Low latency and high throughput dynamic network infrastructures for high performance datacentre interconnects

Small or medium-scale focused research project (STREP)

Co-funded by the European Commission within the Seventh Framework Programme

Project no. 318606

Strategic objective: Future Networks (ICT-2011.1.1)

Start date of project: November 1st, 2012 (36 months duration)

Deliverable D4.1

The LIGHTNESS network control plane architecture

Due date: 31/07/2013

Submission date: 13/09/2013

Deliverable leader: **NXW**

Author list: Giacomo Bernini (NXW), Paolo Cruschelli (NXW), Nicola Ciulli (NXW), Roberto Monno (NXW), Salvatore Spadaro (UPC), Davide Careglio (UPC), Shuping Peng (UNIVBRIS), Alessandro Predieri (IRT), Matteo Biancani (IRT), Chris Liou (Infinera), Ifthekar Hussain (Infinera)

Dissemination level

☒	**PU**	Public
☐	**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)

This page is left intentionally blank

Abstract

This document provides the architecture of the unified network control plane that will be implemented in the context of the LIGHTNESS WP4 research activities.

The LIGHTNESS control plane is conceived to provide on-demand and flexible connectivity services inside the hybrid optical data centre network with the final aim of matching the requirements of data centre applications, also supporting mechanisms and procedures for dynamic modification, optimization, recovery and monitoring. This set of data centre services are offered through a cloud-to-network interface exposed towards data centre management and orchestration entities.

This document presents the LIGHTNESS control plane functional split, based on the Software Defined Networking architecture, providing a description of the architecture model and supported procedures. It also includes a possible deployment scenario where the LIGHTNESS control plane cooperates with end-to-end cloud orchestration and control functions for the provisioning of inter data centre network services in distributed cloud scenarios.

Together with next Deliverables D4.2 (“The LIGHTNESS network control plane protocol extensions”) and D4.3 (“The LIGHTNESS network control plane interfaces and procedures”), this document will drive the subsequent software development activities of the LIGHTNESS control plane in WP4.

Table of Contents

Figure Summary

Table Summary

0. Executive Summary

This deliverable describes the LIGHTNESS network control plane architecture, as a result of the architectural investigations and design activities carried out in the context of Task 4.1. The architecture specified in this document will first drive the identification of control plane protocol extensions in Task 4.2 on the one end, and the detailed specification of the control plane interfaces and procedures in Task 4.3. All these specifications together will then form the basis for the software development in Task 4.4.

LIGHTNESS aims at developing a high-performance network infrastructure for data centres, where innovative optical switching technologies are combined and deployed to meet the emerging needs of distributed cloud and data centre applications. In particular, LIGHTNESS defines and implements a scalable Data Centre Network (DCN) architecture for ultra-high bandwidth, dynamic and on-demand network connectivity services. In this context, the LIGHTNESS control plane is conceived to support automated procedures for the provisioning of flexible and highly reliable connectivity services inside the data centre, with the final aim of overcoming the limitation of currently deployed static and semi-automated data centre network management and control frameworks.

The LIGHTNESS control plane allows the provisioning of the hybrid optical data centre network resources through a set of integrated and innovative functionalities and procedures, implementing connectivity services setup and teardown, dynamic service modification, path and flow computation, data centre network resiliency and monitoring, dynamic and automated resource optimization. The LIGHNTESS control plane is also equipped with an open, flexible and extensible interface at the northbound for cooperation with management and orchestration entities. In addition, an open, standard, vendor independent and technology agnostic southbound interface allows to configure and monitor the underlying Optical Circuit Switching and Optical Packet Switching physical devices composing the hybrid data centre network.

This deliverable follows this structure:

- Chapter 2 provides an overview of network control plane technologies and architectures currently available in the state-of-the-art, along with a brief analysis of their applicability to the data centre environment.

- Chapter 3 describes the LIGHTNESS control plane requirements, in terms of functionalities to be supported to meet the data centre applications and services needs.
- Chapter 4 describes the LIGHTNESS control plane architecture, identifying and detailing its functional decomposition and external interfaces.
- Chapter 5 presents the LIGHTNESS control plane procedures, focusing on the mechanisms supported to handle dynamic and on-demand connectivity services, path computations and enhanced functionalities for data centre network resiliency, optimization and monitoring.
- Chapter 6 summarizes the main conclusions and achievements of this document, including also a brief analysis of the future work to be carried out in the upcoming WP4 activities.

1. Introduction

1.1. Motivation and scope

Data centres represent the key element of the cloud, where storage and compute resources reside. The value of cloud services, as they have emerged in the last years, lies in the cost-effective and on-demand deployment of new applications that can be easily accessed by users. Data centre and cloud providers are therefore able to meet their business needs by broadly instantiate services and applications very quickly with simplified operations. Server virtualization played a fundamental role to introduce dynamic resource sharing and distribution inside the data centres. Today, Virtual Machines (VMs) can be easily created and ran on servers that in the past were dedicated to single environments and applications.

In this context, the Data Centre Networks (DCNs) that interconnect storage and compute resources must operate in coordination with the services running on top of them to deliver cloud-based applications even more efficiently: the network connectivity services need to follow the dynamics of the IT resources in the data centre. Unfortunately, this is not the case of today's data centres. Applications running inside VMs are able to come up very quickly (minutes), but they must wait long time (hours, days) for dedicated network connectivity services to be established. Indeed, current DCN infrastructures require multiple levels of static and human driven configurations to be operated, resulting in rigid and slow deployment procedures. These static, configuration-driven processes bring delays in deploying new applications and services, highly affecting the users' experience while increasing operational costs for data centre and cloud providers.

This means that for data centre and cloud services to be more effective, the networks within and among data centres need to evolve to become as dynamic and flexible as the IT infrastructures. From data centre and cloud providers' perspective, in order to increase the value of their services for the customers, network resources inside data centres must be managed, deployed and handled on-demand with dynamic flexible procedures that follow both applications and users needs.

LIGHTNESS proposes a new data centre infrastructure, where the hybrid optical DCN is controlled and operated by a network control plane that promises to provide on-demand, flexible and programmable connectivity services to meet the requirements of current data centre and cloud services. In particular this document focuses on the specification of the LIGHTENSS control plane high-level architecture and functional decomposition.

2. Reference control plane technologies and architectures

This chapter provides a brief survey of the most relevant reference technologies, architectures and standards that will be evaluated as the basis for the specification of the LIGHTNESS control plane architecture. This initial study will lead to the identification of the architectural models, protocols and functional modules to be included in the LIGHTNESS control plane. Thus it drives the specification of those building blocks to be adopted or extended from the state-of-the-art on the one end, and those to be designed from scratch on the other.

This reference analysis carried out in this chapter is organized in three main streams. First, the DCN control and management frameworks currently adopted by data centre operators to operate their network infrastructures are briefly analyzed, highlighting pros and cons with respect to the emerging needs of highly distributed cloud applications. Then, a quick overview of Internet Engineering Task Force (IETF) GMPLS and PCE reference standards, widely implemented and adopted by network operators to provide control plane procedures for automated provisioning of network connectivity services in transport networks, is carried out to evaluate their applicability in data centre environments. Finally, the main SDN concepts for data centre solutions currently under investigation and standardization in the Open Networking Forum (ONF), and other standardization bodies as well, are briefly summarized.

2.1. Current data centre network control and management

Today, DCNs are mostly managed by data centre and cloud operators as independent pools of resources to be bound to IT services and applications running in the servers. The control and management of network resources inside the data centre are commonly performed with static and semi-automated procedures and tools. Indeed, most of data centre and cloud providers currently deploy operator-driven control and management platforms to provision and monitor the DCN infrastructure. Most of the actions, such as those related to the DCN security (e.g. for Authentication, Authorization and Accounting

- AAA), are therefore performed by semi-automated procedures which need a human supervision and validation.

Current DCN control and management platforms focus on supporting efficient operations and management of data centre fabrics, by providing specific functions to trigger the semi-automated provisioning of secured connectivity services, and proactively monitor the network infrastructure to detect performance degradation and failure conditions. The aim is to provide visibility and control of the DCN infrastructure through a single management point, and ease the diagnosis and troubleshooting of data centre outages. In addition, DCN implementations over the last decade are simply not sustainable in the cloud services era. Traditionally, static network elements and devices management utilities have been stitched together requiring lots of integration resulting in complex operation and maintenance procedures. This rigid control methodology has imposed hard architectural constraints to data centre operators as they try to deliver and deploy cloud services in a cost-effective and scalable way.

On the contrary, modern data centres are becoming increasingly complex and massive, mainly due to the emerging of new virtualization technologies which add further levels of complexity while enabling higher workloads to be accommodated in the DCN. The management of such virtualized infrastructures (at both IT and network level) is requiring novel automated control frameworks able to provide dynamicity, flexibility and high availability of DCN connectivity services. As a further requirement, network services to be delivered inside the data centres for cloud applications are restricted on three key dimensions: the ease of their definition, the speed of their instantiation and the ease of their modification. These three concepts combined together will allow data centre and cloud providers to deploy for their customers new applications very quickly and with dynamic and flexible network performances and characteristics (e.g. bandwidth, latency, jitter, etc.), actually bound to the user and application needs.

2.2. GMPLS and PCE reference standards

The **Generalized Multi-Protocol Label Switching** (GMPLS) framework [RFC3945] is defined within the IETF Common Control and Measurement Plane Working Group (CCAMP WG), as an extension of the MPLS specification [RFC3031]. The GMPLS architecture provides control plane procedures for automated provisioning of network connectivity services with functions for Traffic Engineering (TE), network resource management, and service recovery [RFC4426].

The GMPLS architecture is designed to operate over multiple switching technologies (packet, Layer-2, Time Division Multiplexing –TDM-, fibre and wavelength switching). Extensions in the GMPLS framework, signalling (RSVP-TE [RFC3473]) and routing (OSPF-TE [RFC4203]) protocols to support specific technologies like Wavelength Switched Optical Networks (WSON), G.709 Optical Transport Networks (OTN) and Flexi-Grid Networks are currently under specification ([draft-ietf-ccamp-wson-signaling], [draft-ietf-ccamp-gmpls-g709-framework], [draft-ogrcetal-ccamp-flexi-grid-fwk]). Mechanisms to operate Multi-Layer and Multi-Region Networks (MLN-MRN), comprising multiple data plane switching layers or types under a single instance of GMPLS control plane, are also specified ([RFC4206], [RFC5212]).

The IETF PCE WG defines architectures and protocols for path computation. The **Path Computation Element** (PCE) framework in [RFC4655] identifies two main functional entities: a Path Computation Client (PCC) and a Path Computation Element (PCE). The PCC is the initiator of a path computation request, while the PCE is the entity in charge of computing network paths with a given set of constraints and objective functions. A PCE operates on topology and Traffic Engineering information describing its network domain, while end-to-end inter-domain path computation can be performed through the cooperation of multiple PCEs. The PCE WG specifies different models for inter-PCE cooperation. The Backward-Recursive PCE-based Computation (BRPC) follows a peer-to-peer approach with direct interactions between PCEs located on adjacent domains along the domain path ([RFC5441]). On the other hand, the hierarchical model specified in [RFC6805] introduces the concepts of parent and child PCEs: the parent PCE is in charge of coordinating the end-to-end path computation operating on an abstract view of the inter-domain topology and cooperating with child PCEs responsible for the intra-domain path computation in the different segments. The PCE communication Protocol (PCEP) is the protocol regulating the interaction between PCC and PCE, or between different PCEs. It is initially defined in [RFC5440], and extended in several RFCs and IETF Drafts in support of advanced features (intra-domain confidentiality [RFC5520], Global Concurrent Optimization [RFC5557], path computation in MLN/MRN [RFC5623] and GMPLS [draft-ietf-pce-gmpls-pcep-extensions] networks among the others).

In the latest period, the PCE WG is focusing on extensions for the PCE framework and the PCEP in support of **stateful and active PCEs** ([draft-ietf-pce-stateful-pce]), where the PCE is able to consider not only the network TE information but also the existing Label Switched Paths (LSPs) and, in the active mode, can be delegated to actively operate on the LSPs modifying some of their parameters. These features enables the PCE to efficiently support new functionalities [draft-zhang-pce-stateful-pce-app], like LSPs optimization and re-optimization, auto-bandwidth adjustment [draft-dhody-pce-stateful-pce-auto-bandwidth], bandwidth scheduling, and recovery, that are particularly promising for **intra data centre networks**. Moreover, the stateful active PCE concept can be further extended to allow the PCE to autonomously initiate the creation and deletion of LSPs, as described in [draft-crabbe-pce-pce-initiated-lsp]. The applicability of this approach is relevant in **Software Defined Networking** (SDN) architectures, for example in support of service provisioning triggered from application requirements in very dynamic environments like the internal network of a data centre. Architectures for cooperation between application and network layers are proposed in [draft-farrkingel-pce-abno-architecture], which presents the Application-Based Network Operations (ABNO) model. In this framework, an ABNO controller is responsible of managing the behaviour of the network, according to real-time application requirements and network conditions, and interoperates with a stateful active PCE for network service provisioning and modification.

In the context of data centre networking, extensions to the GMPLS and PCE frameworks have also been investigated in some EU projects in reference to **inter data centre connectivity**. The FP7 GEYSERS project [GEYSERS] has proposed a GMPLS-based Network Control Plane (NCP) architecture for on-demand provisioning of inter DCN services over virtual optical infrastructures. The GEYSERS NCP implements a cloud-to-network interface and internal procedures and protocol extensions to provide connectivity offers and quotations, exchange of information related to resource availability in the data centres at the network edges, and joint selection of IT end-points and network paths. The FP7 CONTENT project [CONTENT] proposes solutions for inter-data-centres and (mobile) user-to-data-centre connectivity over virtual infrastructures where multiple data centres are interconnected through a metro network based on sub-wavelength switching technologies. This metro domain is operated through a GMPLS and PCE-based control plane, derived from the control plane designed in the FP7 MAINS project [MAINS], that is extended to support Time Shared Optical Network (TSON) technologies. While the inter-data-centre connectivity is out of scope for the LIGHTNESS project, it will investigate about possible architectures to interact with the network operators' domains to integrate the inter data centre connectivity. A possible research line may involve the exploitation of the novel interfaces exposed for the provisioning of enhanced network services tailored to cloud service requirements and allowing procedures for cross-stratum optimization between data centres and network domains.

2.3. Software Defined Networking

SDN [onf-sdn-wp] is defined as a control framework that supports programmability of network functions and protocols by decoupling the data plane and the control plane (as depicted in Figure 2.1), which are currently integrated vertically in most network equipments. The separation of control plane and data plane makes the SDN a suitable candidate for an integrated control plane supporting heterogeneous technologies employed in different layers (e.g. optical layer, Ethernet, and IP layer) within data centres. SDN can abstract the heterogeneous technologies employed in the data plane and represent them in a unified way. Proper open standard, vendor and technology agnostic protocols and interfaces are needed for the communications between the SDN controller and the devices in the data plane. OpenFlow (OF) [openflow] is a suitable candidate for the realization of SDN: it is based indeed on flow switching with the capability to execute software/user-defined flow based routing, control and management in the SDN controller which is located outside the data path.

SDN based control plane is an attractive solution for control management of intra data centre network because of the following main features:

- *Abstraction and unification*: SDN abstracts data plane and therefore can be used to build a unified control plane for hierarchy of network technologies i.e. various optical technologies, layer2 and layer 3 within data centre.
- *Application level programmability*: an SDN control plane utilizes a generic network controller that can be used to host any proprietary control and management application.
- *Co-existence*: there are well-defined network protocols and algorithms from which the intra data centre network can benefit, such as PCEP. An SDN based control plane allows rapid and seamless deployment of these protocols as a network application over the SDN controller.

- *Network application level programmable interfacing*: the SDN controller can be equipped with flexible and extendable interfaces to interact with other control and management system. These interfaces can be particularly used by intra data centre network applications to interact dynamically with inter data centre network control and management system.

- *Network function virtualization as well as coordinated virtualization of Network and IT*: SDN allows creation of multiple co-existing and independent network slices (virtual networks) where slicing and traffic isolation is achieved by enforcing routing policy on abstracted data plane switching entity (e.g. flow). This allows dynamic and programmable network virtualization (policy enforcing), which can be controlled by server/IT virtualization mechanisms to achieve coordinated IT + network virtualization. Furthermore, SDN data plane abstraction allows creation of generic network functions (firewall, routing, etc.) that can be virtualized i.e. moved, cloned and deployed over different controllers, controlling different virtual infrastructures.

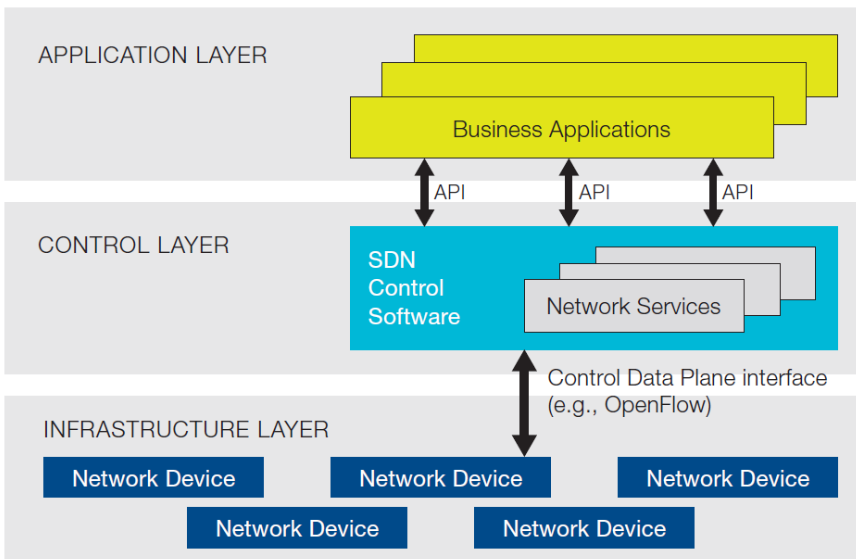

Figure 2.1 SDN split architecture and layering (source: ONF white paper)

In the intra data centre environments, the number of VMs created by different server virtualisation technologies has continuously been increasing over the last decade and it will be doubled in the next few years. These virtual servers together with their virtual access switches need to be tightly integrated with the data centre infrastructure. Furthermore, data centre infrastructures are becoming increasingly multi-tenant, hosting heterogeneous types of tenants by adopting a service model where each tenant is provided with its own virtual data centre infrastructure (VDC or DC as a Cell). In such an environment network virtualization becomes very important and when is combined with IT virtualization allows data centre operators to create multiple co-existing isolated VDC infrastructures for their tenants. The SDN technology allows data centre operators to manipulate logical maps of the network and create multiple co-existing network slices (virtual networks) for each VDC independent of underlying transport technology and network protocols.

Several standardisation bodies are actively working on SDN: *ONF, IRTF, IETF and OMG*.

One of the important industry standards bodies addressing standardization and promotion of SDN is the **Open Networking Foundation (ONF)**. Founded in 2011, the ONF focuses on a number of activities pertaining to an architecture whereby business applications are decoupled from the network control plane and the network control plane is decoupled from the underlying network hardware. This is a departure from the more conventional approach, whereby services and control plane capabilities are tightly coupled with vendor-specific hardware solutions, leading to monolithic systems that are neither open nor programmable at a fine level of granularity, challenging to configure (particularly in multi-vendor environments), and costly to deploy at scale. Central to this approach is the notion of centralized intelligence, where a globalized view of the virtualized network is maintained, and intelligent decisions can be made about policies and flows. From a technical perspective, the ONF organizes a number of working groups to address specific areas of SDN. The currently approved working groups include Architecture & Framework, Configuration & Management, Extensibility (where the OpenFlow protocol is developed), Forwarding Abstractions, Market Education, Testing and Interop, and Optical Transport.

The initial efforts of ONF focused on packet systems within the data centre environment and networking Layers 2 and 3, and through these efforts, the market has seen a number of commercial developments, including commercial and open-source SDN controllers as well as OpenFlow enabled hardware solutions. The interest from the industry around open and programmable networking has now expanded beyond the data centre and beyond Layers 2 and 3. Efforts are now underway to develop standards around Layers 4-7 as well as extending SDN concepts towards the transport domain.

The ONF Optical Transport Working Group was recently formed in 2013, and is focused on the applicability of OpenFlow and SDN to the optical transport layers, inclusive of Layer 0 (optical transmission), Layer 1 (primarily OTN switching/grooming), and Layer 2/2.5 (packet/optical integrated technologies). The current focus of the working group is on defining a common set of use-case scenarios, both enterprise and service-provider models. As the transport layer is predominantly “circuit” based (as opposed to connectionless packet based) when looking at data plane connectivity, the working group is examining both a direct model, whereby the individual transport systems are directly programmed, as well as an abstract model, where an abstraction of the optical transport network is supported. The three key use cases being addressed include private enterprise optical transport, network virtualization and data-centre interconnection for service providers, and packet/optical integration. These latter two use cases have the most relevance to the LIGHTNESS project, and include key concepts for further research, including:

- Virtualization of the transport network enables multi-tenancy with dedicated partitions, which can be an important tool for facilitating server-to-server or application-to-application connectivity
- Dynamically programmable virtual network topologies between applications within data centres over the optical transport network, when the data centre and network are owned by the same provider as well as when the data centre and network are owned by different entities.

- Orchestration of both data centre resources along with Wide Area Network (WAN) transport bandwidth
- Multi-layer integration and optimization of packet/optical solutions for greater network efficiency, including networks that employ converged transport technologies, combining Layers 0, 1, and 2/2.5

At the Internet Research Task Force **(IRTF) and IETF** several working groups are active in the areas of SDN, including the SDN Research Group within IRTF [IRTF-SDN]. General network programmability approaches in IETF have been considered for a long time by the FORCES WG [FORCES], and recently by the I2RS WG [I2RS]. The NVO3 WG addresses the association of heterogeneous network and data centre infrastructures under common virtual data centre containers, as well as multi-tenancy large-scale DCNs and how these issues can be addressed using an overlay-based network virtualization approach. Network modelling and configuration are considered within the successful YANG model [YANG], and applied in NETMOD WG and NETCONF WG [netconf], while the ALTO protocol [ALTO] and its extensions allow defining infrastructure abstractions to be exposed for networked applications.

Finally, the Object Management Group (**OMG)** has recently created an SDN WG (March 2013) [SDN-OMG] with the aim of developing a draft of recommended Northbound APIs and services topology based on an Service Oriented Architecture (SOA) vision of the SDN space, as well as draft two or more case studies relevant to the SDN space to help drive standards requirements and priorities.

3. LIGHTNESS control plane requirements

This chapter describes the requirements for the LIGHTNESS control plane, mainly providing a summary of the high-level needs for data centre services and applications that have impact on the DCN control procedures and mechanisms. The detailed identification, analysis and specification of both functional and non-functional control plane requirements is carried out in the deliverable D2.2 [del-d22]: therefore the aim of this chapter is to summarize them and lay the foundation for the LIGHTNESS control plane high-level architecture specification. The main functionalities offered by the LIGHTNESS control plane that satisfy such requirements are also described. In addition, the positioning of the control plane in the overall LIGHTNESS architecture is provided, through a description and analysis of its main roles, responsibilities and interactions with external entities.

3.1. Positioning in the LIGHTNESS architecture

The aim of this section is to highlight the role of the control plane in the LIGHTNESS intra data centre scenario, as depicted in Figure 3.1. To replace the traditional multi-tier DCNs, LIGHTNESS implements an ultra-high capacity and highly-scalable network, based on combining optical circuit and packet switching technologies. The final aim is to provide high-bandwidth, low latency and energy efficient connectivity services among servers/Top of the Rack (TOR) switches.

In current data centres, connectivity services are provisioned through static and semi-automated procedures; this means that the bandwidth provisioning process is quite slow and the utilisation of the DCN resources is not optimal. As an example, for fault tolerance reasons typically both primary and backup paths between servers are provisioned. This implies that data centre operators must overprovision the DCN, resulting in a waste of resources with very high cost in both equipment and power consumption. In contrast, current data centre workloads highly fluctuate on time and the introduction of virtualization techniques (e.g., VMWare) enables hundreds of VMs in each multi-core CPU server.

Data centre operators may need to perform dynamic VMs migration for multiple reasons, including load-balancing, optimised usage of servers, energy-efficiency, maintenance, hardware and software failures, etc. This results in additional complexity to allow the efficient management of the data centres resources.

Future data centres also require to implement continuous monitoring of the data centre resources in order to gather performance and network usage statistics for Service Level Agreement (SLA) monitoring as well as re-planning and optimization functions.

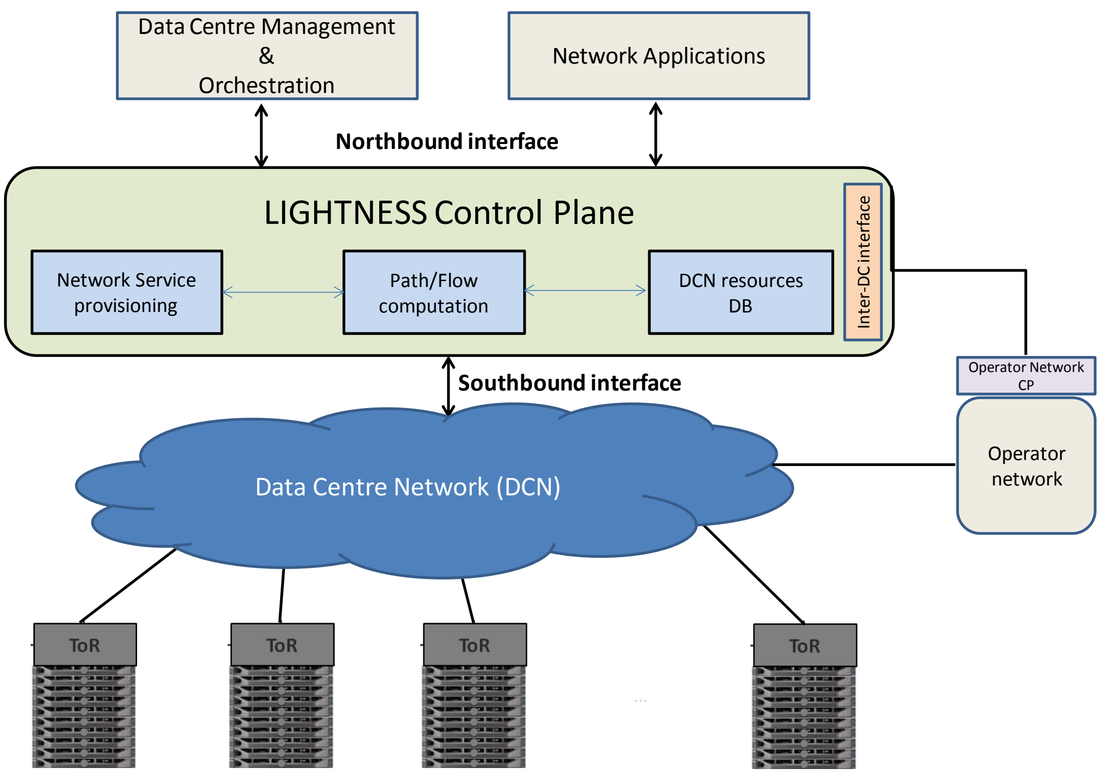

Figure 3.1: High-level LIGHTNESS intra-data centre scenario

In such a context, the introduction of a control plane allows to implement and automate a range of control functions (e.g., flow computation, traffic engineering, VM migration, service and network monitoring, recovery, network optimization). It extends functionalities offered by current data centre management frameworks to support on-demand provisioning of connectivity services between the different entities involved in the data centre distributed system.

However, the design of a control plane for data centre environments need to consider some challenges:

1) it must simplify the management of the applications running inside data centres;
2) it must be able to provision connectivity services quickly in order to meet the SLA requirements of the intra data centre dynamic traffic patterns

3) it must handle failure conditions that may occur, managing alarm notifications coming from the underlying data plane entities (i.e. physical devices) and recovering the involved connectivity services;
4) it should be scalable; its operation should not be affected by the dimension of the data centre itself, in terms of number of servers and data plane entities controlled;
5) it should be extensible; it must have external interfaces that allow to be easily extended to cope with additional features and technologies;
6) it must enable isolation of the logical resources, secure accessibility and traffic segregation for multi-tenancy;
7) it must expose network resource information to both data centre management (for orchestration purposes) and applications;
8) it should not impact the performance (e.g., latency) of the DCN.

3.2. Control plane requirements

The LIGHTNESS network control plane operates in a data centre environment, and therefore is responsible for implementing mechanisms and procedures for the provisioning of enhanced network connectivity services for data centre applications. The associated control plane requirements depend and are strictly related to the application and cloud service requirements and characteristics, e.g. in terms of QoS attributes, service dynamicity and flexibility, etc. In addition, the LIGHTNESS control plane operates over the hybrid OCS/OPS DCN and thus needs to implement mechanisms and procedures for configuration and monitoring of optical network resources in the data plane.

The full set of LIGHTNESS control plane requirements is defined and deeply analyzed in deliverable D2.2 [del-d22]. This section provides a summary of that specification to let this document be self contained also in terms of requirements.

The LIGHTNESS network control plane requirements are listed and analysed in Table 3.1. Each requirement has a severity level assigned:

- *Essential* requirements are related to mandatory functionalities to be supported by the LIGHTNESS system
- *Preferred* requirements are related to those functions representing an added value; when not met, the LIGHTNESS system provides a reduced set of functionalities.
- *Nice to have* requirements are used for optional requirements which does not affect the essential and fundamental operations of the LIGHTNESS system when they are not met.

Requirement Name	Description and Severity
Dynamic on-demand network connectivity	The deployment of a network control plane for automated provisioning of DCN resources is a key innovation in the data centre environment. The LIGHTNESS control plane must implement dedicated mechanisms to allocate network resources inside the data centre dynamically and on-demand, providing the QoS guarantees required by the applications running in the servers. Unidirectional and bidirectional point-to-point and point-to-multi-point connectivity services must be provided. Unidirectional services are intended to serve asymmetric traffic between servers inside the data centre (e.g. big data transfer for replication purposes). Bidirectional services are conceived for symmetric traffic between servers. This requirement is: ***Essential***
Support and integration of multiple optical switching technologies	The integrated and converged control of multiple optical switching technologies allows data centre and cloud operators to operate their network resources efficiently on their infrastructures, fully exploiting the peculiar characteristics of different switching technologies. The LIGHTNESS control plane must support the hybrid OCS/OPS DCN and implement procedures to integrate different optical switching technologies in the connectivity services. This requirement is: ***Essential***
Scheduled connectivity services	The support of advance resource reservations allows data centre and cloud operators to enable an efficient control and sharing of network resources inside data centres. Since it represents an added value for the LIGHTNESS data centre infrastructure, the LIGHTNESS control plane must support the scheduling of connectivity service provisioning inside the data centre and enable the specification of the time constraints like service start, stop time and duration. This requirement is: ***Essential***
On-demand network connectivity modification	The on-demand modification of connectivity services active and installed inside the data centre enables the LIGHTNESS infrastructure to cope with the dynamicity and flexibility of distributed data centre applications. In particular, the LIGHTNESS control plane should allow dynamic modifications of pre-established network services for their adaptation

	according to changing needs of data centre and cloud applications, e.g. in terms of bandwidth, QoS guarantees, recovery strategy, etc. This requirement is: ***Preferred***
Optimization of resource utilization	The deployment of a network control plane in a data centre infrastructure enables the implementation of specialized routing procedures and computation algoritms for resource usage optimization purposes. This provides an innovative functionality with respect to the current DCN management frameworks. Therefore, the LIGHTNESS control plane must provide path computation functions to implement enhanced optimization of network resource usage inside the data centre. This requirement is: ***Essential***
Dynamic re-optimization of network services	The automated re-optimization of network resource provisioning allows data centre and cloud operators to limit the implementation of rigid and human driven monitoring and modification actions on their infrastructures. In this context, the LIGHTNESS control plane should provide automated procedures for dynamic re-configuration of network services and autonomous re-optimization of DCN resources utilization. This requirement is: ***Nice to have***
Control plane scalability	The control plane scalability is a key requirement for any network deployment and infrastructure. In a data centre environment the extremely high number of data flows and connectivity services to be supported make the scalability a fundamental requirement. Therefore, the LIGHTNESS control plane operation and procedures must scale in terms of number of controlled devices and traffic load in the DCN, so that the network service provisioning and recovery performances are not affected by the dimension of network and traffic. This requirement is: ***Essential***
Network connectivity service recovery	The automated recovery of DCN services upon failure conditions allows to reduce the impact of service interruption for data centre and cloud applications. In this context, the LIGHTNESS control plane must provide automated procedures to detect and recover failure conditions that affect established network services. Specialized protection and recovery strategies for data centre environments will therefore be implemented in the LIGHTNESS control plane. This requirement is: ***Essential***

Monitoring	The escalation of monitoring information allows data centre management and orchestration entities to be aware of the actual status and performances of the data centre infrastructure. The LIGHTNESS control plane should therefore provide mechanisms to gather updated monitoring information about the DCN behaviour and let data centre management and orchestration entities to modify and update the network services provisioning accordingly. This requirement is: ***Preferred***
Inter data centre connectivity services	A significant part of the network traffic inside data centres (about 10%, according to recent statistics from Cisco [cisco-gci]) refers to applications running among multiple geographically distributed data centres, e.g. to support scheduled data migrations or databases synchronization for replication purposes. Therefore the implementation of inter data centre connectivity services allows a DCN control plane to broaden its scope to a wider range of applications. In this context, the LIGHTNESS control plane must support procedures and mechanisms to stitch intra data centre services to inter data centre ones, enabling the communication among geographically distributed data centres. This requirement is: ***Essential***

Table 3.1 LIGHTNESS control plane requirements

3.3. Overview of control plane functionalities

The network control plane is a key component in the overall network architecture proposed by LIGHTNESS: it is responsible for the on-demand, dynamic and flexible provisioning of network services in the hybrid OCS/OPS data centre infrastructure. This section presents the network control functionalities supported by the LIGHTNESS control plane with the aim of satisfying the requirements identified in section 3.2. The detailed specification of the LIGHTNESS control plane architecture, in terms of functional decomposition and procedures, is carried out in the next chapters.

The main innovative functionalities provided by the LIGHTNESS control plane can be summarized as follows:

- On-demand and dynamic provisioning of DCN resources
- Optimization of DCN resource utilization
- Dynamic re-configuration of DCN services

These functionalities are briefly described in Table 3.2.

<table>
<tr><th>Functionality</th><th>Description</th></tr>
<tr><td>Automated setup and tear down of network services</td><td>The LIGHTNESS control plane offers automated and on-demand provisioning of network services within the DCN. Service establishment can be either immediate or scheduled: the LIGHTESS control plane will therefore maintain a dedicated calendar where scheduled reservations are stored and processed in terms of start date and planned duration.

Each network service is provisioned to satisfy a given set of QoS constraints (e.g. minimum bandwidth, average bandwidth, maximum latency, packet loss, and jitter). The following types of services are supported:
<u>Point-to-point (P2P)</u>: for connections between two end-points (e.g. between two servers inside the data centre)
<u>Point-to-multi-point (P2MP)</u>: connections between a single source end-point and multiple destination end-points (e.g. for data replication purposes)
The automated setup and tear down of network services also includes the inter data centre connectivity, with the aim of supporting data centre applications or services running in geographically distributed data centres. This kind of services spans across multiple data centres and refers to joint and coordinated network resource provisioning in different network segments/domains:
Network resource provisioning in source DCN
Network resource provisioning in the network(s) interconnecting the involved data centres
Network resource provisioning in destination data centre(s)
In this inter data centre scenario, the LIGHTNESS control plane is responsible only for the intra data centre resource provisioning.</td></tr>
<tr><td>Dynamic modification of network services</td><td>The LIGHTNESS control plane offers the possibility of modifying already established network services. The service modification may be required for different reasons and triggered either by a change of needs of data centre applications or by data centre management or orchestration actions. It is not an automated procedure implemented by the LIGHTNESS control plane: it is decided, generated and triggered by the upper layers. The following type of</td></tr>
</table>

	modifications are supported: • <u>QoS constraints</u>: this type of modification can refer to either an upgrade or downgrade of the installed service, e.g. in terms of bandwidth, latency, etc. According to the specific modification request, a re-computation and re-configuration of the entire service may be implemented by the control plane; • <u>Recovery strategy: </u>modification of the recovery mechanism associated to the given network service. For instance, an unprotected network service may be upgraded to be protected by a new network service to be installed, or to be restored on-the-fly with a dynamic failure recovery procedure • <u>Destination in P2MP network services</u>: modification of destination end-points due to changes in data centre applications requirements. The deletion, addition or modification of an existing destination may be requested, • <u>Monitoring mechanisms</u>: modification of the current mechanisms in place, either in terms of type (e.g. synchronous vs asynchronous) or frequency, for exchange of monitoring information for the given network service
Data centre network resiliency	The LIGHTNESS control plane implements automated mechanisms and procedures to recover from failure conditions in the hybrid optical DCN and re-establish all the involved network services. The detection of the network failures in the DCN, that is carried out in cooperation with underlying data plane entities, dynamically triggers the restoration actions. Several recovery strategies are offered by the LIGHTNESS control plane for a given network service: • <u>Unprotected</u>: when a network failure is not automatically recovered at all; • <u>Protected</u>: when one or more network services are available (i.e. already established and configured, or at least planned and booked in terms of resources) for usage in case of network failure • <u>On-the-fly restoration</u>: when upon a network failure a disjoint network path is dynamically computed and established. No additional resources are therefore configured or booked in advance

	In case of impossible recovery of the network service (e.g. when alternative network paths are not available), the failures are notified to the upper-layer data centre management or orchestration, so that recovery strategies can be applied at the service layer.
Optimization of data centre network resources	The LIGHTNESS control plane implements automated procedures to optimize the resource usage inside the data centre and accordingly re-organize the traffic generated by the network services. These optimization procedures rely on the monitoring functions of the control plane, which offer the possibility of processing and combining real-time and predicted DCN resource usage and performances. The final objective of the DCN optimization function is to dynamically optimize the resource utilization in the whole data centre according to the characteristics, performances and even status of running or expected services. This function is completely automated in the LIGHTNESS control plane and does not depend on any external trigger, e.g. coming from the data centre management or orchestration entities.
Data centre network monitoring	The LIGHTNESS control plane offers and supports a set of monitoring functions at different levels. First of all, it allows the upper layers (i.e. DCN management and orchestration entities, or even network applications) to collect monitoring information about status and performance of physical resources and established network services inside the data centre. This enables the network services modification described above. In addition, the control plane also implements internal monitoring functions to support both restoration and optimization automated procedures. The generic monitoring information is exchanged in terms of failures, performance and status of both established network services and physical network resources (i.e. not yet associated to any given service).

Table 3.2 Overview of LIGHTNESS control plane functionalities

4. LIGHTNESS control plane high-level architecture

The main rationale behind the deployment of a network control plane in a data centre environment is the need to overcome the current limitations of the existing DCN management and control frameworks. Indeed, they mainly operate as stand-alone overlaid infrastructures with semi-automated and static procedures to provision connectivity for the data centre IT services and applications. The LIGHTNESS network control plane is conceived to implement automated procedures for setup, monitoring, recovery and optimization of network connections spanning multiple optical technologies and matching QoS requirements for data centre services and applications. The LIGHTNESS control plane operates on top of the hybrid OCS/OPS DCN using the Southbound interface (that will be specified in deliverable D4.3 [del-d43]) in order to control and provision the optical network resources. As a consequence, the LIGHTNESS control plane is decoupled from the underlying DCN data plane equipment. It will include mechanisms to retrieve the status and the availability of the resources, and to receive notifications and dynamically configure the specific DCN nodes. In this context, the LIGHTNESS control plane aims at providing dynamic and flexible procedures to provision and re-configure DCN resources, by implementing optimization functions according to performances and network usage statistics gathered from the underlying data plane equipments.

As a fundamental concept, the LIGHTNESS control plane integrates functionalities offered by current DCN management frameworks to support on-demand provisioning of connectivity services, while reducing as much as possible human actions and validations with automated procedures. Dynamicity and flexibility in the control of the DCN also means that the LIGHTNESS control plane provides enhanced functionalities for combination and orchestration of short non-persistent data flows (suited for optical packet switching) with long persistent ones (more suited for optical circuit switching), according to specific data centre applications and services requirements and data centre operator policies.

This section introduces the control plane high-level architecture by describing its fundamental concepts and main functionalities. In particular, the adopted SDN-based approach and its functional decomposition are described. Moreover an high-level description of the interfaces exposed towards other building blocks of the LIGHTNESS architecture is provided, also considering the support of inter data centre connectivity services.

4.1. SDN-based control plane functional decomposition

The LIGHTNESS control plane aims at fulfilling the requirements defined in deliverable D2.2 [del-d22] and summarized in section 3.2, mainly in terms of dynamic and flexible control of network resources inside the DCN for data centre operators and providers. The identification and definition of the LIGHTNESS control plane high-level architecture have been carried out as a joint activity between WP2 and WP4, in particular Task 2.1 and Task 4.1. A set of control plane approaches and potential solutions have been proposed and analyzed in the context of the WP2 control plane architecture task force: these approaches are detailed and compared in the deliverable D2.2 [del-d22], in terms of qualitative analysis of pros and cons for each solution. A summary of that evaluation, providing the comparison between a distributed GMPLS based approach and a centralized SDN (with OpenFlow support) approach, is reported in Table 4.1 for sake of completeness of this document.

	Functionalities									
Control Plane approaches	*NB interface support*	*Inter data centre provisioning*	*OCS support*	*OPS support*	*Data plane abstraction*	*Network control functions programmability*	*NB API availability*	*Service provisioning/ reconfiguration time*	*Network slicing support*	*Service and network monitoring*
GMPLS/PCE	No	Yes [1]	Yes	No	No	No	No	High	No	No
SDN	Yes	To be designed	Yes [2]	No	Yes [3]	Yes	To be designed	Low	Yes [4]	Yes

[1] H-PCE architecture applied for the interconnection among geographically distributed data centres through GMPLS-based operator (backbone) networks

[2] OF protocol extensions for OCS networks are available

[3] e.g., for OF switches, FlowVisor performs abstraction

[4] Particularly important for multi-tenancy support in the data centre

Table 4.1 LIGHTNESS control plane potential approaches: summary

As specified in D2.2 [del-d22], a Software Defined Networking (SDN) based solution has been selected as the most suitable approach for the LIGHTNESS control plane architecture. Indeed, conventional distributed control frameworks are not well suited for supporting on-demand and dynamic data centre services. They normally implement complex time consuming control functions and procedures for signalling, routing, provisioning and recovery of network services that do not meet the agility and flexibility requirements of a data centre environment. On the other hand, SDN [onf-sdn-wp] is emerging as a extensible and programmable open way of operating DCN resources. Its main concept is the decoupling of forwarding and control functions, centralizing network intelligence and state information, while providing to the upper layers an abstracted and vendor independent view of network resources through open Application Programming Interfaces (APIs). In other words, SDN allows data centre and cloud providers to build more scalable, agile and easily manageable DCNs. It is conceived to provide a software abstraction of the physical network that allows the network itself to be programmable and therefore closely tied to the data centre applications and services needs.

In addition, the SDN framework and concepts perfectly fit with the novel hybrid OCS/OPS data centre fabric proposed by LIGHTNESS, enabling a valuable solution for handling the complex unpredictable data flows in the data centres. Indeed, the programmability and flexibility provided by an SDN based control plane approach allow to fully exploit the benefits of the LIGHTNESS multi-technology DCN in accommodating both short and long lived traffic flows of data centre applications and services.

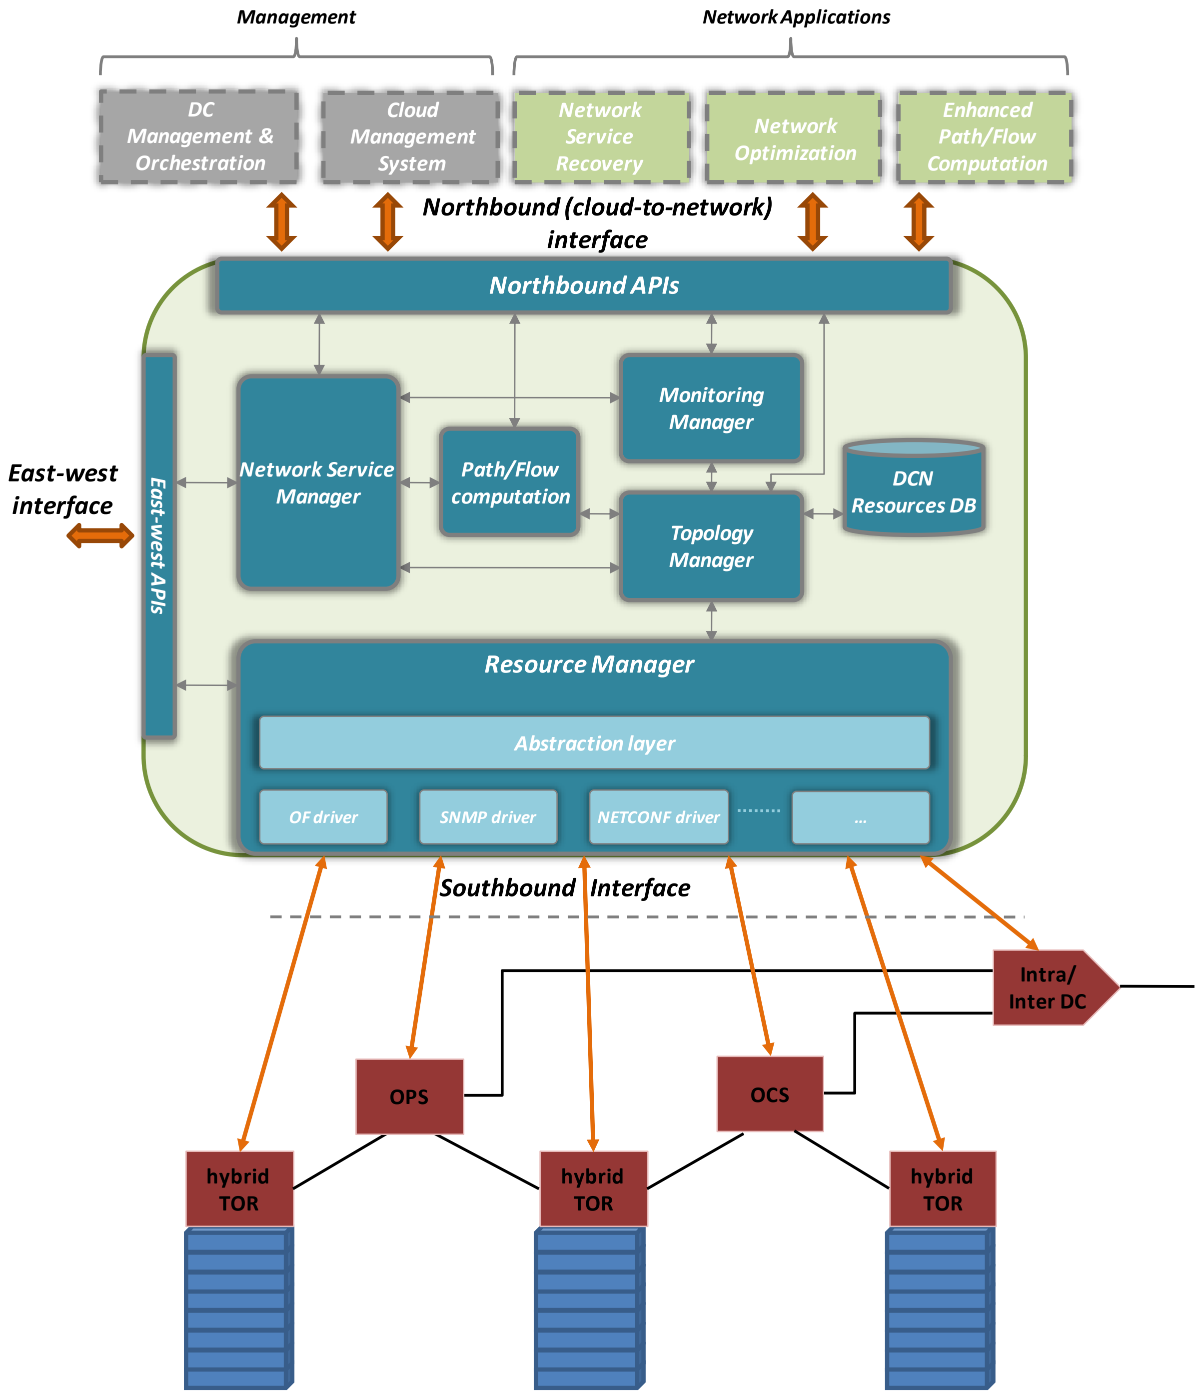

Figure 4.1 LIGHTNESS control plane high-level architecture

The LIGHTNESS control plane functional architecture is depicted in Figure 4.1, that highlights its main functional entities and their interactions. It is based on the SDN architecture proposed by ONF [onf-sdn-wp], which represents the reference for the LIGHTNESS control plane architecture, properly extended in terms of functional decomposition to cope with the specific data centre services and applications requirements defined in the deliverable D2.2 [del-d22]. This means that the control intelligence is logically centralized in an SDN controller (and/or network applications running on top of it) that maintains a full and detailed view of the network.

As shown in Figure 4.1, this LIGHTNESS control plane high-level architecture is composed by an SDN controller natively implementing a restricted set of control plane functionalities that allow to meet the essential requirements for data centre services and applications (i.e. in terms of provisioning, topology discovery and path computation). The enhanced and novel LIGHTNESS DCN control functions are moved outside the SDN controller to let them be implemented and ran as dedicated network applications. Such functional split improves the flexibility of the LIGHTNESS control plane architecture since it provides a generalized control framework, that can be potentially adopted in any data centre environment (in terms of technologies and network devices). This is completely aligned with a fundamental SDN concept [onf-sdn-wp]: to have a simplified, vendor independent and programmable DCN control and operation. It is important to highlight that such a generalized approach is also enabled by the adoption and support (and extension where needed) of open and standard interfaces at the southbound reference point, like OpenFlow [openflow]. In this context, the LIGHTNESS control plane architecture allows data centre operators to implement their own customized network functions and applications according to their specific requirements and needs.

In Figure 4.1 the SDN controller external interfaces are also depicted: the Northbound interface, the Southbound interface, and the East-West interface. The former implements the LIGHTNESS Cloud-to-Network interface, and provides the data centre operator with the management mechanisms and procedures to provision on-demand and flexible network services for its data centre and cloud applications. This interface, implemented in the form of APIs, also enables the implementation of additional and enhanced functions as network applications on top of the SDN controller (e.g. for routing, security, bandwidth management, quality of service, etc.). Moreover, the LIGHTNESS SDN controller could be integrated with Cloud Management Systems, like OpenStack [openstack] and CloudStack [cloudstack] by implementing dedicated plugins (e.g. in the OpenStack Neutron framework [openstack-neutron]) for the Northbound APIs. On the other end, the Southbound interface allows the SDN controller to interact with the underlying DCN network devices, and provides mechanisms and procedures to discover, configure, monitor the network resources. Finally, the East-West interface enables the scalability of the LIGHTNESS control plane: it is conceived to support the cooperation among multiple SDN controllers possibly deployed to control large-scale DCNs organized in clusters. In this scenario, each SDN controller is responsible for the control operations (i.e. network service provisioning, monitoring, recovery, optimization, etc) of a single cluster (or a limited set of clusters), and the East-West interface is used for the information exchange among controllers (e.g. cluster network resource capabilities) and to provision inter-cluster DCN services. More details about the LIGHTNESS control plane external interfaces are provided in section 4.3.

Table 4.2 lists the SDN controller functional modules and for each of them it provides a high-level description of its main functionalities, roles and responsibilities.

SDN Controller Functional Module	Description
Network Service Manager	It implements the SDN controller network service intelligence by handling the on-demand provisioning of connectivity services in the controlled DCN. It uses an abstracted view of the DCN resources provided by the Topology Manager and stored in the DCN Resource Database, and configures the DCN physical devices through the Resource Manager. It enables, through the Northbound APIs, applications running on top of the SDN controller to perform enhanced connectivity functions and gather real time monitoring and network performances information for installed services (as provided by the Monitoring Manager).
Resource Manager	It is the access point to the underlying DCN resources for Network Service Manager and Topology Manager usage, allowing them to synchronize, configure and gather notifications (alarms, failure conditions, performances, etc.) about network resources. The Resource Manager implements the SDN controller southbound interface and includes a specific function for abstracting the vendor-dependent technology details of the DCN resources (i.e. the Abstraction Layer depicted in Figure 4.1). In accordance with the SDN layering and functions proposed in IETF [sdnrg-layer], the Resource Manager is capable to support multiple control and management protocols at the southbound interface through specific control drivers. Potential control and management protocols to be supported are: Openflow [openflow], SNMP [snmp], NETCONF [netconf]. This allows the SDN controller, along with its native control functions and related applications, to be generic and independent from the specific protocols adopted by the potentially different devices in the DCN. In particular the Resource Manager is capable to maintain the technology specific resource granularities of OPS, OCS and TOR switches while providing a uniform resource description.

Topology Manager	It is responsible for maintaining an updated topology view of the underlying network resources in the DCN Resource Database. It is directly interfaced with the Resource Manager for topology discovery and learning purposes. On the other hand it exposes the topology view to the Path/Flow Computation and Network Service Manager for connectivity provisioning purposes. It also provides topology abstraction functions through the Northbound APIs for the network applications running on top of the SDN controller.
Monitoring Manager	It is responsible for handling the notifications about status and performances of the network resources and services in the DCN (e.g. alarms, failures, number of dropped packets, congestion status, etc.). It is conceived to interact with the Resource Manager to gather such information and dispatch them to the proper SDN controller function. It also provides and exports, through the Northbound APIs, the network monitoring information to the network applications running on top of the SDN controllers. These information, e.g. related to failure conditions or performances, enables the support of dedicated applications for network optimization and service recovery.
Path/Flow Computation	This module provides some native routing functionalities for the SDN Controller. In particular it is conceived to implement basic path and flow algorithms in support of the connectivity services to be handled by the Network Service Manager, e.g. such as Dijkstra [dijkstra], All Pairs Shortest Path [floyd-warshall], etc.. It operates on the network topology exposed by the Topology Manager. More sophisticated and enhanced path/flow computation functions could be implemented as network applications on top of the SDN controller (see Table 4.3 and section 5.2).

Table 4.2 LIGHTNESS control plane functional modules description

4.1.1. Enhanced control plane functionalities

By taking advantage of the SDN approach in the LIGHTNESS control plane design, the features and functionalities supported by the control plane can be extended and enhanced during and also after the project lifetime. These features and functionalities can be implemented as applications running on top of the SDN controller. Some general data centre features can be supported, e.g. load balancing, automatic optimization, adaptive and dynamic replanning, data centre virtualization. Considering the unique scenario proposed by LIGHTNESS, that is bringing both OPS and OCS switches into the data centre, the intelligent OPS/OCS handover is also an attractive feature to have. Load balancing can be better supported using the SDN solution, since the controller has a global view of all the underlying available resources (a database containing the information of network resources and network topology, which can be dynamically updated). Based on this view, the traffic load can be routed over the less utilized resources and hence achieve the overall load balancing of the entire data centre. In addition, by leveraging the global database and the communications between the controller and the data plane devices enabled by the protocols, the automatic optimization of the overall resource allocation can be achieved thereby the optimal end-to-end path setup and data delivery. DCN virtualization is another enhanced feature that can be supported by the SDN solution. With this feature, multiple virtual DCN slices can be created and provided to different users to run various applications, which can support the multi-tenant requirement of data centre and also achieve high network utilization ratio.

The SDN controller functional modules described above are conceived to provide a basic set of functionalities for on-demand network connectivity services in the LIGHTNESS hybrid optical DCN. A preliminary set of enhanced network functions candidate to be implemented as applications running on top of the SDN controller has been identified: they aims to fulfil the LIGHTNESS control plane requirements reported in section 3.2, and they are listed in Table 4.3.

SDN Controller Applications	Description
Network Service Recovery	This network application implements recovery functions and mechanisms to let the network services in the DCN be resilient. It interacts, through the Northbound APIs, with the Monitoring Manager to collect and process notifications related to network alarms and failures that could occur for resources involved in established network services. It is responsible for fast re-configuring the failed network services (e.g. establishing new connectivity through the interaction with the Network Service Manager).

Network optimization	This network application provides dynamicity and flexibility in the control of the DCN resources. By interacting with the Monitoring Manager, it gathers real-time information about network performances (latencies, jitter, blocking conditions, loss of packets) and it is responsible to dynamically and automatically modify installed and scheduled network services (e.g. tuning bandwidth, changing data flow path, etc.) with the aim of utilizing the DCN resources efficiently. Multiple network applications may be implemented to support different types of optimization. For instance, a dedicated LIGHTNESS “Technology handover” network application may be provided to dynamically handle the handover of network services between the OPS and OCS technologies in the DCN, e.g. to move long lived data flows initially allocated in the OPS region to be accommodated by the OCS technology.
Enhanced Path/Flow computation	This network application provides enhanced path and flow computation functions, and it is conceived to complement the basic routing functions implemented inside the SDN controller. Either complex computation algorithms or enhanced routing procedures may be implemented here. For instance, the IETF PCE framework [RFC4655] may be adapted to be used as external path and flow computation engine capable to provide stateful and active routing functions [draft-ietf-pce-stateful-pce]. Further details for the LIGHTNESS path and flow computation can be found in section 5.2.

Table 4.3 SDN Controller applications

4.2. Control plane external interfaces

This section provides an overall description of main functionalities supported at the control plane interfaces: Northbound, Southbound and East-West. The detailed specification of these interfaces is currently ongoing in the context of Task 4.3 activities: therefore this section is intended to briefly describe their roles and high-level functions. The full specification of the control plane external interfaces, in terms of procedures, semantics and information models, will be included in deliverable D4.3 [del-d43].

4.2.1. Northbound (Cloud-to-Network) Interface

The Northbound (Cloud-to-Network) Interface is conceived for provisioning dynamic, on-demand and resilient connectivity services in the DCN. It is used at the LIGHTNESS control plane to interact with the various data centre applications, such as coordinated virtualization of IT (i.e. computing and storage) and network resources. Cloud Management such as OpenStack [openstack], and its network module Neutron [openstack-neutron], can also be taken as a plugin to the Northbound interface of the LIGHTNESS control plane. The requests from management and orchestration entities will be sent to the LIGHTNESS control plane through this Northbound interface, specifying the requirements such as the required resources and/or SLA/QoS, etc. In the other direction, highly dynamic, on-demand and resilient DCN services will be provisioned through this interface to the applications. Monitoring functions will also be supported for the mechanisms such as restoration.

Table 4.4 summarizes the functionalities offered by the LIGHTNESS control plane at the Northbound interface.

Functionality	Description
Network Connectivity Service Provisioning	Through the LIGHTNESS Northbound Interface, management and orchestration entities can send their requirements of resources (bandwidth, computing, storage, etc.) and network performance (latency, etc.) to the control plane, and request dynamic, on-demand and resilient network connectivity services.
Data Centre Management	Cloud management systems (e.g. OpenStack) can cooperate with the control plane through proper plugins implementing the LIGHTNESS Northbound Interface, that becomes therefore responsible for operating services and managing resources and data inside data centre. It will ensure the optimal resource utilization and service provisioning.
Monitoring	In order to implement restoration mechanisms and provide resilient connectivity services, performance monitoring (response time, uptime, etc.) also needs to be supported through the Northbound interface. By monitoring, the disaster recovery will be overseen and contingency plans can be initiated, and the practical and up-to-date decisions can be properly made

	regarding of service provisioning.

Table 4.4 Northbound interface main functionalities

4.2.2. Southbound Interface

The Southbound Interface is adopted by the LIGHTNESS control plane to interact with the underlying heterogeneous switches/devices (e.g. OPS, OCS, hybrid TOR) in the data plane. The communications with the gateway towards the core network (i.e. intra-to-intra data centre interface) is also supported by the Southbound interface. Suitable protocols will be chosen, and accordingly extended for the LIGHTNESS purposes, for the communications between the control plane and data plane, such as OpenFlow [openflow]. Messages for realizing different purposes (e.g. monitoring, configuration and modification, collection of features/attributes and statistics) will be exchanged between the control plane and the switches in the data plane through the Southbound interface. According to the messages received, the data plane elements will be configured accordingly.

Table 4.5 provides a brief description of the functionalities supported at the Southbound interface.

Functionality	Description
Collection of attributes	Physical devices have diverse attributes, including both basic (number of ports, number of channels per port, power consumption, etc.) and advanced (nonlinearity, etc.) ones. These attributes/features of heterogeneous devices need to be collected and exposed to the control plane with some levels of abstraction to construct the resource database. In that case, Southbound interface is acting as a uniform interface responsible for the communications between the control plane and the data plane.
Configuration	Through the Southbound interface, the underlying physical devices are configured using the configuration messages of the adopted protocol (e.g. OpenFlow). These messages will include the information for configuring the devices and/or setting up the switches, such as the ports' no., the channels' no., etc.
Collection of statistics	There are some useful statistics (e.g. flow statistics, port statistics,

	queue statistics, etc.), which will be crucial for monitoring the traffic and the utilization of the devices. One of the statistics, counter, could be used for the handover between OPS and OCS in LIGHTNESS.
Monitoring	Monitoring is a key functionality that can guarantee the control plane having the up-to-date information of the data plane, including the network resources and network topology. These information will be used to construct the database which is essential for allocating resources and setting up connectivity.

Table 4.5 Southbound interface main functionalities

4.2.3. East-West Interface

In LIGHTNESS, the East-West interface is defined to address the scalability requirements imposed by the data centre environments. Indeed, the huge number of data flows and services to be supported inside data centres make the scalability a key concept: as a consequence, the LIGHTNESS control plane need to scale in terms of controlled devices and (potentially concurrent) network services to be installed in the DCN, with the aim of providing reliable connectivity services not affected by the dimenstion of the data centre itself. To this purpose, the East-West interface has the purpose to support the control of large-scale data centre organized in clusters of servers by enabling the interconnection of multiple SDN controllers, each one responsible for one cluster or a few clusters.

In such a scenario, the East-West interface allows the SDN controllers deployed on top of each cluster to cooperate and be properly interconnected to either exchange network resource capabilities or provision possible inter-cluster network services. At the time of writing, the East-West interface is conceived to enable peer-to-peer interactions among a set of SDN controllers deployed on top of the DCN infrastructure. Further analysis and investigation will be carried out in the context of Task 4.3 and the detailed specification of this interface will be part of deliverable D4.3 [del-d43].

Table 4.6 gives a summary about main functionalities supported at the East-West interface.

Functionality	Description
Cluster controller interconnection	In a large-scale data centre, the East-West interface could be used to interconnect multiple SDN controllers deployed on top of each cluster. When multiple controllers are used, the East-West

	interface enables a peer-to-peer interconnection among them, with each controller responsible for provisioning resources, recovery and monitoring actions in the cluster under its control.
Cluster information exchange	The East-West interface is responsible for the information exchanges among the multiple clusters, such as the network resources information exchanges and updates, in order to get the global convergence in terms of resource capabilities (but still within one data centre) to facilitate the inter-cluster communications. It also enables the provisioning of inter-cluster network connectivity services.

Table 4.6 East-West interface main functionalities

4.3. Support of inter data centre connectivity

Data centre and cloud service providers typically emplace data centres in geographically distributed sites for high service reliability: in this way, they are able to replicate contents among them periodically, offering enhanced service availability upon failures. In addition, they can deploy and offer services and applications as closest to the user as possible, in particular when the users are spread in different geographical areas. Therefore, inter data centre workload migrations are performed for efficiency reasons and increase data flows among data centres up to daily totals of Terabytes or Petabytes.

In a recent study from Cisco [cisco-gci], an incredible growth of data centre traffic has been observed in the last five years due to the global increase of cloud services demand. Data centre traffic is also expected to quadruple in the next few years, reaching 6.6 zettabytes (1 zetta $=10^{21}$) by 2016. Most of the traffic (about 76%) is forecasted to stay within data centres for server-to-server connectivity purposes, while the rest is expected to leave the DCN. Such a huge intra data centre traffic volume can be considered as the main trigger for deploying ultra-high bandwidth and low-latency optical technologies within data centres: this is aligned with the LIGHTNESS concepts and research towards the hybrid optical DCN infrastructure. On the Internet facing side, the study identifies two main components of traffic leaving data centres: traffic among geographically distributed data centres, and traffic from data centre to end users associated to applications (north-south traffic). Focusing on the first component, the interconnections among data centres are becoming a critical point to provide performance guarantees and high availability for end-to-end cloud services and applications.

A recent study from Telefonica about traffic profile in two data centre hubs in Latin America [telefonica-dc] shows that traffic among data centres is mostly generated in a periodic fashion as result of virtual machines migration and databases synchronization for disaster recovery purposes. Therefore, over-provisioning of inter data centre connectivity based on the expected peak loads brings to dramatic inefficiencies of network resource usage in the end-to-end chain. Even if evidently sub-optimal and inefficient, network providers still offer static connectivity services among data centres, while the traffic characteristics would encourage a more dynamic and flexible provisioning of network resources across the end-to-end data plane (including data centres, metro and core networks). As a consequence, cloud services spanning multiple data centres require new end-to-end mechanisms to provide dynamic, flexible and convergent configuration of intra data centre and metro/core networks, thus avoiding as much as possible over provisioning and static configurations. On the one end, data centre operators need to improve dynamicity and flexibility of their infrastructure, by supporting automated procedures for orchestration, provisioning and optimization of DCN services and resources. On the other end, network providers that offer inter data centre connectivity can meet the above cloud service requirements by leveraging on cost effective and adaptive metro/core network technologies for packet and circuit services, coupled with dedicated control plane mechanisms that can dynamically configure inter data centre network resources.

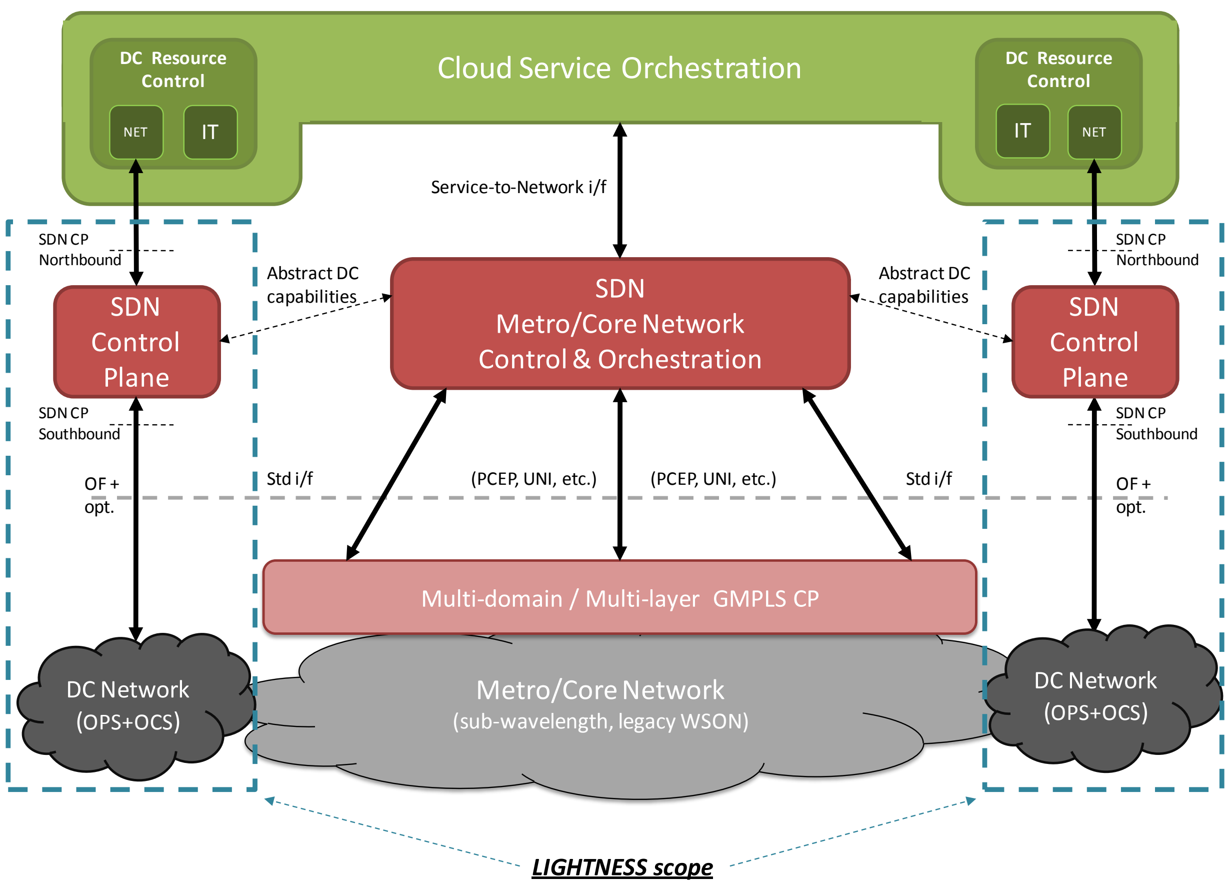

Figure 4.2 SDN framework for orchestration of inter data centre cloud and network services

LIGHTNESS proposes an advanced DCN architecture for ultra-high bandwidth, dynamic and on-demand network connectivity that integrates OCS and OPS technologies inside the data centre. Although the main focus of the project is concentrated on the provisioning of intra data centre network services, the proposed architecture (in particular the control plane) is designed to also support the interfacing (at least at the Northbound) with external cloud orchestration entities and network control planes providing inter data centre connectivity, in order to achieve a consistent host-to-host network configuration. Figure 4.2 shows an SDN based network architecture for the end-to-end orchestration and provisioning of network services in cloud scenarios, where the LIGHTNESS control plane is deployed inside the data centres to dynamically and on-demand configure and optimize the hybrid optical DCN.

The physical infrastructure in Figure 4.2 includes several optical technologies across different segments: OPS and OCS within the data centres, and legacy switching technologies, like Wavelength Switched Optical Network (WSON) among data centres. To better adapt to the cloud traffic dynamicity, the metro network could adopt optical technologies based on sub-wavelength switching, which are particularly suitable for inter data centre scenarios due to their granularity with respect to the traditional lightpath services. An example of this type of technology is represented by the Time Shared Optical Network (TSON) [tson] that is a novel multiplexing metro network solution developed in the in the FP7 MAINS project [MAINS]. Its applicability in inter data centre metro networks is investigated in the FP7 CONTENT project [CONTENT]. The interoperability between heterogeneous metro and core networks is enabled through the adoption of standard GMPLS and PCEP protocols and inter domain interfaces that provide routing and signalling functionalities in a multi-domain and multi-layer environments. An SDN-based inter DCN orchestrator operates on top of this multi domain GMPLS control plane through standard protocols and interfaces, like PCEP and UNI. Such interfaces can be used to retrieve monitoring or abstracted topology information from the underlying networks or enforce commands for connectivity service provisioning. Following the ABNO model described in section 2.2 [draft-farrkingel-pce-abno-architecture], the metro/core network orchestrator is a centralized decision point responsible for inter data centre network resources allocation: an enhanced Service-to-Network interface allows the cooperation with the upper-layer cloud orchestration. A dedicated hierarchical [RFC6805], stateful and active PCE [draft-ietf-pce-stateful-pce] deployed in the SDN orchestrator could allow to autonomously optimize the inter data centre network configuration by maintaining a synchronized view of the services established in the inter data centre network. As a consequence, leveraging the adoption of the PCE framework in the LIGHTNESS SDN control plane (as described in section 5.2), the metro/core orchestrator can optionally take into account abstracted DCN capabilities exchanged through standard PCEP procedures.

At the upper layer, a cloud service orchestrator is responsible for the decisions about end-to-end cloud and network resource provisioning, mainly in terms of allocation, up-/downscaling, re-configuration according to the dynamics required by the cloud applications. In the inter data centre segment abstract information about network topology, capabilities or connectivity quotations may be exposed from the SDN metro/core network orchestrator at the Service-to-Network interface (e.g. through an ALTO server [ALTO]) to allow the cloud service orchestrator to take network-aware decisions about virtual machines placement and/or migration. On the other end, the data centre resource control modules are in charge of network configuration within each data centre: their cooperation with the LIGHTNESS SDN control plane is performed through the Northbound interface described in the previous section, that represents the reference point to provision dynamic, on-demand and resilient network services inside the DCNs. To this purpose, the LIGHTNESS control plane must provide at the Northbound interface the primitives to allow the provisioning of network services which represents the intra data centre portion of an end-to-end inter data centre network connectivity. Consequently, the LIGHTNESS control plane must be able to perform, through its Southbound interface, dedicated network configurations at the edge of the data centre to stitch intra and inter network services, e.g. by supporting the configuration of MPLS and GMPLS labels, outgoing wavelengths, etc.

5. LIGHTNESS control plane procedures

This chapter describes the procedures supported by the LIGHTNESS control plane to implement the enhanced data centre functions introduced in the previous sections. Starting from the SDN functional decomposition and specification carried out in chapter 4, the LIGHTNESS control plane procedures for DCN service provisioning (including path and flow computation), recovery, optimization and monitoring are detailed into dedicated sections below.

5.1. Data centre connectivity services

As highlighted in the previous sections, in future data centre scenarios, flexible and highly scalable control and management of intra data centre server-to-server connectivity will be mandatory. The introduction of automation inside the data centre by means of a control plane entity allows to match connectivity requirements imposed by the running applications. While basic application requirements are limited to certain throughput and latency metrics, the advent of a control plane inside the data centre can include additional features in the computation of the workload flow routes. These features, such as certain reliability levels, link congestion avoidance (e.g., through load balancing), or even energy consumption awareness, can directly impact not only on the particular workload exchange under establishment, but also on the overall data centre performance and operational costs.

The LIGHTNESS control plane proposed solution mainly focused on providing connectivity services among servers/racks in intra data centres scenarios. These service are conceived to provide data connectivity within the DCN, and involve the configuration and provisioning of the data plane entities: OCS, OPS and TOR switches, and possibly the servers' Network Interface Cards (NICs). The control of establishment and teardown of these services is performed by the SDN controller described in the previous chapter.

Basically the following two types of connectivity services to be supported and provided in the DCN have been identified in LIGHTNESS:

- *Point-to-point connectivity*: this service refers to the simplest scenario where a network connectivity is established between two end-points in the DCN, identified by either TOR switches ports or servers ports;
- *Point-to-multipoint connectivity*: this service refers to a more complex connectivity paradigm where a source end-point (either TOR switch or server port) has to be connected to multiple destination end-points, to support specific requirements of data centre applications (e.g. content or data base replication towards multiple servers, gaming applications involving IT resources in several servers, etc.)

In both cases, the main requirements for such connectivity services are:

- *High throughput*: Most of the applications require very high bandwidth connections among servers.
- *Reliability* (*high-availability*): fault tolerance and graceful performance degradation is considered extremely important; in current data centre infrastructure, this is provided by redundant connections to be used upon link failure occurrence, that increase the overall data centre CAPEX.
- *Quality of service (QoS) guarantees*: the connectivity among servers must match requirements in terms of latency, jitter, and packet losses.
- *Energy consumption awareness*: Especially in large data centres for cloud computing (hosting up to 100K servers), the use of power efficient connections is of paramount importance.

Specific procedures have been defined in support of the point-to-point and point-to-multipoint connection paradigms. The following sub-section describes these procedures in a step-wise fashion to highlight main actions to be carried out for automated connectivity establishment inside the DCN. In addition, a dedicated sub-section for dynamic connectivity service modification is also included to cover this enhanced function provided by the LIGHTNESS control plane and describe the related procedures. Finally, the control of scheduled connectivity services is also briefly analyzed and described.

The procedures for path and flow computation are here summarized with a single action at the internal Path/Flow computation module of the SDN controller: the detailed procedures and possible options for these computations are deeply described in section 5.2.

5.1.1. Data centre connectivity service setup and teardown

The procedures for establishment and teardown of point-to-point and point-to-multipoint connectivity are similar from a control plane point perspective, in particular in terms of main interactions at external and internal interfaces. For this reason, the procedures for both types of services are described below in the same table. In particular, the connectivity service setup and teardown procedures are depicted in Figure 5.1a and Figure 5.1b: a step-wise description of main actions and interactions is also provided in Table 5.1.

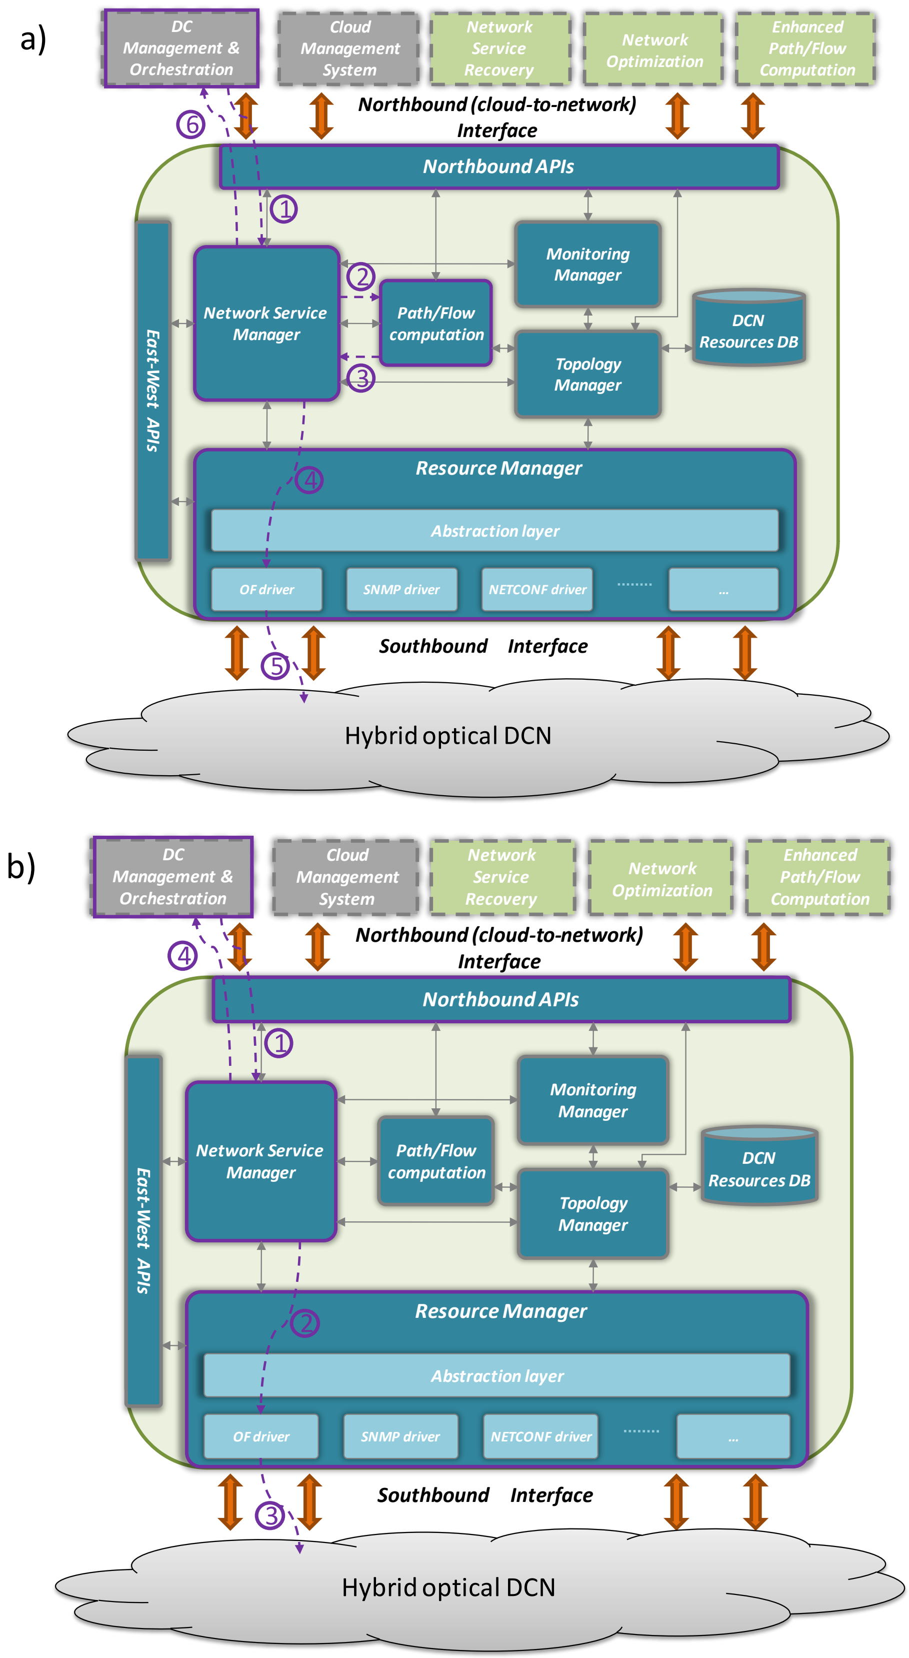

Figure 5.1 Data centre connectivity service setup & teardown

Table 5.1 includes both setup and teardown phases, and differentiates (where needed) the description of point-to-point and point-to-multipoint connectivity services procedures.

Figure 5.1 shows a scenario where the data centre connectivity service is requested by the data centre management entity through the Northbound interface. Moreover, at the Southbound interface the usage of OpenFlow protocol (properly extended to support the LIGHTNESS OPS, OCS and TOR switches) is assumed, as it is the main candidate (under investigation in Task 4.3 at the time of writing) for the implementation of this interface.

Phase	Step	Description
Setup	1	A new connectivity service is requested by the Data Centre Management entity through the Northbound interface. The request includes at least the following parameters: • End-points of the service • QoS parameters (bandwidth, latency, packet loss, etc.) • Recovery strategy (unprotected, protected, on-the-fly – see section 5.3) • Monitoring mechanism (see section 5.5)
	2	The Network Service Manager module inside the SDN controller handles the connectivity service requests, and performs a preliminary check on the requirements imposed by the Data Centre Management entity, in terms of recovery strategy, monitoring mechanisms. If they can be provided, then the Network Service Manager invokes the Path/Flow Computation module to resolve the end-points in the connectivity request and
	3	The Path/Flow Computation module performs the computation of the network path that satisfies the QoS constraints requested, according to the procedure described in section 5.2. If a suitable path is found, it is sent back to the Network Service Manager. Depending on the type of connectivity service, different behaviours are defined: • Point-to-point service: a single path is computed that connects the two end-points specified in the request • Point-to-multipoint service: a single point-to-multipoint path is computed if the data plane entities are able to perform multicast actions at transport level. Otherwise independent point-to-point paths

		(one for each destination) are computed
	4	The Network Service Manager processes the path(s) computed and composes the connectivity service to be installed in the DCN. It then invokes the Resource Manager to provision the service and reserve the network resources in the data plane entities. In case of point-to-multipoint connectivity, the service may be composed by a sub-set of independent connectivity services implementing the point-to-multipoint paradigm if the data plane entities are not able to perform multicast actions at the transport plane. In case, the Network Service Manager is in charge to bundle these independent services and handle them as a single network service to the upper layers.
	5	The Resource Manager converts the provisioning request received from the Network Service Manager into a format suitable for the Southbound interface. This is carried out by the Abstraction Layer, that in this case invokes the OpenFlow driver to interact with the underlying data plane entities and provision the network resources.
	6	Once the provisioning of the network service is carried out, the Network Service Manager notifies the Data Centre Management entity about the successful establishment of the connectivity. This notification may be either synchronous or asynchronous, depending on the Northbound interface procedures and semantics that will be defined in deliverable D4.3 [del-d43] (out of the scope of this document)
Teardown	1	The Data Centre Management entity requests for the cancellation of the connectivity service previously established. This action is performed through the Northbound interface.
	2	The Network Service Manager processes the teardown request, retrieve the connectivity service to be cancelled and identifies the related network resources to be de-provisioned. Then, the Resource Manager is invoked to actually free the resources involved in the given service. In case of point-to-multipoint service, the service may be established as a sub-set of independent connectivity services. In this case, the Network Service Manager is in charge to split the cancellation request into atomic actions

		towards the Resource Manager to free all the involved resources
	3	The Resource Manager converts the de-provisioning request received from the Network Service Manager into a format suitable for the Southbound interface. This is carried out by the Abstraction Layer, that in this case invokes the OpenFlow driver to interact with the underlying data plane entities and free the network resources.
	4	Once the de-provisioning of the network service is carried out, the Network Service Manager notifies the Data Centre Management entity about the successful cancellation of the service. This notification may be either synchronous or asynchronous, depending on the Northbound interface procedures and semantics that will be defined in deliverable D4.3 [del-d43] (out of the scope of this document)

Table 5.1 Step-by-step data centre connectivity procedure

5.1.2. Dynamic connectivity service modification

The dynamic and on-demand modification of already established services is an advanced feature supported by the LIGHTNESS control plane that allows the management and orchestration entities on top of the data centre infrastructure to trigger service modifications upon change of needs and requirements of applications running in the data centre. The trigger for this kind of procedure is therefore external at the LIGHTNESS control plane: however, as per connectivity setup and teardown, the modification procedure is fully automated and does not need any human validation for upgrading/downgrading the given installed service.

As described in section 3.3, the following dynamic modifications may be applied to installed network services:

- QoS constraints: to upgrade or downgrade the installed service in terms of bandwidth, latency, jitter, etc.;
- Recovery strategy: to modify the recovery mechanism associated to the installed service, according to the paradigms defined in section 5.3;
- Destination end-point in point-to-multipoint connectivity: to modify, add or delete a destination end-point due to changes in data centre applications requirements;
- Monitoring mechanism: modification of the current mechanisms in place, according to the different paradigms defined in section 5.5

The modification of the monitoring mechanism is the simplest among the others. Indeed, upon a request from the management and orchestration entities through the Northbound interface, the Network Service Manager only needs to re-configure and update the monitoring behaviour through a simple interaction with the Monitoring Manager performed internally at the SDN controller. On the other end, for what concern QoS, recovery strategy and destination end-point modifications, dedicated re-computation and re-provisioning of the given network service may be needed to let the control plane implement the requested modification, thus involving more modules inside the SDN controller (mainly Network Service Manager, Path/Flow Computation, Resource Manager) and also requiring cooperation with the underlying data plane through the Southbound interface. However, QoS, recovery strategy and destination end-point modification procedures are similar in terms of actions to be performed and internal/external control plane interactions, since they mostly reuse part of the setup and teardown procedures described in the previous section.

Indeed, for QoS modifications, a bandwidth upgrade may need the provisioning of a new network service (with de-provisioning of the original one) to meet the modification requirements (resulting in a modification procedure composed by both setup – of new service - and teardown – of original one – phases described in section 5.1.1). Similarly, in case of upgrade of recovery strategy from unprotected to protected, the control plane needs to establish a new protection network service to be used as backup in case of network failures, following the same procedure described in section 5.1.1 for the setup phase. Finally, also the addition of a destination end-point in a point-to-multi-point network service requires a modification procedure that implements the same steps defined in section 5.1.1 for the setup phase of a new service. The Network Service Manager first invokes the Path/Flow Computation to calculate the new network service segment (or new independent network connectivity in case of data plane not capable to perform multicast actions) needed to connect the given new end-point, and then requests the Resource Manager to provision the new network resources in the data plane through the Southbound interface.

5.1.3. Scheduled connectivity services

A LIGHTNESS connectivity service can be scheduled to allow the advance reservation of data centre services and associated network resource reservation for usage in a given future time slot. During the setup phase of a new connectivity service with advance reservation, the involved network resources (already computed) are reserved for their usage in the specified time interval. At the specified network service start time (included in the connectivity service request performed through the Northbound interface, along with the service duration), they are automatically activated (i.e. actually provisioned in the data plane) by the control plane through the Southbound interface. These scheduled connectivity services are handled at the LIGHTNESS control plane by maintaining a set of calendars that describe the network resource availabilities on a time basis, thus taking into account time slots. These calendars are stored in the DCN Resource DB inside the SDN controller and are accessed by the control plane path and flow computation functions (either internal or external – i.e. network applications) to calculate and take into account the actual network resource availability for the given time slot.

The procedure for the setup of a scheduled data centre connectivity service is similar to the one described in section 5.1.1. The only differences can be summarized as follows:

- Path and flow computation must take into account the start time and the service duration specified in the connectivity request;

- Network resources are not immediately configured at the connection setup phase inside the DCN ; a start time timer is activated and when it expires the procedure for DCN resource provisioning is triggered (i.e. steps 4 and 5 in Table 5.1 for connectivity service setup) .

5.2. Path and flow computation

In LIGHTNESS, the path and flow computation mechanisms are conceived to facilitate the provisioning of network services in support of data centre and cloud services. This path and flow computation can be considered as a fundamental functionality of the LIGHTNESS control plane, that controls the hybrid optical DCN and offers connectivity services such as point-to-point and point-to-multipoint. As described in section 4.1, the LIGHTNESS SDN controller natively implements a set of path and flow computation functions to let the Network Service Manager support basic network services. Additional routing functionalities, such as enhanced computation engines or algorithms, can be provided as network applications running on top of the LIGHTNESS SDN controller.

Figure 5.2 highlights the LIGHTNESS path and flow computation components in the control plane architecture. The key component is the stateful and active PCE [draft-ietf-pce-stateful-pce], which is deployed as an external network application for the SDN controller. The stateful and active PCE is based on the IETF PCE framework [RFC4655] and provides two additional functionalities. First, as a stateful PCE it maintains and stores status (and possibly performances) of network services installed in the DCN. On the other end, as an active PCE, it is able to trigger autonomously the modification of already installed network services or trigger the setup of new connections. The combination of these two enhanced functions makes the stateful active PCE particularly suitable for the adoption in data centre scenarios, since it can natively enable dynamicity and flexibility in the control of DCN resources. This means that the stateful active PCE is able to provide network optimization functionalities in the LIGHTNESS control plane by gathering updated information about performances and status of the DCN from the SDN controller Monitoring Manager, and accordingly triggering, when needed, network services re-configurations and modifications (i.e. interacting with the SDN controller Network Service Manager through the Northbound APIs).

In addition, the adoption of the IETF PCE framework eases the interoperation with PCEs deployed in other network segments and domains (e.g. metro, core/backbone) for abstracted network information exchange, inter data centre computations and network connectivity. Indeed the IETF PCE natively supports standard interfaces and protocols, like the PCEP [RFC5440], that allow the stateful active PCE in charge of path and flow computations inside the data centre to cooperate and exchange data centre capabilities with other PCEs used for inter data centre orchestration and control purposes, as mentioned in section 4.3.

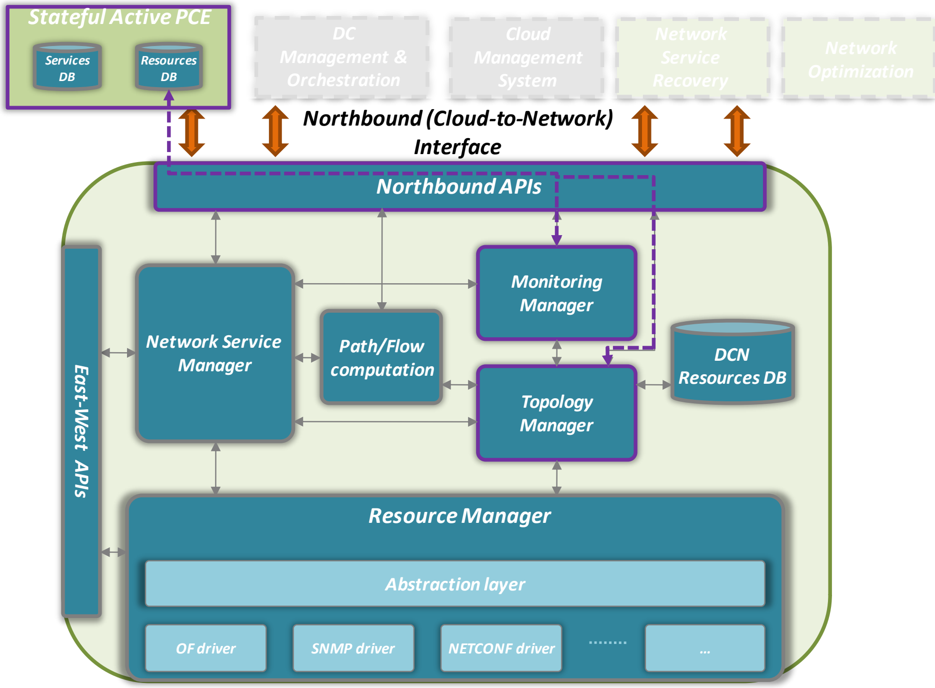

Figure 5.2 LIGHTNESS path computation scenario.

As shown in Figure 5.2, the stateful active PCE hosts two different databases to store on the one hand the status and performances of network services installed in the DCN, and on the other hand the status and capabilities of the whole DCN resources. In particular, this latter database is conceived to store resource status and capability information coming from the DCN Resources database in the SDN controller: however, it may contain additional information (e.g. statically or dynamically configured by the data centre management and orchestration entities) related to Traffic Engineering parameters, IT resources at the edges of the DCN, energy consumption characteristics, that allow to implement enhanced path computation algorithms for DCN services. The synchronization of this stateful active PCE database with the DCN Resources database in the SDN controller can occur following two different approaches:

1. Active and continuous synchronization enabled by the interaction with Topology and Monitoring Managers through the Cloud-to-Network interface, in order to gather all the most updated DCN resources availabilities and capabilities
2. Asynchronous polling of Topology Manager each time a computation request has to be performed. In this case the stateful active PCE database could also be an ad-hoc network graph built on top the topology view gathered from the Topology Manager (augmented with the additional information mentioned above), created and released (i.e. deleted) for each computation.

Two different path computation procedures have been identified: the first one considers the stateful active PCE actually deployed and running on top of the SDN controller, while the second one relies on the only Path/Flow computation module inside the SDN controller (i.e. without stateful active PCE). Indeed, since the stateful active PCE depicted in Figure 5.2 is a network application of the LIGHTNESS control plane, it could be disabled (i.e. not running) or not implemented, depending on the specific data centre or cloud operator requirements and needs. To differentiate between these two scenarios, we assume that the Path/Flow computation module in the SDN controller is aware of the availability of the stateful active PCE and acts as a decision point for its usage.

Table 5.2 describes, step-by-step, the LIGHTNESS path computation procedure when the external PCE is running as a network application on top of the SDN controller. Such procedure is also shown in Figure 5.3.

Step	Description
1	Network Service Manager receives a new request for a connectivity service to be configured in the DCN (e.g. a point-to-point service through the Northbound APIs), with a set of constraints and data flow characteristics to be matched, including network end-points, bandwidth, duration of the service, QoS, etc. This request is internally processed and the Path/Flow computation is invoked through the related internal SDN control interface to find a data path that matches the network service constraints.
2	The Path/Flow computation module receives the request from the Network Service Manager and processes it. It first recognizes that an external PCE is available for this computation and, according to configured policies, forwards to it the request. The communication between the Path/Flow computation and the external PCE could be based on the PCEP protocol: in this case the Path/Flow computation module could include a light Path Computation Client (PCC) to properly manage the request and reply PCEP messages.
3	The external PCE receives the computation request (e.g. through its PCEP interface) and processes it. Based on the network information stored in its databases (for network services and DCN resources) it calculates the optimal path running a computation algorithm according to the constraints received. If such a path exist and is successfully found, the external PCE sends back the computation reply to the Path/Flow computation module.
4	The Path/Flow computation module receives the computation reply from the external PCE and forwards it to the Network Service Manager to proceed with the network

	service setup.
5	The Network Service Manager receives the computed path from the Path/Flow computation module and binds it to the corresponding pending network service. The path computation procedure is now closed and the Network Service Manager is ready to provision the connectivity service by interacting with the Resource Manager.

Table 5.2 LIGHTNESS path computation procedure: external PCE deployed

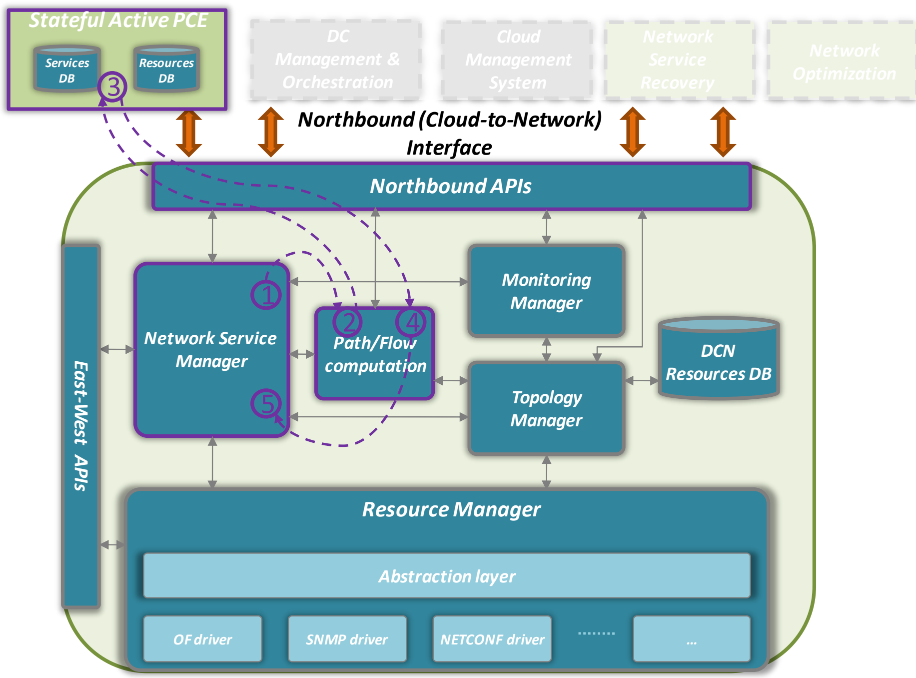

Figure 5.3 LIGHTNESS path computation procedures with external PCE.

On the other hand, a step-by-step description of the path computation procedure without an external PCE is provided in Table 5.3 and also depicted in Figure 5.4.

Step	Description
1	Network Service Manager receives a new request for a connectivity service to be configured in the DCN (e.g. a point-to-point service through the Northbound APIs), with a set of constraints and data flow characteristics to be matched, including network end-points, bandwidth, duration of the service, QoS, etc. This request is internally processed and the Path/Flow computation is invoked through the related internal SDN control interface to find a data path that matches the network service constraints.
2	The Path/Flow computation module receives the request from the Network Service Manager and processes it. According to configured policies and computation request requirements, it decides to internally handle such computation.
3	The Path/Flow computation module interacts with the Topology Manager to get an updated view of the DCN topology, along with resources status, capabilities and availabilities. Based on this topology view, it also builds a network graph to run a computation algorithm on top of it.
4	The Path/Flow computation module calculates the optimal path running a computation algorithm according to the constraints received by the Network Service Manager. If such a path exists and is successfully found, it sends back the computation reply to the Network Service Manager.
5	The Network Service Manager receives the computed path from the Path/Flow computation module and binds it to the corresponding pending network service. The path computation procedure is now closed and the Network Service Manager is ready to provision the connectivity service by interacting with the Resource Manager.

Table 5.3 LIGHTNESS path computation procedure: external PCE not deployed

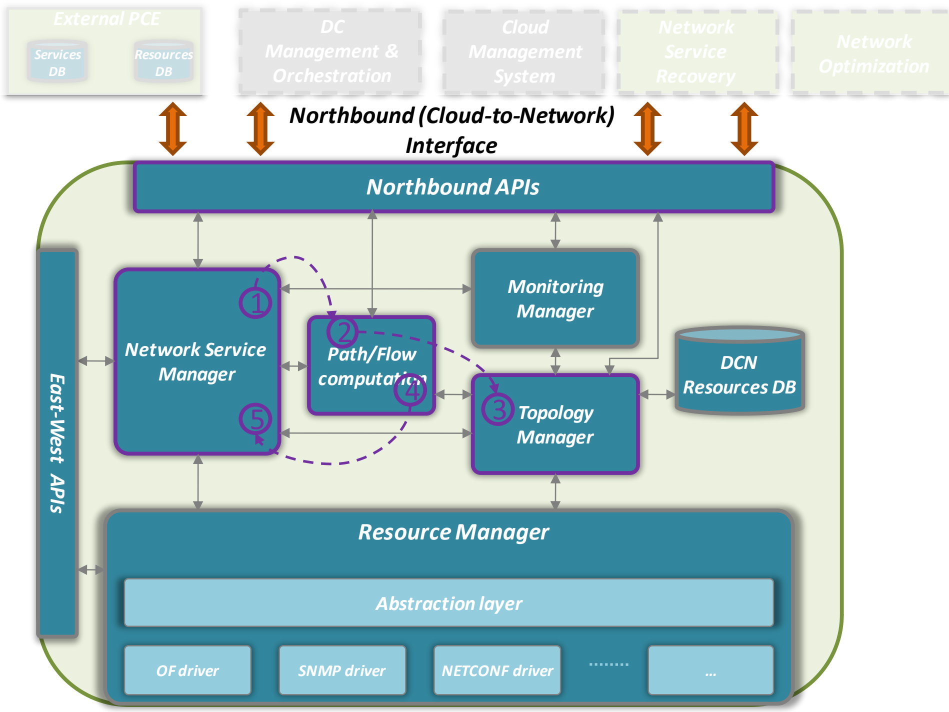

Figure 5.4 LIGHTNESS path computation procedures without external PCE.

5.3. Data centre network resiliency

The deployment of a control plane on top of the DCN infrastructure allows the provision of a set of additional and innovative functionalities that are not currently implemented inside data centres. In particular, besides the dynamic, on-demand and flexible DCN connectivity services provisioning described in previous sections (including enhanced computation functions), a key functionality offered by the control plane is represented by the DCN resiliency. In LIGHTNESS, automated procedures for network failures detection and recovery are implemented provide reliable network service reliability to applications running inside data centres. The final objective is to implement an autonomous DCN infrastructure where the control plane is able to automatically react to network failure conditions that may affect the performances of applications and services running inside the data centre.

The recovery procedures described in this section relies on the implementation of monitoring functions in the LIGHTNESS control plane: this means that each network service installed in the data centre can be monitored and checked in terms of performances, status of involved resource and potential network failures. Further details about monitoring procedures and mechanisms are provided in section 5.5.

The LIGHTNESS control plane offers different types of recovery strategies for the network services to be established inside the data centre. The strategy to be implemented for a given network service is specified at the Northbound interface when the new service is requested to the control plane. Therefore the Northbound interface (that will be specified in deliverable D4.3 [del-d43]) will offer the possibility to select the preferred recovery strategy as a constraint for the network service to be established.

The following recovery strategies are offered by the LIGHTNESS control plane:

- **Unprotected**: this strategy refers to network services that does not need any recovery action in case of network failures. This may be applied to network services associated to not critical data centre applications that are not affected by network connectivity interruptions;
- **Protected**: one or more protection network services are pre-established (or at least booked without actual configuration in the data plane physical devices) by the control plane when the given network service is requested. In this context, two types of protection are offered by the LIGHTNESS control plane:
  - 1:1 protection: the given (primary) network service has one dedicated protection network service (either pre-established or booked);
  - N:M shared protection: this protection refers to N primary network services that share M protection (i.e. backup) network services. These M backup services are differentiated by a priority

  This recovery strategy imply a kind of overprovision of resources, since backup services are immediately established: for this reason its usage can be limited to those network services associated to critical data centre applications highly affected by potential network failures and outages (e.g. banking);
- **On-the-fly restoration**: this strategy does not provide any pre-establishment of backup network services: no additional DCN resources are configured or booked in advance. Upon a network failure a disjoint network service is dynamically computed and established. To speed up the recovery procedure, the control plane implements a make-before-break mechanism: this means that first the new disjoint network service is established, and then the failed one is released.

A step-wise description of the LIGHTNESS restoration procedure is provided in Table 5.4 below. Please refer to Figure 5.1 for the control plane functional entities mentioned.

Step	Description
1	Failure Detection. Once the network failure is detected by the data plane entities, a notification is sent to the SDN controller in the LIGHTNESS control plane through the Southbound interface.
2	Failure Localization. The network failure notification coming from the Southbound interface is processed by the Resource Manager inside the SDN controller, which interacts with the Topology Manager to update the status of the involved network resources. Then the Monitoring Manager is triggered to check if network service is affected by the network failure.
3	Failure Notification. The Monitoring Manager correlates the network failure notification received from the Resource Manager with all the involved network services established inside the data centre. After that, it notifies any network application running on top of the SDN controller that is registered for these type of events. In this step-wise description we assume the presence of a dedicated Network Service Recovery application implemented and running on top of the SDN controller.
4	Recovery. The Network Service Recovery application receives the network failure notification from the SDN controller through the Northbound interface. This is the trigger for the core part of the recovery procedure. Depending on the recovery strategy associated to the given failed service, the Network Service Recovery application interacts with the Network Service Manager in the SDN controller (again through the Northbound interface) to take specific decisions and actions: • **Unprotected**: nothing to do • **Protected:** the switch to one of the backup network services is performed by the Network Service Manager (by interacting with the data plane through the Southbound interface) in case of pre-established 1:1 or N:M protection scheme. On the other end, if the backup network services are booked, the Network Service Manager proceeds with the activation of the highest priority protection service • **On-the-fly-restoration:** the Network Service Manager requests for a disjoint network service computation to the PCE (either external or internal to the SDN controller) and then provision the new service following the make-before-break paradigm

5	Recovery escalation. In case of failure in some step of the recovery mechanisms (e.g. if alternative disjoint network paths not available), the network failures are notified to the upper-layer data centre management or orchestration entities, and the recovery strategies are moved and applied at the service layer, i.e. out-of-the-scope of the control plane responsibilities.

Table 5.4 LIGHTNESS network service recovery procedure

5.4. Dynamic data centre network resources optimization

The network optimization features supported by the LIGHTNESS control plane allow the dynamic modification of connectivity services already provisioned (or potentially just scheduled) on the DCN in automated reply to degradation of service performance or network utilization. As shown in Figure 5.5, this function is enabled through the cooperation of internal modules within the SDN controller and dedicated external applications. In particular, the detection of inefficient network configuration is based on algorithms running on top of the monitoring information provided by the SDN controller Monitoring Manager. These algorithms are implemented in a dedicated external application (indicated as Network Optimization in the picture) that is also responsible of triggering and managing the entire sequence of the re-optimization procedure. When conditions requiring network re-optimization are detected, the new resource allocations are elaborated through dedicated path computation mechanisms and then enforced on the data centre through the Network Service Manager operation. It should be noted that path computation for re-optimization is usually based on dedicated objective functions and adopts Global Concurrency Optimization (GCO) [RFC5557] algorithms that could take into account advanced data on the network usage beyond the traditional topology and TE information (e.g. stateful data about running or planned services). For this reason, this type of path computation may be not performed directly by the internal SDN controller functions, but it may be delegated to the external PCE application.

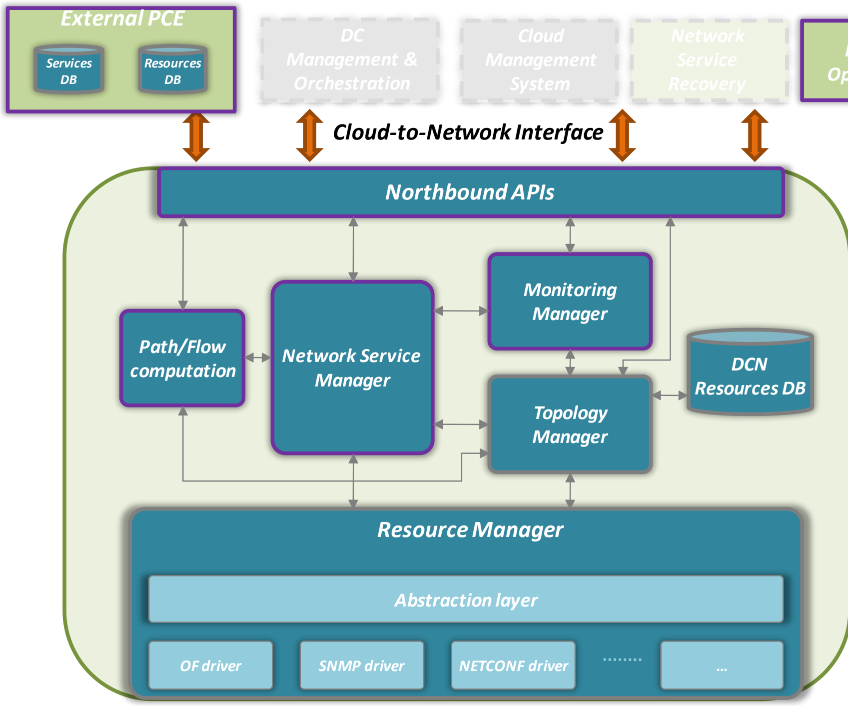

Figure 5.5 Architecture modules involved in LIGHTNESS network re-optimization.

A possible workflow for network re-optimization is described in Table 5.5, while the interaction between the different architecture components is shown in Figure 5.5.

Step	Description
1	<u>Monitoring data collection</u>. Information about the status of the network and the performance of the established services are collected from the Monitoring Manager within the SDN controller and forwarded to the Network Optimization application. The interaction between the two entities can follow different approaches, e.g. based on polling or subscription/notification mechanisms (see section 5.5 for further details).
2	<u>Detection of performance degradation</u>. The Network Optimization application, based on the monitoring information retrieved from the SDN controller, detects some conditions that could reflect a degradation of the network performance, e.g. in terms of inefficient and unbalanced network resource utilization, continued failures in the allocation of new network services, poor performance of existing services. This detection can be based just on real-time data, or could involve some predictions about future services and traffic.

3	<u>Initiation of network re-optimization</u>. Once a performance degradation is detected in the DCN, the Network Optimization application may decide to initiate a new re-optimization procedure to re-configure all or part of the existing or scheduled services. This decision should be based on configurable policies implemented within the application.
4	<u>Computation of re-optimized network services</u>. The new network service configuration must be computed through dedicated path computation procedures. The SDN controller internal path computation function, based on specific policies, may decide to delegate this computation to external PCEs, specialized with algorithms dedicated to global concurrent optimization. The new computed flows are then notified back to the Network Optimization application.
5	<u>Enforcement of re-optimized network services.</u> In case a positive solution has been found for the re-optimization problem, the Network Optimization application requests the enforcement of the suggested service modifications on the DCN. This procedure is handled by the Network Service Manager through the traditional mechanisms for DCN configuration (i.e. through the Resource Manager).

Table 5.5 LIGHTNESS network re-optimization procedure

5.5. Data centre network monitoring

The enhanced functionalities offered by the LIGHTNESS control plane and described in the above sections rely on the implementation of a set of monitoring functions. In particular, while the DCN resiliency and the dynamic network resource optimization refer to automated procedures enabled by control plane internal exchange of monitoring information, the dynamic modification of established network services is enabled by monitoring information exported by the control plane towards external management entities. This means that the monitoring functions in the LIGHTNESS control plane are provided at different levels (i.e. external and internal) and enhance the cooperation with upper layer actors, such as data centre management and orchestration entities.

The main purpose of these control plane monitoring functions is to make the network services established inside data centres a fully manageable set of resources in the overall data centre workflows. Indeed, a synchronization of their status, performances and availability allow either the control plane itself or the external entities to apply their own procedures on the network services (i.e., dynamic modification, re-optimization, recovery).

The monitoring information exchanged at the different levels described above is related to either already established network services inside the data centre, or single network resources not associated to any network service. In particular, the information is expressed in terms of status, failures and performances of these resources. In general, we can assume that the information related to network services are exchanged through the Northbound interface, while the ones related to single network resources are exchanged through the Southbound interface at first, and then through internal control plane interfaces as well.

Three different modes for exchanging the monitoring information (at both Northbound and Southbound interfaces) are defined in LIGHTNESS:

- **Polling**:
  - *Northbound*: network applications running on top of the SDN controller, data centre management and orchestration entities poll the control plane to get monitoring information whenever they need;
  - *Southbound*: the SDN controller polls the underlying data plane entities to retrieve status and performances of DCN resources;
- **Synchronous notifications**:
  - *Northbound*: a monitoring state is maintained at the control plane for each installed network service, and notifications about their status and performances are sent to the registered entities (e.g. network applications running on top of the SDN controller, data centre management and orchestration, etc.) on a periodic basis;
  - *Southbound*: periodic notifications for DCN resources status and performances are received by the SDN controller from the underlying data plane entities, and processed by the Monitoring Manager;
- **Asynchronous notifications**:
  - *Northbound*: a monitoring state is maintained at the control plane for each installed network service, and notifications in case of changes of status and performances are sent upwards based on thresholds and conditions as agreed through the interface itself;
  - *Southbound*: spontaneous notifications are sent from the data plane entities to the SDN controller when failures, changes of status or performances occur at the DCN resources.

The modes for exchanging monitoring information defined above should be configurable at the involved interfaces, i.e. the Northbound and Southbound: such configuration can be defined at either a global control plane level or on a per-network service basis. In particular, each entity involved in the monitoring procedures should be allowed to customize the monitoring behaviour. In LIGHTNESS, the features that can be customized at the control plane interfaces include:

- Monitoring mode (polling, synchronous notification, asynchronous notification).

- Specific parameters of the selected monitoring mode (e.g. synchronous notification frequency, asynchronous notification conditions and thresholds, etc.).

- Depth of monitoring information, in terms of set of parameters to be included in monitoring notifications (resources/services IDs, status, performance parameters, failure reasons, etc.).

6. Conclusions

This document has provided the LIGHTNESS control plane architecture, and presented the enhanced functionalities in support of dynamic, on-demand and flexible provisioning of DCN resources.

An SDN based solution has been proposed for the LIGHTNESS control plane architecture with the aim of operating the DCN resources in a flexible and programmable way, allowing data centre and cloud operators to build more scalable, agile and easily manageable DCNs. The functional modules composing the control plane split architecture have been identified and described. Furthermore, the detailed specification of the enhanced functionalities and automated procedures for on-demand provisioning and dynamic modification of highly resilient network services, path computation and dynamic data centre resource optimization has been also provided. The external control plane interfaces have been also identified ad briefly detailed in terms of supported functions and procedures: this latter part represents a fundamental input for the activities carried out in Task 4.3 for the specification of the LIGHTNESS control plane interfaces.

The relevant set of control plane functionalities described in this document will be then identified and developed as software modules in Task 4.4. The resulting proof-of-concept prototype will be then integrated in the overall LIGHTNESS system in the context of WP5 activities, and will be experimentally validated over the project test-bed to be finally used for trial purposes.

7. References

[ALTO] ALTO Protocol - http://tools.ietf.org/html/draft-ietf-alto-protocol-10

[cisco-gci] Cisco, Global Cloud Index, http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns1175/Cloud_Index_White_Paper.html, 2012

[cloudstack] Apache CloudStack: http://cloudstack.apache.org/

[CONTENT] "Convergence of Wireless and Optical Network an IT Resources in Support of Cloud Services", European Commission's FP7 FP7 CONTENT project, http://content-fp7.eu

[dc-traffic] Theophilus Benson, Aditya Akella, and David A. Maltz, "Network Traffic Characteristics of Data centres in the Wild", ACM IMC, Nov. 2010.

[del-d22] "Design document for the proposed architecture", LIGHTNESS deliverable D2.2, September 2013.

[del-d43] "The LIGHTNESS network control plane interfaces and procedures", LIGHTNESS deliverable D4.3, in progress.

[dijkstra] Dijkstra, E. W., "A note on two problems in connexion with graphs". Numerische Mathematik 1: 269–27.

[draft-ietf-ccamp-wson-signaling] G. Bernstein, S. Xu, Y. Lee, G. Martinelli, H. Harai, "Signaling Extensions for Wavelength Switched Optical Networks", IETF Draft, work in progress, July 2013

[draft-ietf-ccamp-gmpls-g709-framework] F. Zhang, D. Li, H. Li, S. Belotti, D. Ceccarelli, "Framework for GMPLS and PCE control of G.709 Optical Transport Networks", IETF Draft, work in progress, June 2013

[draft-ogrcetal-ccamp-flexi-grid-fwk] O. Gonzales de Dios, R. Casellas, F. Zhang, X. Fu, D. Ceccarelli, I. Hussain, "Framework and Requirements for GMPLS based control of Flexi-grid DWDM networks", IETF Draft, work in progress, June 2013

[draft-ietf-pce-gmpls-pcep-extensions] C. Margaria, O. Gonzales de Dios, F. Zhang, "PCEP extensions for GMPLS", IETF Draft, work in progress, July 2013

[draft-ietf-pce-stateful-pce] E. Crabbe, J. Medved, I. Minei, R. Varga, "PCEP extensions for stateful PCE", IETF Draft, work in progress, June 2013

[draft-zhang-pce-stateful-pce-app] X. Zhang, I. Minei, "Applicability of Stateful Path Computation Element (PCE)", IETF Draft, work in progress, May 2013

[draft-dhody-pce-stateful-pce-auto-bandwidth] D. Dhody, U. Palle, "PCEP extensions for MPLS-TE LSP Automatic Bandwidth Adjustment with stateful PCE", IETF Draft, work in progress, July 2013

[draft-crabbe-pce-pce-initiated-lsp] E. Crabbe, I. Minei, S. Sivabalan, R. Varga, “PCEP extensions for PCE-initiated LSP setup in a stateful PCE model”, IETF Draft, work in progress, July 2013

[draft-farrkingel-pce-abno-architecture] D. King, A. Farrel, “A PCE-based Architecture for Application-based Network Operations”, IETF Draft, work in progress, July 2013

[floyd-warshall] Floyd, Robert W., "Algorithm 97: Shortest Path". Communications of the ACM 5 (6): 345

[FORCES] Forwarding and Control Element Separation (FORCES) Working Group ttps://tools.ietf.org/wg/forces/

[GEYSERS] “Generalized Architecture for Dynamic Infrastructure Services”, European Commission’s FP7 ICT Work Programme, http://www.geysers.eu

[I2RS] Interface to the Routing System (I2RS) Working Group http://datatracker.ietf.org/wg/i2rs/charter/

[IRTF-SDN] IRTF Software Defined Networking Research Group (SDNRG) http://trac.tools.ietf.org/group/irtf/trac/wiki/sdnrg

[MAINS] “Metro Architectures enabling sub-wavelengths”, European Commission’s FP7 MAINS project, http://www.ist-mains.eu

[netconf] R. Enns, “NETCONF Configuration Protocol”, IETF RFC4741, December 2006.

[ONF] Open Networking Foundation: https://www.opennetworking.org/

[onf-sdn-wp] ONF, “Software-Defined Networking: The New Norm for Networks”, Open Networking Foundation white paper, April 2012

[openflow] “OpenFlow Switch Specification”, version 1.3.1, https://www.opennetworking.org/, Open Networking Foundation, , September 2012

[openstack] OpenStack Cloud software: https://www.openstack.org/

[openstack-neutron] OpenStack Networking (Neutron): https://wiki.openstack.org/wiki/Neutron

[RFC3945] E. Mannie, “Generalized Multi-Protocol Label Switching (GMPLS) Architecture”, IETF RFC 3945, October 2004

[RFC3031] E. Rosen, A. Viswanathan, R. Callon, “Multiprotocol Label Switching Architecture”, IETF RFC 3031, January 2001

[RFC4426] J. Lang, B. Rajagopalan, D. Papadimitriou, “Generalized Multi-Protocol Label Switching (GMPLS) Recovery Functional Specification”, IETF RFC 4426, March 2006

[RFC4206] K. Kompella, Y. Rekhter, “Label Switched Paths (LSPs) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)”, IETF RFC 4206, October 2005

[RFC5212] K. Shiomoto, D. Papadimitriou, J.L. Le Roux, M. Vigoureux, D. Brungard, “Requirements for GMPLS-based Multi-Region and Multi-Layer Networks (MRN/MLN)”, IETF RFC 5212, July 2008

[RFC4655] A. Farrel, J.P. Vasseur, J. Ash, “A Path Computation Element (PCE)-Based Architecture”, IETF RFC 4655, August 2006

[RFC5441] J.P. Vasseur, R. Zhang, N. Bitar, J.L. Le Roux, “A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths”, IETF RFC 5441, April 2009

[RFC6805] D. King, A. Farrel, “The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS”, IETF RFC 6805, November 2012

[RFC5440] J.P. Vasseur, J.L. Le Roux, “Path Computation Element (PCE) Communication Protocol (PCEP)”, IETF RFC 5440, March 2009

[RFC5520] R. Bradford, J.P. Vasseur, A. Farrel, “Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism”, IETF RFC 5520, April 2009

[RFC5557] Y. Lee, J.L. Le Roux, D. King, E. Oki, “Path Computation Element Communication Protocol (PCEP) Requirements and Protocol Extensions in Support of Global Concurrent Optimization”, IETF RFC 5557, July 2009

[RFC5623] E. Oki, T. Takeda, J.L. Le Roux, A. Farrel, “Framework for PCE-based inter-layer MPLS and GMPLS Traffic Engineering”, IETF RFC 5623, September 2009

[sdnrg-layer] E. Haleplidis et al., “SDN Layers and Architecture Terminology”, IRTF SDNRG I-D, draft-haleplidis-sdnrg-layer-terminology-00, July 2013.

[SDN-OMG] Software-Defined Networking Working Group in OMG: http://sdn.omg.org/

[snmp] J. Case et al., “A Simple Network Management Protocol (SNMP)”, IETF RFC1157, May 1990.

[telefonica-dc] C. Liou, O. Turkcu, V. López, J.P. Fernández-Palacios, "An Economic Comparison of Cloud Network Architectures," in Proc. PTC 2013

[tson] G. S. Zervas, J. Triay, N. Amaya, Y. Qin, C. Cervelló-Pastor, and D. Simeonidou, “Time Shared Optical Network (TSON): a novel metro architecture for flexible multi-granular services,” Opt. Express 19(26), B509–B514 (2011)

[YANG] YANG, A Data Modeling Language for the Network Configuration Protocol, IETF RFC6020http://tools.ietf.org/html/rfc6020

8. Acronyms

API	Application Programming Interface
CCAMP	Common Control and Measurement Plane
DCN	Data Centre Network
GMPLS	Generalized Multi Protocol Label Switch
IETF	Internet Engineering Task Force
IRTF	Internet Research Task Force
LSP	Label Switched Path
MLN	Multi-Layer Network
MPLS	Multi Protocol Label Switch
MRN	Multi-Region Network
NIC	Network Interface Card
OCS	Optical Circuit Switching
OPS	Optical Packet Switching
OF	OpenFlow
ONF	Open Networking Forum
OMG	Object Management Group
OTN	Optical Transport Network
PCC	Path Computation Client
PCEP	PCE Communication Protocol
PCE	Path Computation Element
QoS	Quality of Service

RG	Research Group
SLA	Service Level Agreement
SDN	Software Defined Networking
SOA	Service Oriented Architecture
TDM	Time Division Multiplexing
TE	Traffic Engineering
TOR	Top-of-the-Rack
TSON	Time Shared Optical Network
VM	Virtual Machine
WAN	Wide Area Network
WG	Working Group
WSON	Wavelength Switched Optical Network