

23-07-2013

Deliverable DJ1-2.1: Technology Investigation of OpenFlow and Testing



Deliverable DJ1-2.1

Contractual Date: 28-02-2013
Actual Date: 23-07-2013
Grant Agreement No.: 238875
Activity: JRA1, JRA2
Task Item: JRA1 T1, JRA2 T5
Nature of Deliverable: R (Report)
Dissemination Level: PU (Public)
Lead Partner: SWITCH
Document Code: GN3-13-003

Authors: Part 1 – JRA1 T1: M. Channegowda (University of Bristol), P. van Daalen (SURFnet), N. Efstathiou (University of Bristol), H. Thstrup Jensen (NORDUnet), R. Nejabati (University of Bristol), I. Torus (NORDUnet), H. Trompert (SURFnet)
Part 2 – JRA2 T5: G. Androulidakis (GRNET), Ch. Argyropoulos (GRNET), K. Baumann (SWITCH), L. Bergesio (i2CAT/RedIRIS), S. Buscaglione (DANTE), J. A. García-Espín (i2CAT/RedIRIS), J. Jofre (i2CAT/RedIRIS), A. Sevasti (GRNET), M. Suñé (i2CAT/RedIRIS)

© DANTE on behalf of the GÉANT project.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7 2007–2013) under Grant Agreement No. 238875 (GÉANT).

Abstract

This deliverable has been jointly prepared by JRA1 T1 and JRA2 T5. It documents the results of JRA1 T1's study of the state of the art of Software-Defined Networking and OpenFlow, and presents JRA2 T5's design and deployment of the GÉANT OpenFlow Facility.

Table of Contents

Executive Summary	1
Part 1	3
1 Introduction	3
1.1 Background	3
1.2 Problem Definition	3
1.3 Software-Defined Networking	4
1.4 In Part 1	5
2 OpenFlow	6
2.1 Overview	6
2.2 Concept of Flows	8
2.3 Capabilities	8
2.3.1 OpenFlow Version 0.9	9
2.3.2 OpenFlow Version 1.0	9
2.3.3 OpenFlow Version 1.0 Draft 0.3 Optical Extensions	10
2.3.4 OpenFlow Version 1.1	10
2.3.5 OpenFlow Version 1.2	10
2.3.6 OpenFlow Version 1.3	11
2.4 Unified Control Plane	12
2.4.1 OpenFlow-Based Unified Control Plane [OF-FirstDemo]	12
2.4.2 FlowVisor	14
2.4.3 Slice-based Facility Architecture	16
2.4.4 Network Services Framework	16
2.5 Project Examples	18
2.6 Hardware and Control Software Maturity	18
3 Investigation of OAM Mechanisms with OpenFlow	21
3.1 Introduction	21
3.2 Ethernet OAM	21
3.3 Integration of OAM in OpenFlow – Ethernet OAM Scenario	23

4	Review of Vendors' Approaches	26
4.1	Ciena	26
4.1.1	Current Status and View of OpenFlow / SDN	26
4.2	Juniper	28
4.2.1	Introduction of OpenFlow / SDN	28
4.2.2	Current Status and View of OpenFlow / SDN	28
4.2.3	Junos Software Development Kit	29
4.2.4	Juniper Controller	29
4.2.5	Juniper OpenFlow-Enabled Equipment	29
4.2.6	Collaborations	30
4.2.7	Future Plans	30
4.3	Cisco	30
4.3.1	Introduction of OpenFlow / SDN	30
4.3.2	Cisco ONE	30
4.3.3	Cisco OpenFlow-Enabled Hardware	31
4.3.4	Future Plans	31
4.4	Brocade	31
4.4.1	Introduction of OpenFlow / SDN	32
4.4.2	Current Status of and Strategy on OpenFlow / SDN	32
4.4.3	Brocade SDN-Ready Product Family	35
4.4.4	Brocade OpenFlow Demonstrations	37
4.4.5	Collaborations	38
4.4.6	OpenFlow Switches Operating at Line-Rate Speeds	38
4.4.7	Brocade Acquiring Vyatta	39
4.4.8	Summary and Conclusions	39
4.5	Big Switch Networks	39
4.5.1	Introduction of OpenFlow / SDN	40
4.6	Arista Networks	42
4.6.1	Introduction of OpenFlow / SDN	42
4.6.2	Collaborations	43
4.6.3	Summary and Conclusions	44
4.7	IBM	44
4.7.1	Introduction of OpenFlow / SDN	44
4.7.2	Programmable Network Controller (PNC)	45
4.7.3	RackSwitch G8264	47
4.7.4	Collaborations	48
4.8	HP	49
4.8.1	Introduction of OpenFlow / SDN	49

4.8.2	HP Controller – HP Virtual Application Networks SDN Controller	50
4.8.3	HP Switches	51
4.8.4	HP APIs and Applications	52
4.8.5	Future Plans	52
4.9	Dell	52
4.10	NEC	53
4.10.1	Introduction of OpenFlow / SDN	53
4.10.2	NEC ProgrammableFlow Family of Products	54
4.10.3	Collaborations	56
4.10.4	Future Plans	56
4.11	Pica8	57
4.11.1	Pica8 Implementation of OpenFlow	57
4.11.2	XorPlus OS	58
4.11.3	Pica8 Switches	59
4.11.4	Future Plans	60
4.12	ONF	60
5	Conclusions	61
Part 2		64
6	Introduction	64
6.1	In Part 2	64
7	SDN/OpenFlow in Research Networks	65
7.1	The Case of GÉANT	66
8	GÉANT OpenFlow Facility	67
8.1	Objectives	67
8.2	Facility Design	67
8.3	GÉANT OpenFlow Control Framework	70
8.4	Facility Deployment Issues and Considerations	71
9	Operational Environment of the Facility	72
10	Use Cases and International Collaboration	76
10.1	Use Cases	76
10.2	International Collaboration	77

11	Conclusions	79
Appendix A	Definition of Terms	80
Appendix B	GÉANT OpenFlow Facility: General Server Specifications	81
References		82
Glossary		89

Table of Figures

Figure 2.1: Idealised OpenFlow switch – a remote controller controls the flow table over a secure SSL channel [OF-EnablingInnovation]	7
Figure 2.2: OpenFlow header field of a flow table	7
Figure 2.3: Packet and circuit flow definition	8
Figure 2.4: Architecture of an optical network control plane based on OpenFlow [OF-FirstDemo]	13
Figure 2.5: FlowVisor is logically placed between the physical network and the slice controllers [ALIEN, ALIEN D2.1]	14
Figure 2.6: FlowVisor Architecture [FlowVisor-NVL]	15
Figure 2.7: Network Service Framework concept [NSFv1.0]	17
Figure 3.1: A hierarchy of maintenance sessions according to IEEE 802.1ag specification [DJ1.1.2]	22
Figure 3.2: OpenFlow Switch forwarding IEEE 802.1ag messages to the OpenFlow controller	24
Figure 4.1: Juniper SDN architecture	28
Figure 4.2: Juniper EX Series Ethernet Switches	29
Figure 4.3: Brocade's view of the SDN stack	33
Figure 4.4: Brocade support for standards-based overlay technologies	34
Figure 4.5: Brocade's change from traditional physical infrastructure to the VCS Fabric-based SDN [BrocadeNetworkTrans]	35
Figure 4.6: Brocade MLXe and NEC PFC demonstration	37
Figure 4.7: Big Switch's three-tier view of SDN	40
Figure 4.8: Floodlight controller in SDN environment	40
Figure 4.9: Big Switch commercial target: controller platform and application development on network virtualisation	41
Figure 4.10: Arista 7150 Switch series	42
Figure 4.11: IBM SDN approach	45

Figure 4.12: IBM and NEC switching architecture for OpenFlow [IBMOFNG]	48
Figure 4.13: HP OpenFlow and SDN activity milestones [HPOF]	49
Figure 4.14: The HP SDN Architecture [HPVAN]	50
Figure 4.15: HP Virtual Application Networks SDN Controller [HPVAN]	51
Figure 4.16: ProgrammableFlow Management Console – centralised monitoring	55
Figure 4.17: Pica8 solution for Data Centers [Pica8startup]	57
Figure 4.18: PicOS architecture [Pica8OSforOpen]	58
Figure 8.2: A generic view of the GÉANT OpenFlow Facility showing control and management plane	69
Figure 9.1: A three-tier architecture for operations	73
Figure 9.2: Demarcation points for the network connectivity services	74

Table of Tables

Table 2.1: Inter-domain flow mapping rules [OF-FirstDemo]	13
Table 4.1: Arista 7150 switches – comparison of specifications	43
Table 4.2: IBM PNC specifications [IBMPNC]	47
Table 4.3: IBM RackSwitch G8264 performance [IBMG8264]	47
Table 9.1: Operational tasks across the tiers and assignments to operations entities	75
Table A.1: Definition of terms	80
Table B.2: GÉANT OpenFlow Facility: general server specifications	81

Executive Summary

This document has been jointly prepared by GN3 Joint Research Activity 1 Future Network, Task 1 Carrier Class Transport Network Technologies (JRA1 T1) and JRA2 Multi-Domain Network Service Research, Task 5 Network Factory (JRA2 T5). Part 1 documents the results of JRA1 T1's study of the state of the art of Software-Defined Networking (SDN) and OpenFlow, to be used by GN3 colleagues, especially as input to the work of JRA2 T5; Part 2 presents JRA2 T5's design and deployment of the GÉANT OpenFlow Facility.

SDN is a new network architecture/approach which allows a network to be programmed as if it were a computer. OpenFlow is a communications protocol that gives access to the forwarding plane of a switch or router over the network. It differs from traditional network technologies in its ability to separate the control plane from the forwarding plane, supporting the concept of SDN: OpenFlow, or any Application Programming Interface that provides a layer of abstraction from the physical network to the control element, allows that network to be configured or manipulated through software, which then opens it up to further innovation – by researchers, operators, application/service providers and third parties as well as by network equipment vendors.

Part 1 describes OpenFlow, covering the concept of flows, the capabilities of the different versions, the unified control plane, project examples, maturity, and Operation, Administration and Maintenance (OAM) mechanisms. It also presents a review of vendors' approaches, considering, for twelve vendors, aspects such as introduction of OpenFlow/SDN, current status, collaborations and future plans.

The study confirms that SDN and OpenFlow are one of the current key topics for all network vendors, representing both an opportunity and a requirement, since no network vendor can afford to neglect a development with which the main part of the market is getting involved, and about which customers are continually asking. The network vendors' market can be divided into three with regard to SDN/OpenFlow: the "heavy" vendors with strong network strategies and long-term roadmaps, such as Cisco or Juniper; vendors such as IBM or HP who are finding the technology a good opportunity to re-enter, or enter more deeply, the networks market; and start-ups that can easily benefit from OpenFlow without significant investments, such as Big Switch or Pica8. The study has shown that a wide variety of OpenFlow equipment is available and OpenFlow software will converge to standard facilities. For the application of OpenFlow on large-scale networks or Wide Area Networks, it is recommended to limit the different equipment to a few brands in order to provide a stable environment. Meanwhile the OpenFlow software applied in a data centre environment is stable, reliable and can be implemented successfully.

Part 2 describes how OpenFlow is being used by Research and Education Networks (RENs) in general, and by GÉANT in particular to deliver a Testbed as a Service (TaaS). It then summarises the objectives, design and

deployment of the GÉANT OpenFlow Facility, describes the facility's operational environment, and provides examples of use cases and international collaborations.

SDN/OpenFlow has obtained significant recognition among RENs in recent years as a means to serve the specialised needs of their users across a range of use cases. The work presented in Part 2 of this document focuses on a particular category of use cases that confirms the suitability of OpenFlow as a technology offering Testbed as a Service (TaaS) capabilities. TaaS delivers to researchers a slice composed of an isolated L2 network environment with attached computation resources, and gives them full control of the resources in their slice through a set of management utilities.

The GÉANT OpenFlow Facility has been successfully implemented on top of the GÉANT backbone production environment. It combines GÉANT substrate capabilities, software-based OpenFlow-enabled network elements, computing resources and a control framework for managing the facility and delivering slices to users. It is the first testbed facility implemented on top of GÉANT that offers “vanilla” L2 slices, and is one of the pioneering OpenFlow-enabled facilities globally offering WAN TaaS. A three-tier architecture for operations has been successfully designed and delivered for development as part of the evolution of the facility offerings in a service context.

A slice contest held by JRA2 T5 resulted in several interesting proposals for experimentation utilising the specific capabilities of the facility. This has confirmed the interest of the research community in experimentation on top of the GÉANT wide-area OpenFlow Facility and its particular value in supporting applications and Future Internet research. It has also provided a strong indication of returns on the investment. As the GN3 project concludes, a roadmap of facility enhancements as well as a plan for the handover of its operations set the foundation for an operational TaaS service within GN3plus.

While no standalone technology testing of SDN/OpenFlow has taken place within either JRA1 or JRA2, as originally envisaged for this deliverable (i.e. comparable to the transport network technology testing documented in [DJ1.1.2]), JRA2 T5's OpenFlow Facility design and deployment has encompassed practical testing, and the experiments carried out using the facility will also contribute to testing both the technology and the facility itself.

Part 1

1 Introduction

1.1 Background

OpenFlow is a communications protocol that gives access to the forwarding plane of a switch or router over the network [OF-EnablingInnovation]. OpenFlow differs from traditional network technologies in its ability to separate the control plane from the forwarding plane, supporting the concept of Software-Defined Networking (SDN). This enables features like virtualisation, advanced forwarding and programmability of networks, and allows researchers to test their ideas on real networks by isolating the production traffic.

In 2012, a sub-task of GN3 Joint Research Activity 1 Future Network, Task 1 Carrier Class Transport Network Technologies (JRA1 T1) conducted a study of the state of the art of SDN and OpenFlow, to be used by GN3 colleagues, especially as input to the work of JRA2 Multi-Domain Network Service Research, Task 5 Network Factory (JRA2 T5). Part 1 of this report documents the results of that study, which were also presented to JRA2 T5 in two workshops held on 12 June 2012 in Utrecht and during the GN3 Symposium in Vienna in October 2012. It focuses on trends in SDN and OpenFlow; JRA1 T1 did not carry out any practical, hands-on research and development as part of the sub-task. The information presented was current at the time of the study (2012); further research will be carried out in GN3plus, affording the opportunity to assess new developments and bring existing details up to date.

This section briefly defines the problems presented by current network deployments and introduces the SDN concept.

1.2 Problem Definition

The major hurdles faced by telecoms researchers using current network architectures and deployments include:

- Networks are rigid and hard to manage, e.g. network administration.
- Innovation is slow and comes in patches in the form of new protocols or algorithms.

- Supporting complex requirements is a challenge, e.g. Access Control List (ACL), Virtual Local Area Networks (VLANs), Deep Packet Inspection (DPI).
- OpEx is higher in large networks.
- The solution should be platform agnostic – a multi-vendor cross-domain solution.

In an effort to address such basic but key problems, a layer-based (Open Systems Interconnection (OSI), TCP/IP) generalisation approach is currently employed. However, a lower level control plane abstraction is needed. For example, in the optical domain, Generalised Multi-Protocol Label Switching (GMPLS) separates the control plane from the data plane. Traditional networking-based approaches to control plane abstraction lead to either:

- Defining new protocols, e.g. routing.
OR
- Reconfiguring existing mechanisms, e.g. traffic engineering.
OR
- Manual configuration, e.g. access control.

However, there are other, better alternatives, one of which is Software-Defined Networking (SDN).

1.3 Software-Defined Networking

“OpenFlow or more broadly Software-Defined Networking is a new network architecture/approach that enables innovations by researchers, operators, application/service providers and third parties as well as by network equipment vendors.” [POMI-OF-SDN]

SDN allows a network to be programmed as if it were a computer. OpenFlow, or any API that provides a layer of abstraction from the physical network to the control element, allows that network to be configured or manipulated through software, which then opens it up to further innovation.

The features of SDN that address researchers’ increasingly complex network requirements include:

- Abstraction of the forwarding function.
Separation of data plane from control plane which gives freedom to service providers to try different topologies and protocols without many restrictions from the physical infrastructure (OpenFlow/FlowSpace).
- Distributed state abstraction.
Global network view (logical and virtual) providing different virtualised topologies to support experimentation on production networks.
- Abstraction of distributed network control.

Network Operating System (NOS) in global service environment providing a single virtualised operating system.

Areas where SDN is already making an impact and delivering benefits include:

- Network virtualisation.
- Dynamic network access control.
- Load balancing.
- Energy-efficient networks.
- Adaptive network monitoring and counting.

1.4 In Part 1

Part 1 of this document has been structured as follows:

- Chapter 2 *OpenFlow* – describes OpenFlow, covering the concept of flows, the capabilities of the different versions, the unified control plane, project examples and maturity.
- Chapter 3 *Investigation of OAM Mechanisms with OpenFlow* – covers Ethernet Operation, Administration and Maintenance (OAM) and its integration in OpenFlow.
- Chapter 4 *Review of Vendors' Approach* – for twelve vendors, considers aspects such as introduction of OpenFlow/SDN, current status, collaborations and future plans.
- Chapter 5 *Conclusions* – provides an overall assessment of the study results.

2 OpenFlow

2.1 Overview

OpenFlow [OpenFlow] is envisioned as the solution for virtualised programmable networks that will drive research innovations in the networking arena. OpenFlow is a protocol that allows administrators/researchers to define flows in resources to control and isolate traffic. This separation of control allows the flow of packets to be determined by the software, which acts as the controller.

As shown in Figure 2.1, OpenFlow consists of a controller, which connects to an OpenFlow-enabled device via a Secure Sockets Layer (SSL) channel. After successful channel setup, the OpenFlow controller has access to the device to define flow-table rules to control the path of the traffic flow. Every packet processed by the device is matched against the flow-table entries populated by the controller. If a packet matches an entry, the corresponding action (dedicated by controller) is performed or else the packet is sent to the controller to make the necessary decision. A packet flow is a group of packets defined by the subset of a header field of a flow-table entry. The format of the header field is defined in Figure 2.2.

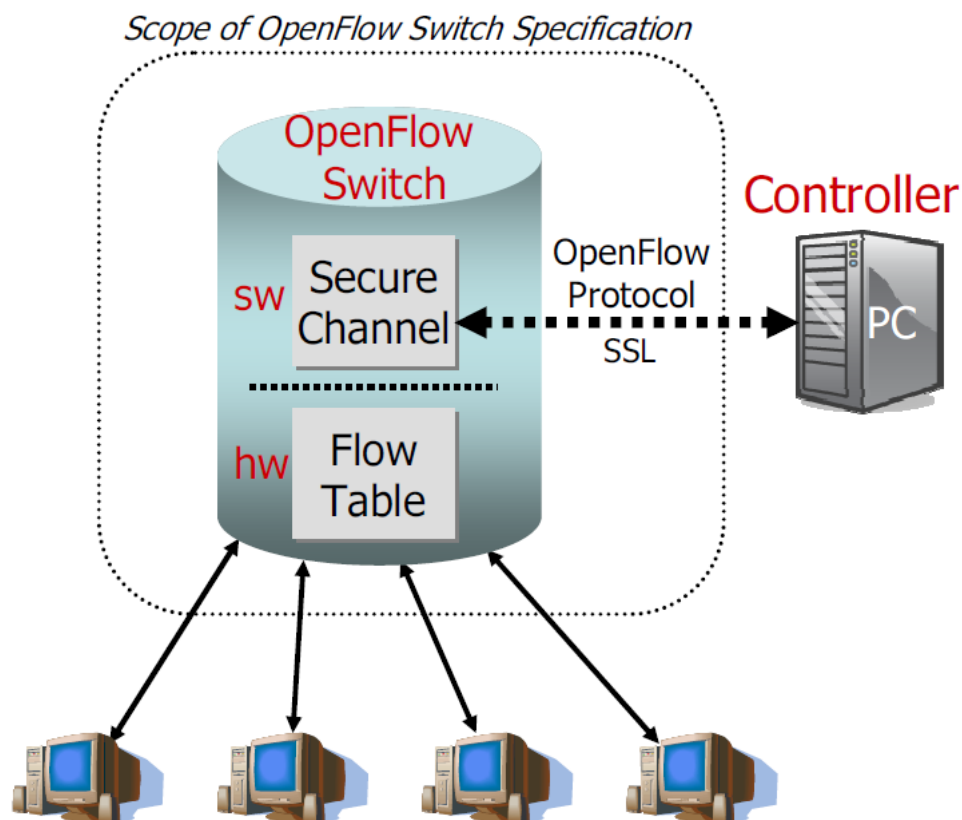


Figure 2.1: Idealised OpenFlow switch – a remote controller controls the flow table over a secure SSL channel [OF-EnablingInnovation]

Input Port	Ethernet			VLAN		IP				TCP/UDP	
	SA	DA	Type	ID	Priority	SA	DA	Protocol	TOS	Src port	Dst port

Figure 2.2: OpenFlow header field of a flow table

The protocol supports three message types – asynchronous, symmetric and controller-to-switch – to establish connection between a controller and an OpenFlow-enabled device. The specification [OF-SwitchSpec] also lists the various message exchanges needed to support the communication.

Due to the attention that OpenFlow has attracted recently, there have been substantial new developments. Arrays of controllers like NOX [NOX], Beacon [Beacon] and the controller framework Trema [Trema] have been implemented. Many vendors have started supporting OpenFlow on their products and an association, the Open Networking Foundation (ONF) [ONF] has been set up for standardisation. One such innovation is the FlowVisor, which is a controller providing network virtualisation.

2.2 Concept of Flows

Using a flow, it is possible to define a sequence of packets from a source to a destination that have common characteristics. In the OpenFlow concept, a flow is defined as a combination of the 12-tuple L2/L3/L4 headers (Figure 2.3), which aggregate those attributes into a single flow definition and push it in the flow table of the OpenFlow switch. Except from the packet domain, a flow can be defined in the circuit domain by a 5-tuple field, which identifies an optical flow. However, in the flow definition, actions and statistics are attached to the flow so the OpenFlow switch will know what to do with those kinds of packets and collect statistics about that particular flow.

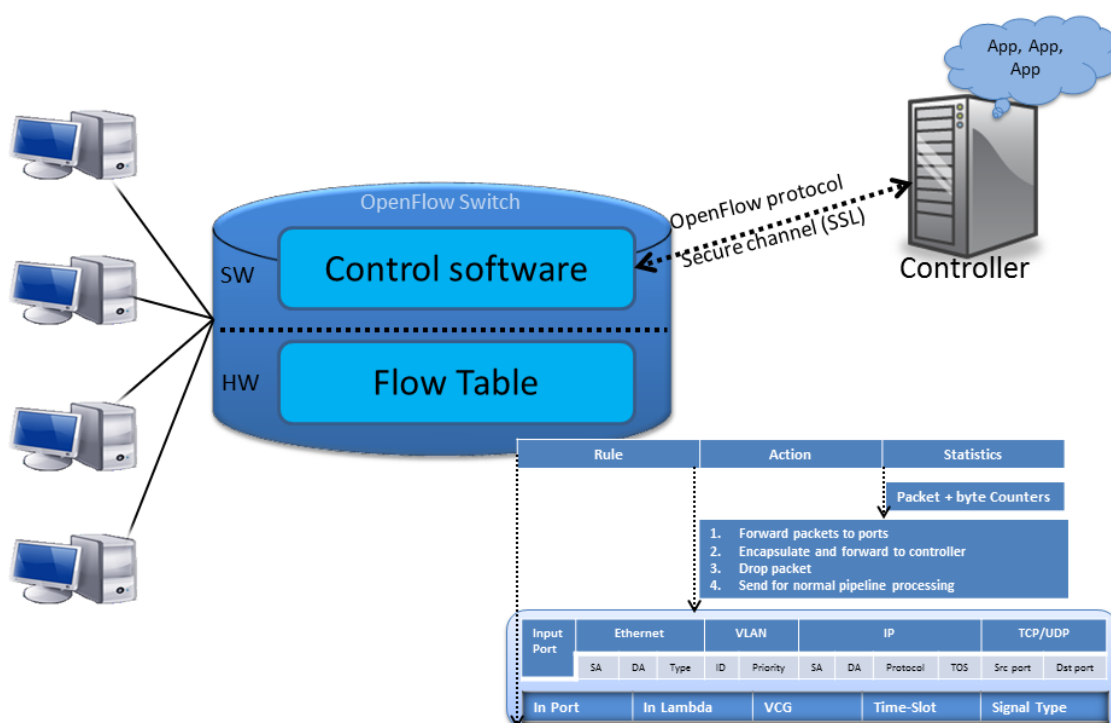


Figure 2.3: Packet and circuit flow definition

2.3 Capabilities

In 2008, Stanford University released the first version of the OpenFlow protocol and in the following years released more versions with significant updates, along with version 1.0, which is the version that has been implemented by hardware vendors in commercial networking products. In 2011, Deutsche Telekom, Facebook, Google, Microsoft, Verizon, and Yahoo! formed a non-profit consortium called the Open Networking Foundation (ONF), with the goal of promoting and commercialising SDN and OpenFlow as an approach to networking [ONFAbout]. To date, ONF has standardised OpenFlow version 1.2 and OF-Config 1.0, which is the first version of the configuration and management protocol for OpenFlow switches [ONFConfig1.0]. In the near

future, OpenFlow 1.3 and OF-Config 1.1 will be standardised, since they are under the ratification period [ONFConfig1.2].

This section summarises the OpenFlow versions and provides a brief list of the supported features.

2.3.1 OpenFlow Version 0.9

OpenFlow 0.9 [OF-V0.9] incorporates a large number of bug fixes from previous versions, clarifications and new features including:

- Failover mechanism to a redundant controller.
- Emergency flow entries.
- VLAN priority matching.
- Support for Type of Service (ToS) field re-writing.

2.3.2 OpenFlow Version 1.0

OpenFlow 1.0 [OF-SwitchSpec, OF-V1.0] adds a number of key features and is the first stable and most deployed version to date. The largest additions are:

- Slicing, where multiple queues per output port have been added and provide the ability to configure the bandwidth per queue.
- Quality of Service (QoS) mechanism, where IP ToS / Differentiated Services Code Point (DSCP) bits can be matched.
- User-specifiable datapath description, where a datapath description field has been added to the switch description, which allows a switch to return a string specified by the switch owner to describe the switch. This can be used to provide more specific data and parameters such as delay, buffer capacity, policing and forwarding parameters directly to the user.
- IP addresses in Address Resolution Protocol (ARP) packets.
- Flow cookies to identify flows.
- Selective port statistics.
- Improved flow duration resolution in statistics/expiry messages.

The release also includes a large number of small changes to the specification and bug fixes in the reference implementation.

Hardware-accelerated OpenFlow 1.0-capable switches are available on the market with a variety of vendors supporting it. On the controller side, the reference controller supports 1.0 and a variety of OpenFlow controllers support version 1.0, e.g. NOX [NOXRepo], Beacon [Beacon], Floodlight [Floodlight], Open VSwitch [OpenVSwitch], etc.

2.3.3 OpenFlow Version 1.0 Draft 0.3 Optical Extensions

OpenFlow 1.0 draft 0.3 [OF-PAC.C, OF-SwitchSpecAdd] covers the components and basic functions of circuit switches based on switching time slots, wavelengths and fibres. It also covers hybrid switches with both packet and circuit interfaces and/or switching fabrics. In addition, it specifies OpenFlow protocol changes required to manage such OpenFlow switches from a remote controller. The circuit switch addendum [OF-SwitchSpecAdd] introduces the following changes to the OpenFlow protocol specification:

- Extensions to the capabilities field in `ofp_switch_features` to account for circuit switch capabilities and to include circuit ports.
- Creation of structure (struct) `ofp_phy_cport` to account for circuit port characteristics.
- Two additional OpenFlow messages to enumeration (enum) `ofp_type`:
 - `OFPT_CFLOW_MOD` to create virtual ports and to specify cross-connections.
 - `OFPT_CPORT_STATUS` with new “reasons” added to `ofp_port_reason`.
- Definition of struct `ofp_connect` for specifying the cross-connection in circuit switches.
- Addition of two action types, `OFPAT_CKT_OUTPUT` and `OFPAT_CKT_INPUT`, to account for adapting packet flows to circuit flows and extracting packet flows from circuit flows.
- Addition of error messages for problems in circuit flow modifications.
- Extension of `ofp_port_state` to report feature changes of the circuit port.

2.3.4 OpenFlow Version 1.1

The goal of OpenFlow 1.1 [OF-SwitchSpec110] was to produce a software prototype based on Open VSwitch to test the new specification features. Those features were:

- Multipath, where a flow can be sent over several paths.
- Tags and tunnels support, where tag support has been extended to include MPLS shim headers. VLAN tag handling has been modified to support packets with multiple VLAN tags (Q-in-Q encapsulation) and tunnels support via virtual ports.
- Multiple tables, where an OpenFlow switch can now expose multiple tables to a controller.

2.3.5 OpenFlow Version 1.2

The Open Networking Foundation (ONF) released OpenFlow 1.2 [ONF-OF-V1.2]. It encompasses the switch specification, an evolution from OpenFlow 1.1, 1.0, and previous versions and describes the formats and protocols by which an OpenFlow switch receives, reacts to, and responds to messages from an OpenFlow Controller.

The OpenFlow 1.2 Switch Specification builds significantly upon previous releases in many ways, including the following improvements:

- It adds support for IPv6. In addition to the previous support for IPv4, MPLS, and L2 headers, OpenFlow 1.2 now supports matching on IPv6 source address, destination address, protocol number, traffic class, ICMPv6 type, ICMPv6 code, IPv6 neighbour discovery header fields, and IPv6 flow labels.
- It adds support for extensible matches. By employing a Type Length Value (TLV) structure, the protocol allows far greater flexibility for the treatment of current and future protocols. TLV support was added in the matching function, in packet-in and in error messages.
- It allows experimenter extensions through dedicated fields and code points assigned by ONF.

2.3.6 OpenFlow Version 1.3

The Open Networking Foundation (ONF) released OpenFlow 1.3 [ONF-OF-V1.3], which is the latest version to date. It includes a more flexible framework to express capabilities, something that was missing from the previous versions.

The main change is the improved description of table capabilities. Those capabilities have been moved out of the table statistics structure into their own request/reply message, and encoded using a flexible TLV format. This enables the additions of next-table capabilities, table-miss flow entry capabilities and experimenter capabilities.

Other changes include renaming the “stats” framework into the “multipart” framework to reflect the fact that it is now used for both statistics and capabilities, and the move of port descriptions into their own multipart message to enable support of a greater number of ports.

A summary of all the changes is given below.

- More flexible table miss support.
- IPv6 Extension Header handling support.
- Per flow meters, which adds the ability to measure and control the rate of packets sent to the controller.
- Per connection event filtering, which improves the multi-controller support by enabling each controller to filter events from the switch it does not want.
- Auxiliary connections, which enable a switch to create auxiliary connections to supplement the main connection between the switch and the controller to carry packet-in and packet-out messages.
- MPLS Bottom of Stack (BoS) matching, which indicates if other MPLS shim headers are in the payload of the present MPLS packet; matching this bit can help to disambiguate cases where the MPLS label is reused across levels of MPLS encapsulation.
- Provider Backbone Bridging (PBB) tagging, which enables OpenFlow to support networks based on PBB and Provider Backbone Bridging with Traffic Engineering PBB-TE.
- Rework tag order, where each tagging operation adds its tag in the outermost position.
- Tunnel-ID metadata.
- Cookies in packet-in, where a cookie field was added to the packet-in message, which takes its value from the flow and sends the packet to the controller, enabling it to classify more efficiently the packet-in messages.

- Duration field added for most statistics, which enables packet and byte rate to be calculated more accurately from the counters included in those statistics.
- On demand flow counters, which enable or disable packet and byte counters on a per-flow basis.

2.4 Unified Control Plane

With separation of data and control, and the treatment of packets as flows, together with the introduction of circuit-flow features in the OpenFlow protocol, a unified architecture becomes realisable for converged packet-circuit networks. OpenFlow abstracts each data-plane switch as a flow table. It allows the definition of a flow to be any combination of L2-L4 packet headers for packet flows, as well as L0-L1 circuit parameters for circuit flows.

2.4.1 OpenFlow-Based Unified Control Plane [OF-FirstDemo]

Software-Defined Networking (SDN) and flexible Dense Wavelength-Division Multiplexing (DWDM) grids are two technologies that enable network operators to adapt their infrastructure to changing application requirements, thereby increasing network efficiency and minimising network cost. An OpenFlow-based software-defined control plane, which integrates packet and optical networks over multiple domains with heterogeneous transport technologies, can decrease the current deployed control planes' complexity, OpEx and CapEx and provide control across different technologies.

Figure 2.4 shows an architecture block diagram of the OpenFlow-based unified control plane for multi-domain packet over fixed/flexible grid optical networking, where the principal components are the extensions to the existing OpenFlow (OF) protocol circuit specification to support flexi-grid transport technology along with its integration with fixed DWDM grid and layer 2 packet switching. The required OF extensions of this architecture are mainly focused on the Switch_Features and CFlowMod messages. Switch_Features extensions include the support of central frequency, spectrum range, and granularity of Bandwidth Variable Transponder (BVT) and Bandwidth Variable Optical Cross-Connect (BV OXC); number of ports and wavelength channels of WDM OXC; peering connectivity inside and across multiple domains. Also the CFlow_Mod message supports a flexible grid domain based on the ITU-T G.694.1 recommendation [ITU-T G.694.1]. The allowable granularity of flexible grid equipment can be determined using the Switch_Feature messages. To control a BVT or BV OXC only grid parameters (central frequency and slot widths) are exchanged between the controller and a Network Element (NE) via CFlow_Mod messages.

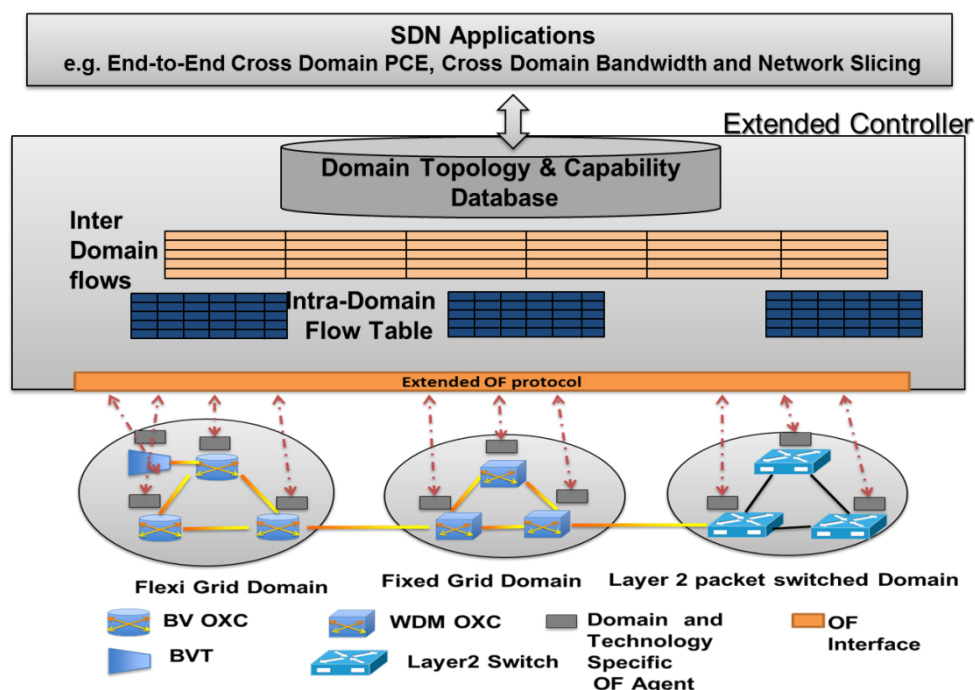


Figure 2.4: Architecture of an optical network control plane based on OpenFlow [OF-FirstDemo]

Table 2.1 shows the flow specification for fixed and flexible grid optical networks. In a fixed grid optical network, a flow can be identified by a fixed flow identifier comprising port, wavelength and signal type (e.g. bitrate, optical signal format, protocol and mapping structure) fields associated with a specific optical switch. In a flexible grid optical network, a flow is identified by a flexible flow identifier comprising port, centre frequency (CF), frequency slot bandwidth (FSB) and type of signal fields associated with that switch.

An extended OF controller is required in order to support the required networking functions for multi-domain/multi-technology operation such as peering and flow mapping between multiple domains and/or technologies. It utilises intra-domain and inter-domain flow tables, where an intra-domain flow table holds flow identifiers and associated actions for each NE within a particular domain. For optical NEs, an action is defined as a cross connection associated with one or more flow identifiers. For a flexible grid NE, the action also includes CF and FSB. Inter-domain flow tables hold flow identifiers and associated actions for NEs that interconnect between neighbouring domains. Actions in flow tables that are associated with two heterogeneous technology domains must comply with multi-domain mapping rules described in Table 2.1.

Multi-Domain Scenario	Flow Mapping Rule
Flexi Grid \leftrightarrow Fixed Grid	Each centre frequency and bandwidth must be compatible with WDM Grid
Flexi Grid \leftrightarrow Packet Switched	One or multiple packet flow identifiers must be mapped to each centre frequency + bandwidth
Fixed Grid \leftrightarrow Packet Switched	One or multiple packet flow identifiers must be mapped to each channel

Table 2.1: Inter-domain flow mapping rules [OF-FirstDemo]

This work has led to the published document “First Demonstration of an OpenFlow based Software-Defined Optical Network Employing Packet, Fixed and Flexible DWDM Grid Technologies on an International Multi-Domain Testbed” [OF-FirstDemo].

The application of both the frontier technologies, SDN with OF and flexible elastic networks, provides an extensible control framework for packet over optical transport embracing existing and emerging technologies. This enables dynamic, flexible and sliceable networking with virtualisation.

2.4.2 FlowVisor

FlowVisor is a special-purpose OpenFlow controller which implements a network virtualisation layer [FlowVisor] between the OpenFlow-enabled switches in an environment where there are multiple controllers. FlowVisor allows multiple logical networks, each with different addressing and forwarding schemes, to co-exist on the same physical infrastructure. FlowVisor “slices” the OpenFlow network by intercepting all the OpenFlow-protocol messages sent by the switches to the controllers and forwarding them to the correct controller according to pre-defined policies (a process called network slicing). FlowVisor allows the configuration of several logical topologies (also called slices) with links and nodes that are a subset of the physical topology. FlowVisor generally hosts multiple guest controllers, one controller per slice (see Figure 2.5) and it ensures that a controller can observe and control its own slice only, while isolating one slice from another. For example, a slice can be defined as the set of OpenFlow messages coming from specified subsets of switches and containing flow headers with destination and source IP addresses belonging to a subnet. Although a switch can belong to multiple slices, each slice only has control over its own flows. When a controller tries to control flows outside the slice to which it is assigned, i.e. its flowspace, FlowVisor can either modify the control message or simply reject that message.

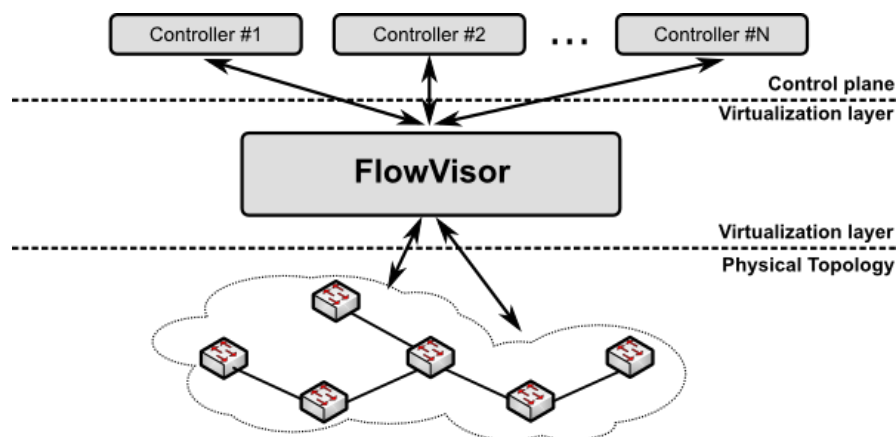


Figure 2.5: FlowVisor is logically placed between the physical network and the slice controllers [ALIEN, ALIEN D2.1]

Referring to Figure 2.6, FlowVisor maintains a policy engine, translation and forwarding mechanisms to fulfil the following virtualisation goals:

- Isolation.
Multiple virtual segments are created and allocated to user controllers within the same physical substrate. Each virtual network segment is independent and hence gives users a secure isolation.
- Transparent.
FlowVisor is transparent to both the user controller and the OpenFlow-enabled device. This solves two problems:
 - It provides flexibility by decoupling the control from the forwarding plane.
 - Since both physical infrastructure and control software are unaware of the virtualisation layer, design of the controller and FlowVisor are not dependent on each other.
- Scalability and flexibility.
FlowVisor supports creation of highly diverse virtual networks keeping in mind the allocated resources such as bandwidth, traffic, CPU, etc.

FlowVisor has been successfully applied as a network virtualisation layer in several wired and wireless scenarios. It has been applied using different network technologies in several different deployment experiments including: GENI [GENI] and OFELIA [OFELIA].

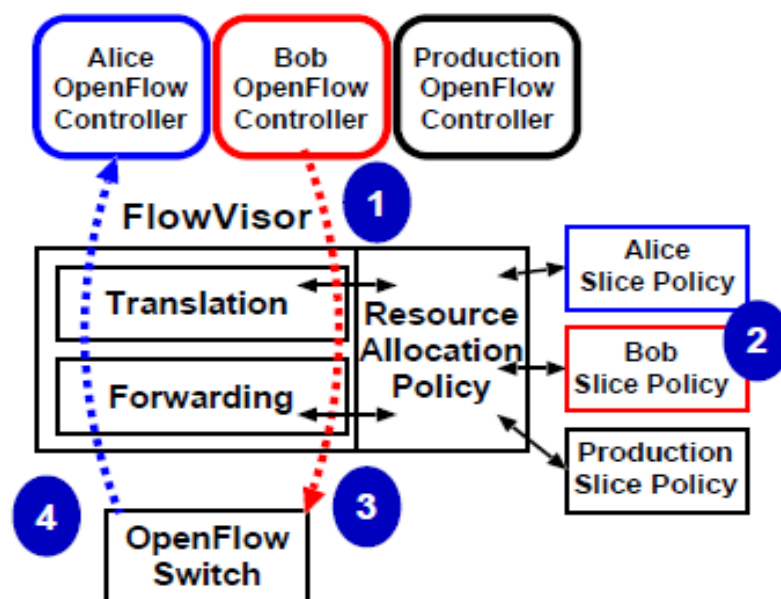


Figure 2.6: FlowVisor Architecture [FlowVisor-NVL]

In Figure 2.6, the FlowVisor intercepts OpenFlow messages from guest controllers (1) and, using the user's slicing policy (2), transparently rewrites (3) the message to control only a slice of the network. Messages from switches (4) are forwarded to guests only if it matches their slice policy [FlowVisor-NVL].

2.4.3 Slice-based Facility Architecture

The goal of Slice-based Facility Architecture (SFA) is to provide a minimal interface, a narrow waist that enables testbeds of different technologies and/or belonging to different administrative domains to federate without losing the control of their resources. This will allow researchers to combine all available resources and run advanced networking experiments of significant scale and diversity.

However, the coexistence of a highly heterogeneous set of experimentation practices and information coming from a wide variety of sources and administrative domains is a rather challenging task due to the different semantics and terminology, different usage models (best effort vs. reservations), different authorisation policies, political issues, and more.

To achieve its ambitious objective, SFA defines a distributed and secure API that allows researchers affiliated with federated administrative domains to browse all the available resources and allocate those required to perform a specific experiment, respecting the agreed federation policies. Note that SFA is more of a standard specification than a specific implementation (there are currently different implementations for PlanetLab, ProtoGENI, and OpenFlow testbeds), which is being constantly updated based on the requirements that arise during its utilisation.

2.4.4 Network Services Framework

The Network Services Framework (NSF) concept is presented in Figure 2.7.

Network Service Framework concepts

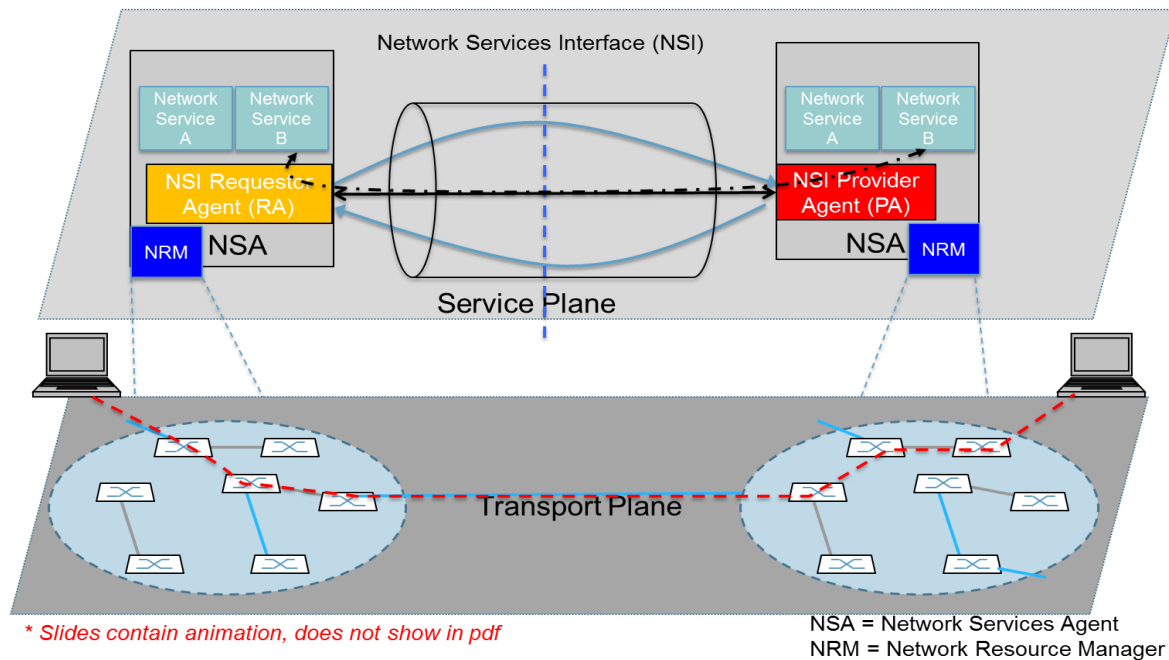


Figure 2.7: Network Service Framework concept [NSFv1.0]

The key components are as follows:

- **Service Plane.**
Abstraction of multi-layer, multi-domain, network capabilities for users, applications, network administrators.
- **Network Services Interface (NSI).**
Base interface between requestor agent and provider agent to request and get network services.
- **Composable Services.**
Ability to create a higher-layer, customised service with multiple network services to meet an application need.
- **Connection Service.**
First network service being defined carried by NSI.
- **Topology Service.**
Candidate for the next service to be carried by NSI .

2.5 Project Examples

OpenFlow projects all over the globe have been started including initiatives from USA, EU, Brazil, Japan, etc. Two of these are outlined below.

The OpenFlow in Europe: Linking Infrastructure and Applications (OFELIA) project [OFELIA] is a three-year project in the EU FP7 programme started in October 2010. The project creates a unique OpenFlow experimental facility offering virtualisation and control of the network environment through secure and standardised interfaces, enabling the researchers to experiment, to control and extend the network itself precisely and dynamically. Initially, ten partners, including universities and industry, formed OFELIA, and after the first open call another two partners joined. New partners will join from the second open call.

The Future Internet Testbeds Experimentation between Brazil and Europe (FIBRE) project [FIBRE] is a collaboration between the European partners and Brazil to create a shared Future Internet (FI) experimental facility to support joint experiments in network infrastructure and distributed applications. It will be based on the federation and enhancement of the OFELIA and OneLab projects. The FIBRE project includes 15 partners in total, where 5 partners are from the EU, 9 from Brazil and 1 from Australia.

2.6 Hardware and Control Software Maturity

More and more hardware vendors are releasing new OpenFlow-enabled commercial products supporting various features [SDNProducts]. A detailed vendor review is provided in Chapter 4 of this report; this section presents a brief overview.

- Broadcom as a switching silicon manufacturer produces OpenFlow-enabled chip sets, e.g., Trident/Trident+, supporting 10 GE / 40 GE, which are used by hardware vendors as their switching fabric [Broadcom].
- Brocade provides a variety of OpenFlow-enabled products offering various features. The Brocade MLX series [BrocadeMLX] and the Brocade NetIron CER 2000 series [BrocadeNICER] routers deliver scalability, flexibility, performance and high reliability for service providers and enterprise networks, supporting high-performance Ethernet edge routing and MPLS applications in Metro Ethernet, data centre and campus networks. The Brocade NetIron CES 2000 series [BrocadeNICES] of switches provides IP routing and advanced Carrier Ethernet capabilities. Based on the Brocade Multi-Service IronWare software, these 10 GbE-capable 1U switches provide Carrier Ethernet service delivery at the network edge and for data centre top-of-rack server access.
- The Extreme BlackDiamond X8 series [ExtremeBD] offers 2.56 Tbps per slot capacity and more than 20 Tbps total switching capacity with 768 ports of line rate 10 GbE or 192 ports of line rate 40 GbE in a single chassis. It has the ability to automatically track and report 128,000 virtual machines (VMs) and automatically migrates Virtual Port Profiles (VPPs) in a touch-less manner through XNV™ (ExtremeXOS® Network Virtualization). OpenFlow support will be available in the near future.

- HP offers a variety of OpenFlow-enabled products from edge switches to high performance and capacity switches. However, recently it was reported that OpenFlow-enabled HP products do not operate at line rate. HP products vary from campus networks switches (HP 3500 series [HP3500]) to chassis-based (HP 5400 zI Switch series [HP5400]), and modular switches (HP 8200 zI Switch series [HP8200]), supporting various combinations of interfaces e.g. 24-port and 48-port 1 GE and versatile 10 GbE connectivity (CX4, X2, and SFP+) with the integrated L2 to L4 intelligent edge features and HP AllianceONE solutions support.
- The IBM RackSwitch G8264 and IBM RackSwitch G8264T [IBMG8264] are built to accommodate the needs of data centres and cloud networks delivering line-rate, high-bandwidth switching, filtering, and traffic queuing. The RackSwitch G8264 is optimised for applications requiring high bandwidth and low latency. It supports IBM Virtual Fabric, to help clients reduce the number of I/O adapters to a single dual-port 10 Gb adapter, and OpenFlow. In addition, it can support 48 1/10 GbE SFP+ ports, 4 40 GbE QSFP+ ports in a 1U form factor. The RackSwitch G8264T is suitable for flexible and low-cost connectivity for 10 Gb environments supporting distances up to 100 m with 48 10GBase-T ports and 4 40 GbE QSFP+ ports in a 1U form factor.
- The NEC ProgrammableFlow PF5820 [NECPF5820] and PF5240 [NECPF5240] switches along with the NEC's ProgrammableFlow controller provide full functionality on OpenFlow-enabled networks with line rate multi-layer switching capability. Both NEC switches deliver extensive network virtualisation, multipath, security, and programmability capabilities. The PF5820 supports 1.2 Terabits per Second non-blocking bi-directional throughput in a 1U form factor including 48 1G/10G (SFP+) and 4 40G (QSFP) interfaces. The PF5240 supports 176 Gigabits per Second non-blocking bi-directional throughput in a 1U form factor including 48 1 GE ports and 4 1 GE / 10 GE SFP+, being able to maintain up to 160,000 flows.
- NetGear, with support from Broadcom's chip set and Indigo firmware, provides the GSM7352Sv2 ProSafe Gigabit Ethernet L3 Managed Stackable Switch [NETGEARGSM7352Sv2] in various combinations of 24 and 48 1 GE ports and 2 10G SFP+ ports, supporting high-speed, high-density workgroups at the edge of the network and in the backbone.
- Last but not least, the recently founded Pica8 [Pica8] offers a variety of OpenFlow products with various combinations of interfaces ranging from 48-port 1 GE to 16-port 10 GbE QSFP+ with XORPlus OS, Indigo and Open VSwitch software solutions.

Regarding control software and commercial OpenFlow controllers, various alternatives are available, either free or licensed.

With the free option, there is a wide range of OpenFlow controllers coded in different programming languages, e.g. Java, C++ or Python, that provide the basic OpenFlow communication and functionality, e.g. NOX and POX [NOXRepo], Beacon [Beacon], Floodlight [Floodlight] and Maestro [Maestro].

On the other hand, hardware vendors and software companies have developed and released commercial OpenFlow-capable platforms that support a wide range of features, e.g. centralised or distributed control, network virtualisation, load balancing, etc. Some of them are listed below.

- The NEC ProgrammableFlow Controller and Management Console [NECPF] provide network-level virtualisation, with centralised control that provides easy deployment, management, control and visibility of all end-to-end flows.
- The Nicira Network Virtualisation Platform (NVP) [NiciraNVP] creates an intelligent abstraction layer between end hosts and an existing network. NVP consists of a virtual switch component (OVS) and a clusterable controller system that allows programmatic management of large numbers of isolated virtual networks.
- The Embrane Heleos [Embraneheleos] is a multi-service, distributed software platform for powering on-demand elastic network services such as load balancers, firewalls, virtual private networks (VPNs) and WAN optimisation. Heleos delivers agility, rapid procurement and provisioning, multi-tenancy at scale, programmability and cost savings users expect with the cloud.
- The ConteXtream Grid [ContXtream] solution is an SDN-based network virtualisation platform that forms a software abstraction layer on top of traditional networks, creating an any-to-any network grid.

3 Investigation of OAM Mechanisms with OpenFlow

3.1 Introduction

The reason for investigating the impact of the Operation, Administration and Maintenance (OAM) mechanism in an OpenFlow-driven environment is the possibility of offloading the processing of OAM messages from the vendor hardware and moving this to a centralised OpenFlow controller.

At the same time, OpenFlow is meant to focus only on forwarding and be transparent to most OAM mechanisms, while the aspect of how the OAM mechanisms interact with OpenFlow represents an open discussion.

3.2 Ethernet OAM

As Ethernet represents a Layer 2 technology defined for operating on Local Area Networks (LANs), it does not include specifications for OAM capabilities to allow it to be used as a carrier-grade transport technology.

In order to compensate for this drawback, a number of Ethernet OAM standards have been developed to enable Ethernet to be converted into Carrier Ethernet. The IEEE 802.1ag standard, also known as Connectivity Fault Management (CFM), defines a protocol for running the Ethernet OAM aspects.

The standard describes the new entities – maintenance domains, maintenance points – the relationship between maintenance domains and also the procedures used by the maintenance points to manage the connectivity faults inside a maintenance domain.

The management of the Ethernet OAM messages is related to a couple of key elements depicted in Figure 3.1:

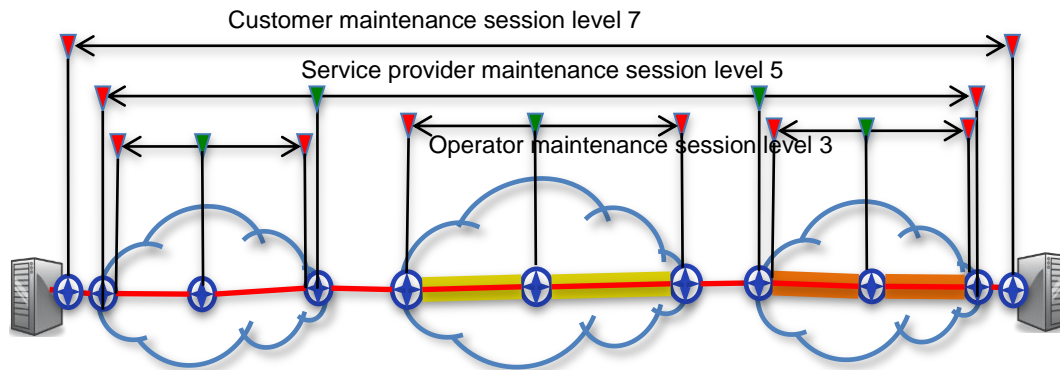


Figure 3.1: A hierarchy of maintenance sessions according to IEEE 802.1ag specification [DJ1.1.2]

- Maintenance Domain (MD).
 - Represents a part of the network that is owned or operated by a single entity. There is a hierarchical relationship between the different domains, which is represented as levels; the larger the domain, the higher the level value assigned to it. The levels are defined from 0 to 7, ranging from the smallest, representing the Operator Domain, to the Service Provider Domain and to the largest one – the Customer Domain.
- Maintenance Association/Session (MA).
 - Consists of a number of Maintenance Points (MPs) all configured with the same Maintenance Association Identifier (MAID) and MD level.
- Maintenance End Point (MEP).
 - Represents the Maintenance Point placed at the edge of the domain.
 - Sends and receives CFM frames
 - Drops all CFM frames coming from the link with an MD level similar to or lower than its own.
- Maintenance Intermediate Point (MIP).
 - Represents the Maintenance Point placed inside the domain (not at the edge).
 - Receives CFM frames from MEPs and MIPs, classifies and forwards them further.
 - Drops all CFM frames with a lower MD Level than its own.

Regarding the actions the MPs are taking in relation to the Maintenance Levels (Figure 3.1), the packets with a higher level are forwarded; the ones with the same level are processed, while the packets with a lower level are dropped.

Three types of messages are sent by the MPs associated to the Ethernet ports, as also described in [vdPol]: Loopback Messages, Link Trace Messages and Continuity Check Messages.

- Loopback Messages (or L2 ping).
 - Used for service troubleshooting by checking for reachability.
 - Request/Reply.

- Sent manually from MEPs via the Command-Line Interface (CLI.)
 - Processed by MIPs/MEPs.
 - Unicast request, unicast reply.
- Link Trace Messages (or L2 traceroute).
 - Used for service troubleshooting by requiring path information.
 - Request/Reply.
 - Sent manually from MEPs via the CLI.
 - Processed by MIPs/MEPs along the path.
 - Multicast request including Time To Live (TTL), unicast replies.
- Continuity Check Messages (or periodic hello messages).
 - Used for service status monitoring.
 - Detect loss of connectivity.
 - Periodic hello messages sent by MEPs.
 - Processed by MEPs only.
 - Continuity Check (CC) frames sent to multicast group; no replies are required.

3.3 Integration of OAM in OpenFlow – Ethernet OAM Scenario

As described in [vdPol], an Ethernet OAM-enabled OpenFlow controller solution was proposed by Ronald van der Pol, from SARA, which used the OpenFlow switches (as in Figure 3.2) to forward the IEEE 802.1ag packets, used to monitor links, to the OpenFlow controller. Further on, the OAM packets are sent to the daemon of the open source implementation of IEE 802.1ag (dot1ag-utils) and are stored in a Round Robin Database (RRD). perfSONAR is then used to publish the current link status from the RRD file.

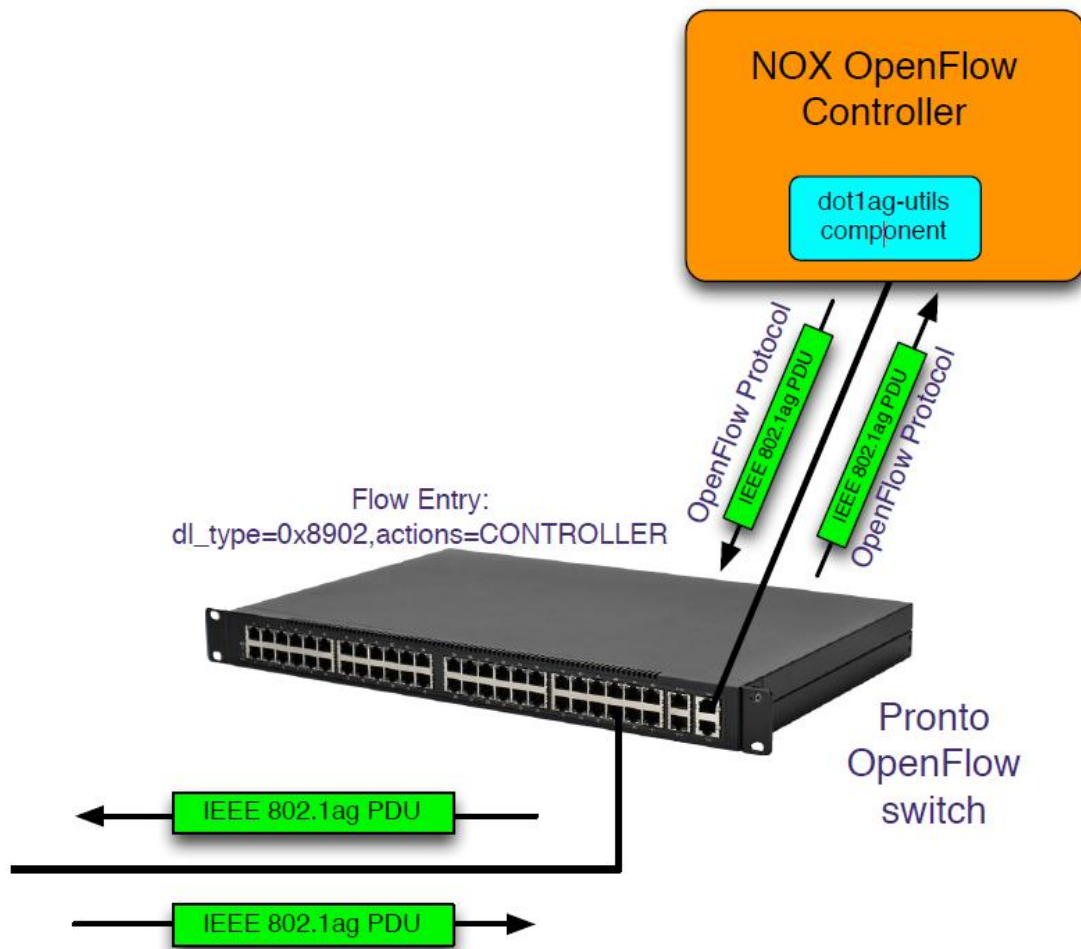


Figure 3.2: OpenFlow Switch forwarding IEEE 802.1ag messages to the OpenFlow controller

A significant benefit of implementing the IEEE 802.1ag protocol in the OpenFlow controller is that the OAM provided by this protocol can then be used by any switch that supports the OpenFlow API.

There is still ongoing work in the project presented in [vdPol], and the planned development is to define the forwarding for the 802.1ag frames, fully implement MEP functionality and also ITU Y.1731 frame loss and frame delay measurements.

In [CloudToad], Cloud Toad proposes to define the OAM mechanism in the OpenFlow specification, with the purpose of enabling all vendors to support the same specific functionality. The OAM messages are intended to have specific flow criteria embedded (e.g. L2 overhead) so that the receiving MIP can match it against its lookup structure. The results would be that, following a TRACE between an ingress and an egress node, the OpenFlow controller can discover if the path followed by the OAM message is not the same as the one specified in the forwarding tables and also if there are possible delays along the path.

As far as it has been presented to date in the literature, it is still an open debate whether the OAM functionality is going to be provided by the switches or routers, and whether vendors will include some solutions for the OpenFlow controller to access specific OAM features in the hardware or whether the controller will handle the OAM mechanism entirely.

4 Review of Vendors' Approaches

This chapter considers the OpenFlow/SDN approaches of twelve vendors, covering aspects such as introduction of OpenFlow/SDN, current status, collaborations and future plans. The twelve vendors are:

- Ciena.
- Juniper.
- Cisco.
- Brocade.
- Big Switch Networks.
- Arista Networks.
- IBM.
- HP.
- Dell.
- NEC.
- Pica8.
- ONF.

4.1 Ciena

4.1.1 Current Status and View of OpenFlow / SDN

In October 2012, Ciena publicly announced for the first time the support of OpenFlow on their equipment [CienaOFStripes]. Even though OpenFlow protocol implementation is just in pre-production phase at Ciena, they demonstrated interoperability of the 5410 Service Aggregation Switch with other software and equipment from the ONF members at the ONF's second Plugfest event.

Ciena not only demonstrated OpenFlow support on packet equipment, but also emphasised their focus on "openness at both API levels" (API for applications to network control software and APIs enabling network control software to network physical equipment). An important goal for Ciena in attending the Plugfest was to

get useful feedback, to be used for defining the final aspects of the production-ready OpenFlow-enabled equipment.

In November 2012, Ciena announced their SDN strategy as well as latest plans for OpenFlow support, separating this into two dimensions: equipment and software [CienaSDNStrat].

4.1.1.1 *Ciena SDN Strategy on Hardware*

Ciena are committed to providing OpenFlow support on their packet portfolio, in close alignment with their “OPⁿ vision” for open interfaces and optical packet networks.

According to [CienaSDNStrat], Ciena are planning to release OpenFlow support on their latest generation of packet equipment:

- The 5400 product family, including the 5430 Reconfigurable Switching System and the 5410 Service Aggregation Switch.
- Two platforms in the 5100 product family (the 5160 and 5142 Service Aggregation Switch).

By providing OpenFlow support on this equipment, Ciena want to enable their customers to explore, trial, validate and deploy the OpenFlow equipment in real case scenarios in both data centre and metro/access infrastructures.

4.1.1.2 *Ciena SDN Strategy on Software*

According to [CienaSDNStrat], Ciena are working in close collaboration with customers and partners to establish a deployable “carrier grade and scale” SDN network control layer foundation framework.

Ciena are also adapting key functional kernels related to the automated and northbound API-based interactions with business applications from the V-WAN Hypervisor product while adding to the OneControl-based southbound equipment layer control with OpenFlow southbound support. Based on this approach, Ciena are planning to deliver control plane applications that will enhance the control plane capabilities (with the goal of expanding the utility of the SDN software- based operational framework for their customers).

Ciena are also focusing on developing software-based solutions that can be enabled by the implementation of the SDN architecture. One such example is related to real-time, big data analytics [CienaSDNStrat].

4.2 Juniper

4.2.1 Introduction of OpenFlow / SDN

Juniper Networks announced in June 2012 their first SDN strategy, initially focusing on data centre solutions as a safe approach. This was probably due to the fact that SDN was still a growing, not a mature, technology, whereas the data centre is a more reliable environment where a failure can be handled more easily than in an open environment.

According to [JuniperSDN], Juniper's main focus regarding SDN strategy is on a northbound API between the network devices and the orchestration layer in the SDN environment, while reducing the use of the OpenFlow protocol (but not its role). Juniper see an advantage in the fact that they are using the same operating system (OS), JunOS, on all their devices (while Cisco are using three OSs), which is a premise for a faster SDN integration (due to the fact that it's easier to make a controller that is only controlling one type of OS).

4.2.2 Current Status and View of OpenFlow / SDN

Juniper's SDN strategy is presented in Figure 4.1:

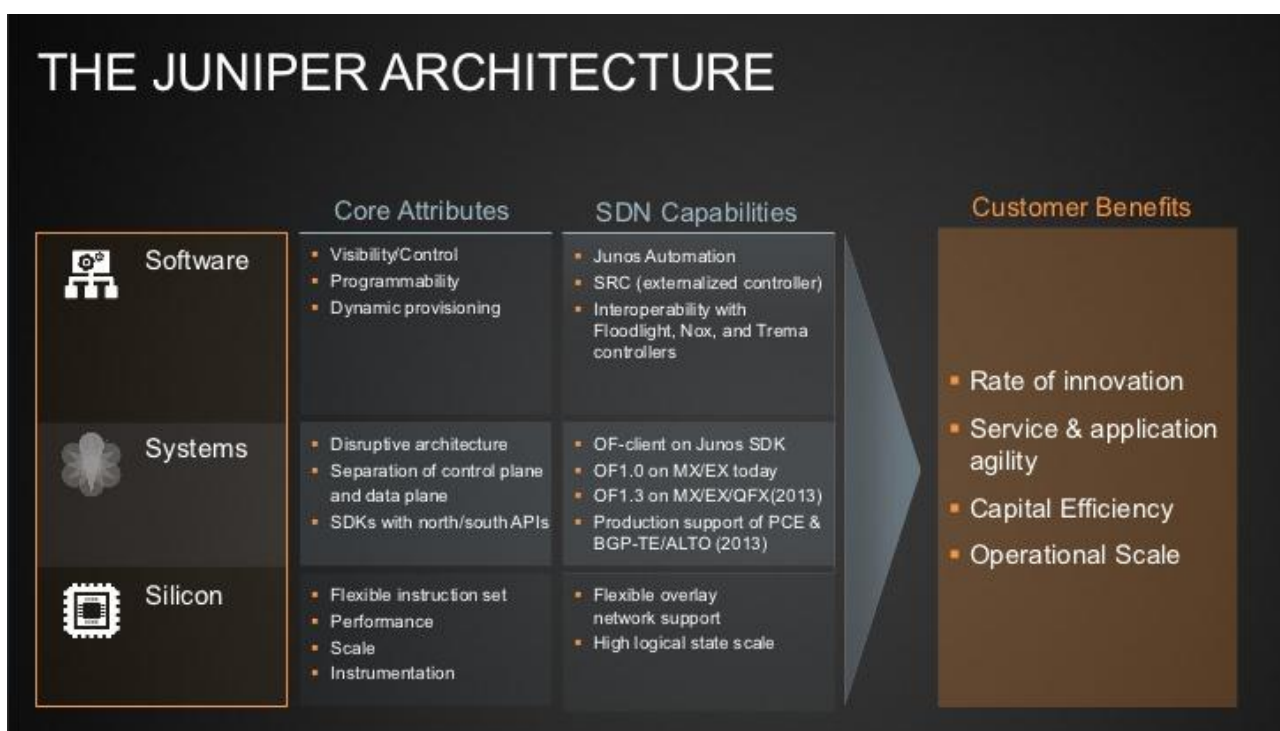


Figure 4.1: Juniper SDN architecture

4.2.3 Junos Software Development Kit

Junos Software Development Kit (SDK) allows developers to build custom applications on top of JunOS, to expand or to create additional functionality for Juniper equipment.

OF-Client 1.0 is enabled on Junos SDK, which makes Juniper's routers and switches programmable through software.

4.2.4 Juniper Controller

According to [JuniperOpenSource], Juniper are involved with other industry players in developing an open source-based controller in support of SDN that would offer an alternative to the proprietary solutions from VMware, Cisco and others. As Bob Muglia, executive vice president of Juniper Software Solutions Division, says, Juniper's goal is to have an open source SDN controller that is a de facto standard widely supported by the industry, in a similar manner to Linux or Apache, which have become the standard for open source operating systems or web servers.

It is still unclear whether Juniper will brand their own SDN controller or source another one. However, [JuniperOFCntrl] indicates that a decision is imminent as to whether they will have their own or will deploy it in cooperation with possible partners such as Big Switch or existing open-source controller projects. Juniper insist on the fact that while the controller will support the OpenFlow protocol, it will not be based solely on it; at the same time, Juniper believe that control doesn't have to be done on a per flow basis [JuniperOpenSource].

4.2.5 Juniper OpenFlow-Enabled Equipment

Juniper MX routers and EX switches (Figure 4.2) are OpenFlow 1.0 enabled.



Figure 4.2: Juniper EX Series Ethernet Switches

4.2.6 Collaborations

According to [JuniperSDNVisio], in 2011 Juniper and Big Switch demonstrated an OpenFlow implementation on a carrier-grade router (Juniper MX-series) connected to a Big Switch controller.

4.2.7 Future Plans

OpenFlow 1.3 will be enabled on the MX, EX and Qfabric QFX lines in 2013.

4.3 Cisco

4.3.1 Introduction of OpenFlow / SDN

In June 2012 Cisco introduced their official SDN architecture, also called Cisco Open Network Environment (Cisco ONE). However in Cisco's view of SDN architecture, OpenFlow is just an enabler of the programmability of the hardware from the centralised control plane, and is not the only protocol that should be considered for doing this. Cisco claim that the whole idea of SDN is more important than just the method of communicating from the separated control plane to the forwarding planes.

4.3.2 Cisco ONE

Cisco ONE is designed to be a flexible and programmable solution that allows customers to customise it easily.

Cisco are taking a phased approach to the development of Cisco ONE [CiscoSDNsuspense], meaning that customers will implement the architecture in phases, which makes it easier for Cisco to analyse the impact and decide on progress.

Cisco ONE contains APIs, agents and controllers, as well as overlay network technologies that enable the programmability of the different layers based on the customer's needs. The goal, in Cisco's view, is to open up all the layers of the network and go beyond the decoupling of control plane and forwarding plane via OpenFlow, and to be able to use a greater variety of southbound API protocols in the SDN environment and not be limited to OpenFlow. Examples of such additional protocols include DevroFlow [DevroFlow], NETCONF, OpenStack's Quantum application, and other vendor-dependent solutions [NorthboundAPI]. The way Cisco ONE is able to provide more benefits than OpenFlow alone, according to David Ward, Chief Technology Officer of Cisco's Service Provider division, is the use of the northbound API which is able to provide counters, analytics and statistics from the routers to the programming applications [CiscoOF] and through taking more intelligent decisions already from the application level.

Cisco ONE also includes the One Platform Kit (onePK), which provides APIs for developers working with Cisco's operating systems (IOS, IOS-XR, NX-OS) [CiscoSDNsuspense]. The Cisco onePK will be supported on different Cisco platforms in phases; initially, the ASR 1000 and ISR G2 routers will be onePK enabled.

Cisco are looking into five different target markets for their Cisco ONE programmable network architecture: academia/research, enterprise, service provider, cloud service provider and data centre. Compared to their competitors, e.g. Juniper, this is clearly a wider focus (Juniper are mainly looking into data centre environments).

4.3.3 Cisco OpenFlow-Enabled Hardware

Cisco have announced that they will provide "proof-of-concept" OpenFlow v1.0 agents on the Catalyst 3750-X and 3560-X series switches [CiscoSDNsuspense], but generally they were slow in adding OpenFlow support in their switches. However, so far the only purpose of supporting OpenFlow on Cisco hardware is to enable research and academia to test it on Cisco equipment.

Cisco also planned to support OpenFlow on their Nexus data centre switches but it is still uncertain whether this will happen soon as Cisco still see OpenFlow deployment as being of more interest in the research community rather than in data centre production environments.

According to [CiscoCatalyst], the Catalyst 6500 series might be one of the next Cisco OpenFlow-enabled hardware families, and will also get software-defined networking support working as an OpenFlow aggregation node that interfaces with conventional network architecture.

4.3.4 Future Plans

According to [CiscoThreat], Cisco have adopted an SDN strategy that is broader or more complete than just focusing on OpenFlow. Executives of Cisco call this a "wait-and-see attitude" to adopting the OpenFlow technology, meaning that Cisco are following the trend and waiting for enough research to be done around OpenFlow before deciding how this can be integrated and what the benefits in the long term could be.

4.4 Brocade

According to [BrocadeOFswitch], Brocade's Software-Defined Networking strategy focuses on wide area networks (WANs) and service provider networks rather than data centre and campus networking. This is the reason why Brocade's initial OpenFlow support is available on the MLX router.

4.4.1 Introduction of OpenFlow / SDN

In June 2010, according to [BrocadeTechDay], Brocade first presented their vision of virtualisation, cloud, convergence, building the virtual data centre, application optimisation and new opportunities for service providers, and also introduced OpenFlow as one of the first major networking vendors looking into software-defined networking.

4.4.2 Current Status of and Strategy on OpenFlow / SDN

In May 2012, as presented in [Brocade100GE], Brocade outlined their strategy for software-defined networking. In support of this, Brocade announced they they had integrated hardware-based OpenFlow support in the Brocade® MLX® Series of routers and related Brocade NetIron® platforms, enabling their customers to deploy SDN at wire-speed of 100 GbE performance.

Brocade's focus for their SDN investments is on network virtualisation, automation and simplification in large-scale data centres, as well as traffic engineering and flow management on high-speed networks.

Brocade's strategy provides an upgrade path to SDN, while letting the operators keep their existing services and slowly switch to the new software-defined services in a way that minimises the risks – the so-called hybrid solution.

Brocade's view of network layering and the implementation of the OpenFlow-based SDN stack, based on the ONF reference model, is presented in Figure 3 [BrocadeExploreSDN]:

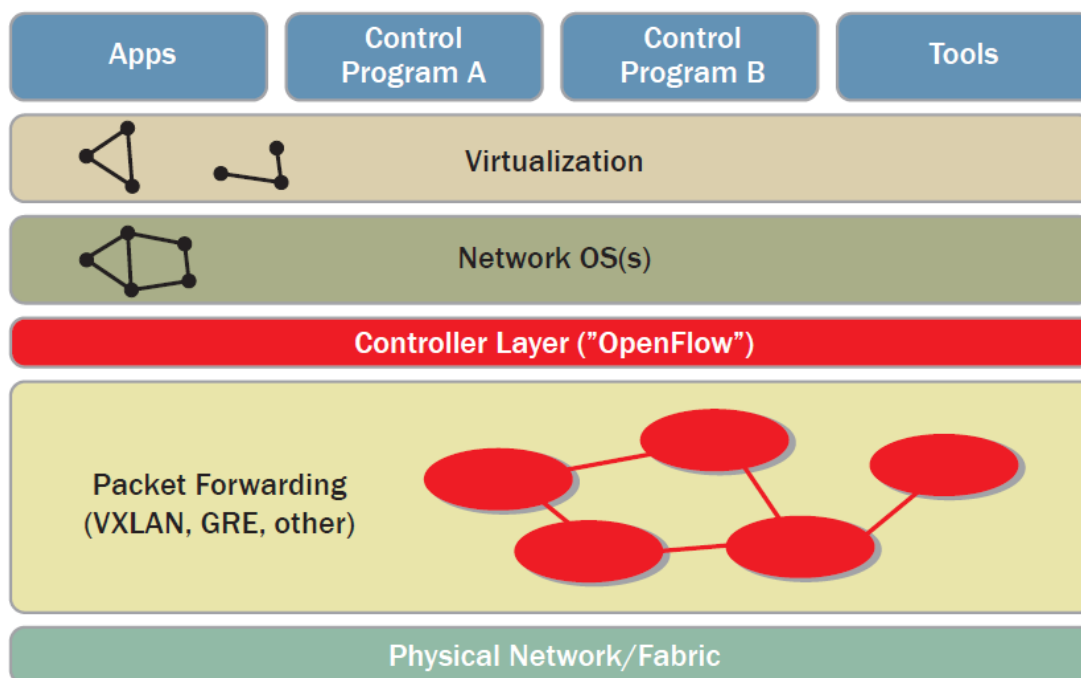


Figure 4.3: Brocade's view of the SDN stack

Brocade present in [Brocade100GE] a number of key technologies and capabilities that compose the SDN strategy and are intended to reach the goal of offering a highly flexible cloud-optimised network solution.

According to Ken Cheng, Vice President of Service Provider Products at Brocade, SDN will transform networking infrastructure into a platform for innovation, enabling Brocade customers to deliver new services and applications faster and to scale. This is done through OpenFlow and open service API support (in Brocade's routing and data centre platforms). Brocade is the first in the industry to enable OpenFlow support in high-performance 100G networks.

According to [Brocade100GE], "With programmatic control of network infrastructure through partnerships with a broad set of OpenFlow controller vendors, Brocade enables service innovation via open APIs. Specifically, Brocade is delivering industry-standard OpenFlow for Layer 2/3 forwarding and the Brocade OpenScript™ engine for Layer 4/7 switching to unlock the network to increase service velocity for highly customised services."

4.4.2.1 Tunnel-Agnostic Solutions

Brocade will support extensive overlay technology to enable network virtualisation. Brocade's approach to overlay networks will be tunnel-technology agnostic. According to [Brocade100GE], Brocade products have a tunnel-technology-agnostic design that supports overlay technologies (Virtual Extensible LAN (VXLAN) and Network Virtualisation using Generic Routing Encapsulation (NVGRE) as well as emerging standards (Figure 4.4)).

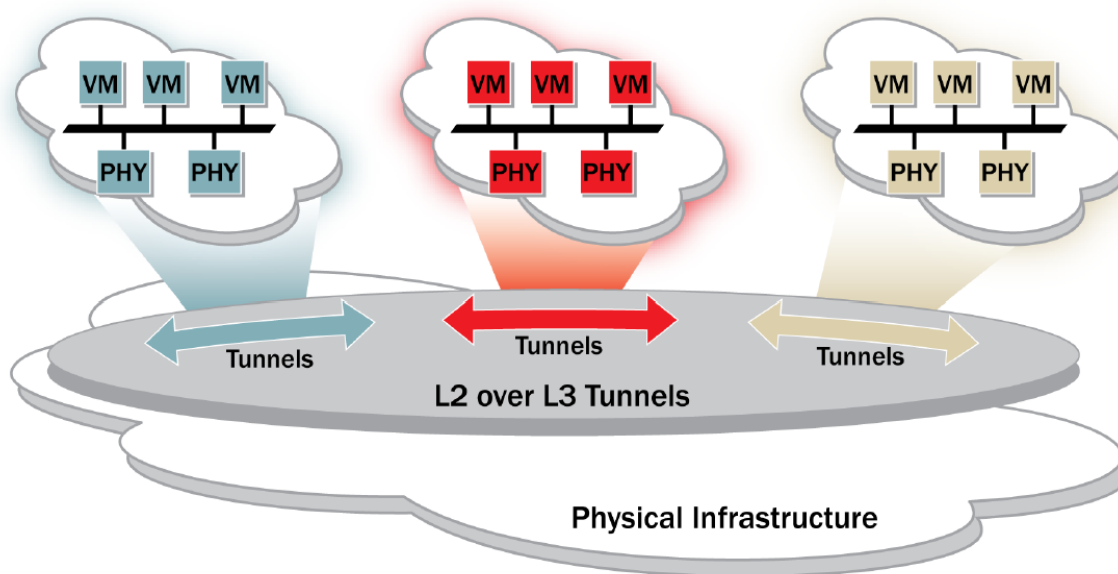


Figure 4.4: Brocade support for standards-based overlay technologies

4.4.2.2 Cloud Management Capabilities

Ken Cheng says that Brocade will provide rich cloud management capabilities through the Open Stack framework with standards-based plug-ins and RESTful interfaces.

According to [Brocade100GE]], “Brocade is providing a single common cloud management and orchestration interface through northbound standards-based plug-ins and standards-based RESTful interfaces. The plug-ins to cloud management frameworks – including Cloudstack, Microsoft System Center, OpenStack, and VMware vCenter/vCloud™ Director – provide comprehensive orchestration capabilities for cloud service delivery and automated operations.”

4.4.2.3 Partnerships to Deliver High-Value Network Applications

Ken Cheng says that Brocade will partner with industry technology leaders and network operators to deliver high-value network applications including advanced traffic engineering, advanced flow management and large-scale data centre management.

“To enable SDN deployments, Brocade supports OpenFlow 1.0 on its carrier-grade products, and Brocade is committed to interoperating with any controller that supports the OpenFlow 1.0 standard. Brocade believes that flexibility and interoperability enable the most efficient and valuable SDN deployments for network operators. With this in mind, Brocade enables OpenFlow in its hardware. This means that OpenFlow on Brocade routers and switches runs at line rate for any interface speed, from 1 GbE to 100 GbE.” [BrocadeNetworkTrans]

4.4.2.4 SDN Optimised by Ethernet Fabric

Brocade believe that Ethernet fabric-based networks are critical to realising the full potential of SDN. Brocade Virtual Cluster Switching (VCS) fabric technology offers key features and functionality to enable infrastructure scalability, virtualisation and management simplification, making it ideally suited for SDN-based cloud deployment.

According to [Brocade100GE], “Resilient and auto-forming Ethernet fabrics will enhance SDN. Brocade VCS® Fabric technology is optimised for virtualised applications through its active-active topology and virtual machine (VM)-aware intelligence that enables seamless VM mobility without manual configuration, including the network policies associated with each VM. In addition, multiple switches that make up the fabric can be managed and programmed as one logical switch, dramatically reducing operational complexity and improving SDN controller scalability. Finally, VCS fabrics, with a self-forming and self-healing architecture, are highly resilient to help increase availability.” This transformation is shown in Figure 4.5.

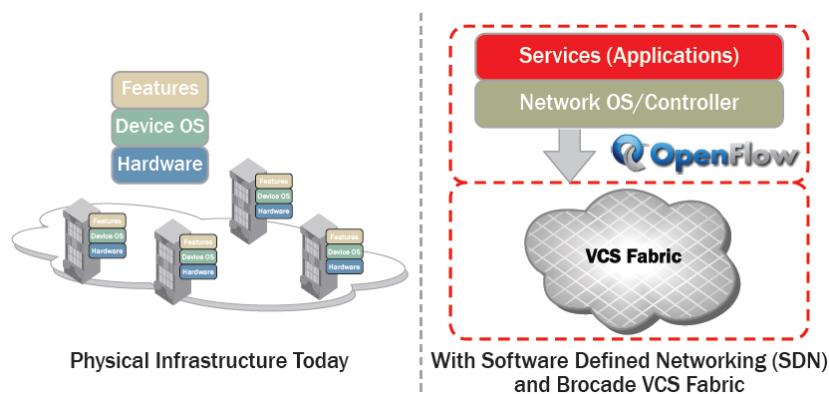


Figure 4.5: Brocade's change from traditional physical infrastructure to the VCS Fabric-based SDN [BrocadeNetworkTrans]

4.4.2.5 OpenFlow in Hybrid Mode

According to [Brocade100GE], “Brocade will be the industry's first network vendor to deliver OpenFlow in hybrid mode¹. With Brocade Hybrid Mode, customers can simultaneously deploy traditional Layer 2/3 forwarding with OpenFlow. This unique capability enables network operators to integrate OpenFlow into existing networks, giving them the programmatic control offered by SDN for specific flows while the remaining traffic is handled as before. The Brocade hardware support for OpenFlow enables customers to apply these capabilities at line rate in 10 GbE and 100 GbE networks.”

4.4.3 Brocade SDN-Ready Product Family

Brocade offer an SDN-ready product family which includes:

¹ “Hybrid mode” in this context has a specific meaning and does not denote IP plus optical circuit.

- Brocade VDX® family of data centre switches.
- Brocade ADX® Series application delivery switches.
- Brocade MLX Series of 100 GbE routers.
- Brocade NetIron CER Series routers.
- Brocade NetIron CES Series switches.

Brocade NetIron® 5.4 OS provides:

- Hybrid-mode OpenFlow support.
- Layer2/Layer3 forwarding.
- High-performance 100 GbE.

“The upcoming Multi-Service IronWare software release 5.4 for NetIron Family of Routers for MLX, XMR, CES and CER will support OpenFlow version 1.0. These will be the first carrier-grade systems to support OpenFlow. OpenFlow will be supported as part of a normal software release. That is, OpenFlow is just another feature to be supported on these routers along with all other supported features, such as Ipv4-6 routing, MPLS, VPLS, PBB, etc. This is not an experimental feature or prototype. OpenFlow will be available as a fully supported feature on these routers.” [BrocadeOFComing]

The Brocade products that come with software already enabling OpenFlow, or that get a software upgrade, can enable OpenFlow on a per-port basis. This allows the equipment to run some of the ports controlled by OpenFlow/Controller while the others can continue working on the existing configuration (the typical hybrid-switch configuration).

Brocade prioritised OpenFlow development on its 100 GbE Brocade MLXe® Series routers, along with the Brocade NetIron® CES and CER edge platforms. This allows users to create logical networks for multi-tenancy purposes via L2 over L3 tunnelling, according to [BrocadeNetworkTrans].

“Brocade is also building an ecosystem of partners for SDN through strategic investments, solution designs and interoperability testing. Brocade recently unveiled an OpenFlow lab in Japan designed to demonstrate Brocade and partner technologies to aid in the development of tested and validated solutions. Brocade is also a member of the Indiana Centre for Network Translational Research and Education (InCNTRE), which seeks to advance development, increase knowledge, and encourage adoption of OpenFlow and other standards-based SDN technologies. Brocade has also participated in a number of industry multi-vendor demonstrations with NEC, including at the Open Networking Summit (ONS) in Santa Clara, CA in April 2012, and Interop Las Vegas in May 2012.” These demonstrations are described below.

4.4.4 Brocade OpenFlow Demonstrations

4.4.4.1 NetIron CES

Brocade's first demonstration based on OpenFlow took place at Stanford University at the 1st Open Networking Summit (ONS) in October 2011. There, Brocade demonstrated a NetIron CES running an early OpenFlow implementation as part of the network virtualisation demonstration given by NEC using ProgrammableFlow Controller.

4.4.4.2 MLXe and CER #1

At the 2nd ONS in April 2012, Brocade demonstrated 10G and 1G multipoint flows running at wire speed. The demonstration focused on the actions of sending copies of the packets to multiple destination ports, while supporting VLAN tag modifications on a per-destination port basis. A schematic of the demonstration is shown in Figure 4.6.

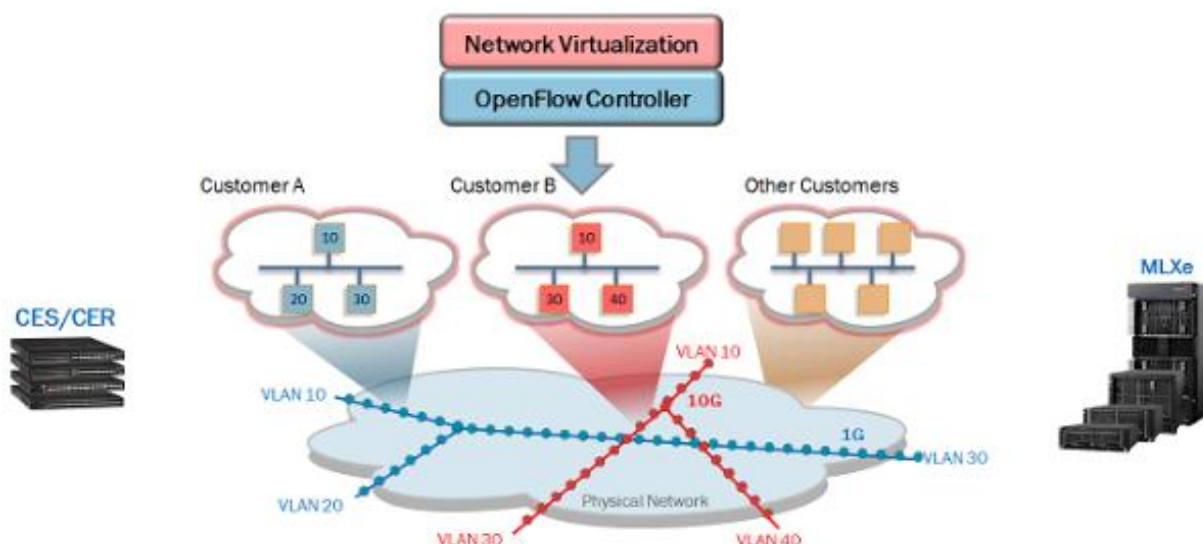


Figure 4.6: Brocade MLXe and NEC PFC demonstration

As usual, the equipment proposed for testing by Brocade was an MLXe and a CER router while the chosen controller was the NEC ProgrammableFlow Controller.

4.4.4.3 MLXe and CER #2

In May 2012, as part of Interop Las Vegas, Brocade again demonstrated the capability of MLXe and CER routers to run OpenFlow, while the MLXe was part of a NEC demonstration of a multi-vendor network fabric via OpenFlow.

In NEC's demonstration the aim was to show an interoperable OpenFlow network solution featuring Brocade, Extreme Networks, IBM, Intel, Microsoft and Radware controlled by the NEC ProgrammableFlow controller version 2.0.

4.4.4.4 MLXe and CER #3

During the 1st Open Networking Foundation (ONF) OpenFlow plugfest in March 2012, Brocade again presented their MLXe and CER routers interworking with NEC ProgrammableFlow Controller, Indiana University ON-SS (based on NOX), NTTData, Big Switch (FloodLight) and FlowVisor.

4.4.5 Collaborations

4.4.5.1 Brocade and NEC

As stated in [BrocadeNEC], Brocade and NEC Corporation announced on 22 May 2012 a collaboration intended to bring to market combined solutions based on SDN and OpenFlow switching, including network virtualisation. The solution will be looking into network virtualisation, large-scale data centre infrastructure management, traffic engineering and WAN flow management, based on open standards.

In the collaborative solution, Brocade contributes with the OpenFlow-enabled products while NEC introduces the ProgrammableFlow Controller. The NEC ProgrammableFlow Controller PF6800, based on OpenFlow standardised technology, used in combination with Brocade OpenFlow-enabled switches or routers, eliminates most of the configuration management in traditional networks.

4.4.6 OpenFlow Switches Operating at Line-Rate Speeds

According to [BrocadeOFswitch], a common issue with the OpenFlow implementation in the vendors' equipment is related to the way it is implemented: in dedicated hardware / Application-Specific Integrated Circuits (ASICs) or in software.

There is a problem with some of the switches which just implemented OpenFlow in their firmware but not running through the ASICs. It can also be the case where OpenFlow is implemented in dedicated ASICs but, due to hardware limitations, the OpenFlow controller can overwhelm the ASICs by sending more commands to a switch than the switch is capable of processing in its flow table. In such a case, the extra commands/rules have to be handled in software instead of hardware, which makes the switch speeds slower than line rate.

As presented in [BrocadeOFswitch], Brocade claim that their implementation of OpenFlow in the MLX routers does not suffer from this possible issue. It is, however, expected that this issue will start to disappear by the end of this year when the merchant silicon vendors start supporting OpenFlow on their own chips.

4.4.7 Brocade Acquiring Vyatta

As stated in [BrocadeVyatta1], Brocade announced in November 2012 the acquisition of Vyatta, a California-based provider of software-based network operating systems.

The main resources that make Vyatta an asset for Brocade are the network virtualisation, SDN and cloud computing platforms, which would be integrated with Brocade's Ethernet fabric architecture. As shown in [BrocadeVyatta2], the main advantages to Brocade of this acquisition are:

- Obtain software expertise.
- Obtain new customer opportunities in VMware environments (due to Vyatta's relationship with VMware).
- Reduce Brocade's reliance on OpenFlow for their strategy and the reliance on open-source controllers for SDN deployment.

According to Miky Klayko, CEO of Brocade, "the acquisition complements Brocade's R&D investments in Ethernet fabrics and SDN, as well as Brocade's broad industry and solutions-level partnerships that enable Brocade to pursue new market opportunities in data centre virtualisation, public cloud, enterprise virtual private cloud, and managed services." [BrocadeVyatta3]

4.4.8 Summary and Conclusions

Brocade first introduced support for OpenFlow in June 2010. They deliver OpenFlow on their switching and routing products, enabling flow control up to 100 GbE speeds.

In Brocade's view, OpenFlow is an enabler of SDN, which helps Brocade to meet the goal of having an open cloud-optimised network, built on a resilient, automated fabric foundation.

According to [BrocadeExploreSDN], Brocade believe that OpenFlow will initially be of greatest interest to service providers and large enterprises with early-adopter IT cultures and active cloud deployments. Virtually all of the ONF board members fit this profile, and it is likely that the first standardised OpenFlow applications will be of greater interest to IT organisations of this type than to the typical enterprise.

Brocade enable OpenFlow version 1.0 on the equipment shipped with an OS version 5.4 and also on older equipment if the software is upgraded to OS 5.4. OpenFlow works on a per-port basis, meaning that Brocade equipment can work in a hybrid mode. Also, all the ports can be OpenFlow-enabled, which allows OpenFlow to run over ports running from 1G to 100G line cards.

4.5 Big Switch Networks

Big Switch Networks is a start-up created by former Stanford MScs/PhDs and network engineers from Cisco, Juniper, Arista and VMware with the purpose of exploiting the new software-defined networking opportunities.

Another similar start-up, Nicira, was acquired by VMware, leaving Big Switch as the main independent start-up in the SDN market.

4.5.1 Introduction of OpenFlow / SDN

Big Switch's view of OpenFlow is shown in Figure 4.7:

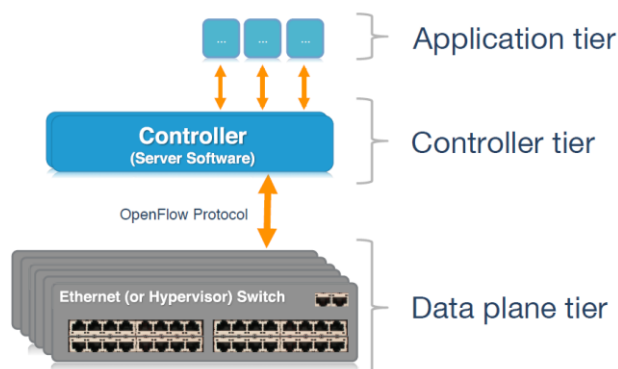


Figure 4.7: Big Switch's three-tier view of SDN

Big Switch are developing their own open source version of OpenFlow controller, as announced in autumn 2012. It is based on the open SDN controller Floodlight. Floodlight, as shown in Figure 4.8, is the only available Apache-licensed controller and is developed by an open community of developers.

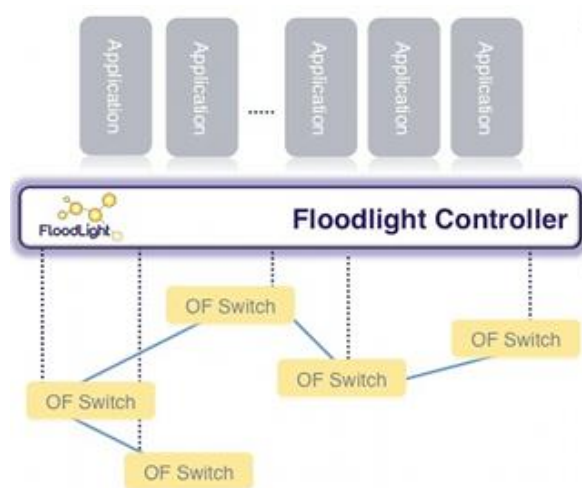


Figure 4.8: Floodlight controller in SDN environment

Based on their three-tier view of software-defined networking, Big Switch are planning to provide to the public domain their own open source controller for the controller tier. In this way the code is exposed and rapidly

hardened, and becomes more attractive for the clients and partners by letting them choose the deployment of application tier from different vendors.

Also, according to [BigSwitchSDNPrimer], Big Switch's version of the Floodlight controller is in use by more than 15 companies in Fortune magazine's Fortune 500. Their controller is trying to differentiate from the open source Floodlight by providing a combination of physical and virtual environment support, integration with the existing networking standards, abstraction of the underlying network infrastructure complexity, and a platform that is open for third parties in the ecosystem to build their own applications internally.

The commercial target for Big Switch is the application tier, where Big Switch plan to build applications ranging from rudimentary ones, such as multi-switch forwarding models and topology discovery, to more advanced services such as load balancing or firewalls. Also, as shown in Figure 4.9, Big Switch will help the partners to build their own applications on top of the Big Switch commercial controller.

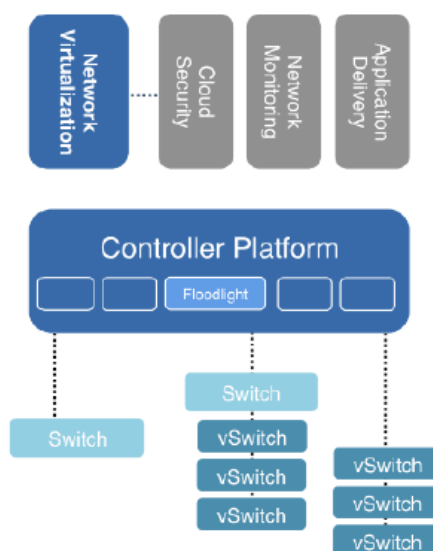


Figure 4.9: Big Switch commercial target: controller platform and application development on network virtualisation

In [BigSwitchFloodlight], Big Switch announced in August 2012 that they had delivered over 6000 downloads of the Floodlight OpenFlow controller and APIs. This emphasises the fact that the possible SDN customers that are evaluating the Big Switch SDN solutions can rely on a large number of OpenFlow physical and hypervisor switches and third-party applications from which to choose.

4.6 Arista Networks

4.6.1 Introduction of OpenFlow / SDN

Arista is one of the young companies that found an opportunity in the market due to the new merchant silicon improvements. As merchant silicon (e.g. Broadcom, Marvell or Fulcrum) is lately able to provide similar performance at a lower price compared to the ASICs provided by vendors such as Juniper and Cisco, commodity hardware is a good opportunity for Arista to add their software on top of the cheaper boxes and differentiate themselves from the big vendors without needing big investments.

Arista's main focus is towards data centre switching rather than the enterprise switching market.

According to a presentation given by Anshul Sadana, Senior Vice President of Customer Engineering at Arista Networks, the 7150 Switch series (shown in Figure 4.10) was designed in support of SDN via the Arista Extensible Modular Operating System (EOS), which is controller-friendly software, with the APIs talking to the controllers, and uses OpenFlow.

For SDN, Arista see many models emerging in the market, while some of the use cases require new features on the hardware which generally take 2-3 years to get into the next generation of ASICs. The 7510 switch has the advantage of having a flexible forwarding path and can add new frame formats, or forwarding behaviour, by only changing EOS. This allows the support of new features like VXLAN and operations at wire speed, again by just changing the EOS software.



Figure 4.10: Arista 7150 Switch series

Arista are implementing OpenFlow 1.0 in the 7150 Series.

According to Arista's specifications, the main advantages of the 7150S switch are:

- Architected for leading-edge applications including Big Data, Cloud Networks, Financial Trading, HPC and Web 2.0 environments.
- Accurate and predictable performance in a range of densities (deterministic high performance, low latency and jitter for all traffic).

- Balanced resources and deployment flexibility for “Any Application” suitability.
- Network-wide virtualisation platform for support of next-generation cloud-bursting (i.e. flexibility and versatility in the cloud) with wire-speed VXLAN hardware-based Tunnel Endpoint termination.
- Support of sub-microsecond Network Address Translation (NAT).
- Deterministic latency and high-performance 10 GbE / 40 GbE switch and IEEE 1588 platform.
- Advanced forensics and monitoring capabilities – redefining instrumentation, automation and analysis of high-end infrastructure.

As described in Arista’s specifications, the 7510 switch series comprises three models with 24, 52 and 64 10 Gb ports, while all three models support wire-speed performance at both L2 and L3. The 7150S-24 and 7150S-52 switches support SFP+ on all ports, while 7150S-64 has 48 SFP+ ports and 4 QSFP+ ports (the difference is that SFP+ ports support 100 Mb, 1G and 10G operation with 100/1000-TX transceivers, while QSFP+ ports support 40 GbE and 4x10 GbE modes) as described in Table 1.

Aspect	7150 Switches		
	7150S-24	7150S-52	7150S-64
Description	7150 Switch 24-Port SFP+	7150 Switch 52-Port SFP+	7150 Switch 48-Port SFP+ 4 QSFP+
Total Ports	24	52	64
SFP+ Ports	24	52	48
L2/3 Throughput	480 Gbps	1.04 Tbps	1.28 Tbps
L2/3 PPS	360 Mbps	780 Mbps	960 Mbps
Latency	350ns	380ns	380ns
Typical Power Draw	191 W	191 W	224 W

Table 4.1: Arista 7150 switches – comparison of specifications

4.6.2 Collaborations

According to [AristaSDN], at Interop, Arista planned to show the autoprovisioning mechanism on VMware virtualisation software by using the VMware vCloud SDN controller with the Arista switches.

Arista also planned to demonstrate integration with the Big Switch controller by using OpenFlow to connect to their switches, and to integrate the Arista switches with the Nebula controller by using the OpenStack Nova APIs.

Arista's goal for these collaborations is to show that all three controllers are able to program Arista switches through three different APIs and protocols.

4.6.3 Summary and Conclusions

Arista do not see OpenFlow as a key element in enabling network programmability. Arista are looking at (proprietary) alternatives for enabling network programmability by using their novel EOS operating system together with the Arista Field-Programmable Gate Array (FPGA) switches.

Even though Arista mention the 7150 switch as their SDN-specialised switch, this is not from an OpenFlow perspective, as OpenFlow 1.0 Agent is also enabled on other switches such as the 7050 series. Rather it refers to the novel architecture of the switching hardware and EOS operating system.

4.7 IBM

4.7.1 Introduction of OpenFlow / SDN

IBM seem to be returning to focusing on the networking industry with their development of the Programmable Network Controller (PNC) and their strong support for the SDN and OpenFlow standardisation activity.

IBM's view of SDN is more data-centre driven than a WAN approach. According to [IBMSDN], "OpenFlow holds great promise for smarter computing by enabling organisations to more easily modify, control and manage today's dynamic physical and virtual networks". IBM's vision of SDN is shown in Figure 4.11.

Integrated: Software-Defined Networking Enables User Control

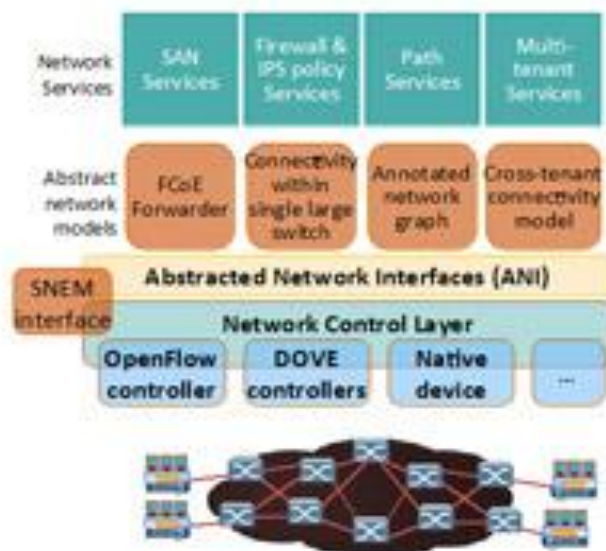


Figure 4.11: IBM SDN approach

As a consequence of their data-centre focus, IBM are missing a suitable core switch for a complete wide-area SDN solution, and are reliant on partners such as Brocade and Juniper Networks. For the specific application of OpenFlow in the (concentrated) data centre, however, IBM have developed two strong main components: the IBM Programmable Network Controller and the IBM Rack Switch G8264. Each of these is described below.

4.7.2 Programmable Network Controller (PNC)

According to [IBMOFController], IBM became the second large IT vendor to introduce its own OpenFlow controller in October 2012 by announcing the availability of its Programmable Network Controller (PNC). The controller provides control of network flows and unlimited virtual machine (VM) mobility – implemented in enterprise-class software, as described in [IBMSDN].

IBM's OpenFlow controller software is proprietary and runs on an x86 server. IBM license the software in configurations to control 1, 10 or 50 switches and plan soon to offer a licence to enable the management of 100 switches. The controller features web-based, RESTful APIs for the development of network applications on top of it and it also enables, by default, topology discovery, fault detection and redundant configurability as presented in [IBMOFController].

IBM's OpenFlow controller isn't cheap, as it is competing with open source controllers like Floodlight and NOX. The IBM PNC controller licensed for one controlled switch costs 92k US dollars, while for controlling more than one switch, additional licences are needed for at least 1,7k US dollars per device.

The detailed IBM PNC specifications are shown in Table 2.

Specification	Details
OpenFlow standard	OpenFlow 1.0 compliant
Verified OpenFlow Switches	Aggregation Switch: IBM RackSwitch G8264, NEC PF5240 Edge Switch (Edge of OpenFlow domain): IBM RackSwitch G8264, NEC PF5240
OpenFlow Vendor Extensions	Bit masking In band Broadcast/Multicast for wire speed forwarding
Virtual Tenant Network (VTN)	vRouter (L3) vBridge (L2) vFilter
North bound API	Web API
Number of VTN	1000 (Extended VLAN mode: 10.000 VLANs)
Number of Flows	300.000
Redundancy Features	IBM PNC Active/Standby
Switch and Link Discovery	Topology discovery
IP	IPv4 IPv6 (L2 forwarding)
ARP	ProxyARP
Routing options	Shortest Hop ECMP (L2/L3) Avoid switch routing
QoS	Type of Service / Class of Service (ToS/CoS) marking Policing
Link Aggregation Group (LAG) features	Multi-Chassis Link Aggregation Group (MCLAG)
Policy Setting	ACL (Pass, Drop) Redirect Path-Policy NET-WATCH
Monitoring	Failure Status Link traffic monitoring

Specification	Details
Visualisation	Topology Station (Host, VM) Path information Flow statistics sFlow
OS	Red Hat Enterprise Linux RHEL 6.1 (x86_64)

Table 4.2: IBM PNC specifications [IBMPNC]

4.7.3 RackSwitch G8264

IBM's first switch involved in OpenFlow deployment is the 1.28 Tbps RackSwitch G8264 top-of-rack switch. The switch offers 48 SFP/SFP+ 10 GbE ports and 4 QSFP 40 GbE ports that can be split out to an additional 16 x 10 GbE ports. The switch supports OpenFlow 1.0.0 and can handle a maximum of 97,000 flow entries.

According to [IBMG8264], IBM offers two versions of the RackSwitch G8264 and G8264T; only G8264 is OpenFlow enabled.

The G8264 switch offers:

- Optimisation for applications requiring high bandwidth and low latency.
- Support for Virtual Fabric and OpenFlow.
- 48 x 10 GbE SFP+ ports and 4 x 40 GbE QSFP+ ports in a 1U form factor.
- Future-proofed with 4 x 30 Gb QSFP+ ports.
- Performance details are summarised in Table 4.3.

	IBM RackSwitch G8264
Performance	100% line rate performance 880 nanoseconds latency 1.28 Tbps non-blocking switching throughput (full duplex) 960 Mpps

Table 4.3: IBM RackSwitch G8264 performance [IBMG8264]

4.7.4 Collaborations

4.7.4.1 IBM and NEC

IBM have teamed up with NEC to provide a novel OpenFlow-enabled network solution, dedicated to data centre and cloud environments, which has been successfully deployed by various customers according to [IBMNEC]. For example:

- Stanford University uses this solution in a parallel network to test OpenFlow's applicability to the university's production environment.
- Tervela, provider of market-leading, distributed data fabric, uses this solution for delivering a breakthrough in dynamic networking to ensure predictable performance of Big Data.
- Selerity, provider of ultra-low latency event data, uses the solution to accelerate real-time decision-making for global financial markets.

The solution that IBM and NEC propose, as shown in Figure 4.12, consists of the NEC's ProgrammableFlow Controller: NEC PF5240 1/10 Gigabit Ethernet Switch and NEC PF5820 10/40 Gigabit Ethernet Switch combined with IBM's OpenFlow-enabled RackSwitch G8264 10/40 GbE top-of-rack switch.

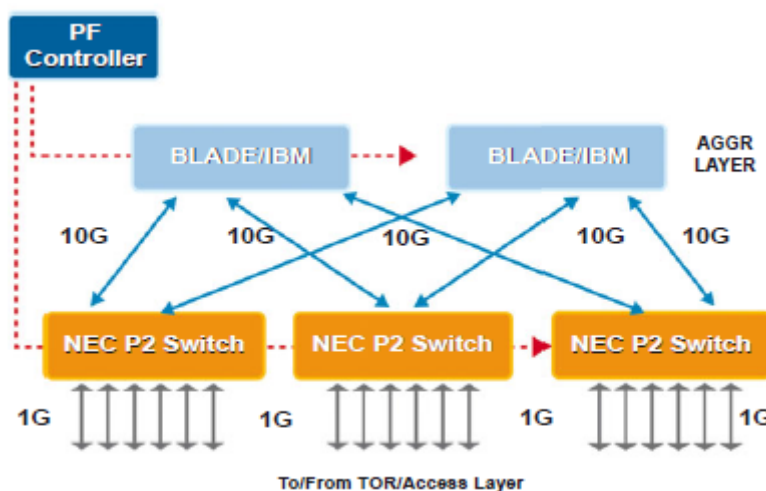


Figure 4.12: IBM and NEC switching architecture for OpenFlow [IBMOFNG]

4.8 HP

4.8.1 Introduction of OpenFlow / SDN

HP's involvement in OpenFlow started in 2007 with their collaboration with Stanford University for Ethane (the predecessor of OpenFlow), according to [HPOF]. Their OpenFlow-related activity until today is shown in Figure 4.13:

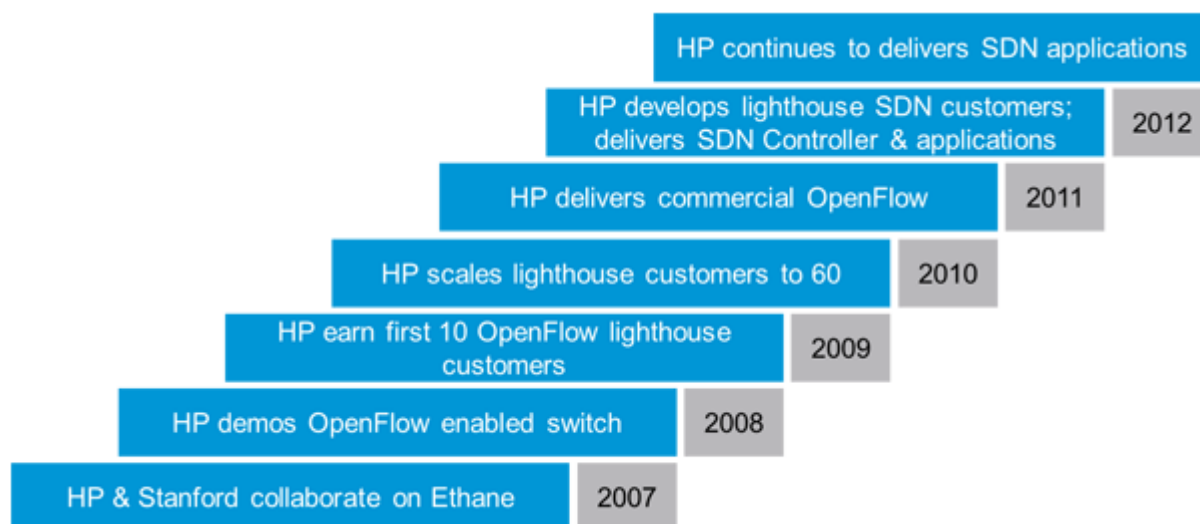


Figure 4.13: HP OpenFlow and SDN activity milestones [HPOF]

According to [HPSDN], at the beginning of October 2012 HP announced their software-defined network (SDN) technology development including an SDN controller and SDN applications, services and solutions for virtualisation strategies.

HP are dividing their SDN solution based on three layers, as shown in Figure 4.14:

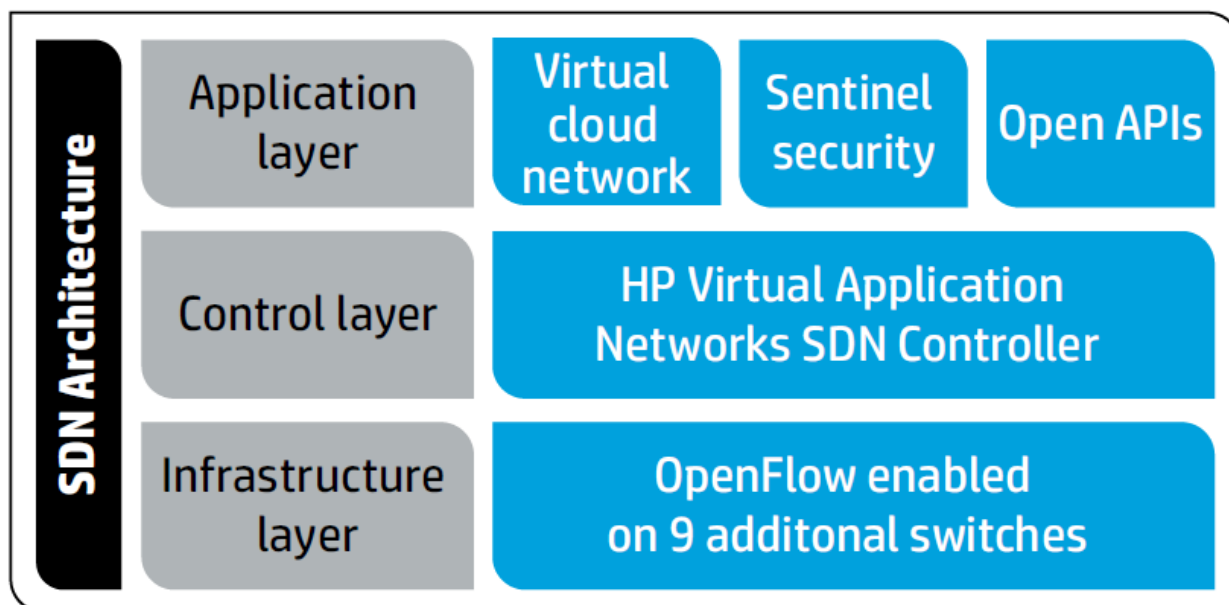


Figure 4.14: The HP SDN Architecture [HPVAN]

- Infrastructure layer – by providing open programmable access to their devices through OpenFlow. HP are adding the new HP 3800 switch series to the existing 16 models in the product portfolio of OpenFlow-enabled equipment.
- Control layer – by providing the new HP Virtual Application Networks SDN Controller (proprietary HP controller) used to abstract the physical layer and to provide a centralised view with an automatic network configuration of all devices in the infrastructure. Also, APIs are provided to third-party developers to integrate custom enterprise applications.
- Application layer – by providing open programmable interfaces to automate applications across the network. In this case HP's main focus is on the virtualisation application HP Virtual Cloud Networks and on security suite application – HP Sentinel Security.

According to Bethany Mayer, HP's senior vice president and general manager of networking, the HP SDN solution "is end-to-end, not just a data centre solution, and it goes from the data centre to campus and branch, and even across the wide area network if you need that".

4.8.2 HP Controller – HP Virtual Application Networks SDN Controller

The HP Virtual Application Networks SDN Controller will be shipped both as pure software and as an integrated appliance based on HP Proliant server hardware, according to [HPSDNPortfolio]. The controller is based on HP intellectual property, but it also includes some elements of open source code.

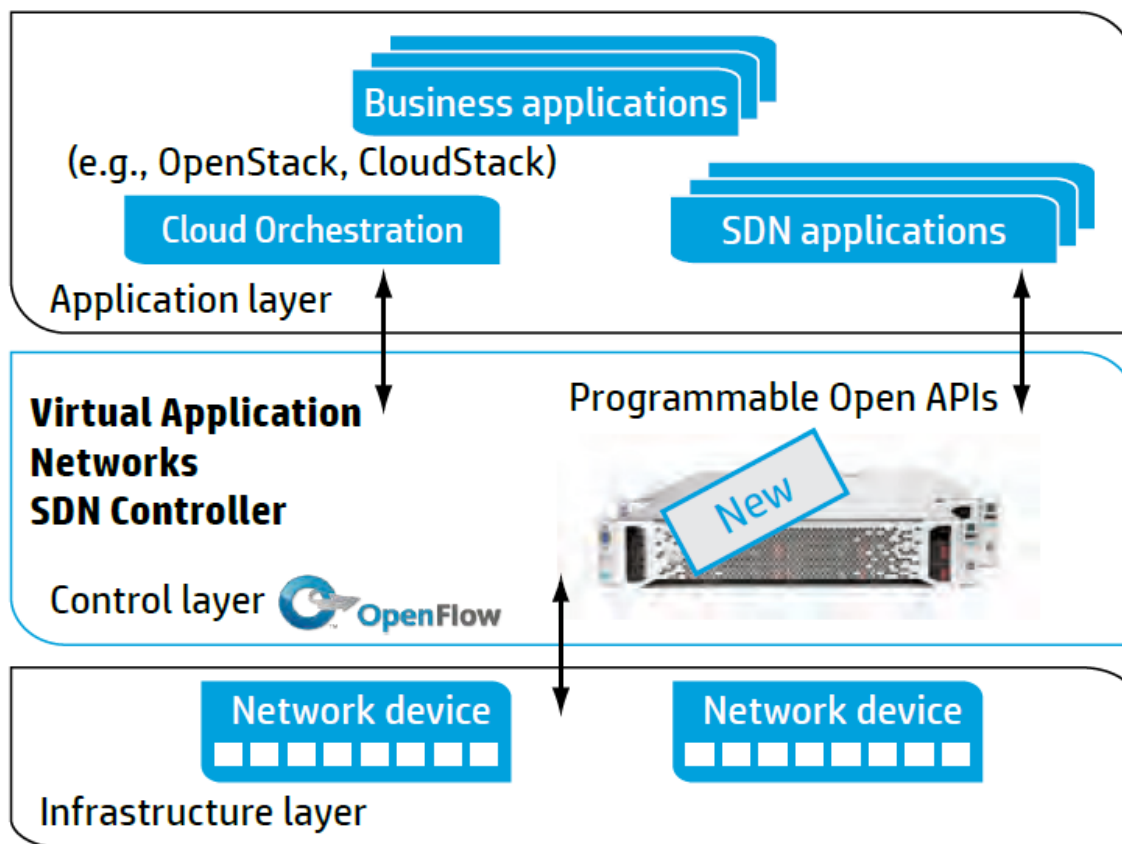


Figure 4.15: HP Virtual Application Networks SDN Controller [HPVAN]

As shown in Figure 4.15, the HP controller is the centrepiece of the SDN architecture, and according to [HPVAN] it provides:

- Availability as either software or appliance.
- Full support for the OpenFlow protocol.
- Open APIs to enable third-party SDN application development.
- HP SDN applications and OpenFlow-enabled infrastructure for an end-to-end solution.

4.8.3 HP Switches

According to [HPOF16switch], on February 2012 HP announced that OpenFlow support is generally available on 16 different switches within their HP 3500, 5400 and 8200 series.

Also, during the announcement in October, HP stated that they are extending the number of switches that support OpenFlow from 16 to 25 by adding the HP 3800 switch line. HP also announced the expansion of OpenFlow support to their entire switching portfolio by the end of 2012, according to [HPSDNPortfolio].

4.8.4 HP APIs and Applications

As detailed in [HPSDNPortfolio], the HP controller's northbound APIs, based on the Representational State Transfer (REST) architecture, represent the main point of differentiation for the HP SDN solution.

There are three network services applications that HP are proposing as a starting point for emphasising the benefits of the SDN architecture:

- The network virtualisation application – enabling cloud network automation, multi-tenancy and public-private cloud integration on top of the HP OpenFlow controller.
- The distributed load-balancing application – implemented by HP for the Swiss research institute CERN with the purpose of distributing network traffic across multiple devices, including firewalls and servers, increasing simplicity and reducing cost and bandwidth bottlenecks.
- The SDN security application – protecting the customer by looking at his DNS queries and making sure that the customer doesn't go to unwanted URLs. This solution, called HP Sentinel Security, can solve various issues without adding new hardware and it is used to protect HBO's entire network infrastructure.

4.8.5 Future Plans

According to [HPSDN], HP's short-to-medium-term plans for deploying the SDN solution are:

- HP Virtual Application Networks SDN Controller expected to be available worldwide by the second half of 2013.
- HP Virtual Cloud Networks Application expected to be available worldwide by the second half of 2013.
- HP Sentinel Security Application is currently available as an early access program to selected customers.
- OpenFlow support is currently available on HP Networking switches as a software upgrade.
- HP SDN services have been available worldwide since early 2013.

4.9 Dell

As presented in [DellSDNevolve], Dell think that SDN will turn enterprise switches into an inexpensive commodity not earlier than 3-5 years from now. Even though the idea of moving most of the higher-layer network functions into software is working, and in the long term would probably determine a decrease in the price of the switches while also making them interoperable between different vendors, operators are looking into running a hybrid architecture. This means that they want to keep the traditional architecture, where the switch logic is kept inside the switch and implemented in hardware, while slowly moving towards taking the logic into centralised controllers.

DELL acquired Force10 in 2011 in order to meet competition in networking from Cisco Systems, Hewlett-Packard and IBM, while being able to provide all the elements of a data centre.

Dell are now planning to extend SDN beyond its current state in order to provide more flexibility and management capability to the organisations that are deploying it. As OpenFlow to date is used as an SDN standardisation between a controller and the data plane or switch forwarding, Dell is looking into providing APIs that are able to reach all the way up to the hypervisor. The plan also includes the optimisation of Dell's software for the key hypervisors, which include VMware ESX, Microsoft Hyper-V and the open-source OpenStack framework while it remains hypervisor-agnostic.

Dell's longer term plan is to be able to connect, in a very simple way, any of the hypervisors into any controller and any switch. In this way, if Dell's customers are moving from ESX to Hyper-V or OpenStack, they can keep using Dell equipment.

According to [DellForce10], from the SDN perspective, Dell are not just providing an OpenFlow-enabled switch, they are including OpenFlow in the whole Dell Virtual Network Architecture (VNA) solution. In this way customers are allowed to implement SDN, for instance, by using an OpenFlow controller (e.g. from Big Switch) by interfacing it to the connection abstraction that VNA provides.

At Interop 2012, Dell demonstrated a common physical network fabric composed of Force10 S60 switches and a common network virtualisation abstraction, in collaboration with the Big Switch controller, showing how tenant networks can span multiple orchestration technologies (Dell AIM and Dell OpenStack) [DellForce10Interop].

4.10 NEC

4.10.1 Introduction of OpenFlow / SDN

NEC was one of the first vendors to enter the SDN market with an OpenFlow network solution, called ProgrammableFlow. Version 2 of the NEC ProgrammableFlow Controller, the PF6800, won the Grand Prize and the Management, Monitoring and Testing category at Best of Interop 2012. The ProgrammableFlow software separates out the physical network from the control layer, providing customers with new levels of flexibility and control to support the deployment of new network services and gain the benefits of SDN.

Even though NEC has significant switching hardware in their portfolio, their plan is to interoperate their controller with as many switching vendors as possible [NECSDN].

The SDN technology is delivered by NEC under the umbrella of the ProgrammableFlow family of products, described below.

4.10.2 NEC ProgrammableFlow Family of Products

4.10.2.1 *ProgrammableFlow Controller*

The NEC ProgrammableFlow Controller enables virtualisation in enterprise networking and is defined to allow data centres to provide network management based on high levels of automation while reducing operating costs and the time required to deliver network services, according to [NECOF]. In this centralised way (i.e. from the controller configuration interface), network tenants can have their own set of policy, management rules and security, with no interference from other tenants.

There are a number of features specific to the ProgrammableFlow Controller, according to the specifications:

- OpenFlow network control:
 - Topology discovery.
 - Policy-based route control.
 - Shortest hop, flow count and link weight.
- Virtual networking:
 - Virtual bridging/routing.
 - Physical to virtual mapping.
 - MAC/VLAN, Port/VLAN and VLAN.
- Custom filtering/waypoint routing:
 - Source/Destination MAC.
 - VLAN tag.
 - Source/destination IP address, DSCP, protocol.
 - Source/destination TCP port.

4.10.2.2 *ProgrammableFlow PF5240 Switch*

PF5240 is an Ethernet-based switch which can integrate into Ethernet environments while simultaneously functioning as an OpenFlow Switch. According to [NECOF], it was the first OpenFlow Switch to be generally available to customers worldwide and it offers:

- 176 Gbps non-blocking bi-directional throughput in a 1U form factor.
- 48x UTP ports that operate at 10 Gb / 100 Gb /1000 Gb speeds.
- 4x SFP/+ ports that operate at 1 GbE or 10 GbE.

It also supports line-rate multi-layer switching, and maintains up to 160,000 network flow entries. According to the specifications, the switch can work in hybrid OpenFlow mode, connecting OpenFlow networks to L2/L3 networks.

The PF5240 switch provides a virtual switch instance for running OpenFlow and distributed protocols on the same equipment, while the implemented OpenFlow Agent is version 1.0.

4.10.2.3 ProgrammableFlow PF5820 Switch

PF5820 is the latest member of the ProgrammableFlow family and it supports:

- 1.2 Tbps (1280 Gbps) non-blocking bi-directional throughput in a 1U form factor.
- 48x SFP+ ports that operate at 10 GbE.
- 4x QSFP+ ports that operate at 40 GbE or as 16 additional 10 GbE ports.

According to the PF5820 specs, the switch provides a full line rate performance while the OpenFlow implemented is Version 1.0; it can support more than 80,000 Layer 2 flows.

4.10.2.4 ProgrammableFlow Management Console

The ProgrammableFlow Management Console (Figure 4.16) provides:

- End-to-end visualisation of the physical network.
- Virtual tenant network view.
- End-to-end flow monitoring.
- Real-time fault monitoring.

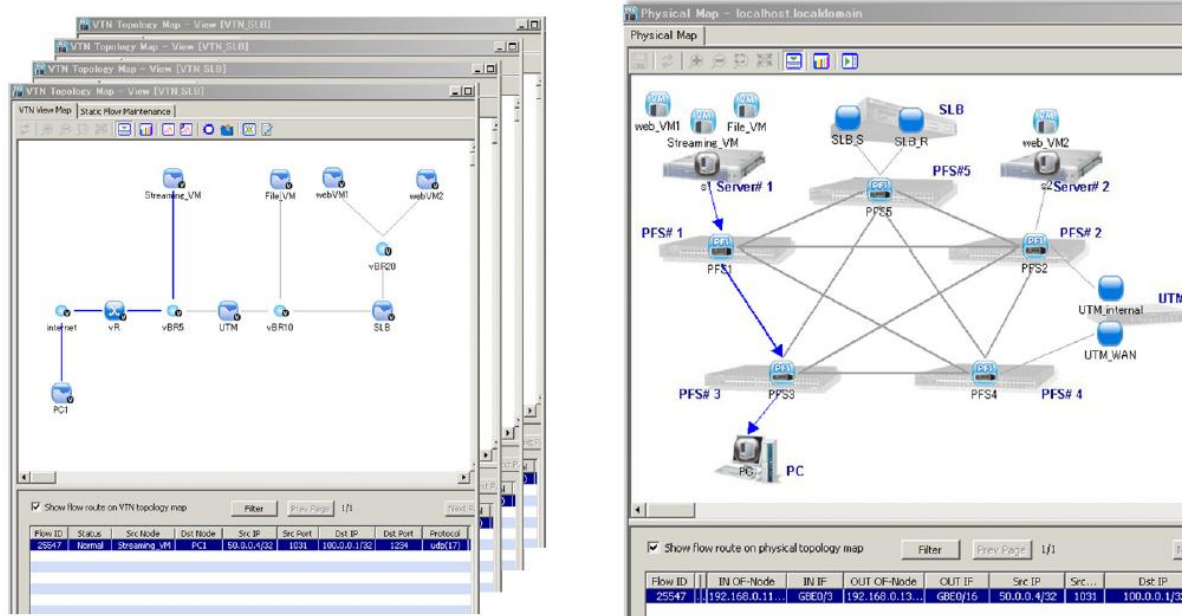


Figure 4.16: ProgrammableFlow Management Console – centralised monitoring

4.10.3 Collaborations

4.10.3.1 *NEC and Brocade*

As described in Section 4.4.5.1.

4.10.3.2 *NEC and IBM*

As described in Section 4.7.4.1.

4.10.3.3 *NEC and Microsoft*

According to [NECPF], in June 2012 NEC and Microsoft demonstrated the NEC ProgrammableFlow Network suite SDN operating in a windows-based virtualised environment – Windows Server 2012 Hyper-V, in Booth 18 at Microsoft TechEd North America 2012.

4.10.3.4 *NEC and Radware*

As announced in [NECRadware], Radware and NEC are cooperating to create a joint solution that integrates Radware's Attack Mitigation System into NEC's ProgrammableFlow OpenFlow-based switches and controller. The goal of this cooperation is to extend the advantages of the SDN architecture in the area of network security.

The joint solution will provide an application security-aware network that dynamically assigns security protection resources that can be customised per need according to varying levels of detected threats and traffic volumes.

The main benefits of the joint solution between NEC and Radware can be summarised, as in [NECRadware], as follows:

- Reduced network complexity.
- Service-centric attack mitigation protection.
- Better service-level guarantees.
- High availability.

4.10.4 Future Plans

According to [NECSDN], NEC continue to extend the interoperability tests with as many vendors as possible, while also contributing to the development of northbound applications for network management support. At the same time, integration of the ProgrammableFlow fabric with cloud computing platforms (e.g. OpenStack) is of high interest because it would allow engineers to provision virtual compute, storage and networks solutions on-demand.

NEC announced a complete automation of SDN (the controller) in the next version of the ProgrammableFlow Controller [NECSDN].

4.11 Pica8

Pica8 is a three-year-old start-up company with a focus on competing with big networking vendors such as Cisco by enabling networking based on merchant silicon available from Broadcom, Marvell or Fulcrum [Pica8startup]. This is motivated by the fact that the merchant silicon market is lately able to provide similar performance and a lower cost for their chips, compared to the custom ones (ASICs) provided by Cisco, Juniper or other vendors.

Pica8's plan is to provide cheaper boxes, also called "white boxes", and also to build the software to run on top of the boxes, opening this up as a whole solution to the industry (as shown in Figure 4.17). By contrast, one of Pica8's competitors, Arista, who are also using the advantage of merchant silicon, are placing their own software on top of their boxes, which makes Arista's solution "less open".

Scalability of Data Centers

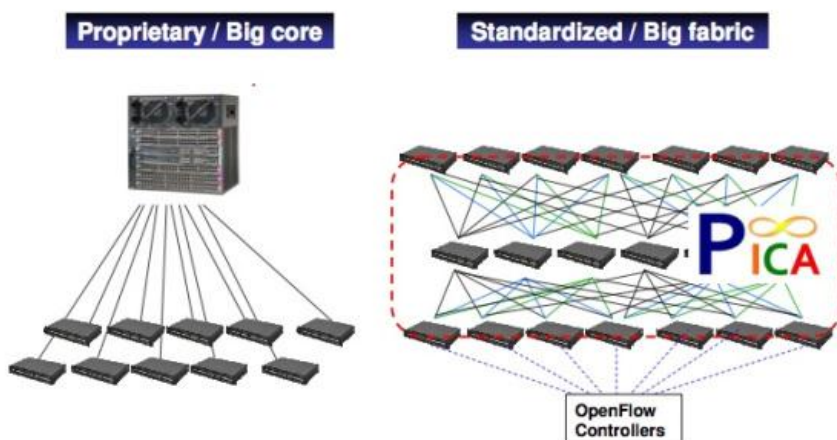


Figure 4.17: Pica8 solution for Data Centers [Pica8startup]

4.11.1 Pica8 Implementation of OpenFlow

According to [Pica8OF], Pica8 has released three different OpenFlow implementations:

- Indigo – an OpenFlow implementation run by Big Switch Networks, a strategic partner of Pica8. Indigo supports OpenFlow 1.0 and Big Switch also plan to implement future OpenFlow standards.

- Pica8 OpenFlow – a pre-installed component within XorPlus OS 1.0 and 1.1 images. This implementation was based on porting the Indigo User Mode to the Pica8 hardware drives but Pica8 decided to stop this implementation in favour of the direct Indigo implementation.
- Pica8 OVS (Open vSwitch) – implies porting the open source OVS into XorPlus 1.3. The latest OVS supports OpenFlow 1.0 spec and most of the OpenFlow 1.2 features; Pica8 has an agreement with Nicira to jointly develop OVS for future 1.2 and 1.3 OpenFlow standards and make this a clear direction for future development.

In the end, as described in [Pica8OF], Pica8 decided to use OVS as the default OpenFlow image for their switch platforms (also called Pronto).

The current release of PicOS (shown in Figure 4.18) is version 1.5.1, released on October 2012; it supports OVS version 1.2.2, which is OpenFlow 1.0 compatible.

PicOS Architecture

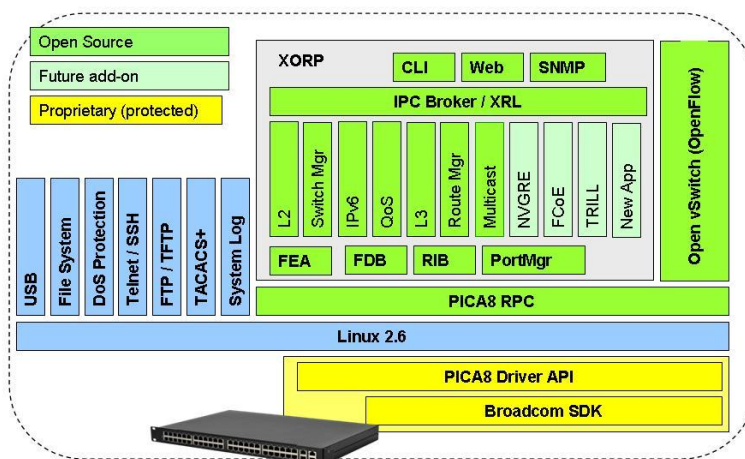


Figure 4.18: PicOS architecture [Pica8OSforOpen]

In order to open source the protocol stack in PicOS and also protect the proprietary code from the chip suppliers, Pica8 separated their code into two components: PicOS and Pica8 Driver [Pica8OSforOpen].

4.11.2 XorPlus OS

Pica8 XorPlus OS is a free-licence network OS requiring no subscription fee and allowing free use of any new features and free bug-fixing releases.

Pica8 XorPlus is a switching software supported by the open community focusing on performance, scalability and stability issues for data centres while providing [Pica8XorPlus]:

- L2 features.
- L3 features.
- OpenFlow (OVS 1.1 release, compliant to OpenFlow 1.0 standard).
- Management.

According to [Pica8XorPlus], the XorPlus OS is independent of underlying switch chips and it allows the protocol stack to run on a different platform from the driver and the switching hardware.

4.11.3 Pica8 Switches

Pica8 switches are provided under the brand Pronto.

4.11.3.1 *Pica8 3290*

The Pronto 3290 switch provides [Pica83290]:

- 48x 1 GbE ports.
- 4x 10 GbE uplinks (SFP+).
- 1U height switch for standard rack.
- Support layer 2/3 switching capabilities.
- Supports full suite of PicOS and consequently OpenFlow 1.0 capabilities.

4.11.3.2 *Pica8 3295*

Pica8 Pronto Switch 3295 provides similar specs as 3290 switch but with differences in number of management ports (1 instead of 2) and better failure times.

4.11.3.3 *Pica8 3780*

Pica8 Pronto Switch 3780 provides also similar specs as 3920:

- 48x 1 GbE SFP+ ports.
- 1x management GE service port.
- 1x RJ45 serial console port.
- 1U height switch for standard rack.
- Support layer 2/3 switching capabilities.

4.11.3.4 *Pica8 3920*

Pica8 announced in March 2012, according to [Pica83920], the general availability of their Pronto 3920 switch , which enables 40 GbE interfaces together with OpenFlow and provides:

- 48x 10 GbE SFP+ ports.
- 4x 40 GbE QSFP ports.
- 1U height switch for standard rack.
- Support layer2/3 and OpenFlow capabilities.

4.11.4 Future Plans

Pica8 are planning to work with other partners (e.g. Google, and Deutsche Telekom) to develop and qualify a full specification of OpenFlow 1.2, according to [Pica83920].

Another plan Pica8 are working on is for their new Pica8 Pronto 3980 switch.

4.12 ONF

OpenFlow is being promoted by the Open Networking Foundation (ONF), founded by six companies that own and operate some of the largest networks in the world: Deutsche Telekom, Facebook, Google, Microsoft, Verizon and Yahoo, as well as around 50 other member companies, major equipment vendors, software suppliers and IC technology providers.

The main benefits that ONF provides to the members are:

- Real-time participation in the evolution of concepts and programs.
- Collaboration with inventors, market leaders and start-ups driving SDN and OpenFlow.
- Interoperability testing.
- Royalty-free access to OpenFlow.
- Frequent speaking engagements that can improve a member organisation's market visibility.

5 Conclusions

SDN and OpenFlow are one of the current key topics for all network vendors. For many years SDN/OpenFlow was a promising technology and a marketing buzz word; since 2012 it represents an opportunity which is also, in part, a requirement, since no network vendor (of hardware or software) can afford to neglect a development with which the main part of the market is getting involved, and about which customers are continually asking – what SDN/OpenFlow is, what it does, why they should use it, when and how they can use it.

SDN and OpenFlow are creating a virtual division of the network vendors' market, placing on one side the “heavy” vendors with strong network strategies and long-term roadmaps, such as Cisco or Juniper. On another side are vendors such as IBM or HP who are finding the technology a good opportunity to re-enter, or enter more deeply, the networks market. On a third side are start-ups that can easily benefit from OpenFlow without significant investments, such as Big Switch or Pica8.

Looking to the classic network vendors such as Cisco and Juniper, the adoption of OpenFlow is slow. Cisco are deploying a proprietary platform, called Cisco ONE, and see SDN as a good opportunity and a natural transition. At the same time they see OpenFlow as one small piece in the SDN picture, but a piece that can be recreated in other ways or with other protocols. Cisco claim they can offer a higher degree of programmability for their customers at all layers compared to the “simple” view of OpenFlow, which focuses on the southbound interface of the programmability of the hardware from a centralised controller. However, Cisco are not ruling out OpenFlow, simply following a “wait-and-see” approach before committing to a long-term decision.

Juniper are taking a similar approach to Cisco, looking rather into the northbound APIs and providing useful information from the network to the application layers, and do not see OpenFlow as a dependency. However, unlike Cisco, Juniper have enabled OpenFlow agents (version 1.0) on some of their switch and router series, via the JunOS SDK. Juniper are also looking at SDN/OpenFlow controllers, but it is still to be decided whether they will brand a proprietary SDN controller or source another one.

In Q3 2012, both HP and IBM announced their SDN strategy and new OpenFlow-enabled equipment deployment. Both vendors are trying to cover the whole stack of SDN (OF-enabled hardware, controllers and applications) and use the software competence and the wide range of customers as an advantage for a proprietary deployment of SDN. HP are offering the HP Virtual Application Network SDN Controller (either in software or hardware implementation and OF enabled), 25 OF-enabled switches (probably the widest variety of OF-enabled switches in the market) and also a number of predefined SDN applications that are already in production. However, the HP switches are running OF as a firmware upgrade, making this unlikely to be able to control flows at full line-rate speed (as in the case of hardware/ASIC OF implementation).

IBM have a similar approach to HP, offering their IBM Programmable Network Controller (PNC) and the IBM Rack Switch G8264 OpenFlow enabled.

NEC are also providing SDN and OpenFlow architecture under the umbrella of their ProgrammableFlow family of products. NEC offer the ProgrammableFlow controller as well as two switch series (PF 5240 and 5820 – both running OF 1.0 agents) and a ProgrammableFlow Management Console, making SDN a core business for NEC and ProgrammableFlow a complete solution for SDN customers.

Also in Q3 2012, Ciena announced their strong interest in SDN and OpenFlow, and they are committed to providing OpenFlow support on their packet portfolio. Ciena's plans are to release OpenFlow support for the 5400 and 5100 switch series and enable their customers to test it, and also to develop software-based solutions and new control plane capabilities based on SDN/OpenFlow.

Dell are getting involved in the SDN market via their Dell Virtual Network Architecture (VNA) and they see OpenFlow as being integrated into their VNA solution and not just as OpenFlow-enabled switches. From a controller perspective, Dell have been interworking with the Big Switch controller (for example, at Interop 2012).

Brocade are showing a high interest in SDN/OpenFlow deployment and are focusing on WAN and service provider networks (instead of on data centres, as most other vendors are). Brocade enabled OpenFlow 1.0 on their MLX routers on 100 GbE interfaces running at full line rate. Brocade also plan to deliver OpenFlow 1.0 on the new equipment shipped with the NetIron OS 5.4, as they allow a per-port usage of OpenFlow as well as hybrid mode. Brocade also acquired Vyatta (a software-based network operating system) in Q3 of 2012, which shows their growing interest in deploying SDN-enabled solutions.

SDN and OpenFlow were not (at least in the first phase) a benefit for the “big” vendors only, but also allowed start-ups like Nicira (subsequently acquired by VMware) or Big Switch to find a niche market where they could grow and differentiate themselves from the other vendors. Nicira and Big Switch were both focused on providing SDN solutions and new network services, and while Nicira have been acquired by VMware, Big Switch still represent one possible acquisition for the other vendors considering Big Switch's open source controller and the network application development on top of their controller – which is currently a long-term business opportunity for Big Switch.

In recent years the merchant silicon market (e.g. Broadcom, Marvell, Fulcrum, etc.) has been able to provide similar performance at significantly lower costs compared to the main proprietary chip/ASICs vendors such as Cisco or Juniper. This was a good opportunity for companies such as Arista Networks or Pica8 to be able to compete with the main vendors and to try to offer similar networking solutions at lower cost. Arista are also focusing on SDN architecture based on their proprietary Extensible Modular Operating System (EOS), which is controller-friendly software, with the APIs talking to the controllers, and uses OpenFlow. However, Arista do not see OpenFlow as a key element for SDN and are instead looking into proprietary deployment of network programmability with their EOS operating system, even though Arista provide their 7150 Switch with the OpenFlow 1.0 Agent.

An interesting presence in the SDN market is Pica8, a three-year-old start-up company, also based on the merchant silicon advantage but, in contrast to Arista Networks, trying to build open source protocol stack/OS for

their architecture instead of proprietary OS. Pica8 are also offering three switches that are OpenFlow 1.0-enabled by implementing OVS in their PicOS.

Observing the network vendors' deep market involvement in SDN and OpenFlow, and also seeing the Open Networking Foundation (ONF) organisation, which is pushing ahead the development, deployment and interoperability of OpenFlow-enabled SDN components, it is clear that a transition is taking place, if not technological then at least cultural, in the way vendors and operators are providing their services and the way they are making money out of it.

The focus of the SDN/OpenFlow study was the application layer; OpenFlow's optical/transport support was not included. While a number of suppliers are producing programmable optical equipment, their work was very much at the research stage at the time the study was undertaken (Y4 Q3). For Infinera, in particular (now a key supplier to the GÉANT backbone, although this had not been announced at the time of the study), implementation of SDN was at an early stage. Similarly, while new activities in the area of optical extensions for OpenFlow were underway at the time of the study, the rate and number of developments were too great to draw meaningful, useful conclusions. Such work and developments will be considered in GN3plus, in which NSI is also a major point of interest.

To conclude, the study has shown that a wide variety of OpenFlow equipment is available and OpenFlow software will converge to standard facilities. For the application of OpenFlow on large-scale networks or Wide Area Networks, it is recommended to limit the different equipment to a few brands in order to provide a stable environment. Meanwhile the OpenFlow software applied in a data centre is stable, reliable and can be implemented successfully.

Part 2

6 Introduction

This part of the document presents the design and deployment of the GÉANT Software-Defined Networking (SDN) / OpenFlow Facility.

6.1 In Part 2

Part 2 is organised as follows:

- Chapter 7 *SDN/OpenFlow in Research Networks* – describes how OpenFlow is being used by Research and Education Networks, including GÉANT, to deliver a testbed as a service.
- Chapter 8 *GÉANT OpenFlow Facility* – describes the design principles of the wide area GÉANT SDN/OpenFlow Facility, covering:
 - Objectives of the facility.
 - Facility design.
 - The GÉANT OpenFlow Control Framework (GOCF).
 - Facility deployment issues and considerations.
- Chapter 9 *Operational Environment of the Facility* – describes the OpenFlow Facility operational environment, focusing on the three-tier architecture for operations and including a list of operational tasks at all tiers, with the corresponding assignments to the different operations entities.
- Chapter 10 *Use Cases and International Collaboration* – describes the OpenFlow Facility use-case categories and examples, and international collaborations.
- Chapter 11 *Conclusions* – provides an overall assessment of the design and deployment of the GÉANT OpenFlow Facility.

7 SDN/OpenFlow in Research Networks

Software-Defined Networking (SDN), as an architectural concept, and OpenFlow, as a technology to realise it, have been conceived by the academic community to provide “a way for researchers to run experimental protocols in the networks they use every day” [OF-EnablingInnovation2]. Although originating from campus environments, SDN/OpenFlow has obtained significant recognition among Research and Education Network (RENs) in recent years.

There are several cases of adoption of OpenFlow among RENs. The Global Environment for Network Innovations project (GENI) has adopted OpenFlow [OF-GENI], encompassing backbone OpenFlow-enabled resources from the US national backbones of NLR and Internet2, to deliver experimental resources to its user community. ESnet has integrated OpenFlow capabilities with their On-Demand Secure Circuits and Advance Reservation System (OSCARS) [OF-OSCARS], to deliver an end-to-end network virtualisation solution across both the local area network (LAN) and wide area network (WAN). Internet2 has included OpenFlow capabilities in its Advanced Layer 2 Service [OF-I2]. JGN-X, the nation-wide network testbed of Japan, has been enhanced with OpenFlow capabilities to deliver RISE, a large-scale OpenFlow testbed for researchers, students, engineers, and operators [OF-RISE]. Overall, RENs are gradually researching, piloting and adopting SDN/OpenFlow to serve the specialised needs of their users.

Together with dynamic circuit or bandwidth-on-demand (BoD) services, OpenFlow capabilities are expected to be an integral part of REN infrastructures in the near future. Recent research [FederatingSDN, DCN-OF] is also demonstrating possibilities for BoD and OpenFlow to complement each other for advanced service offerings to REN users. Dynamic circuit capabilities interconnecting or extending SDN clouds are currently under investigation.

The introduction of OpenFlow to RENs presents many different types of use cases. These include [APAN-FIT-OF-Wkshop]:

- Security applications (Network Access Control (NAC), Intrusion Detection System / Intrusion Prevention System (IDS/IPS), Remote Packet Capture and Injection).
- Virtual Machine (VM) mobility.
- Load balancing.
- n-casting.
- Policy-based firewalling.
- Flow-based network provisioning.

- Extensions to dynamic circuit provisioning.
- Optical layer extensions.

On top of these, OpenFlow is a technology with recognised Testbed as a Service (TaaS) capabilities [CanProdNetTest]. Utilising its isolation and programmability characteristics, resource and management/control sharing of network elements (switches) and links is possible. Multiple user support, data plane slicing through Layer 2 switching, control plane slicing and isolation, programmability and real-time configuration of OpenFlow-enabled slices² offer the infrastructure and functional elements for the delivery of TaaS capabilities to end users. The GENI developments mentioned above validate this finding further.

7.1 The Case of GÉANT

GN3 operates GÉANT, the next-generation pan-European network, and related services that meet the communication needs of its R&E community across Europe. Service offerings are currently under investigation to provide user-accessible network environments in the form of testbeds on top of the production environment.

To prevent the production traffic of GÉANT's connectivity services [GÉANTServices] from being disrupted by high-bandwidth applications and experiments, it makes sense to separate them. This also enables researchers to modify the behaviour of infrastructure elements, such as traffic routing, which could not be realised on the production infrastructure.

SDN is considered a valid technology for delivering logical networks on top of the GÉANT production environment, so that experiments can be conducted on OpenFlow-enabled network slices without affecting production GÉANT services, albeit utilising GÉANT backbone resources. Chapter 8 presents the GÉANT OpenFlow Facility, as designed and implemented to provide a proof of concept along these lines.

The goal is to provide TaaS utilising OpenFlow, by giving researchers full control of the resources in their slice through a set of management utilities for the slice. This translates to delivery of seamless slice request submission, instantiation, management and decommissioning functionalities to end users, so that their core experimental work is not disrupted in any stage and the GÉANT production environment remains unaffected.

It has to be made clear that OpenFlow-based TaaS in GÉANT is not expected to accommodate all needs for partitioned physical network and IT resources that provide secure and isolated application-specific infrastructures. For example, requirements for all-optical slices for experimentation (slicing at the physical layer) cannot be fulfilled by the facility presented in Chapter 8.

Beyond TaaS, SDN/OpenFlow capabilities can further enable GÉANT to extend its operational capabilities and service offerings in the areas mentioned above as recognised opportunities for RENs. However, research and implementation in these areas is beyond the scope of this document.

² For a definition of terms, please refer to Appendix A.

8 GÉANT OpenFlow Facility

8.1 Objectives

As mentioned in Chapter 7, the GÉANT OpenFlow Facility has been designed and piloted to demonstrate the potential of OpenFlow in delivering TaaS capabilities over GÉANT. The facility is deployed on top of the GÉANT backbone production environment and is so far a single-domain environment. Multi-domain extensions are possible in subsequent lifecycles.

As a proof of concept, the GÉANT OpenFlow Facility supports:

- Announcement of the facility's network resources (Layer 2 links and OpenFlow-enabled switches) that can be used for the OpenFlow-controlled slice topologies interconnecting computing resources.
- Announcement of the available computing resources per Point of Presence (PoP) of the facility that are attachable to the slice topologies.
- A reservation mechanism for the network and computing resources.
- Management and control plane functionalities for the reserved resources.

8.2 Facility Design

Computing resources are offered as Virtual Machines (VMs) upon dedicated physical servers using Xen [Xen]³ hypervisor-based virtualisation. Network resources are offered utilising software-based OpenFlow switches based on Open vSwitch (OvS) [Open vSwitch] and network links that interconnect the OpenFlow software switches.

Two (2) general-purpose servers are installed on each of the five (5) GÉANT PoPs hosting the GÉANT OpenFlow Facility PoPs (for the GÉANT OpenFlow Facility General Server Specification, see Appendix B). Each server at a PoP (see Figure 8.1) is either:

³ Xen hypervisor is a well-known hypervisor that provides computing resources' virtualisation. It was chosen due to its compatibility with the OFELIA Control Framework (OCF), realising the management plane of the GÉANT OpenFlow Facility.

- The host of a software-based OpenFlow switch (Open vSwitch) on top of a native Linux Debian distribution, or
- The host of a Xen hypervisor for the instantiation of multiple VMs allocable to user slices.

The data plane topology of the GÉANT OpenFlow Facility is a full mesh graph, so that every OpenFlow switch has direct connectivity with all four (4) of the other OpenFlow switches. This is achieved by:

- A back-to-back connection of each OpenFlow switch to the local GÉANT production MX router, through four 1 Gigabit Ethernet ports.
- Interconnecting the GÉANT OpenFlow Facility PoPs over the GÉANT backbone using pseudowires (L2MPLS VPNs) configured on the MX routers. This setup eliminates the need for VLAN switching on GÉANT equipment and its associated limitations to the OpenFlow Facility's slicing capabilities.

The network slicing technique used by the GÉANT OpenFlow Facility dynamically allocates one or a range of VLAN IDs to each user slice. In this way, the network is partitioned per Open vSwitch interface (physical or logical) and user traffic is distinguished by the VLAN ID. VLANs can be involved in experimentation within the slice when a set of VLAN IDs is allocated to a slice. Thus the control logic of a slice is defined by the slice controller permitting the involvement of the reserved VLAN IDs to control logic decisions.

A simple example is that of an experimenter who can use his own range of VLAN IDs for routing purposes on top of the OpenFlow topology, forcing the OpenFlow switch to forward Ethernet frames based on non-standard algorithms, rather than having them handled by a legacy broadcast domain.

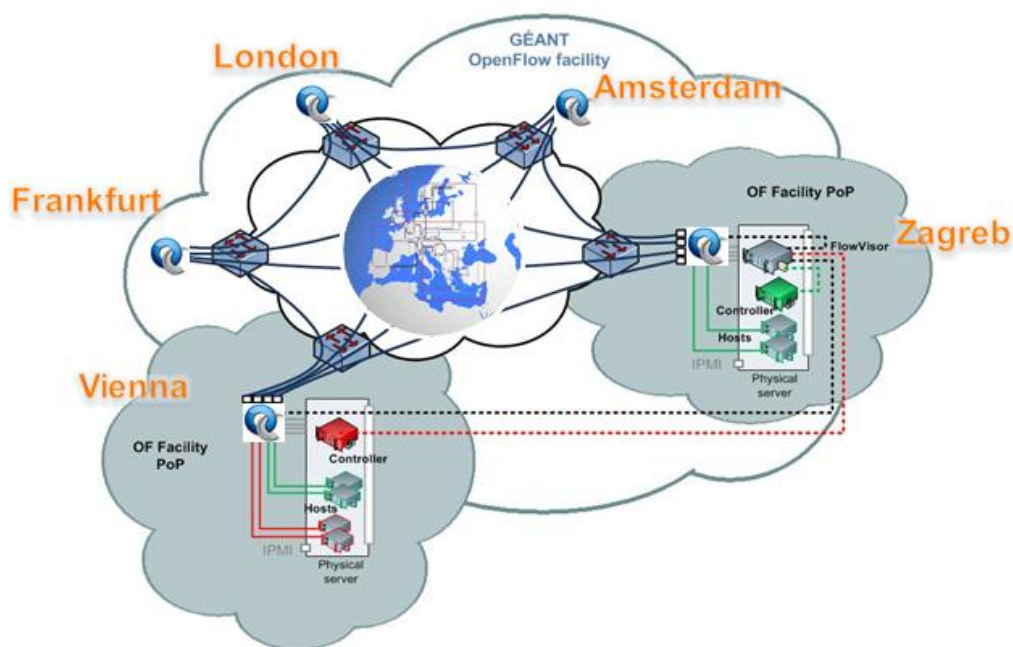


Figure 8.1: Overview of the GÉANT OpenFlow Facility

The GÉANT OpenFlow Facility PoPs are collocated with five of the GÉANT PoPs in Vienna, Zagreb, London, Amsterdam and Frankfurt. The GÉANT OpenFlow Facility overview is provided in Figure 8.1.

Additionally, in the Frankfurt PoP, the facility's management and control plane elements/software are hosted, namely the GÉANT OpenFlow Control Framework (GOCF) and the FlowVisor software. The GÉANT OpenFlow Facility's control and management plane are depicted in Figure 8.2.

End users may connect to the GÉANT OpenFlow Facility in order to transmit/receive traffic with three options⁴:

- (a) Using VMs, hosted on servers offered by the facility.
- (b) Installing their own equipment within the OpenFlow Facility PoPs.
- (c) Connecting their remote labs to the OpenFlow Facility through an (emulated or native) Ethernet circuit provided by their local NREN or an ISP.

However, in option (a) the traffic that the VMs exchange with the facility has to comply with Ethernet standards. This means that the VMs can be used as traffic producers/consumers but they cannot be used to extend the OpenFlow Facility. For such functionality, the users have to resort to options (b) or (c).

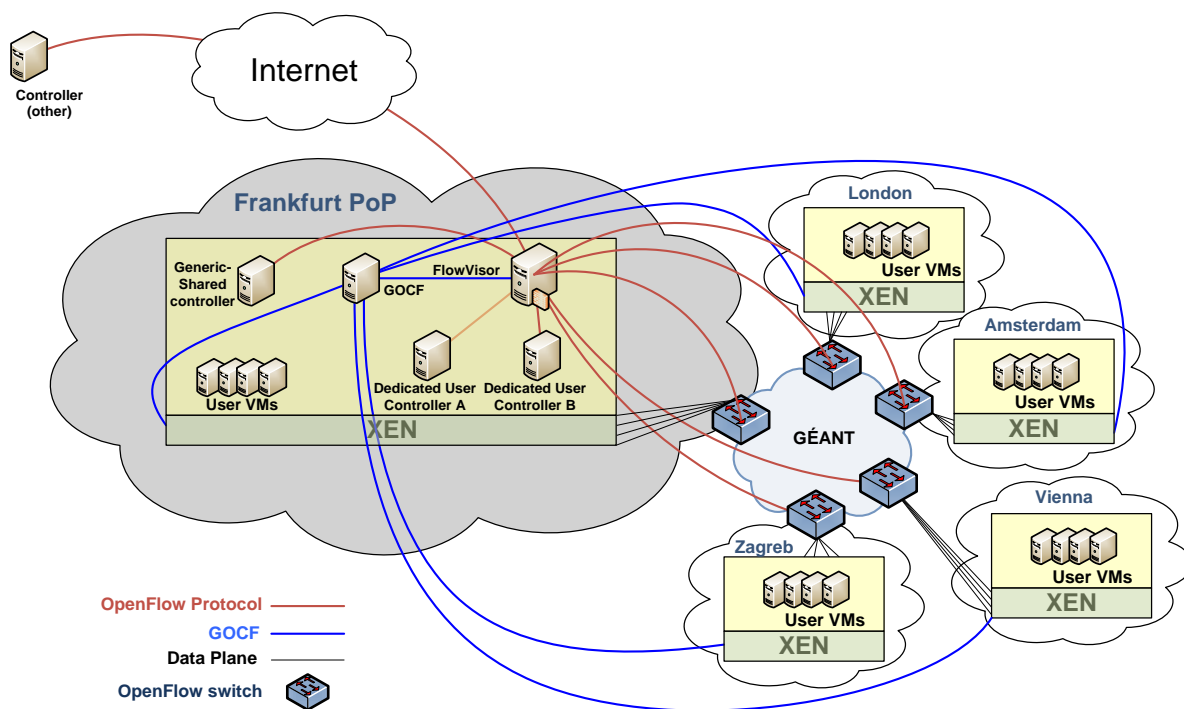


Figure 8.2: A generic view of the GÉANT OpenFlow Facility showing control and management plane

⁴ Not all of the options are available in the early phases of the facility's lifetime.

8.3 GÉANT OpenFlow Control Framework

The orchestration software deployed on top of the GÉANT OpenFlow Facility to implement management plane functionality is the OFELIA Control Framework (OCF) [OCFRepos]. OCF is developed within the OFELIA FP7 project [OFELIA], and it has been adapted to fulfil the GÉANT OpenFlow Facility requirements.

The main functionalities of GOCF are:

- Resource allocation and instantiation. The GOCF supports resource allocation, instantiation and de-allocation for any type of resource (e.g. an OpenFlow slice or a virtual machine out of a server farm).
- Experiment-based resource allocation. The resource allocation/de-allocation is performed per slice, with a slice being the smallest indivisible entity composed by the resources necessary to carry out an experiment. Slices are isolated from each other, even though they might share the same infrastructure substrate.
- Authentication and Authorisation (AA) and policy framework. GOCF supports the necessary mechanisms for authentication and authorisation (in several “scopes” or layers).
- Usability. End users/experimenters have access to comprehensive and easy-to-use user interface(s). In this sense, the main focus of the development is towards a web-based user interface.

An OpenFlow-aware network control plane is built on top of the facility's OpenFlow proxy controllers so as to slice the network flowspace of the entire data plane topology. FlowVisor [FlowVisor] is the selected proxy controller supporting the facility's network slicing and decoupling of forwarding (switches) and control (controller) elements. FlowVisor implements the network flowspace sharing and allocation per user controller logic. A single instance of a FlowVisor is deployed at the facility, controlling the OpenFlow switches and acting as a proxy between them and the experimenters' OpenFlow controllers. A second FlowVisor may be added in the future for redundancy.

Each experimenter's OpenFlow controller determines the control logic of a certain flowspace or slice. A controller can be either dedicated to a user/slice or shared (see Figure 8.2).

The controllers, FlowVisor and OpenFlow switches communicate utilising the public Internet over the management ports on the servers hosting them. Thus the control plane of the facility is implemented “out-of-band” at the IP layer, over the GÉANT network.

Authentication and Authorisation is handled by the Expedient tool [Expedient]. Expedient is a pluggable web user interface (UI) and clearinghouse originally developed by Stanford University, and adapted and extended to be a part of the OCF.

The Expedient clearinghouse functionality addresses user management, authentication and authorisation together with the projects'/experiments' administration. It can be integrated with an LDAP server in order to support authentication against a users' directory.

For the purposes of the GÉANT OpenFlow Facility, the GOCF Expedient instance has been integrated with Active Directory (provided by DANTE Operations), so that users in the GÉANT community can be approved

and obtain access to the facility using their credentials. Authorisation requires the definition of the appropriate user groups in Active Directory. The GÉANT OpenFlow Facility distinguishes between an admin/management user and end users/researchers.

- Admin/Management user. This user is able to create projects and accept slice requests, and handles the general operation and management of the infrastructure.
- End User/Researcher. End users/researchers can submit through the UI projects, create slices and perform experiments on top of the OpenFlow Facility. They have to belong to a project, which can be shared by multiple researchers.

8.4 Facility Deployment Issues and Considerations

The current data plane setup of the facility and the underlying substrate results in a quite complicated environment, with 30 data plane links, 5 software switches controlled by the OpenFlow protocol and the specialised handling of VLAN tagged packets. Recurrent testing of the GÉANT OpenFlow Facility data plane is required in order to ensure a reliable experimental environment. A set of diagnostic tests has been defined for that purpose.

Geographically dispersed OpenFlow-enabled devices introduce flow establishment delays. Proactive or reactive flow recording and correlation with flow number constraints are under constant surveillance. Long distances between OpenFlow-enabled devices and the OpenFlow controller in wide area environments introduce delay to the flow establishment process. Flows can be established using either proactively distributed or reactively distributed flow entries. At this point, the facility itself consists of an experiment for wide area OpenFlow controlled network environments.

9 Operational Environment of the Facility

The facility is supported by several GÉANT services and therefore can be considered as one of the multiple customers of GÉANT. A set of access points to the GÉANT services can be defined as the demarcation points between the facility and underlying GÉANT.

One layer higher, facility slices, administered by several end users/experimenters, can operate concurrently. An administration interface between the facility's operator and the experimenters is defined. End users/experimenters are able to manipulate their own functional subset of the facility's resources taking advantage of the provided OpenFlow facility services. (For more details on the capabilities delivered to end users, please refer to "GÉANT OpenFlow Facility Design" [GÉANTOFFDesign].)

A three-tier architecture for operations (see Figure 9.1) is the result of using the OpenFlow facility on top of GÉANT to deliver TaaS.

- The lowest level (Tier-1) is the view of the GÉANT backbone production environment operators. On this level, operators are responsible for the physical infrastructure (bare-metal servers of the facility, physical connections towards the GÉANT backbone routers, management LANs at the PoPs) as well as GÉANT connectivity and other services supporting the facility (L2MPLS VPNs across the facility PoPs, IP connectivity across the facility's control plane elements (controllers), perimeter firewalling).
- On the second level (Tier-2), the OpenFlow facility operators are responsible for the operations of the services provided to the end users/experimenters (such as monitoring, internal firewalling), including the required resources such as server hypervisors, virtual machines, the OpenFlow proxy controller, control and management capabilities of the OpenFlow slices.
- On the third level (Tier-3), the experimenter himself operates a subset of OpenFlow resources, such as the capabilities of his allocated virtual machines and a set of network flows through his own OpenFlow controller, to fulfil the concepts of user-controlled slices and TaaS.

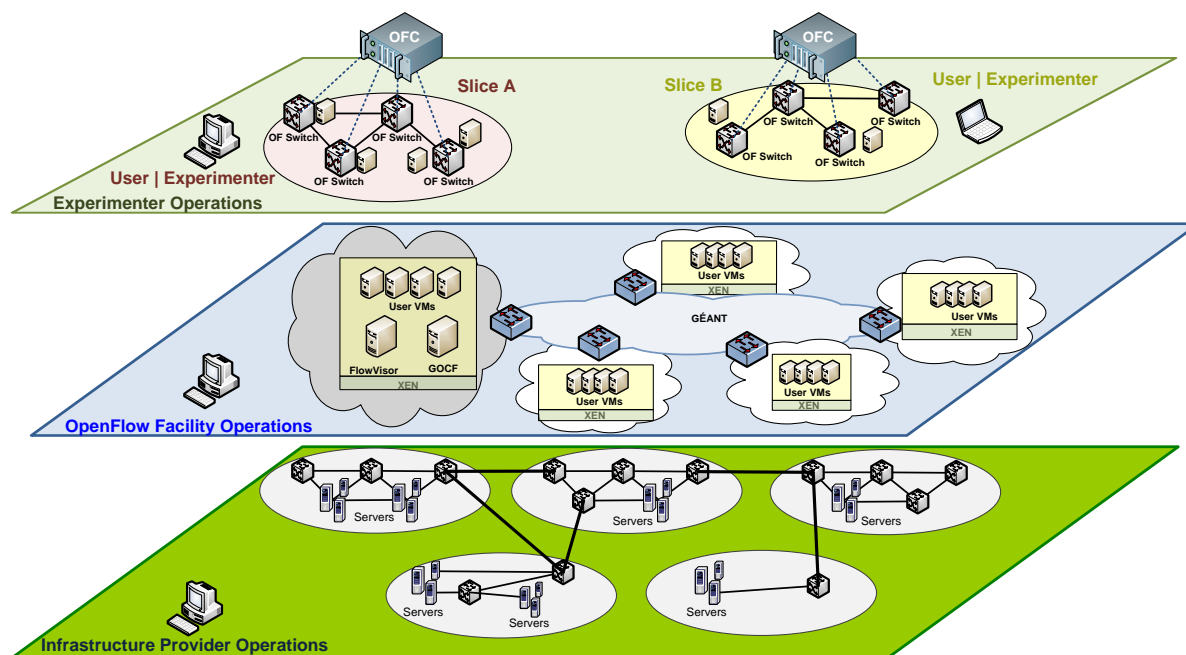


Figure 9.1: A three-tier architecture for operations

Based on this architecture, a list of services at different layers, operational tasks and corresponding responsibilities is provided in Table 9.1 below.

In order to illustrate further the concepts behind the three-tier architecture, Figure 9.2 presents a closer look at part of the services delivered from Tier-1 to Tier-2 and their demarcation points. Built-in interfaces, labelled with “BltIn/1-eth0”, are used by the control and management plane of the OpenFlow Facility. The service delivered to these interfaces by Tier-1 (the substrate) is IPv4/IPv6 Internet connectivity over the GÉANT backbone. The interfaces acting as OpenFlow data plane ports (highlighted with circles in Figure 9.2) participate in the facility’s data plane. For these, point-to-point L2 Ethernet connectivity is required, delivered by the L2MPLS pseudowires over GÉANT shown in Figure 9.2.

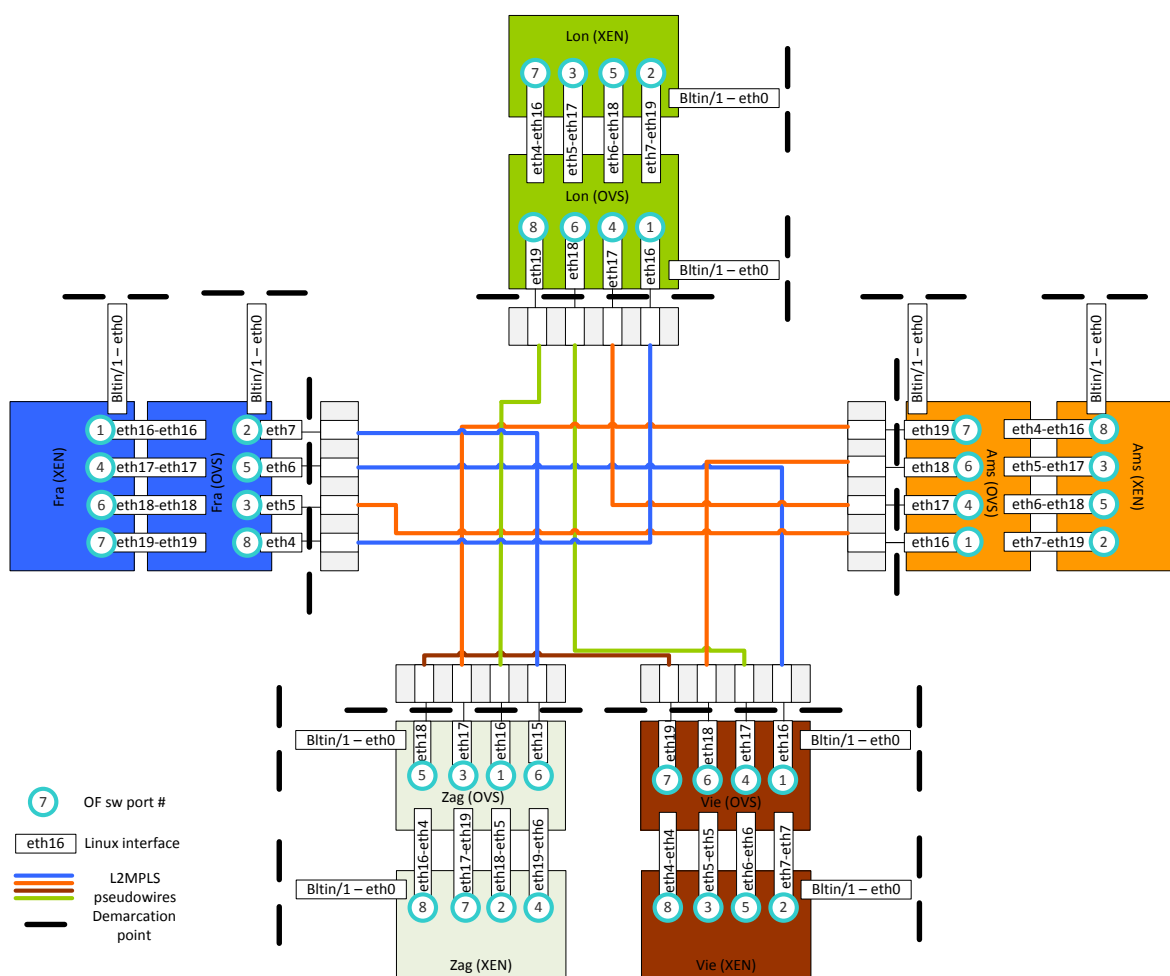


Figure 9.2: Demarcation points for the network connectivity services

Services running outside the scope of Tier-1 (the substrate), such as the GOCF, are closely related to the experimental scope of the facility and introduce non-standard operations for Tier-2.

Table 9.1 presents the operational tasks at all tiers and the corresponding assignments to the different operations entities. Generally speaking, production GÉANT services that are required by the OpenFlow Facility are operated by the GÉANT Network Operations Centre (NOC) at Tier-1, while facility-specific components, exclusively developed for the facility operations, are operated by Subject Matter Experts (SMEs) from JRA2 T5 at the moment of writing⁵. Table 9.1 also highlights slice operations at Tier-3 exercised by slice users.

⁵ In the long term, Tier-2 operations should be devolved to the operations entity defined for the facility.

Operational Task	Service Level/Provider	Operations Entity
Cabling	Substrate Provider	GÉANT NOC
Power Supply	Substrate Provider	GÉANT NOC
Air Conditioning	Substrate Provider	GÉANT NOC
Physical Installation/Maintenance	Substrate Provider	GÉANT NOC
IPv4/IPv6 Network Connectivity (for the control and management plane)	Substrate Provider	GÉANT NOC
IPv4/IPv6 Address Space Allocation	Substrate Provider	GÉANT NOC
Domain Name System	Substrate Provider or OpenFlow Facility provider	GÉANT NOC
Perimeter Firewalling	Substrate Provider	GÉANT NOC
L2 Data Plane Connectivity	Substrate Provider	GÉANT NOC
Operating System Administration (Xen and OVS servers)	OpenFlow Facility provider	JRA2 T5
Open vSwitch Administration	OpenFlow Facility provider	JRA2 T5
OpenFlow Proxy Controller (FlowVisor) Administration	OpenFlow Facility provider	JRA2 T5
Xen Virtual Machine Administration	OpenFlow Facility provider	JRA2 T5
GOCF Software	OpenFlow Facility provider	JRA2 T5
Experimenters Support	OpenFlow Facility provider	JRA2 T5
Experimenters Registration	OpenFlow Facility provider	JRA2 T5
Virtual Machine Operations (create, delete, start, shutdown)	User/Experimenter	User/Experimenter
OpenFlow Controller Operations	User/Experimenter	User/Experimenter
Slice Topology (create, edit, delete)	User/Experimenter	User/Experimenter

Table 9.1: Operational tasks across the tiers and assignments to operations entities

As part of the future work on the facility and foreseen service offerings, thorough documentation of the operational processes at all three tiers, vertical communication across operational entities from the different tiers (e.g. Tier-1 to Tier-2 interactions), ticketing, request fulfilment and troubleshooting are essential.

10 Use Cases and International Collaboration

10.1 Use Cases

The GÉANT OpenFlow Facility aims to support applications and user communities requiring specific TaaS functionality that is not offered by the GÉANT service portfolio at the moment of writing.

Due to the technical scope and the short timeframe of the work presented in this report, user-community requirements have not been systematically defined and addressed.

However, use cases in particular are essential in further development of the GÉANT OpenFlow Facility and its path towards service offerings. Based on the literature but also on some basic analysis of use cases observed in the R&E community, two major categories of use cases have been identified:

1. Using OpenFlow as a traffic-engineering mechanism to manage programmatically the backbone capacity and paths to serve specialised applications and protocols at the end systems. Typical examples of such use cases include virtual machine live migration across remotely located data centres over a WAN and efficient big data transfer utilising OpenFlow-enabled traffic engineering [MPTCP].
2. Using OpenFlow to deliver “vanilla” Layer 2 slices for Layer 2 (and above layers) research and experimentation on the actual network data and control plane technologies. Research on new protocols or capabilities such as that of [EthOAMOF] fall into this category.

At the moment of writing, a user manual has been delivered [GÉANTOFFUM] and the Multipath TCP (MPTCP) use case has been invited onto the facility. While the experimentation itself, as reported in [MPTCP], has already been carried out in a user-implemented testbed, the user experience with the GÉANT OpenFlow Facility has been reported as very positive.

At the same time, a slice contest [OFComp] has been held by the JRA2 T5 team. During the contest, an invitation for a researcher or team to experiment on the facility for one month with their own research/experimentation project was announced. Several interesting proposals were received and evaluated based on the relevance of the experimentation proposed to a wide-area OpenFlow facility, such as that of GÉANT. The winner has been invited to collaborate with the GN3 JRA2 T5 team during the slice setup and experimentation time, by providing feedback on the experience and the outcomes of their research utilising the facility, and to contribute to testing the facility itself.

The winner of the contest was the MTA-BME Future Internet Research Group in the Budapest University of Technology and Economics (Department of Telecommunications and Media Informatics). The proposed research [OFCompMTA-BME] focuses on network resiliency research at the transport layer, utilising OpenFlow capabilities to provide edge-disjoint routes on top of a network core. It is therefore a category 2 use case, as defined above.

The research aims to eliminate the need for concurrent resiliency mechanisms at the physical, transport and network layer and associated complexities/costs. Using the GÉANT OpenFlow Facility for this research is considered to be particularly relevant, since the facility provides a realistic environment to measure resilience mechanisms and convergence times exploiting the OpenFlow capabilities offered.

The proposal [OFCompGWDG] from the eScience research group of the Society for Scientific Data Processing mbH Göttingen (GWDG), a corporate facility of the Georg-August University of Göttingen and the Max-Planck-Gesellschaft, was the runner-up in the contest. It focuses on research into OpenFlow itself, proposing testing and evaluation of a specialised SDN controller developed by the group, with a focus on the use case of inter-data centre federation. A theoretical framework is suggested to confirm performance findings as a result of testing on top of the GÉANT OpenFlow Facility. The facility is considered particularly useful for the group's work as the use case requirements include a wide-area environment with disjoint paths and the GÉANT OpenFlow Facility can also be used to test the controller capabilities.

Other proposals received include that of Synchromedia Laboratory, Ecole de Technologie Supérieure, University of Quebec and the GreenStar Network (CANARIE) from Canada [OFCompSynch]. The proposal introduces a "follow the wind, follow the sun" paradigm utilising OpenFlow capabilities for migrating virtual data centres and providing QoS for cloud-based applications. The use case is dependent upon features such as those of the GÉANT OpenFlow Facility, as it requires a global-scale cloud federation environment across Canada and Europe to test and deploy cloud-based power management for reduction of greenhouse gas (GHG) emissions. This is clearly a category 1 use case, as defined above.

Another proposal of category type 2 was that from the i2CAT Foundation, Distributed Applications and Networks Area [OFCompi2CAT]. The proposal builds upon the concept of Network Function Virtualisation (NFV) and utilises the properties of the GÉANT OpenFlow Facility to experiment with and research upon a routing function virtualisation solution. The proposal is endorsed by an industry partner (Telefonica I+D) with potential for dissemination towards the Network Function Virtualisation Industry Specification Group (NFV-ISG) under the auspices of the European Telecommunications Standards Institute (ETSI).

This short-term, limited-scope contest has confirmed the interest of the research community in experimentation on top of the GÉANT wide-area OpenFlow Facility and its particular value in supporting applications and Future Internet research. It has also provided a set of very interesting use cases and a strong indication of returns on the investment.

10.2 International Collaboration

Regarding international collaborations, the team has followed the Internet Research Task Force (IRTF) Software Defined Networking Research Group (SDNRG), closely collaborated with OFELIA partners on the

facility design and deployment issues, as well as providing feedback on the OCF releases, and discussed with Internet2 possible synergies with their Advanced Network Services and Network Development and Deployment Initiative / Open Science, Scholarship and Services Exchange (NDDI/OS³E) initiatives.

Foreseen future work includes multi-domain extensions of the GÉANT OpenFlow Facility offerings through relevant NRENs' and global-wide initiatives as well as engagement in relevant international initiatives, liaison with FP7 projects, and the presence standardisation activities of the Open Grid Forum (OGF) and Internet Engineering Task Force / Internet Research Task Force (IETF/IRTF) [SDNIETF].

11 Conclusions

Although originating from campus environments, SDN/OpenFlow has obtained significant recognition among RENS in recent years. Overall, RENS are gradually researching, piloting and adopting SDN/OpenFlow to serve the specialised needs of their users. The introduction of OpenFlow to RENS presents many different types of use cases. The work presented in Part 2 of this document focuses on a particular category of use cases, providing solid evidence of the suitability of OpenFlow as a technology offering Testbed as a Service (TaaS) capabilities. TaaS utilising OpenFlow delivers to researchers a slice composed of an isolated L2 network environment with attached computation resources, and gives them full control of the resources in their slice through a set of management utilities.

The GÉANT OpenFlow Facility has been successfully implemented on top of the GÉANT backbone production environment. It combines GÉANT substrate capabilities, software-based OpenFlow-enabled network elements, computing resources and a control framework for managing the facility and delivering slices to users. It is a complex environment, which at the same time preserves simplicity of access and use for the end users. It is the first testbed facility implemented on top of GÉANT that offers “vanilla” L2 slices. It is also one of the pioneering OpenFlow-enabled facilities globally offering WAN TaaS. A three-tier architecture for operations has been successfully designed and delivered for development as part of the evolution of the facility offerings in a service context.

As part of the dissemination and user outreach activities carried out in the last months of GN3, JRA2 T5 invited proposals for experimentation utilising the specific capabilities of the facility; several interesting proposals were received. This has confirmed the interest of the research community in experimentation on top of the GÉANT wide-area OpenFlow Facility and its particular value in supporting applications and Future Internet research. It has also provided a strong indication of returns on the investment. Working with some of the user groups has resulted in essential feedback for the evolution of the facility and associated service offerings as well as in concrete facility-enabled research results. As the GN3 project concludes, a roadmap of facility enhancements as well as a plan for the handover of its operations set the foundation for an operational TaaS service within GN3plus.

Appendix A Definition of Terms

ID	Description
OpenFlow	A standardised communication interface acting between the respective control and data-forwarding planes of an SDN architecture [ONFSDN]
OpenFlow Facility	A physical infrastructure of interconnected data forwarding and control plane elements adhering to and communicating in accordance with the OpenFlow specification
OpenFlow Slice	A functional subset of an OpenFlow Facility's resources allocated for exclusive use by a specific End User/Experimenter
OpenFlow Switch	A software or hardware implementation of an OpenFlow-specification-compliant switch
OpenFlow Controller	An OpenFlow-specification-compliant agent implementing control plane functionality for one or more OpenFlow Switches by managing their flow tables
OpenFlow Proxy Controller	A special-purpose controller, permitting the OpenFlow Facility slicing to different End Users by handling OpenFlow control messages from End User controllers to the OpenFlow Facility switches and vice versa, acting as a proxy and policy-enforcing entity
End Users/Experimenters	OpenFlow Facility users provided with access to the facility resources and allocated with a facility slice for use/experimentation. NRENs as well as academic and research institutions and research projects or groups/consortia connected to/through NRENs are potential End Users of the OpenFlow Facility
OpenFlow Facility substrate	The set of physical resources upon which OpenFlow Facility slices are implemented, including CN-provided backbone links, switches, servers

Table A.1: Definition of terms

Appendix B **GÉANT OpenFlow Facility: General Server Specifications**

Item	Quantity/Duration
Server	Ten (10)
Rack Units	≤ 2
Number of CPUs	≥ 2
Number of cores per CPU	≥ 4
CPU Cache Size	≥ 6 Mbyte
CPU Frequency	≥ 2.60 GHz
Memory Size	≥ 16 Gbyte
RAID Controller	RAID 1/5 with SAS HDDs & ≥ 4 HDDs support
HDDs	≥ 2 x SAS 146 GB
Network Interfaces	≥ 12 x Gigabit
Out-of-Band Server Management Card	1
Redundant Power Supply	1
Optical Drive	1
4-Hour Mission Critical Support	1 year

Table B.2: GÉANT OpenFlow Facility: general server specifications

References

Part 1

- [ALIEN D2.1]** ALIEN Deliverable D2.1 - Report on Hardware abstraction models [Online]
<http://fp7-alien.eu/files/deliverables/D2.1-ALIEN-final.pdf>
- [ALIEN]** ALIEN Project. [Online]
<http://www.fp7-alien.eu>
- [AristaSDN]** “Arista Networks Showcases Software Defined Cloud Networking With SDN Controllers from Vmware, Big Switch Networks and Nebula”, Amanda Jaramillo, Arista Networks Press Releases, May 2012
<http://www.aristanetworks.com/news/pressrelease/467-pr-20120502-01>
- [Beacon]** <https://openflow.stanford.edu/display/Beacon/Home>
- [BigSwitchFloodlight]** “Big Switch Networks delivers over 6000 Downloads of the Apache-Licensed Floodlight Controller and APIs creating the largest SDN ecosystem”, Press Releases, Big Switch Networks, August 2012
<http://www.bigswitch.com/press-releases/2012/08/20/big-switch-networks-delivers-over-6000-downloads-of-the-apache-licensed>
- [BigSwitchSDNPrimer]** “Software-Defined Networking Primer + Deep dive into Big Switch Networks”, Brian Marshall, Stephen Patel, International Strategy & Investment Group INC., August 2012
http://www.bigswitch.com/sites/default/files/sdn_resources/isi-research.pdf
- [Broadcom]** <http://www.broadcom.com/products/features/BCM56846.php>
- [Brocade100GE]** “Brocade Delivers 100 Gigabit Ethernet Solutions for Software-Defined Networks”, Brocade Newsroom, 22 May, 2012,
<http://newsroom.brocade.com/press-releases/brocade-delivers-100-gigabit-ethernet-solutions-fo-nasdaq-brcd-0890051>
- [BrocadeExploreSDN]** “Exploring Software-Defined Networking (SDN), its expected role in cloud-optimized networks, and the Brocade plans for SDN support”, Brocade White Paper, 2012
http://www.brocade.com/downloads/documents/white_papers/exploring-sdn-wp.pdf
- [BrocadeMLX]** <http://www.brocade.com/products/all/routers/product-details/netiron-mlx-series/system-options.page>
- [BrocadeNEC]** “Brocade and NEC Expand Efforts to Drive Market Adoption of Software-Defined Networking”, Brocade Newsroom, 22 May, 2012
<http://newsroom.brocade.com/press-releases/brocade-and-nec-expand-efforts-to-drive-market-ado-nasdaq-brcd-0890053>
- [BrocadeNetworkTrans]** “Network Transformation with Software-Defined Networking and Ethernet Fabric”, Brocade White Paper, 2012

	http://www.brocade.com/downloads/documents/positioning-papers/network-transformation-sdn-wp.pdf
[BrocadeNICER]	http://www.brocade.com/products/all/routers/product-details/netiron-cer-2000-series/index.page
[BrocadeNICES]	http://www.brocade.com/products/all/switches/product-details/netiron-ces-2000-series/index.page
[BrocadeOFComing]	“OpenFlow Coming to MLX and CES/CER”, Norrival, Community Blog Brocade, May 2012 http://community.brocade.com/community/brocadeblogs/service_providers/blog/2012/05/21/openflow-coming-to-mlx-and-cescer
[BrocadeOFswitch]	“Brocade: Some OpenFlow switches can’t operate at line-rate speeds”, Shamus McGillicuddy, TechTarget, May 2012 http://searchsdn.techtarget.com/news/2240150716/Brocade-Some-OpenFlow-switches-cant-operate-at-line-rate-speeds
[BrocadeTechDay]	“Brocade Technology Day”, John McHugh, June 9, 2010, Gotham Hall, New York City
[BrocadeVyatta1]	“Brocade bolsters its SDN game, buys Vyatta”, Larry Dignan, ZDNet News, November 2012 http://www.zdnet.com/brocade-bolsters-its-sdn-game-buys-vyatta-7000006913/
[BrocadeVyatta2]	“Brocade buying Vyatta to add to SDN capabilities”, Matthew Palmer, SDNcentral.com, November 2012 http://www.sdncentral.com/sdn-blog/breaking-news-brocade-buying-vyatta/2012/11/
[BrocadeVyatta3]	“Brocade to buy Vyatta, Software-Defined Network play”, Eric Savits, Forbes, November 2012 http://www.forbes.com/sites/ericsavitz/2012/11/05/brocade-to-buy-vyatta-software-defined-network-play/
[CienaOFStripes]	“Ciena Shows Its OpenFlow Stripes at ONF Plugfest”, Chris Janz, Head of Market Development, Ciena Insights Blog, October 2012 http://www.ciena.com/connect/blog/Ciena-Shows-Its-OpenFlow-Stripes-at-ONF-Plugfest.html
[CienaSDNStrat]	“Detailing Ciena’s SDN Strategy”, Chris Janz, Head of Market Development, Ciena Insights Blog, November 2012 http://www.ciena.com/connect/blog/Detailing-Cienas-SDN-Strategy.html
[CiscoCatalyst]	“Catalyst 6500 upgrade: OpenFlow, LISP, next-gen VSS”, Shamus McGillicuddy, Techtarget, October 2012 http://searchnetworking.techtarget.com/news/2240169262/Catalyst-6500-upgrade-OpenFlow-LISP-next-gen-VSS [membership required]
[CiscoOF]	“Cisco on why OpenFlow alone doesn’t cut it”, Jim Duffy, NetworkWorld, June 2012 http://www.networkworld.com/community/blog/cisco-why-openflow-alone-doesnt-cut-it
[CiscoSDNsuspense]	“Cisco ends the SDN suspense”, Jim Duffy, NetworkWorld, June 2012 http://www.networkworld.com/news/2012/061312-cisco-sdn-260138.html
[CiscoThreat]	“Cisco Systems faces threat to core router business”, Drew Fitzgerald, MarketWatch, November 2012 http://www.marketwatch.com/story/cisco-systems-faces-threat-to-core-router-business-2012-10-24
[CloudToad]	Cloud Toad, Packet Pushers, “More OpenFlow Thoughts... Then I’m Done, I Promise” http://packetpushers.net/more-openflow-thoughts-then-im-done-i-promise/
[ContXtream]	http://www.contextream.com/products/products.html
[DellForce10]	“Dell Builds on Force10 for Virtual Networking”, Sean Michael Kerner, Enterprise Networking Planet, April 2012

- <http://www.enterprisenetworkingplanet.com/datacenter/dell-builds-on-force10-for-virtual-networking.html>
- [DellForce10Interop]** “Interop 2012 OpenFlow Roundup”, Blogs, OpenFlowHub, May 2012
<http://www.projectfloodlight.org/blog/2012/05/10/interop2012/>
- [DellISDNevolve]** “SDN won’t fully evolve for 3-5 years, Dell says”, Stephen Lawson, Computerworld news, 30 August 2012
http://www.computerworld.com/s/article/9230798/SDN_won_t_fully_evolve_for_3_5_years_Dell_says
- [DevoFlow]** A. Curtis et al., “DevoFlow: Scaling Flow Management for High-Performance Networks”
<http://conferences.sigcomm.org/sigcomm/2011/papers/sigcomm/p254.pdf>
- [DJ1.1.2]** “Deliverable DJ1.1.2: Transport Network Technologies – Study and Testing”, JRA1, GÉANT Project, May 2012
http://geant3.archive.geant.net/Media_Centre/Media_Library/Media%20Library/GN3-12-100_DJ1-1-2_Transport-Network-Technologies-Study-and-Testing.pdf
- [Embraneheleos]** <http://www.embrane.com/products>
- [ExtremeBD]** <http://www.extremenetworks.com/products/blackdiamond-x.aspx>
- [FIBRE]** FIBRE Project. [Online]
<http://www.fibre-ict.eu/>
- [Floodlight]** <http://floodlight.openflowhub.org/>
- [FlowVisor]** R. Sherwood et. al., “FlowVisor: A Network Virtualization Layer”, October 2009
<http://www.openflow.org/downloads/technicalreports/openflow-tr-2009-1-flowvisor.pdf>
- [FlowVisor-NVL]** Rob Sherwood, Glen Gibby, Kok-Kiong Yapy, Guido Appenzellery, Martin Casado, Nick McKeown, Guru Parulkary, “FlowVisor: A Network Virtualization Layer” OPENFLOW-TR-2009-1
<http://www.openflow.org/downloads/technicalreports/openflow-tr-2009-1-flowvisor.pdf>
- [GENI]** GENI Project. [Online]
<http://www.geni.net>
- [HP3500]** http://h17007.www1.hp.com/us/en/products/switches/HP_3500_zl_Switch_Series/index.aspx
- [HP5400]** http://h17007.www1.hp.com/us/en/products/switches/HP_5400_zl_Switch_Series/index.aspx
- [HP8200]** http://h17007.www1.hp.com/us/en/products/switches/HP_8200_zl_Switch_Series/index.aspx
- [HPOF]** “OpenFlow: Enabling technology for software-defined networking”, Teri Shiozaki, Social Media Program Manager, HP Networking, September 2012
<http://h17007.www1.hp.com/uk/en/solutions/technology/openflow/index.aspx>
- [HPOF16switch]** “HP Networking makes OpenFlow generally available on 16 switch models”, Shamus McGillicuddy, TechTarget, February 2012
<http://itknowledgeexchange.techtarget.com/networkhub/hp-networking-makes-openflow-generally-available-16-switch-models/>
- [HPSDN]** “HP Advances Software-Defined Networks with Integrated Infrastructure, Controller and Application Solutions”, Samantha Singh, HP News release, October 2012
<http://www8.hp.com/us/en/hp-news/press-release.html?id=1300459>
- [HPSDNPortfolio]** “New HP SDN Portfolio: Controller, switches and network apps”, Shamus McGillicuddy, TechTarget, October 2012
<http://searchsdn.techtarget.com/news/2240164400/New-HP-SDN-portfolio-Controller-switches-and-network-apps>

- [HPVAN]** “Virtual Application Networks”, Fact Sheet, HP Networking, October 2012
<http://www.tsa.com/documents/news/TSA%20Networking%20for%20Web2.pdf>
- [IBMG8264]** “IBM System Networking RackSwitch G8264 and RackSwitch G8264T: 10 Gb/40 Gb Ethernet Switch designed for the data center”, Data Sheet, IBM Systems and Technology, July 2012
http://www.hamburgnet.de/products/ibm/G8264_Datasheet.pdf
- [IBMG8264]** <http://www-03.ibm.com/systems/networking/switches/rack/g8264/>
- [IBMNEC]** “IBM and NEC Team Up to Enable Industry Innovators Tervela and Selerity to Transform their Networks with OpenFlow”, IBM News room, January 2012
<http://www-03.ibm.com/press/us/en/pressrelease/36566.wss>
- [IBMOFController]** “IBM’s OpenFlow controller: Big Blue eyes SDN’s application layer”, Shamus McGillicuddy, TechTarget, October 2012
<http://searchsdn.techtarget.com/news/2240164654/IBMs-OpenFlow-controller-Big-Blue-eyes-SDNs-application-layer>
- [IBMOFNG]** “OpenFlow: The next generation in networking interoperability”, Systems and Technology – Technical White Paper, IBM, May 2011
<http://public.dhe.ibm.com/common/ssi/ecm/en/qcw03010usen/QCW03010USEN.PDF>
- [IBMPNC]** “IBM Programmable Network Controller”, Data Sheet, IBM Systems and Technology, October 2012
<http://public.dhe.ibm.com/common/ssi/ecm/en/qcd03018usen/QCD03018USEN.PDF>
- [IBMSDN]** “Software Defined Networking”, Thought Leadership White Paper, IBM Systems and Technology, October 2012
<http://public.dhe.ibm.com/common/ssi/ecm/en/qcw03016usen/QCW03016USEN.PDF>
- [ITU-T G.694.1]** G.694.1 “Spectral grids for WDM applications: DWDM frequency grid” [Online]
<http://www.itu.int/rec/T-REC-G.694.1/recommendation.asp?lang=en&parent=T-REC-G.694.1-201202-1>
- [JuniperOFCntrl]** “When will Juniper ship an OpenFlow Controller?”, Roy Chua, Co-Founder, SDN Central, October, 2012
- [JuniperOpenSource]** “Juniper plotting open source alternative to Cisco, Vmware SDN controllers”, Jim Duffy, News, Network World, September 2012
<http://www.networkworld.com/news/2012/091412-juniper-muglia-262469.html>
- [JuniperSDN]** “Juniper confines SDNs to data center”, Jim Duddy, News, Network World, June 2012,
<http://www.networkworld.com/news/2012/061912-juniper-sdn-260342.html>
- [JuniperSDNVision]** “Juniper’s Vision for SDN Enables Network Innovation”, Michael Leonard, Technical Marketing Manager at Juniper, Juniper Forum, July 2012
<http://forums.juniper.net/t5/Data-Center-Directions/Juniper-s-Vision-for-SDN-Enables-Network-Innovation/ba-p/146972>
- [Maestro]** <http://code.google.com/p/maestro-platform/>
- [NECOF]** “Designing and Building a Datacenter Network: An alternative Approach with OpenFlow”, Rohit Mehra, IDC, January 2012
http://www.nec.com/en/global/prod/pflow/images_documents/Designing_and_Building_a_Datacenter_Network.pdf
- [NECPF]** “ProgrammableFlow Network Suite Brings Software Defined Networking (SDN) to Windows Server 2012 Hyper-V Environments”, Press Center, NEC, June 2012
<http://www.necam.com/about/read.cfm?ID=12196dd8-5e3d-44fa-8f7c-8d3f9a6ba619>
- [NECPF]** <http://www.necam.com/Pflow/>

- [NECPF5240] <http://www.necam.com/SDN/doc.cfm?t=PFlowPF5240Switch>
- [NECPF5820] <http://www.necam.com/Docs/?id=ba0dad4-f253-4a8a-b27a-a791378f9acf>
- [NECRadware] “Radware and NEC Collaborate to take OpenFlow security to the next level”, Press Release, Radware, May 2012
<http://www.radware.com/newsevents/pressrelease.aspx?id=1629520>
- [NECSDN] “Open Networking Summit: NEC waits for no one; SDN ‘ready to go’”, Rivka Gewirtz, Executive Director, TechTarget, April 2012
<http://searchnetworking.techtarget.com/news/2240148711/Open-Networking-Summit-NEC-waits-for-no-one-SDN-ready-to-go>
- [NETGEARGSM7352Sv2] <http://support.netgear.com/product/GSM7352Sv2>
- [NiciraNVP] <http://nicira.com/en/network-virtualization-platform>
- [NorthboundAPI] Greg Ferro, “Northbound API, Southbound API, East/North – LAN Navigation in an OpenFlow World and an SDN Compass”, 21 August 2012
<http://etherealmind.com/northbound-api-southbound-api-eastnorth-lan-navigation-in-an-openflow-world-and-an-sdn-compass/>
- [NOX] N. Gude, et. al., “NOX: Towards an Operating System for Networks”, SIGCOMM CCRewiew , Vol. 38, Issue 3, July 2008
- [NOXRepo] <http://www.noxrepo.org/>
- [NSFv1.0] Guy Roberts, Tomohiro Kudoh, Inder Monga, Jerry Sobieski, John Vollbrecht, “Network Services Framework v1.0”, GFD.173 NSI-WG
<http://www.ogf.org/documents/GFD.173.pdf>
- [OFELIA] OFELIA Project. [Online]
<http://www.fp7-ofelia.eu/>
- [OF-EnablingInnovation] N. McKeown, et al., “OpenFlow: Enabling Innovation in Campus Networks”, March 14, 2008
<http://www.openflow.org/documents/openflow-wp-latest.pdf>
- [OF-FirstDemo] M. Channegowda, R. Nejabati, M. Rashidi Fard, S. Peng, N. Amaya, G. Zervas, D. Simeonidou, R. Vilalta, R. Casellas, R. Martínez, R. Muñoz, L. Liu, T. Tsuritani, I. Morita, A. Autenrieth, J. Elbers, P. Kosteki, and P. Kaczmarek, “First Demonstration of an OpenFlow based Software-Defined Optical Network Employing Packet, Fixed and Flexible DWDM Grid Technologies on an International Multi-Domain Testbed”, in European Conference and Exhibition on Optical Communication, OSA Technical Digest (online) (Optical Society of America, 2012), paper Th.3.D.2.
<http://www.opticsinfobase.org/abstract.cfm?URI=ECEOC-2012-Th.3.D.2>
- [OF-PAC.C] Packet and Circuit Network Convergence (PAC.C) with OpenFlow. Web site:
<http://openflowswitch.org/wk/index.php/PAC.C>
- [OF-SwitchSpec] “OpenFlow Switch Specification” Version 1.0.0 (Wire Protocol 0x01), December 31, 2009
<http://www.openflow.org/documents/openflow-spec-v1.0.0.pdf>
- [OF-SwitchSpec110] “OpenFlow Switch Specification” Version 1.1.0 Implemented (Wire Protocol 0x02), February 28, 2011
<http://www.openflow.org/documents/openflow-spec-v1.1.0.pdf>
- [OF-SwitchSpecAdd] “Extensions to the OpenFlow Protocol in support of Circuit Switching: Addendum to OpenFlow Protocol Specification (v1.0) – Circuit Switch Addendum v0.3”, June, 2010
http://www.openflow.org/wk/images/8/81/OpenFlow_Circuit_Switch_Specification_v0.3.pdf
- [OF-V0.9] <http://www.openflow.org/wp/2009/07/openflow-09-released/>
- [OF-V1.0] <http://www.openflow.org/wp/2009/12/openflow-1-0-released/>

[ONF]	Open Networking Foundation website https://www.opennetworking.org/
[ONFAbout]	https://www.opennetworking.org/about
[ONFConfig1.0]	https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow-config/of-config1dot0-final.pdf https://www.opennetworking.org/standards/of-config-10
[ONFConfig1.2]	https://www.opennetworking.org/standards/1dot2-and-pending
[ONF-OF-V1.2]	https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-spec-v1.2.pdf
[ONF-OF-V1.3]	https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-spec-v1.3.0.pdf
[OpenFlow]	OpenFlow website http://www.openflowswitch.org
[OpenVSwitch]	http://openvswitch.org/
[Pica8]	http://www.pica8.org/
[Pica83920]	"Pica8 launches 3920", James Liao, Pica8 Blogs, March 2012 http://pica8.org/blogs/?p=162
[Pica83290]	http://www.pica8.org/products/p3290.php
[Pica8OF]	"Pica8 & OpenFlow", James Liao, Pica8 Blogs, February 2012 http://pica8.org/blogs/?p=40
[Pica8OSforOpen]	"Pica8's OS for Open Switches" http://www.pica8.com/open-switching/open-switching-overview.php
[Pica8startup]	"Pica8: A startup taking advantage of network commoditization", Stacey Higginbotham, Gigaom News, February 2012 http://gigaom.com/2012/02/17/pica8-a-startup-taking-advantage-of-network-commoditization/
[Pica8XorPlus]	http://www.pica8.org/products/xorplus.php
[POMI-OF-SDN]	http://pomi.stanford.edu/content.php?page=research&subpage=openflow
[SDNProducts]	"Shipping SDN Products", SDN Central website http://www.sdncentral.com/shipping-sdn-products/
[Trema]	http://trema.github.com/trema/
[vdPol]	Ronald van der Pol, SARA, "Ethernet OAM enabled OpenFlow Controller", November 2011 https://noc.sara.nl/nrg/presentations/SC11-SRS-8021ag.pdf

Part 2

[APAN-FIT-OF-Wkshop]	Chris Small, "OpenFlow Workshop", APAN FIT Workshop, Feb 22 2011 http://www.apan.net/meetings/HongKong2011/Session/Slides/fit/10.pdf
[CanProdNetTest]	R. Sherwood, G. Gibb, K. Yap, G. Appenzeller, M. Casado, N. McKeown, and G. Parulkar, "Can the production network be the testbed?", in Proceedings of the 9th USENIX conference on Operating systems design and implementation (OSDI'10). USENIX Association, Berkeley, CA, USA, 1-6.
[DCN-OF]	Jeff Boote, et al., "Dynamic Circuit Network Services and OpenFlow Integration" http://archstone.east.isi.edu/twiki/pub/ARCHSTONE/OpenFlow/sc-2011-srs-internet2-v2.pdf
[EthOAMOF]	R. van der Pol, S. Boele, F. Dijkstra, "OpenFlow Demo, IEEE 802.1ag Ethernet OAM", Nov 12-18, 2011, SCInet Research Sandbox (SC11, Seattle, USA)

[Expedient]	http://www.openflow.org/wk/index.php/Expedient
[FederatingSDN]	“Federating SDN-enabled islands with an extended NSI Framework” https://tnc2013.terena.org/core/presentation/29
[FlowVisor]	https://github.com/OPENNETWORKINGLAB/flowvisor/wiki
[GÉANTOFFDesign]	“GÉANT OpenFlow Facility Design” https://geant3-intranet.archive.geant.net/sites/Research/JRA2/T5/Documents/GEANT%20OpenFlow%20facility%20documentation/OF%20Facility%20Design%20document/GN3-12-331_GEANT-OpenFlow-Facility-Design_v1.7_A.pdf
[GÉANTOFFUM]	“OpenFlow Facility User Manual” https://geant3-intranet.archive.geant.net/sites/Research/JRA2/T5/Documents/GEANT%20OpenFlow%20facility%20documentation/User%20manual/GEANT-OF-Facility%20User_Manual.docx
[GÉANTServices]	http://www.geant.net/Services/ConnectivityServices/pages/home.aspx
[MPTCP]	R. van der Pol, S. Boele, F. Dijkstra, A. Barczyk, G. van Malenstein, J.H. Chen, J. Mambretti, “Multipathing with MPTCP and OpenFlow”, Supercomputing 2012 SCInet Research Sandbox Proposal, June 2012
[OCFRepos]	OFELIA Control Framework GitHub repository http://fp7-ofelia.github.com/ocf/
[OFComp]	http://www.geant.net/Pages/openflow_competition.aspx
[OFCompGWDG]	https://geant3-intranet.archive.geant.net/sites/Research/JRA2/T5/Documents/G%C3%89ANT%20OpenFlow%20Slice%20Contest/GWDG_Proposal_for_GEANT_Competition_Final.pdf
[OFCompi2CAT]	https://geant3-intranet.archive.geant.net/sites/Research/JRA2/T5/Documents/G%C3%89ANT%20OpenFlow%20Slice%20Contest/Geant3-OF-comp-last-ok-130222-1.pdf
[OFCompMTA-BME]	https://geant3-intranet.archive.geant.net/sites/Research/JRA2/T5/Documents/G%C3%89ANT%20OpenFlow%20Slice%20Contest/bme_hu.pdf
[OFCompSynch]	https://geant3-intranet.archive.geant.net/sites/Research/JRA2/T5/Documents/G%C3%89ANT%20OpenFlow%20Slice%20Contest/OpenFlow%20Competition%20Proposal-2.pdf
[OFELIA]	OFELIA FP7 project http://www.fp7-ofelia.eu/
[OF-EnablingInnovation2]	N. McKeown, et al., “OpenFlow: Enabling Innovation in Campus Networks”, March 14, 2008 http://ccr.sigcomm.org/online/files/p69-v38n2n-mckeown.pdf
[OF-GENI]	http://www.openflow.org/wk/index.php/GENI_Deployment
[OF-I2]	http://www.internet2.edu/pubs/IS-advanced-layer-2-service.pdf
[OF-OSCARS]	https://www.es.net/services/virtual-circuits-oscars/
[OF-RISE]	http://www.jgn.nict.go.jp/rise/english/index.html
[Open vSwitch]	http://openvswitch.org/
[SDNIETF]	SDNRG in IRTF http://trac.tools.ietf.org/group/irtf/trac/wiki/sdnrg
[Xen]	http://xen.org/

Glossary

AA	Authentication and Authorisation
AAI	Authentication and Authorisation Infrastructure
ACL	Access Control List
AP	Access Port
API	Application Programming Interface
ARP	Address Resolution Protocol
ASIC	Application-Specific Integrated Circuit
BoD	Bandwidth on Demand
BoS	Bottom of Stack
BP	Backbone Port
BV	Bandwidth Variable
BV OXC	Bandwidth Variable Optical Cross-Connect
BVT	Bandwidth Variable Transponder
CapEx	Capital Expenditure
CC	Continuity Check
CF	Centre Frequency
CFM	Connectivity Fault Management
CLI	Command-Line Interface
CN	Carrier Network
CoS	Class of Service
CPU	Central Processing Unit
DNS	Domain Name System
DPI	Deep Packet Inspection
DSCP	Differentiated Services Code Point
DWDM	Dense Wavelength-Division Multiplexing
EoS	Arista Extensible Modular Operating System
EoSDH	Ethernet over SDH
ETSI	European Telecommunications Standards Institute
EU	European Union
FI	Future Internet
FOAM	FlowVisor OpenFlow Aggregate Manager
FP7	EU's Seventh Framework Programme for Research and Technological Development
FPGA	Field-Programmable Gate Array
FSB	Frequency Slot Bandwidth
GHG	Greenhouse Gas

GMPLS	Generalised Multi-Protocol Label Switching
GOCF	GÉANT OpenFlow Control Framework
GUI	Graphical User Interface
GWDG	Society for Scientific Data Processing mbH Göttingen (Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen)
HDD	Hard Disk Drive
HPC	High Performance Computing
HW	Hardware
ICMPv6	Internet Control Message Protocol version 6 Internet Control Message Protocol version 6
ID	Identifier
IDS	Intrusion Detection System
IETF	Internet Engineering Task Force
InCENTRE	Indiana Centre for Network Translational Research and Education
IP	Internet Protocol
IPMI	Intelligent Platform Management Interface
IPS	Intrusion Prevention System
IRTF	Internet Research Task Force
ISP	Internet Service Provider
JRA	GN3 Joint Research Activity
JRA1 T1	JRA1 Future Network, Task 1 Carrier Class Transport Network Technologies
JRA2 T5	JRA2 Multi-Domain Network Service Research, Task 5 Network Factory
L2	Layer 2
LAG	Link Aggregation Group
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
LDAPS	LDAP (Over SSL)
L_n	Layer <i>n</i>
MA	Maintenance Association
MAC	Media Access Control
MAID	Maintenance Association Identifier
MCLAG	Multi-Chassis Link Aggregation Group
MD	Maintenance Domain
MEP	Maintenance End Point
MIP	Maintenance Intermediate Point
MP	Maintenance Point
MPLS	Multi-Protocol Label Switching
MPTCP	Multipath TCP
MTU	Maximum Transmission Unit
NAC	Network Access Control
NAT	Network Address Translation
NDDI	Network Development and Deployment Initiative
NE	Network Element
NFV	Network Function Virtualisation
NFV-ISG	NFV Industry Specification Group
NIC	Network Interface Card

NLR	National LambdaRail
NOC	Network Operations Centre
NOS	Network Operating System
NSI	Network Services Interface
NVGRE	Network Virtualisation using Generic Routing Encapsulation
NVP	Network Virtualisation Platform
OAM	Operation, Administration and Maintenance
OCF	OFELIA Control Framework
OF	OpenFlow
OFELIA	OpenFlow in Europe: Linking Infrastructure and Applications
OGF	Open Grid Forum
ONE	Open Network Environment (Cisco's SDN architecture)
ONS	Open Networking Summit
OpEx	Operating Expenditure
OS	Operating System
OS³E	Open Science, Scholarship and Services Exchange
OSCARS	On-Demand Secure Circuits and Advance Reservation System
OSI	Open Systems Interconnection
OvS	Open vSwitch
OXC	Optical Cross-Connect
P2P	Point to Point
PBB	Provider Backbone Bridging
PBB-TE	Provider Backbone Bridging with Traffic Engineering
perfSONAR	Performance-focused Service Oriented Network monitoring Architecture
PFC	ProgrammableFlow Controller
PNC	IBM Programmable Network Controller
PoP	Point of Presence
QoS	Quality of Service
QSFP	Quad (4-channel) Small Form-factor Pluggable
R&E	Research and Education
RAID	Redundant Array of Independent Disks
REN	Research and Education Network
REST	Representational State Transfer
RRD	Round Robin Database
SA	GN3 Service Activity
SA1	SA1 Network Build and Operations
SAOS	Service-Aware Operating System
SAS	Serial Attached SCSI
SCSI	Small Computer System Interface
SDH	Synchronous Digital Hierarchy
SDK	Software Development Kit
SDN	Software-Defined Network/Networking
SDNRG	Software Defined Networking Research Group
SFA	Slice-based Facility Architecture
SFP	Small Form-factor Pluggable

SFP+	Enhanced Small Form-factor Pluggable
SME	Subject Matter Expert
SSH	Secure Shell
SSL	Secure Sockets Layer
T	Task
TaaS	Testbed as a Service
TCP	Transmission Control Protocol
TE	Traffic Engineering
TLS	Transport Layer Security
TLV	Type Length Value
ToS	Type of Service
TTL	Time To Live
UI	User Interface
UTP	Unshielded Twisted Pair
VCS	Virtual Cluster Switching
VLAN	Virtual Local Area Network
VM	Virtual Machine
VNA	Virtual Network Architecture
VPLS	Virtual Private LAN Service
VPN	Virtual Private Network
VPP	Virtual Port Profile
VT AM	Virtualisation Aggregation Manager
VTN	Virtual Tenant Network
VXLAN	Virtual Extensible LAN
WAN	Wide Area Network
WDM	Wavelength-Division Multiplexing