Ref. Ares(2015)1042095 - 09/03/2015

Grant Agreement No. 619572

COSIGN

Combining Optics and SDN In next Generation data centre Networks

Programme: Information and Communication Technologies

Funding scheme: Collaborative Project – Large-Scale Integrating Project

Deliverable D2.1

Integration of Open Flow SW interface with POLATIS system and Venture 4x4 OXS

Due date of deliverable: January 15, 2015
Actual submission date:, January 16, 2015

Start date of project: January 1, 2014 Duration: 36 months

Lead contractor for this deliverable:

Project co-funded by the European Commission within the Seventh Framework Programme		
Dissemination Level		
PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Executive Summary

This document surveys the progress made in WP2 relating to the integration of an Open Flow software interface to the two optical switching technologies included in the project. The optical switches make part of the novel data plane architecture included in the COSIGN project. Both switching technologies represent complementary solutions to the problem of optical interconnects in large scale data center networks. The Polatis switch represents a low loss (<2dB) and high port count (>200 ports) optical interconnect with switching time in the millisecond range. The Venture switch represents an ultrafast (~1nsec) optical switch with a low port count (initially 4x4) and slightly higher losses (<4dB for first samples, but improving to 0dB by end of project). Together they make up the building blocks for the COSIGN photonic data plane switch fabric.

OpenFlow is a communications protocol that gives access to the forwarding plane of a network switch or router over the network. As such it builds upon a centralized controller to manage and control the entire network insuring optimal use of resources and improved throughput and latency under certain traffic patterns. In the COSIGN project OpenFlow has been designated as the software agent to be used to enable an SDN layer in the data plane. The different switches (including also the Top of Rack (TOR) switches) are all required to support the OpenFlow protocol.

In this report we summarize the progress made in the process towards the integration of OpenFlow in the optical switches. The result of the effort over the past 12 months are that we have successfully integrated OpenFlow into the Polatis switch and demonstrated the first use of optical circuit switching under the OpenDaylight SDN controller.

The Venture switch is still in the early stages of development and first samples have been delayed as a result of the need for further modelling and process refinement with the bespoke chip. In parallel with resolving these challenges a drive board is being developed with the help of DTU and University of Bristol. It is anticipated that the interface of the Open Flow software with the Venture switch will be demonstrated in May 2015.

Document Information

Status and Version:	Draft 1.6	
Date of Issue:	31/12/2014	
Dissemination level:	Public	
Author(s):	**Name**	**Partner**
	Harm J.S Dorren	TUe
	Oded Raz	TUe
	Tim Durrant	Venture Photonics
	Nick Parsons	Polatis
Edited by:	Oded Raz	TUe
Checked by :	Sarah Ruepp	DTU

Table of Contents

1 Introduction

1.1 Reference Material

1.1.1 Reference Documents

[1]	COSIGN FP7 Collaborative Project Grant Agreement Annex I - "Description of Work"
[2]	S. Das, "Extensions to the OF Protocol in support of Circuit Switching", addendum v0.3, June 2010
[3]	S. Das, G. Parulkar, N. McKeown, P. Singh, D. Getachew, and L. Ong, "Packet and circuit network convergence with OpenFlow," in Proc. OFC/NFOEC 2010, paper OTuG1
[4]	M. Channegowda, P. Kostecki, N. Efstathiou, S. Azodolmolky, R. Nejabati, P. Kaczmarek, A. Autenrieth, J.P. Elbers and D. Simeonidou, "Experimental Evaluation of Extended OpenFlow Deployment for High-Performance Optical Networks," in Proc.ECOC 2012, Tu.1.D.2
[5]	M. Channegowda, http://www.youtube.com/watch?v=wZTMGRfIKks
[6]	M. Channegowda, http://youtu.be/hhHMJ1i6XiQ?list=UUFt-mPr81Z6oVAX_ppIqAmw
[7]	Y. Yoshida et al," First International SDN-based Network Orchestration of Variable Capacity OPS Over Programmable Flexi-Grid EON", in Proc. OFC2014, PDP ThA.2
[8]	S Yan et al., "First Demonstration of All-Optical Programmable SDM/TDM Intra Data Centre and WDM Inter-DCN Communication", in Proc ECOC2014, PD 1.2
[9]	M Channegowda, B Guo, https://www.youtube.com/watch?v=rwvezPUwSK0&feature=youtu.be

1.1.2 Acronyms and Abbreviations

Most frequently used acronyms in the Deliverable are listed below. Additional acronyms can be specified and used throughout the text.

DC	Data Centre
DCN	Data Centre Network
DCIM	Data Centre Infrastructure Management
DMZ	Demilitarized Zone
IaaS	Infrastructure as a Service
IETF	Internet Engineering Task Force
ISO	International Organization for Standardization
IPS	Intrusion Prevention System
MOOC	Massive Open Online Course
MPLS	Multi-Protocol Label Switching
NaaS	Network as a Service
NIST	National Institute of Standards and Technology
OSS	Operational Support Services
PaaS	Platform as a Service
PCI DSS	Payment Card Industry's Data Security Standard
PE	Provider Edge
PoD	Point of Delivery
RAS	Reliability, Availability, Serviceability
RFC	Request For Comments
SaaS	Software as a Service
SDE	Software Defined Environments
SDN	Software Defined Networking
SLA	Service Level Agreement
TCAM	Ternary Content-Addressable Memory

TCO	Total Cost of Ownership
ToR	Top of the Rack
VDC	Virtual Data Centre
VLAN	Virtual LAN
VPN	Virtual Private Network
VRF	Virtual Routing/Forwarding

1.2 Document History

Version	Date	Authors	Comment
00	08/12/2014	See the list of authors	TOC first draft
01	16/1/2015		Final version

2 OpenFlow integration

2.1 Integration of OpenFlow and Polatis switch

2.1.1 Background

Polatis and the High Performance Networks Group at the University of Bristol have developed a strong collaboration on SDN applications of optical circuit switching since early 2012. Although the OpenFlow 1.0 protocol was originally designed to support packet switching, extensions for circuit switching have been proposed [2,3] and further enhanced by Bristol and others [4].

At the ECOC 2013 exhibition in London, Polatis and Bristol created a live demonstration of an OpenFlow-controlled hybrid packet-optical circuit switched datacentre network with automatic detection and routing of persistent (elephant) flows between servers from a packet-switched connection to a direct optical circuit overlay, as shown in Figure 1. An embedded OpenFlow agent (POLOA) was installed on the network interface cards of all 3 Polatis optical switches used in the demonstration. The network included two OpenFlow NEC top of rack switches and was managed via a POX OpenFlow controller. Video clips of the ECOC demonstration and a subsequent more detailed sequence can be found in [5,6].

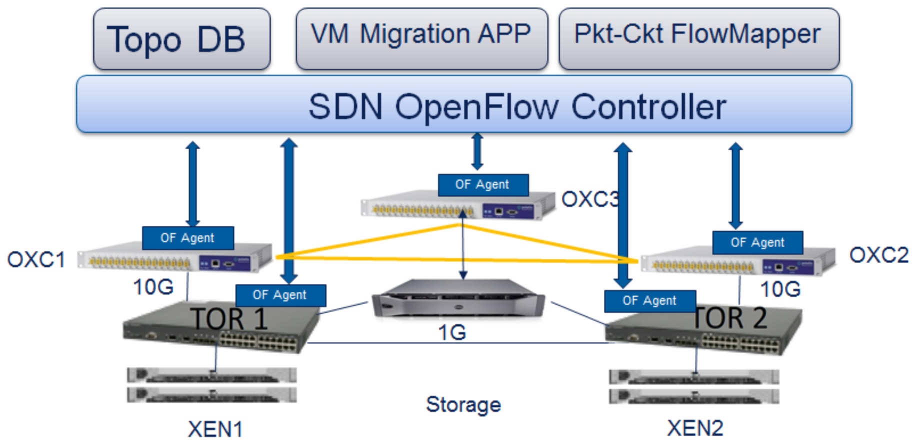

Figure 1 Hybrid packet-optical circuit switched datacentre network demonstration at ECOC 2013

2.1.2 Progress

In February 2014, Polatis delivered a fully-featured 192x192 DirectLight optical switch with input and output optical power monitors to the University of Bristol's High Performance Networks lab (figure 2). This unit, along with 4 other OpenFlow-enabled Polatis optical circuit switches installed in the networks lab, was used in a wide-scale UK-Japan SDN orchestration experiment under the STRAUSS project and presented as a post deadline paper at OFC 2014 [5]. The equipment was also used for the programmable all-optical intra-datacentre network experiment reported in [8].

Figure 2 A standard Polatis 192x192 optical circuit switch with LC/APC connectors: front panel(above); rear panel (below) showing redundant hot-swappable network interface cards, power supply units and fans

During the year, the team has created southbound plugins for the OpenDaylight controller to interface with the Bristol/Polatis OpenFlow agent and were able to demonstrate for the first time OpenDaylight control of optical circuit switches during the COSIGN consortium meeting at Bristol in September 2014 [9].

2.1.3 Plans

On a separate project, the Bristol team are developing a faster and more efficient version of the OpenFlow 1.0+ agent which integrates directly with the Polatis user services API (figure 3), rather than connecting indirectly via the TL1 interface as was done in the first release. The agent size has also been reduced by a factor of ~3 by eliminating redundant code. The agent is currently being ported to the network interface on a 192x192 optical switch and will go through integration testing during January 2015.

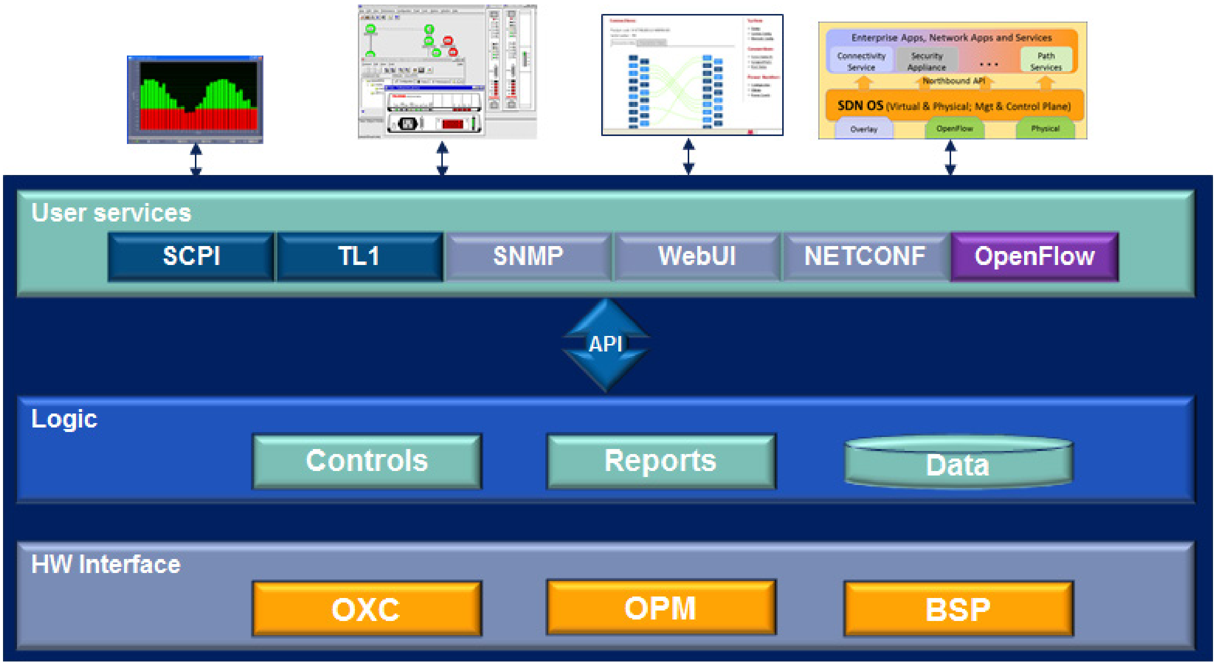

Figure 3 Polatis Optical Switch User Interface Structure

It is intended to develop the embedded OpenFlow agent further over the next 6 months to be compliant with the next stable release of OF 1.4/1.5, which will provide direct support for optical circuit switching, as opposed to the vendor/experimenter-specified extensions used with OF1.0+. Additional user interfaces will also be created to add REST and NETCONF support for configuration management.

2.2 Integration of OpenFlow and Venture switch

The Venture switch is still in the early stages of development and first samples have been delayed as a result of the need for further modelling and process refinement with the bespoke chip.

In order to minimize risk and shorten the time line to completion, a parallel activity to the chip development a drive board is being developed with the help of DTU and University of Bristol. This will facilitate the OpenFlow interface.

It is anticipated that the interface of the Open Flow software with the Venture switch will be demonstrated in May 2015

3 Annex: OpenFlow 1.0+ optical circuit switch extensions

This annex describes the current command extensions used by the embedded OpenFlow agent (POLOA) developed by UNIVBRIS to control the Polatis optical circuit switch. The command set is defined by OpenFlow 1.0 and the Circuit Switch Addendum v0.3 [2] together with additional proprietary extensions as detailed below.

Within the lifetime of the project, it is planned to enhance the embedded OpenFlow agent to support OpenFlow 1.4 (which natively supports circuit switching) and so the following information will be revised.

3.1 Connect input to output port

The circuit switch Flow Table is modified using an OFPT_CFLOW_MOD command, and as noted in the addendum it is required that OFPFC_MODIFY_STRICT and OFPFC_DELETE_STRICT are used to modify and terminate existing connections. The Polatis switch only supports real ports.

3.2 Read power at given port

Uses a proprietary extension to the OFPT_STATS_REQUEST command with type set to OFPST_CPORT = 0x18 and a following datablock laid out as:

```
struct ofp_cport_stats_request {
     uint16_t port_no;
     uint8_t direction;
     uint8_t pad[5];
 };
  OFP_ASSERT(sizeof(struct ofp_cport_stats_request) == 8);
```

The OFPST_PORT message must request statistics either for a single port (specified in port_no) or for all ports (if port_no == OFPP_NONE).
The response appears in the corresponding OFPT_STATS_REPLY, with a proprietary body either once or multiple times:

```
struct ofp_cport_stats {
    uint16_t port_no;
    uint16_t pmonset_lambda;
    uint16_t pmonset_offset;
    uint16_t pmonset_atime;

    uint32_t mes_pmon;
    uint8_t pad[4];          /* Align to 64-bits. */

    uint8_t voaset_mode;
    uint8_t voaset_level;
    uint8_t voaset_ref;
    uint8_t pad1[5];          /* Align to 64-bits. */

    uint64_t port_alarms;
};
OFP_ASSERT(sizeof(struct ofp_cport_stats) == 32);
```

If a value is unsupported, the field is set to all ones.

3.3 Trigger alarm for power below threshold

The port_alarm field of the OFPT_STATS_REPLY structure can be checked to discover if an alarm has been raised.