Layer 2 services

L3-SDX and L2-SDX have been identified as the first step of the GÉANT *softwarization* path. We carried out a thorough analysis of the requirements on these services from the perspective of the GÉANT service provider, summarized in Table 1. The requirements are classified as Functional, Non-functional (i.e. referring to performance or reliability), or Operational (e.g. related to monitoring or logging).

Based on this requirement analysis, we selected ONOS [7] as the controller platform. In particular, the existing SDN-IP ONOS application provided a very good fit with the functional requirement of L3-SDX, and the ONOS resilience and distribution features provided a good match with the identified Non-functional requirements.

Requirement	Service	Type	Priority	Status
L2 virtual circuit between two edge ports or VLANs	L2-SDX	Functional	Must	Completed
MPLS encapsulation of L2-SDX circuits	L2-SDX	Functional	Must	Completed(1)
VLAN and Stacked-VLANs (802.3ad) encapsulation	L2-SDX	Functional	Must	Completed(2)
IP transport between BGP peers	L3-SDX	Functional	Must	Completed
Custom route selection process	L3-SDX	Functional	Desirable	Planned
IPv6 support	BOTH	Functional	Must	Completed
Control plane resiliency	BOTH	Non-functional	Must	Completed
Control plane failure recovery	BOTH	Non-functional	Must	Completed
Network status after control plane failure	BOTH	Non-functional	Must	Completed
BGP control plane resiliency	L3-SDX	Non-functional	Must	Completed
Traffic rerouting after data plane failures	BOTH	Non-functional	Must	Completed
Control BGP attributes for each BGP peer	L3-SDX	Functional	Desirable	Not Needed(3)
Apply separate policies for each BGP peer	L3-SDX	Functional	Desirable	Not Needed(3)
Add, remove or shutdown BGP peers without impacts	L3-SDX	Functional	Must	Completed
Scale up to 100 BGP peers	L3-SDX	Non-functional	Desirable	Planned
Scale up to 100K routes	L3-SDX	Non-functional	Must	Planned(4)
Data plane statistic collection	BOTH	Operational	Desirable	Completed
Export of statistics to standard NMSs (SNMP, IPFIX)	BOTH	Operational	Optional	Ongoing(5)
Logging facilities	BOTH	Operational	Must	Completed
(1) Not fully supported in all switches (2) Stacked-VLANs not fully supported in the switches (3) Realized in the BGP peer (4) Tested up to 15K routes (5) SNMP not supported				

Table 1 - GEANT SDX requirements

ONOS: SDN Network Operating System for Service Providers

ONOS is an open source SDN control plane platform, meeting Service Provider requirements, released in 2014 as an open source project by ON.Lab. ONOS provides a stable implementation of a scalable, highly available and resilient Network Operating System (NOS).

The overall system has been conceived as a distributed system in the form of a cluster, composed of multiple instances, all functionally identical to each other. The architecture (Figure 3) can be structured in three tiers: a protocol-aware SouthBound (SB) layer, a protocol-agnostic distributed core layer, and an application layer. Each tier is a collection of pluggable modules/subsystems realizing specific functionalities that make up the ONOS platform. An API is exposed at each tier, providing isolation and modularity.

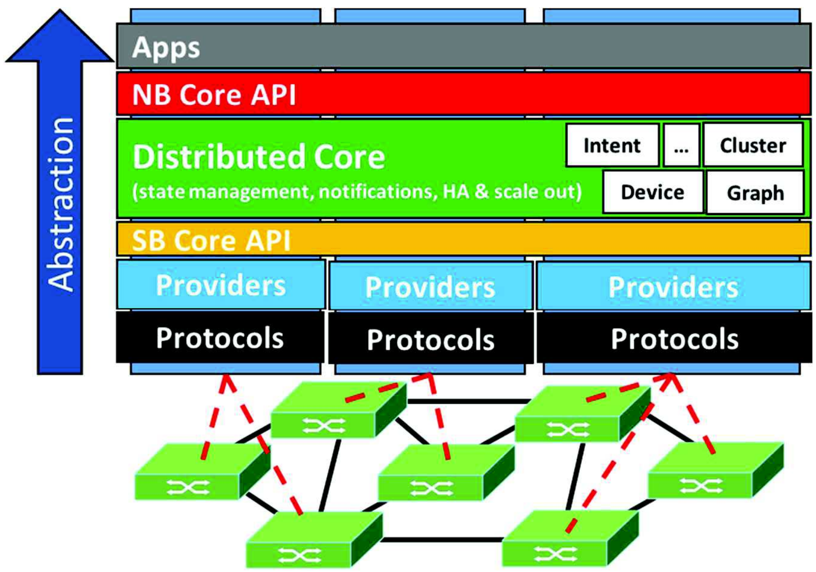

Figure 3 - ONOS Architecture

The distributed core is responsible for synchronization and coordination between the instances in the cluster. It builds a global network view based on information learned on the SouthBound API and offers services to the application layer. In order to achieve scalability and provide resiliency, various distribution mechanisms are available through a set of primitives. Each core subsystem uses these primitives in different ways according to the consistency requirements of the state it is managing. On top of the distributed state, a logically centralized network view is constructed and presented to applications. In addition, work is partitioned amongst the instances in the cluster. For example, each instance is elected to be responsible for managing a subset of the devices in the network, while the other instances are ready to step in if the primary instance fails. In case of data plane failures, built-in mechanisms for traffic rerouting are activated.

The SouthBound layer consists of a collection of software modules called ‘providers’, which interact with data plane devices using different southbound protocols. Providers gather information about network state and pass it to the distributed core, and receive instructions from the core to program the devices.

On the NorthBound side, ONOS presents abstractions to the applications, including Network Topology, Flow Objectives and Intents. Intents provide applications with a network-centric programming abstraction that allows developers to program the network through the usage of high-level policies that capture *what* needs to be done, rather than *how* to do it. The Intent framework determines how to implement an intent based on what other policies are in the system, and abstracts low-level details of this implementation. Intents make network policies configuration easier, speed up management procedures and tend to reduce the occurrence of configuration errors. Intents are backed by a dedicated subsystem that: i) translates Intents into device instructions; ii) coordinates and ensures the installation of the generated instructions; iii) reacts to network changes and modifies paths accordingly; iv) permits optimization across intents translations. The Intent framework has been widely used for developing the L3-SDX and L2-SDX applications.

ONOS is supported by an active open source community. Different ONOS applications have been developed over the years by ON.Lab and by the community as listed in the documentation of the project. For example, the SDN–IP application allows SDN islands to seamlessly interconnect with external networks using the standard BGP protocol. Among the applications, we mention CORD™ (Central Office Re-Architected as Datacenter)[12]. It aims to revolutionize the way Service Provider Central Offices are built and operated. It brings in the principles of SDN, NFV, cloud technologies and disaggregation, thereby making the Central Offices more manageable and agile.

High level architecture for GÉANT SDX

The proposed SDX architecture is based on SDN enabled networking equipment, controlled by a cluster of ONOS controllers. The ISP services, such as L3-SDX and L2-SDX, are designed as NorthBound applications running simultaneously on top of the NOS, offering both Layer 2 and Layer 3 connectivity services. Coexistence of different services in the data plane can be enforced through slicing mechanisms (e.g. VLAN tagging). As for the networking equipment, their integration is possible through open APIs, like OpenFlow, or by vendor-specific APIs implemented in ONOS in the form of pluggable drivers. The use of the so-called *white box* devices is currently under investigation and testing in GÉANT as they could replace traditional equipment to achieve relevant cost savings.

A SDX can span a single location (e.g. replacing an OXP or PoP) or multiple locations (e.g. federating PoPs in a single logical PoP, or creating a distributed OXP). This issue is further discussed in the section about the practical experience. The L3-SDX has been developed on top of an existing ONOS application, called SDN-IP, while L2-SDX has been realized as a new ONOS application.

L3-SDX/SDN-IP in GÉANT

SDN-enabled networks still need to interoperate with traditional networks on the Internet. The ONOS SDN-IP application interconnects an SDN island with external networks leveraging the BGP protocol. The solution allows: a) external ASs to exchange routes and transit traffic through an SDN network; b) the SDN network to advertise routes to the external networks; c) a Service Provider to scale its SDN control plane by segmenting an AS into multiple SDN domains, which communicate through BGP. Besides the technical advantages, the Service Providers also gain benefits in reduced CAPEX and OPEX, since they can use a single set of devices to manage (possibly L0/L1), Layer 2 and Layer 3 connectivity.

The high-level architecture of SDN-IP is shown in Figure 4. The SDN network is composed of different data plane devices controlled by ONOS, which are directly connected to the BGP-speaking border routers of the external ASs. Finally, one or more internal BGP speakers peer with the external routers and act as bridges between the external domains and the SDN-IP application. From the legacy networks perspective, the SDN domain appears as a standalone AS, as though it was running legacy BGP routers at the edges. Within the SDN network, SDN-IP has two main roles. The first is to install flows for BGP traffic between the external routers and the internal BGP speakers, thus allowing BGP sessions to be established. The second is to translate received routes into ONOS intents, which are compiled down into flows on the SDN switches. In order to transport the data traffic in the SDN network, ONOS makes use of *multi-point to single-point* tunnels, avoiding the use of n x n-1 tunnels to connect the endpoints, thus reducing the flow table entries in the data plane.

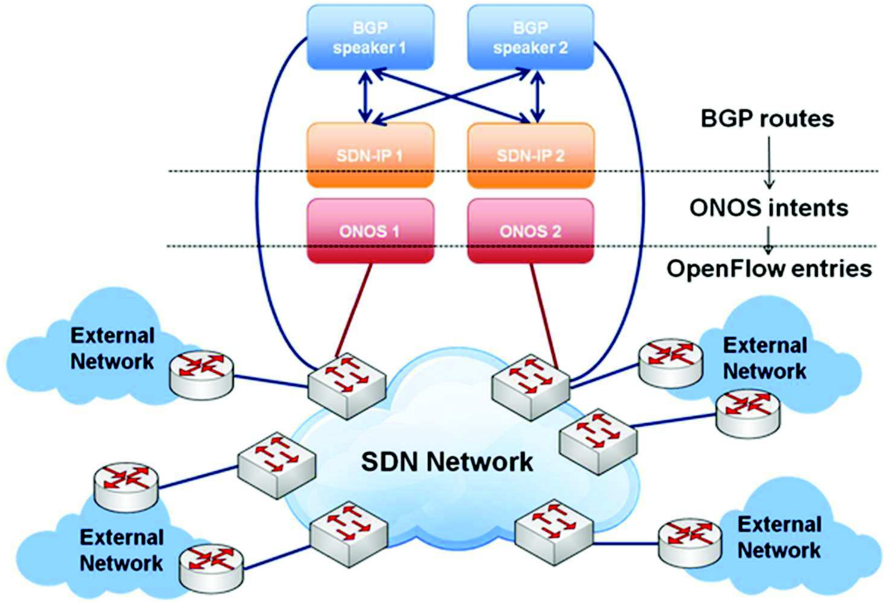

Figure 4 - High-level representation of the SDN-IP architecture

SDN-IP provides a feasible migration path towards the *softwarization* of ISP networks. It can be integrated with networks that already use BGP both externally and internally. From an operational point of view, SDN-IP guarantees flexibility in the covered use case, as it does not make any assumptions on the deployment scenario. The application can run on one or multiple ONOS instances. Moreover, the BGP settings can be changed dynamically with the addition or removal of peers. SDN-IP provides HA within the application itself: the service keeps working seamlessly, as long as there is at least one instance of the SDN-IP application running. In addition, SDN-IP leverages the HA mechanisms provided by ONOS for maintaining a consistent forwarding state in the data plane. SDN-IP provides a scalable solution able to control large-scale of SDN networks by using BGP-based confederations and ONOS clusters of different sizes.

L3-SDX extends SDN-IP adding the support for new deployment scenarios and providing facilities to monitor the BGP and transit traffic in the network. L3-SDX improves the flexibility of SDN-IP, making it is possible to deploy multiple peers belonging to the same AS and interconnected through different connection points controlled by ONOS. The application supports the typical IXP scenario where all the BGP routers as well as the Route Server belong to the same subnet [6]. An integration with ONOS IPFIX application allows exporting the counters related to the BGP sessions and to the L3 tunnels using the standardized IPFIX protocol. This can be used to realize advanced monitoring tools. L3-SDX and SDN-IP are both available under a liberal open source license.

L2-SDX service in GÉANT

L2-SDX is an ONOS application that allows the automated provisioning of layer 2 tunnels between endpoints, which can be physical Ethernet interfaces or VLANs. The offered layer 2 services belong to the class of IP Virtual Leased Line services (IP VLL) or Virtual Private LAN services (VPLS), which are a fundamental part of the service portfolio offered by large-scale ISPs. At the time of writing, only IP VLL has been integrated in the L2-SDX application, but VPLS can be provided with a straightforward extension of the current implementation. From a customer perspective, the L2-SDX appears as a black box that transports his traffic from the source to the destination end-point, as if they were in the same Ethernet LAN. Inside the SDX infrastructure, the L2-SDX application provides the necessary mechanisms for the service provisioning and monitoring. The human operators can manage and monitor the application through a CLI and a GUI, which accepts high-level customer's requests and translates them into ONOS *point-to-point intents*.

The application is fully integrated in ONOS and implemented as callable service. In a next release, its services will be exposed through a REST API, allowing the integration with orchestration platforms. Monitoring is achieved through the ONOS IPFIX application. L2-SDX can run over a single ONOS instance or on a cluster of ONOS instances that share a common state information. The multi-instance deployment is useful to control large-scale SDXs made up of many SDN devices. L2-SDX leverages the high availability mechanisms provided by ONOS in order to maintain a consistent state both in the control and data planes. Failures in the control plane are managed through the redundancy of the ONOS cluster. Instead, data plane failures are automatically resolved, through transparent re-routing mechanisms provided by ONOS.

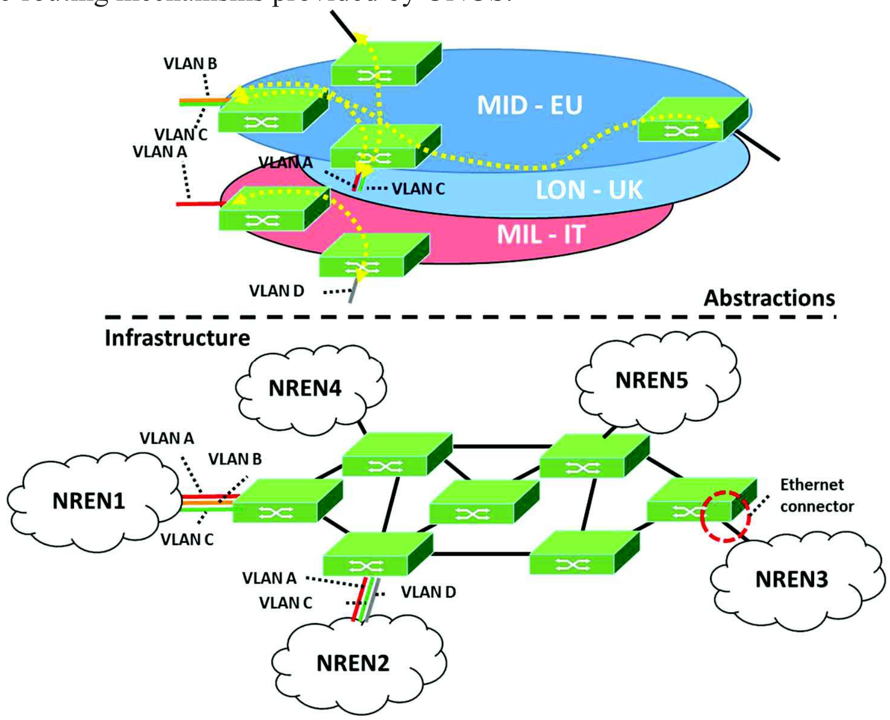

Figure 5 - L2-SDX application and its abstractions

L2-SDX provides users with powerful APIs and abstractions as shown in Figure 5. A virtual SDX (e.g. MID-EU, LON-UK in the figure) contains a number of end-points modeled as edge connectors, which can be interconnected through virtual circuits. Customers manage only the edges of the SDN network controlled by ONOS. The L2-SDX application eases service management and provisioning, e.g. enforcing isolation and avoiding several types of conflicts: i) the resources (ports or VLAN tags) associated with a connector cannot be reused, ii) an edge connector can only be used in a single circuit and iii) a connector in a virtual SDX instance cannot be interconnected with a connector in another virtual SDX. L2-SDX is available under a liberal open source license.

Proof of Concepts and worldwide experimental deployment

We realized two Proof of Concepts (PoCs). The first PoC has been deployed in a laboratory at University of Rome Tor Vergata and used mainly for validation and testing. It is based on Virtual Machines (VMs) and Open vSwitch (9 VMs emulate the data plane, a cluster of 3 VMs composes the ONOS control plane).

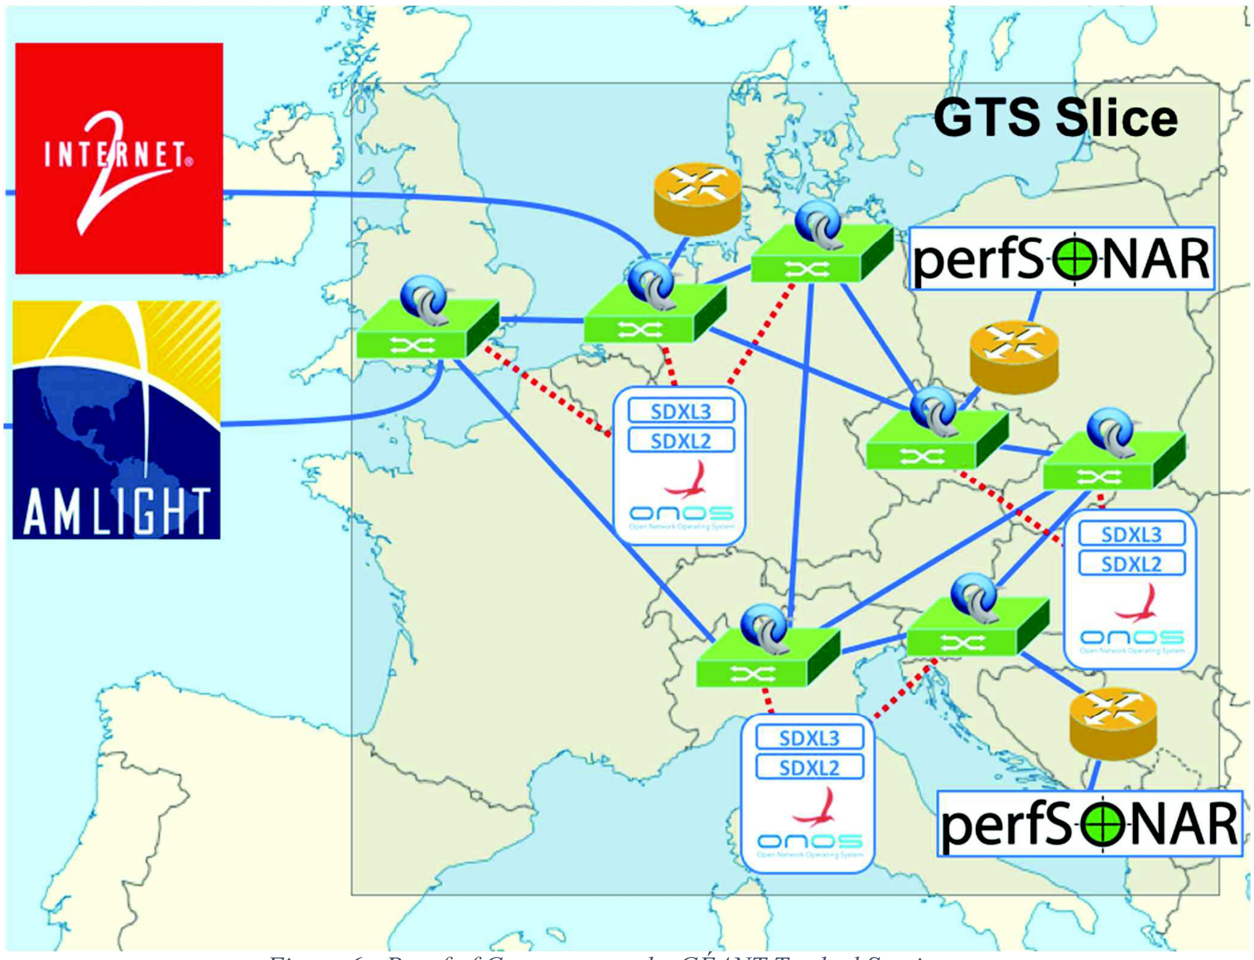

Figure 6 - Proof of Concept over the GÉANT Testbed Service

The second PoC has been realized in the GÉANT Testbed Service (GTS) [8]. The GTS delivers virtual testbeds powered by several facilities, co-located with GÉANT PoPs, offering different type of resources like VMs, SDN devices, virtual circuits. Using GTS, we have built a large-scale PoC with 7 HP OpenFlow switches deployed in 7 PoPs (Figure 6). This data plane is controlled by a cluster of 3 ONOS instances located in three of the PoPs. Three VMs, working as BGP peers, and two stub networks with perfSONar hosts have been deployed. PerfSONAR is a performance measurement and troubleshooting tool for multi-domain scenarios.

The SDX PoC on GTS has been integrated in a worldwide demo hosted at the Open Networking Summit 2016, where ON.Lab has successfully deployed ONOS and SDN-IP creating a global network facility entirely based on SDN. The network spans over 5 continents, interconnecting 9 RENs and more than 30 universities and research centers

Deployment experience and benefits for GÉANT

Overall the SDX PoC on GTS worked according to our expectation, L2-SDX and L3-SDX passed all the functional tests that we have performed. *Status* column in Table 1 reports the coverage assessment of the input requirements. The scalability and efficiency of L2-SDX and L3-SDX are tightly related to ONOS performances, and it has been demonstrated in [7] that the platform can meet carrier grade requirements in specific deployment conditions. L3-SDX and SDN-IP can scale up to 15K routes, achieving the current GÉANT requirement of 12K announced routes.

The SDX deployed in the GTS represents a single geographically distributed SDX, spanning 7 PoPs, with three ONOS instances running in different countries. During the execution of the functional tests, we gained feedback (e.g. the mastership election duration) that drove us not to further stress the SDX under critical events like controller instance failure or data plane failures. Therefore, we believe that having single-location SDXs, spanning a single OXP or PoP, is a safer approach to start the migration toward SDN. OXPs are good candidates for early deployment of SDX due to the complexity of the services that are offered and that makes the introduction of the SDN-based approach attractive. Moreover, with single-location SDXs, the devices can keep their independence and troubleshooting capabilities. In the current GÉANT network, each PoP is seen as a "hop" by the IP traffic, while in a geographically distributed L3-SDX simple troubleshooting tools like *traceroute* would be no longer useful. Incidentally, we observe that there is a gap to be filled with troubleshooting tools for SDN based networks, because layer 3 tools based on ICMP (ping, traceroute) do not work hop-by-hop in a network of SDN controlled switches.

The transition toward geographically distributed SDX could start with federations of nodes controlled by the same NOS instance. We have made some preliminary work on this issue with ICONA [14] an application to interconnect multiple ONOS clusters seamlessly through an "East-West" interface. Initially, OXPs could be interconnected creating a geographically distributed Layer 2 fabric controlled by ONOS. Likewise, geographically close PoPs could be federated in small clusters.

From the point of view of the development costs, it has been possible to release the L2-SDX and L3-SDX in the PoCs with a relatively small effort (in the order of three man months) thanks to the possibility of relying on the ONOS code base and documentation.

Let us now provide a high-level analysis of the benefits achievable by GÉANT with the *softwarization* of the infrastructures, in terms of operational costs like services provisioning and services management. The deployment of a SDX in place of an OXP will automate most of the configuration operations reducing the efforts of the human operators. Currently, in order to set-up a *GÉANT OPEN* service between two access points (ports or VLANs) inside an OXP, the customer has to contact the operators who manually configure the connection [8]. These operations (creation of virtual interfaces, VLAN id selection on both endpoints, VLAN id rewriting) are error prone and require coordination between the interested parties. Any arising issue requires further manual intervention of the operators. A typical target for the provisioning time of these services is 5 days. Using the L2-SDX, it will last minutes instead [15]. Moreover, most failure cases are automatically resolved by L3-SDX/L2-SDX using ONOS built-in mechanisms (e.g, a failure of a controller is solved using redundancy of controller instances, a switch failure is solved by ONOS with re-computation of data plane paths around the faulty switch).

As regards *GÉANT IP* and *GÉANT Plus*, similar improvements in the services provisioning and management are an affordable objective with single-location SDX. They represent a tangible result when compared to the current procedures (5 days to obtain IP connectivity or to set-up a layer 2 circuit).

When considering geographically distributed SDXs, having a centralized view of the network potentially brings further benefits like a more efficient traffic management, but the challenges for a wide area SDN deployment need still to be solved. In particular, a first issue is the impact of the latency and unreliability of the control plane connections between controllers and remote network nodes. A second issue arises when the controller instances are distributed in a geographical way, in order to reduce the latency towards the controlled nodes. The mechanisms used to achieve a consistent view across all the distributed controllers works well in local area networks, where the latency is low and the capacity is high. The performance of these consistency mechanisms can become critical in geographically distributed wide area networks; a careful assessment of these aspects is still needed.

Finally, let us consider an NREN that would like to establish a layer 2 connectivity with a third party (not-GÉANT). Different configurations have to be in place in order to have the connectivity operational: i) layer 2 configuration of the PoP where the NREN is located; ii) layer 2 configuration of the PoP where the target OXP is located; iii) Steering of the NREN flow into a LSP, provided by MPLS/BGP, towards the destination PoP; iv) establishment of layer 2 connectivity inside the OXP. Despite the services being logically similar, the provisioning procedure can require up to 10 days, because they are managed in two completely separated infrastructures, with specific hardware, different policies and configuration mechanism. Moreover, coordination is needed between all the operators (GÉANT, NREN and the third party). Instead, within an SDX environment, the separation is blurred and these services can be managed through a single platform, reducing coordination efforts, manual intervention of the network operators and complexity of the procedures. SDX allows a network engineer to provision a layer 2 circuit without prior technical coordination with the other network teams [15]. Moreover, they are delivered using a unique infrastructure, with reductions in term of OPEX and in future hardware investments.

Conclusions

In this paper, we have first considered the long term *softwarization* path of a Service Provider network like GÉANT. Then we have described the prototypes of two services, called L3-SDX and L2-SDX that have been deployed in a Proof of Concept over the GÉANT Testbed Service. The development has been based on the open source ONOS controller platform for SDN. We have performed a functional evaluation of the PoC and an analysis of the potential benefits, which have been very satisfactory. The developed services can bring considerable savings in the operational costs and can dramatically reduce the service provisioning time, as they automate many tasks that are manually performed. From the point of view of the performance, the SDX solution is ready for a “local” deployment, i.e. considering a single location (even if composed of a large number of nodes). It has to be further assessed if geographically distributed locations can be combined in single logical instances of the SDX.

In order to bring the L2-SDX and L3-SDX services to production, some additional concerns related for example to security and integration with management systems needs to be addressed. The GÉANT team is committed to work on these issues in the near future.

Acknowledgements

This work has been partially funded by the EC under project GN4-1.

References

[1] Open Networking Foundation, “Software-Defined Networking: The New Norm for Networks”, ONF White Paper, April 13, 2012.

[2] A. Galis, S. Clayman, L. Mamatas, J. Rubio Loyola, A. Manzalini, S. Kuklinski, J. Serrat, J. and T. Zahariadis, “Softwarization of Future Networks and Services-Programmable Enabled Networks as Next Generation Software Defined Networks”, IEEE workshop SDN for Future Networks and Services (SDN4FNS), Trento, Italy, November 2013.

[3] M. Kind, R. Szabó,C. Meirosu, F.J. Westphal "Softwarization of carrier networks", Information Technology. Volume 57, Issue 5, Pages 277–284, ISSN (Online) 2196-7032, ISSN (Print) 1611-2776, DOI: 10.1515/itit-2015-0019, October 2015.
[4] J. Mambretti et al., "Software-Defined Network Exchanges (SDXs) and Infrastructure (SDI): Emerging innovations in SDN and SDI interdomain multi-layer services and capabilities", Science and Technology Conference (Modern Networking Technologies - MoNeTeC), 2014 First International. IEEE, pp. 1-6, 2014.
[5] I. Monga, "Software Defined Exchanges: The new SDN?" [Online] http://meetings.internet2.edu/media/medialibrary/2015/03/31/Monga_SDX_SDN.pdf. Accessed Jan. 14, 2017.
[6] I. Pepelnjak – "Could IXPs Use OpenFlow To Scale?" [Online] http://www.menog.org/presentations/menog-12/116-Could_IXPs_Use_OpenFlow_To_Scale.pdf. Accessed Jan. 14, 2017.
[7] ONOS project – ONOS homepage [Online] http://onosproject.org. Accessed Jan. 14, 2017.
[8] GÉANT Connectivity and network management services - Connectivity and netwok management services home page [Online] http://www.geant.org/Services/Connectivity_and_network/Pages/Home.aspx Accessed Jan 14, 2017.
[9] A. Gupta et al., "Sdx: A software defined internet exchange", Proceedings of the 2014 ACM conference on SIGCOMM. ACM, pp. 551-562, 2014.
[10] J. Chung et al., "AtlanticWave-SDX: An International SDX to Support Science Data Applications", 2015.
[11] R. Buyya et al., "Market-oriented cloud computing: Vision, hype, and reality for delivering it services as computing utilities", High Performance Computing and Communications, 2008. HPCC'08. 10th IEEE International Conference on. IEEE, 2008.
[12] CORD project – CORD homepage [Online] http://opencord.org. Accessed Jan 14, 2017.
[13] B. Claise (ed.), B. Trammell (ed.), P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", IETF RFC 7011, September 2013
[14] M. Gerola et al., "ICONA: Inter Cluster ONOS Network Application", Network Softwarization (NetSoft), 2015 1st IEEE Conference on. IEEE, pp 1-2, 2015.
[15] J. Ibarra et al., "Benefits brought by the use of OpenFlow/SDN on the AmLight intercontinental research and education network", Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on. IEEE, pp 942-947, 2015.