



Grant Agreement No. 619572

COSIGN

Combining Optics and SDN In next Generation data centre Networks

Programme: Information and Communication Technologies

Funding scheme: Collaborative Project – Large-Scale Integrating Project

Deliverable D1.4

Architecture design

Due date of deliverable: March. 31th, 2015

Actual submission date: April 21st, 2015

Second submission date: February 23rd, 2016

Start date of project: January 1, 2014

Duration: 36 months

Lead contractor for this deliverable: i2CAT Foundation

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

Executive Summary

On the basis of the defined requirements carried out in Tasks T1.1, T1.2, and the analysis of the proposed use cases, this deliverable reports possible network architecture designs with respect to the data and control plane requirements for future DCN.

The first part of the deliverable focuses on the discussion of the key challenges that COSIGN architecture should address and the benchmarking of COSIGN relevant features and concepts while defining a DCN architecture against several State of The Art (SoTA) projects to identify major innovations that COSIGN architecture may bring to future DCNs. The compilation and analysis of the requirements has been fundamental to define COSIGN architecture; therefore, there has been dedicated a section to identify COSIGN architecture assumptions and the relation of the requirements to the overall ambitions of the architecture. The specification of the full set of requirements (reviewing the ones identified in deliverable D1.1 and including some additional ones) has enabled to identify basic functionalities, layers and interfaces suitable for the COSIGN architecture design. The complete list and analysis has been reported in the Appendix.

Based on identified challenges, the assumptions and the analysis of the requirements, the document presents the proposal of data plane network architecture designs. The document describes a set of physical layer scenarios for the short, medium, and long-term which rely on novel features and optical technologies, that are being developed under the WP2 umbrella. The document illustrates the proposal of a control plane layered plane, defining the functional blocks and interfaces and also explaining the innovation they entail. The complete control plane architecture specification and implementation is in the scope of WP3. In order to complete the full picture, the document also describes a preliminary version of the orchestrator (to be further developed in WP4) and the service and application layers. A first version of the interfaces between the different COSIGN architecture layers is provided too.

This deliverable constitutes the final result of the WP1 Task 3, i.e. the specification of the high-level COSIGN DCN architecture based on the requirements compiled in Tasks T1.1 and T1.2. The COSIGN architecture model is intended to serve as a the reference for the different activities of WP2, WP3 and WP4 while implementing the different functional blocks, layers and interfaces and WP5 while guiding the different scenarios and demonstrators.

This document constitutes the second release of the deliverable. In this new version the following inputs have been added:

- A section (section 4.2) summarizing the requirements, assumptions and ambitions (consistent with the requirements in the Appendix) of the COSIGN Architecture.
- A final section (section 8) to indicate how requirements above are addressed/met by the proposed architecture.
- For each Requirement KPI in the Appendix, there have been added (when possible) numerical measurable references reflecting SoTA and/or specific COSIGN values.

Document Information

Status and Version:	Final Version	
Date of Issue:	23/02/2016	
Dissemination level:	PUBLIC	
Author(s):	Name	Partner
	José Ignacio Aznar	i2CAT
	Amaia Legarrea	i2CAT
	Eduard Escalona	i2CAT
	Josep Batallé	i2CAT
	Adrián Rosello	i2CAT
	Daniel Guija	i2CAT
	Bingli Guo	UNIVBRIS
	Shuping Peng	UNIVBRIS
	Reza Nejabati	UNIVBRIS
	Georgios Zervas	UNIVBRIS
	Dimitra Simeonidou	UNIVBRIS
	Sarah Ruepp	DTU
	Michael Berger	DTU
	Anna M. Fagertun	DTU
	Tim Durrant	VENTURE
	Giada Landi	NXW
	Giacomo Bernini	NXW
	Gino Carrozzo	NXW
	Nicola Ciulli	NXW
	Fernando Agraz	UPC
	Salvatore Spadaro	UPC
	Albert Pagès	UPC
	Matteo Biancani	IRT
	Alessandro Predieri	IRT
	Katherine Barabash	IBM
	Anna Levin	IBM
	Oded Raz	TU/e
Edited by:	José Ignacio Aznar	i2CAT
Reviewed by:	Alessandro Predieri	IRT
	Lars Grüner	OFS
Checked by :	Sarah Ruepp	DTU

Table of Contents

Executive Summary	2
Table of Contents	4
1 Introduction.....	6
1.1 Reference Material	6
1.1.1 Reference Documents	6
1.1.2 Acronyms and Abbreviations	6
1.2 Document History	7
2 Key challenges of COSIGN architecture	9
3 COSIGN Related European projects and initiatives.....	12
3.1 European projects	12
3.1.1 GEYSERS	12
3.1.2 LIGHTNESS	12
3.1.3 INSPACE.....	13
3.1.4 PhoxTroT.....	14
3.1.5 SoTA benchmarking.....	15
4 COSIGN architecture: survey of requirements.....	19
4.1 List of requirements.....	19
4.2 Analysis of the requirements in terms of COSIGN main ambitions and objectives	20
4.2.1 COSIGN architecture design and requirements assumptions.....	20
4.2.2 COSIGN architecture design ambition and relation to the requirements	20
5 COSIGN Layered Architecture and Functional Specification	26
5.1 Physical layer – Data Centre Network	26
5.1.1 Taxonomy of COSIGN physical resources features	27
5.1.2 Foreseen Data Centre physical scenarios/architectures	28
5.1.3 Orchestration, Management and Control aspects of physical resources.....	35
5.2 Control layer.....	36
5.2.1 Preliminary functional blocks of the control layer	37
5.2.2 Functional blocks features and innovations.....	38
5.3 Orchestration layer	40
5.3.1 Preliminary layers and functional blocks of the COSIGN orchestration layer.....	41
5.3.2 Functional blocks features and innovations.....	42
6 Architecture inter-Layer Interfaces Preliminary Specification	44
6.1 Control plane and Orchestration layer interface.....	44
6.2 Control plane and data plane interface	45
6.3 Orchestration layer and Service/application layer interface	46
7 COSIGN service delivery model.....	48
7.1 COSIGN service model actors and roles.....	48
7.1.1 COSIGN roles:.....	48
7.1.2 COSIGN actors:.....	49
7.2 Service model phases	50
7.2.1 Service request.....	50
7.2.2 Service deployment	51
7.2.3 Service Operation	52
7.2.4 Service termination.....	52

8	Matching the requirements to COSIGN architecture	53
9	Summary	63
	Appendix: Complete list of COSIGN and Future DCN architectures requirements	65

1 Introduction

1.1 Reference Material

1.1.1 Reference Documents

[1]	COSING FP7 Collaborative Project Grant Agreement Annex I – “Description of Work”
[2]	T. Benson, A. Akella, and D. A. Matlz, “Network Traffic Characteristics of Data Centers in the Wild”, Proc. IMC 2010
[3]	COSING WP3 Deliverable D3.1 SDN framework functional architecture
[4]	COSING WP4 Deliverable D4.1 COSIGN orchestrator requirements and high level architecture
[5]	COSIGN WP1 Deliverable D1.1 Requirements for next generation intra-Data Centres network design
[6]	LIGHTNESS project, “Release of the design and early evaluation results of OPS switch, OCS switch, and TOR switch”, Deliverable D3.1, November 2013.
[7]	LIGHTNESS project, “Implementation results of the OPS switch, the OCS switch, and the TOR switch”, Deliverable D3.2, June 2014.
[8]	LIGHTNESS project, “The LIGHTNESS network control plane architecture”, Deliverable D4.1, July 2013.
[9]	LIGHTNESS project, “The LIGHTNESS network control plane protocol extensions”, Deliverable D4.2, June 2014.
[10]	SDN architecture, ONF technical report, 2014.
[11]	COSIGN WP1 Deliverable D1.2 Comparative analysis of optical technologies for Intra-Data Centres network.
[12]	COSIGN WP1 Deliverable D1.3 Comparative analysis of control plane alternatives.
[13]	G. Booch, et al., “The UML Reference Manual”. AW, 1999, ISBN 1931777446
[14]	Reference Model of Open Distributed Processing, p.17, def 9.4
[15]	http://www.jencotech.com/documents/PolatisOverviewRev1.pdf
[16]	Lightness Project: Low latency and high throughput dynamic network infrastructures for high performance datacentre interconnects. http://www.ict-lightness.eu/
[17]	GEYSERS Project (Generalized Architecture for Dynamic Infrastructure Services). http://www.geysers.eu/
[18]	INSPLACE Project: Spatial-Spectral Flexible Optical Networking: Solutions for Efficient SDM. www.ict-inspace.eu
[19]	PhoxTroT Project: Photonics for High-Performance, Low-Cost and Low-Energy Data Centers, High Performance Computing Systems. http://www.phoxtrot.eu/
[20]	COSIGN WP3 Deliverable D3.2 SDN framework north-bound and south-bound interfaces specification.

1.1.2 Acronyms and Abbreviations

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

AAA	Authentication, Authorisation, Accounting
A-CPI	Application-Controller Plane Interface
AoD	Architecture on Demand
API	Application Programming Interface
AWG	Array Waveguide Grating
Cloud-SP	Cloud Service Provider
CRUD	Create Read Update Delete
DC	Data Centre
DCN	Data Centre Network

D-CPI	Data-Controller Plane Interface
DC-PIP	Data-Center Physical Infrastructure Provider
DC-VIP	Data-Center Virtual Infrastructure Provider
GMPLS	Generalized Multi-Protocol Label Switch
MCF	Multicore Fibre
MPO	Multipath Push On
NBI	North Bound Interface
NCP	Network Control Plane
NE	Network Element
NIC	Network Interface Card
NMS	Network Management System
OCS	Optical Circuit Switching
OF	Open Flow
OPS	Optical Packet Switching
OSPF-TE	Open Shortest Path First – Traffic Engineering
OXS	Fast Optical Switch
PBGF	Photonic Band Gap Fibre
PCE	Path Computation Element
PCEP	Path Computation Element Protocol
PCI	Protocol Control Information
QoS	Quality of Service
RDB	Resource Data Base
REST	Representational State Transfer
RSVP-TE	Resource Reservation Protocol – Traffic Engineering
SBI	South Bound Interface
SDN	Software Defined Networking
SLA	Service Level Agreement
SMF	Single Mode Fibre
SML	Service Management Layer
SoTA	State of The Art
STREP	Specific Targeted Research Project
TDM	Time Division Multiplexing
ToR	Top of the Rack
VDC	Virtual Data Centre
WSS	Wavelength Selective Switch

1.2 Document History

Version	Date	Authors	Comment
0.1	18/11/2014	See the list of authors	TOC first draft
2.6	23/02/2015		Initial Draft after partners contributions
2.7	02/03/2015		IBM, DTU, UBristol contributions to all sections
2.9	12/03/2015		DTU, UBristol input
3.0	24/03/2015		New version edited by i2CAT
3.1	01/04/2015		New version edited by i2CAT
3.2	10/04/2015		New version edited by i2CAT
3.5	15/04/2015		New version edited by i2CAT
3.6	20/04/2015		New version edited by i2CAT
Final version	20/04/2015		New version edited by i2CAT
4.0 (for re-submission)	23/11/2015		New version edited by i2CAT
5.0 (for re-submission)	17/12/2015		New version edited by i2CAT
6.0 (for re-submission)	26/01/2016		New version edited by i2CAT,

			including additional comments and inputs from IRT, IBM, and WP2 partners.
6.5 (for re-submission)	08/02/2016		New version including the review of IRT and OFS.
7.0 (for re-submission)	26/01/2016		New version edited by i2CAT
7.5 (for re-submission)	20/02/2016		Input from UNIVBRIS
Final Version	23/02/2016		Quality check and final submission.

2 Key challenges of COSIGN architecture

The main objective of the COSIGN project is to design and demonstrate a Data Centre Network (DCN) architecture built over innovative optical technologies combined with SDN based network control and service orchestration for future-proof, dynamic, on demand, low-latency and ultra-high bandwidth intra-data centre applications. In other words, to create the blueprint of the future Data Centre network where highly efficient optical interconnect technologies, controlled by the advanced programmable software control plane, satisfy the diverse and dynamic requirements of modern Data Centre workloads as it was stated in the Description of Work (DOW) document [1].

According to recent studies and publications [2], DC providers are experiencing an annual increase of over 70% in the amount of data centre traffic, of which about 50%-75% is east-to-west (server to server and rack to rack) traffic within data centres. One of the main challenges that the Data Centre network architecture will have to overcome in the coming years will be the ever-increasing data centre traffic volumes and new traffic patterns. COSIGN infrastructure aims to provide Data Centre applications with efficient and reliable provisioning and migration mechanisms, while supporting huge scale and high level dynamics.

In order to reach this target, WP1 partners started from a complete description of the three proposed use cases (namely *Multi-tenant Software Cloud*, *Virtual Data Centre (VDC) as advanced IaaS provisioning* and *Advanced infrastructure orchestration and management*) to retrieve functional requirements that future Data Centre Networks architectures would incorporate and adopt. The outcome was compiled in D1.1 [5]. A brainstorming session was carried out in order to extend, refine and assess previous requirements and analyse the way these would match the architecture. The input from IBM and Interoute COSIGN partners resulted key in the identification of the three use cases and their analysis to retrieve the functional aspects of COSIGN DC network architecture. These challenges that arise due to new traffic patterns are common to both IBM and Interoute, industrial partners of the consortium. In parallel to previous requirements analysis, partners working on the COSIGN DC network architecture carried out a comparative analysis of optical technologies and software solutions suitable for intra-DC networks.

Aiming to satisfy future Data Centre network requirements, **optical technologies** are being proposed, introducing novel optical fibres and optical switches. Specifically, high Radix optically enabled Ethernet switches will be developed as ToR switch using mid-board optics achieving reduced cost, low power consumption and increased front panel bandwidth density. Also, the spatial division multiplexing fibres and low latency fibre (e.g., hollow core photonic band-gap fibres, PBGFs) are object of research in different network scenarios to reduce cabling complexity and end-to-end data transmission latency. Large port count and low loss free space 3D-beam steering optical switches will be developed to interconnect ToR switches building up a flatten DCN architecture. At the same time, fast optical switches with highly modular scalability will be investigated at the optical packet level in a synchronous intra-rack network. COSIGN WP2 partners are focused on the study of these optical technologies and together with WP1 have released D1.2 [11] document, that reports the analysis of the optical technologies comparison analysing alternatives for switches, fibres, interfaces and scenarios suitable for the definition of COSIGN data plane architecture.

WP3 aims to integrate the characteristics of all the aforementioned novel optical technologies. The adoption of SDN technologies targets the abstraction and simplification of the underlying network complexity. Based on an **SDN-based control plane** approach, a network controller customized for optical technologies allows DC operators to manipulate logical maps of the infrastructure and the creation of multiple co-existing virtual networks that are not dependant on the underlying technology and network protocols, enabling in this way application level programmability. Coordinated with WP1, COSIGN WP3 is working on the development of the SDN based control plane. A survey about Software Defined Networking (SDN) control plane solutions and a discussion on network virtualization techniques was carried out and compiled in D1.3 [12] and an initial control plane

functional specification has been recently released D3.1 [3]. Currently, the definition of southbound and northbound interfaces and functionalities is being explored and will be released in D3.2 [20].

COSIGN WP4 focuses on the definition and integration of the services of the orchestrator; D4.1 [4] has been delivered and describes the COSIGN view of architectural trends in the area of services orchestration. COSIGN vision is that a unified orchestration layer is required to enable the Data Centre management stack to efficiently manage the underlying resources while providing the demands of multiple heterogeneous workloads. In particular, WP1 has developed a high level design of the **orchestration layer** and its interfaces. It is in the scope of WP4 to develop such layer and interfaces to other COSIGN components.

Therefore, it is fundamental the continuous monitoring and coordination of activities carried out in WP2, WP3 and WP4 to guarantee the coherence between the expected Hardware and Software technologies considered in each of these WPs and the architecture design specification, as an outcome of WP1. Based on previous activities, the action plan to ensure a suitable COSIGN functional architecture description includes:

- Definition of **high level building blocks**, matching reference model and draft the functional description of the architecture layers (M12).
- **Layers definition** and functionalities allocation (M12).
- **Identification of interfaces** to the overall DC orchestration tools so that future DCN architecture fits into the COSIGN vision of the future programmable virtualized Data Centre (M13).
- **Evaluation of the proposed architecture** from entire workflow perspective (M13 and M14)
- **Architecture specification**: Specification of the architecture and functionalities (M15).

This document summarizes the outcomes of this action plan and specifies possible network architecture designs with respect to the data and control plane requirements for future DCN. The remaining sections of this document are structured as follows:

Section 3 of this document lists and provides an explanation and benchmark of recent related European projects that happened in the past or are currently running in the area of DCN, looking at efficiency and optical control plane optimisation for ToR switches. The compilation and analysis of the requirements has been fundamental to define COSIGN architecture; therefore, section 4 has been dedicated to identify COSIGN architecture assumptions and the relation of the requirements to the overall ambitions of the architecture. The specification of the full set of requirements (reviewing the ones identified in deliverable D1.1 and including some additional ones) has enabled to identify basic functionalities, layers and interfaces suitable for the COSIGN architecture design. The complete list and analysis has been reported in the appendix.

Section 5 provides an overview of the COSIGN architecture and its innovative features, the kind of problems faced by the consortium and the proposed solutions. This section also explains the different levels of the COSIGN architecture: physical, control and orchestration. The physical (or Data Plane) level analyses 4 scenarios that have been proposed for the short, medium long and long terms. The control plane architecture consists of 3 layers: abstraction, infrastructure and network virtualization layer. Functional control blocks are defined from a high level point of view and the relationship between them is also specified. The control layer is responsible for the optical resource virtualization, topology management, path computation, optical services provisioning, fault monitoring, etc.; delivering the DCN services according to policies and SLAs set up by the upper layer entities (e.g. the orchestrator) and/or the DCN administrator. The orchestration layer is also outlined providing a description of its functionalities and its relationship with other layers.

Section 6 of this report provides an explanation about the interfaces between control plane and orchestration layer; together with the interactions that may take place between them, such as: customised virtual network requests, DC dynamic optical connectivity requests and information about status and resource occupation in real-time.

Section 7 is related to the service delivery model. The services running on top of COSIGN design are the reason-being of the whole system. It includes the roles, actors and the service model to deliver those services on top of the proposed COSIGN DCN architecture.

Section 8 of the document validates the architecture design by matchmaking the requirements to the different layers, functional components and interfaces that comprise the COSIGN architecture. The final summary provides with concrete conclusions and findings derived from the definition and proposal of the COSIGN architecture as guidelines for the design of future DC networks.

3 COSIGN Related European projects and initiatives

3.1 European projects

Data Centre (DC) infrastructure technology evolves at a fast pace stressed by continual growth data volumes and the lack of capability of current architectures to cope with fast service provisioning and higher consumption needs. More powerful and efficient DCs are required and this is reflected in the plurality of research projects funded by the EU and US. This complex and multi constraint problem has been addressed in various ways including down-to-earth approaches that combine commodity switches with optics to more ambitious all-optical solutions. Here we explore the relationship between COSIGN and four related EU projects: GEYSERS [17], LIGHTNESS [16], INSPACE [18] and PhoxTrot [19].

3.1.1 GEYSERS

3.1.1.1 Description

The FP7 GEYSERS project [17] (Generalised Architecture for Dynamic Infrastructure Services) run from January 2010 to March 2013 and defined an architecture for the composition and delivery of virtual optical infrastructures combined with pools of Information Technology (IT) resources for virtual operators. On top of each virtual infrastructure, GEYSERS defined a Generalized Multi-Protocol Label Switching (GMPLS) and Path Computation Element (PCE) based Network Control Plane (NCP) for the on-demand provisioning of energy-efficient optical connectivity between DCs and a Service Management Layer for the provisioning of cloud services. The virtualization of the optical resources allowed the GMPLS controllers to adopt a unified south-bound interface to interact with and configure the virtual nodes, abstracting the vendor details of the optical technologies.

GEYSERS identified new connectivity paradigms to enable the cross-layer cooperation between cloud management and network control to achieve a joint, converged optimization between network and IT. The selection of the target datacentres were in some cases delegated to the Network Control Plane, with IT and energy aware algorithms running at the Path Computation Element. Additionally, inter-DC network connectivity was established on demand during the deployment of the end-to-end distributed cloud service. This procedure was triggered by the Service Management Layer (SML), which could also choose the preferred path among several alternatives proposed by the Network Control Plane, depending on price, power consumption or other metrics.

3.1.1.2 GEYSERS - COSIGN comparison

GEYSERS project was one of the first projects to investigate virtualization of optical resources and on-demand delivery of virtual optical infrastructures. These concepts can be applied in COSIGN, even if translated and adapted to a completely different scope, i.e. from the core transport network to the infrastructure of a DC network, addressing also new requirements and challenges especially in terms of service dynamicity.

The abstraction of the virtual optical resources investigated in GEYSERS can potentially influence the COSIGN data plane model at the control level. However, the core control mechanisms in the two projects are completely different: while GEYSERS relied on a GMPLS and PCE-based control plane, with distributed controllers and the wide adoption of signalling and routing protocols (i.e. RSVP-TE, OSPF-TE, PCEP), COSIGN will adopt an SDN based control plane with a centralized controller in charge of configuring and program the whole datacentre network.

3.1.2 LIGHTNESS

3.1.2.1 Description

LIGHTNESS [16] (Low latency and hIgh throuGHpuT dynamic Network infrastructures for high pErformance datacentre interconnectS) is a running STREP project funded by the Seventh Framework

Programme (FP7) of the European Commission. The main goal of the LIGHTNESS project is to demonstrate an SDN-controlled all-optical DCN, which combines Optical Circuit Switching (OCS) and Optical Packet Switching (OPS) technologies to provide tailored transport to data centre applications with different requirements (e.g., latency, throughput, etc.). In the proposed flattened optical network infrastructure, the traffic coming from the servers is classified into long-lived data flows, which are transported by means of OCS technology, and short-lived ones, which take advantage of the fast switching capability provided by OPS. The SDN-controller manages this operation, thus enabling the control and programmability of the DCN.

The main objective of the project is two-fold. First, the design and prototyping of the data plane devices is targeted. This includes the development of the OPS node, an Architecture-on-Demand (AoD), which integrates the OCS and OPS transport technologies by seamlessly forwarding data flows coming from the ToR switch to the appropriated device (i.e. to the OPS or the OCS depending on the requirements of the applications) [6],[7]. Second, the project targets at the design and development of an SDN-controller able to manage the heterogeneous data plane in an integrated way [8]. To this end, the OpenFlow (OF) protocol has been provided with the extensions necessary to control the prototyped devices [9]. In addition, a set of OpenFlow (OF) agents have been developed and integrated into the different optical network devices to enable the OF-based configuration of the data plane from the SDN-controller.

3.1.2.2 *LIGHTNESS - COSIGN comparison*

LIGHTNESS investigates the SDN-based integrated control of a flattened hybrid all-optical DCN, based on the OCS and OPS switching technologies.

Although COSIGN proposes a different data plane architecture, the expertise achieved in LIGHTNESS will leverage the control and management of the new data plane. In addition, the knowledge on the SDN paradigm obtained in LIGHTNESS will be extremely useful to overcome the ambitious challenges committed by COSIGN at both the control and orchestration levels. COSIGN approach is more practical, aiming at realistic DC requirements that include cost, equipment readiness, scale, etc., while LIGHTNESS explores the feasibility of all-optical technologies like OPS.

Although the main objective of LIGHTNESS is to provide tailored transport services to DC applications, the project has also proposed some tasks oriented to virtualize the optical DCN infrastructure. However, LIGHTNESS virtualization strategies aim to be an initial approach to be used within the optical DCN and just focused on virtualizing the LIGHTNESS specific network infrastructure. On the contrary, the objectives posed by COSIGN in this regard are much more ambitious, targeting a cross-layer approach to the virtualization of the whole DC infrastructure including IT and network resources orchestrated by the application layer to provide, for example, virtual data centre (VDC) infrastructures and multi-tenancy in a dynamic and efficient way. This includes a plethora of tasks to be conducted from the orchestration and control layers such as optimizing the optical resource utilization, optical resource virtualization, transport service provisioning, VDC provisioning, tenants management, AAA and SLA management, etc. In any case, the experience gained during LIGHTNESS in the control and management fields is expected to be helpful to overcome COSIGN's challenges.

3.1.3 INSPACE

3.1.3.1 *Description*

INSPACE project [18] (Spatial-Spectral Flexible Optical Networking: Enabling Solutions for a Simplified and Efficient SDM) proposes a novel networking approach by extending the established spectral flexibility concepts to the SDM domain and significantly simplifying the super-channel allocation and control mechanisms, by removing current limitations related with the wavelength continuity and fragmentation issues. The new concept utilises the benefits of the high capacity, next generation, few-mode/multi-core fibre infrastructures, providing also a practical short-term solution, since it is directly applicable over the currently installed multi-fibre cable links. The realisation of INSPACE approach is enabled by the development of novel multi-dimensional spatial-spectral switching nodes, which are fabricated by extending the designs of the existing flexible WSS nodes,

incorporating advance mode/core adapting techniques. The concept is further supported by novel processing techniques that minimise the mode/core interference as well as new network planning algorithms and control plane extensions that are enhanced with the space dimension. The INSPACE consortium forms a strong industry driven research team targeting not only the demonstration of the new network concept and its ability to meet the challenges of delivering exponentially growing content over the next twenty years, but also the full exploitation of its potential towards commercialisation. Specifically, INSPACE is proposed with the following objectives:

- Definition of the next generation flexible optical networking functions and capabilities enabled by the introduction of spatial-spectral switching nodes
- Design and development of spatial-spectral WSS-based switching node
- Development of advanced link designs and MIMO digital signal techniques for MCF/FMF transmission impairments mitigation
- Development of routing and resource allocation algorithms for network planning
- Development of control plane extensions and interfaces with the spatial-spectral routing node
- Integration and performance evaluation of the INSPACE spatial-spectral optical networking system
- Techno-economic and energy efficiency evaluation of the INSPACE network solution considering also the migration scenarios from current deployed networks and offered solutions
- INSPACE technology exploitation, dissemination of project results and contribution to standards

3.1.3.2 INSPACE - COSIGN comparison

INSPACE introduces a novel logical hierarchical structure for next generation multidimensional dynamic and elastic optical networks, based on enabling spatial switching and transmission technologies, which are related to the spatial multiplexing scenario (multi-core or multi-element fibre) in COSIGN.

INSPACE focuses on a general network scenario and also explore other network benefits arising from a combined Spatial/Spectral/Signal-coding domain (e.g., mode, spectrum, transmission format, electronic MIMO processing). In contrast, COSIGN aims to develop appropriate optical technologies for data centres network so that future data centres achieve the required large scale at low cost, as well as enabling multi-tenancy services, flexibility, ease of management and operations. Regarding the data plane devices, as well as employing multi-core/multi-element fibre based spatial transmission media, COSIGN also explores ToR switch with on-board optical interface, large scale optical switch and optical fast switch for different network scenario. In addition, control and management technologies are explored in COSIGN, together with the coordination of the architecture and IT resources.

3.1.4 PhoxTroT

3.1.4.1 Description

PhoxTroT [19] is a large-scale research effort focusing on high-performance, low-energy and cost and small-size optical interconnects across the different hierarchy levels in DC and high-performance computing systems: on-board, board-to-board and rack-to-rack. PhoxTroT will tackle optical interconnects in a holistic way, synergizing the different fabrication platforms in order to deploy the optimal “mix&match” technology and tailor this to each interconnect layer. PhoxTroT will follow a layered approach from near-term exploitable to more forward looking but of high expected gain activities.

The objective of PhoxTroT is the deployment of:

- Generic building block that can be used for a broad range of applications, extending performance beyond Tb/s and reducing energy by more than 50%.
- A unified integration/packaging methodology as a cost/energy-reduction factor for board-adaptable 3D SiP transceiver and router optochip fabrication.
- The whole “food-chain” of low-cost and low-energy interconnect technologies concluding to 3 fully functional prototype systems: an >1Tb/s throughput optical PCB and >50% reduced energy requirements, a high-end >2Tb/s throughput optical backplane for board-to-board interconnection, and a 1.28Tb/s 16QAM Active Optical Cable that reduces power requirements by >70%

3.1.4.2 *PhoxTroT - COSIGN comparison*

The focus of the PhoxTroT project is innovation in optical technologies for interconnecting data centre equipment. While the data plane objectives of COSIGN project are similar to those of PhoxTroT, the scope of COSIGN is broader and includes the control and the orchestration layers of the Data Centre Network, especially the innovation in the software stack. Novel interconnect devices and prototypes developed in PhoxtroT are of a great interest to the COSIGN and must be watched in order to ensure their inclusion into the overall ecosystem of the Next Generation DCN envisioned by COSIGN.

Table 1 shows the previously described architecture as a function of several general concepts relevant to the definition of COSIGN architecture. By comparing the different alternatives, it can be seen that most of them have a clear focus on the definition of a suitable data plane by leveraging different optical technologies. Only COSIGN and LIGHTNESS proposes an SDN based control plane, meanwhile it is only COSIGN the one proposing a full orchestrated architecture considering services which imply a joint provisioning of networking and IT resources. Some other parameters relevant to COSIGN (e.g. virtualization, abstraction, automation mechanisms) are also taken into account.

3.1.5 SoTA benchmarking

It can be observed that all the architectures chase the research and adoption of data plane optical technologies although each project focusses on different technologies to be applied for different goals. On the other hand only COSIGN and LIGHTNESS projects show an important focus on the control plane, standing up for an SDN based approach. In terms of management and orchestration, GESYSERS and COSIGN present dedicated effort to the joint provisioning of IT and Network resources. Nevertheless only COSIGN solution clearly aims to define and implement a full orchestrated architecture considering services which imply a joint provisioning of networking and IT resources, including virtualization and abstraction mechanisms while provisioning the services. Additionally, the benchmarking has also considered several parameters which are relevant to the COSIGN architecture as per the benefit they may bring to the services provided. Per each of the related State of The Art (SoTA) it has been identified if any of them have considered virtualization, abstraction, automation mechanisms, coupled provisioning of IT and Network solutions and if the architecture design considers inter-DC environments and energy consumption features.

ARCHITECTURE	GEYSERS	LIGHTNESS	PhoxTrot	COSIGN	INSPACE
CONCEPT					
MAIN GOAL	To define, develop and validate an end-to-end capable of provisioning Optical Network and IT resources for end-to-end service delivery	To design, implement and carry out the experimental evaluation of a high-performance network infrastructure for data centres. Develop a high-performance network infrastructure for data centres	It focuses on high-performance, low-energy and cost and small-size optical interconnects across the different hierarchy levels in data Center and high-performance computing systems: on-board, board-to-board and rack-to-rack. It also tackles optical interconnects in a holistic way, synergizing the different fabrication platforms in order to deploy the optimal “mix&match” technology and tailor this to each interconnect layer.	To define and implement a flat, scalable Data-Centre-Network architecture, empowered by novel optical technologies and SDN based network control and service orchestration platform for dynamic, on-demand, low-latency and ultra-high bandwidth Intra-DC applications.	To explore the additional degrees of freedom in signal multiplexing (offered by the latest advances in both multi-mode and multi-core fibre systems) and to examine the capabilities of the enabling technologies in support of the multi-dimensional networking concept
DATA PLANE	ROADM and OXC technologies	Optical Circuit Switching (OCS) and Optical Packet Switching (OPS) and hybrid top-of-the-rack (TOR) switches technologies	Optical printed circuit board with a throughput greater than one terabyte per second, a high-end optical backplane for board-to-board interconnection with a throughput greater than two terabits per second, and an active optical cable (AOC).	High Radix optically enabled Ethernet switches multi-core fibre (MCF) TOR switches based on large port count and low loss free space 3D-beam fast optical switches	Multi-dimensional spatial-spectral switching nodes
CONTROL PLANE					

	GMPLS and PCE	Unified SDN-based DC control plane	Not covered	SDN paradigm extended to leverage capabilities from high-performance optical elements while developing technology agnostic protocols for Software/user defined routing and control.	Routing and resource allocation algorithms for the planning and operation of spatial-spectral flexible optical networks
MANAGEMENT AND ORCHESTRATION	Partial E2E provisioning of static IT and Network resources. Not really a full orchestrated solution but a service middleware layer (SML).	Not covered	Not covered	Framework for optical network and IT infrastructure abstraction, Virtualization and end-to-end service orchestration is considered.	Not covered
NETWORK VIRTUALIZATION CONSIDERED	YES	YES	NO	YES	NO
DCN AUTOMATED PROVISIONING MECHANISMS CONSIDERED	YES	YES	NO	YES	NO
ABSTRACTION OF VIRTUAL OPTICAL RESOURCES CONSIDERED	YES: LICL software component is responsible for resource abstraction	Partially covered	Not covered	YES	YES: based on an SDN control plane module present in the SDN controller
INTER-DC					

SOLUTION	YES	NO	YES	NO	NO
COUPLED PROVISIONING OF OPTICAL NETWORK AND IT RESOURCES	YES	NO	NO	YES	NO
ENERGY CONSUMPTION CONSIDERED	YES: Control plane extensions in OSPF and path computation algorithms to achieve it.	YES: energy efficiency of data centre network infrastructures	YES: focuses on high-performance, low-energy and low-cost optical interconnects across the different hierarchy levels in Data Center and high-performance computing systems.	YES: Data plane solutions expect to greatly reduce energy consumption. It is also a challenge that COSIGN provides aware network + IT services in intra-DC environments.	YES: It is carried out an energy efficiency evaluation of the INSPACE network solution

Table 1: Benchmarking of COSIGN related initiatives.

4 COSIGN architecture: survey of requirements

Previous SoTA review of related COSIGN initiatives provides us with a good perspective on the added value that COSIGN architecture is committed to bring starting from the technologies and challenges stated in the DoW document. COSIGN WP1 working group together with WP2, WP3 and WP4 partners have also retrieved a complete list of requirements that it would be desirable to have in future DCN solutions. This survey extends the requirements that were already proposed in D1.1 [5]. To this target, a more extensive analysis of the use cases and a dedicated brainstorming session during the Bristol F2F meeting was carried out. The structure of the survey (business requirements, service requirements and infrastructure requirements) proposed in D1.1 has been preserved.

4.1 List of requirements

It is important to remark, that we aim to define the architecture of future DCNs Network, so that independently of the feasibility to implement the requirements or not in today's DC networking solutions, it is interesting to have them in the survey. The objective is therefore to provide a complete list of requirements that enable the later definition of the functionalities across COSIGN architecture layers.

To ensure a suitable full analysis of the requirements and the implications they may impose to the COSIGN architecture design, there has been provided a template (see Figure 1) aiming to cover as many aspects as possible to make the best of each of them. The proposed template is organized as follows:

1. Requirement description.
2. Identify the KPI(s) that will enable us to measure that we are accomplishing such requirement. Both quantitative and qualitative KPIs have been identified. For the quantitative ones, some COSIGN partners have provided reference SoTA and expected future numerical values (concrete number or ranges) when possible. It must be taken into account that the numbers for some KPIs are not available or constitute private information.
3. Impact: Which impact do these business requirements have? (HIGH, MEDIUM, LOW) for NG-DC design?
4. Novelty: Which is the novelty of the Business requirements? (HIGH, MEDIUM, LOW) for NG-DC design?
5. Required Functionalities: Provide (if possible) suggestions about possible functionalities required to accomplish each requirement and the architecture layer (Data, Control, Management, Application) in which such functionalities could be allocated.
6. Layer: Which layers are supposed to be affected and should take this requirement into account.
7. Related Data Plane (DP) requirements: WP2 needs to know the consequences (physical requirements) that upper layer requirements may impose on the Data plane. This is really an important field of the template that may contribute to overcome and minimize the gap between SW and HW layers while defining interfaces and workflows among them in WP2 and WP3 specially. Not only the requirements coming from the SW layers have a direct impact on the HW layer, thus, each of the requirements has been tagged with an specific colour depending on whether the requirement has a direct impact on the physical HW layer or not. The colour code is as follows:
 - Red: There are no direct HW implications.
 - Yellow: It is not clear if there are direct HW implications or not.
 - Blue: There are HW implications.

<Requirement code>	<Name> <Description>
KPIs	
Impact	
Novelty	
Functionalities	
Layer	
Related DP requirements	

Figure 1 Template for the survey of COSIGN architecture requirements.

The survey includes more than 50 requirements some of which have been already utilised by WP1 (see D1.2 [11]) and WP3(see D3.1 [6]) in the process of defining the specification of the building blocks, interfaces and layers of the HW infrastructure, the control plane and orchestration levels as well as for the specification of the COSIGN architecture. The complete list of requirements can be found at the end of this document as an appendix.

4.2 Analysis of the requirements in terms of COSIGN main ambitions and objectives

The main goal of COSIGN is to define and implement a flat, scalable DCN architecture, empowered by novel optical technologies SDN based network control and service orchestration for future-proof dynamic, on-demand, low-latency, and ultra-high bandwidth intra-data centre applications [1]. To this target, there were identified a number of objectives to accomplish and limitations to overcome.

4.2.1 COSIGN architecture design and requirements assumptions

COSIGN architecture design starts from the **assumption** that a cross-layer interaction among the proposed layers is fundamental to make the most of the proposed optical technologies, the control and management mechanisms to support the new technologies while progressively migrating from legacy to all-optical DCNs. The list of requirements has been derived based on the analysis of the use cases described in D1.1, the technologies brought to the project and the expertise of the COSIGN team. The list aims to address and satisfy the COSIGN proposed reference use cases that were chosen to validate the architecture. The list of, do not only comprise requirements that to be satisfied during the COSIGN project time-being, but also some additional ones that should be able doable in the long term.

4.2.2 COSIGN architecture design ambition and relation to the requirements

The **ambition** of the COSIGN architecture design covers several items around the main objective: In terms of application and service provisioning, the **fast provisioning of new Cloud-based services** (bringing up services in minutes that usually takes months) is a must. Additionally, COSIGN partners took as reference three key flag use cases that have enabled to identify objectives and requirements at the application and service level. From the perspective of DCN orchestration and management the ambition of COSIGN is to **gain service control and management**, moving from manual, cumbersome and error-prone systems towards more **automated and optimized** ones, enabling flexible workload **management**, on-demand resource allocation and scalable and **efficient virtualized connectivity**. The adoption of a unified DCN control plane based on the SDN concept aims to enable a **converged IT and Network** solution. Resource virtualization and enhanced abstraction mechanisms are also in the scope of COSIGN. Finally, at data plane level, it is the ambition of COSIGN to move from traditional fat-tree designs to flattened DC network architectures, based on the adoption of high

performance optical solutions and fibres to enable **the DCN capacity grow, high throughput and low latency infrastructures**.

The following table enumerates the requirements, identifies the impact and importance they entail for future DCN designs and evaluates the relation of each of them against the most relevant COSIGN ambitions and objectives as it was stated in the DoW. The requirements that have been retrieved are consistent and aligned with the different topics that COSIGN aims to improve by proposing the design. From the table, it can be observed that “*Gain service control and management*” and “*increase network performance*” ambitions are represented by a larger number of requirements. “*Gain service control and management*” is consistent with the cornerstone idea of bringing an SDN based network control plane (WP3) that is orchestrated (WP4) with the IT resources. Integrating SDN technology enables a better automation and control of the DC network, whereas the orchestration with the IT part of the solution brings a better resource management while providing services. The basis to “*increase network performance*” is grounded on data plane requirements and the technologies proposed for the data plane (WP2). The SDN control plane and the orchestration layer propagate such enhancement of the network performance towards the services.

	IMPACT	OVERALL COSIGN AMBITION AND OBJECTIVES						
BUSINESS REQUIREMENTS		Automation and optimization of management and control procedures	Increase network performance	Guarantee high capacity and throughput	Fast real time and time-sensitive service provisioning	Low latency	enabling innovation by means of efficient network virtualization	Gain Service control and management
R-BUS-01 - Resource and Services provisioning	High	X					X	
R-BUS-02 - Elasticity	High	X						
R-BUS-03 - Monitoring Tools	High		X	X	X			
R-BUS-04 - Virtual infrastructure control APIs	High							X
R-BUS-05 - Accountability, billing and SLA management								
R-BUS-06 - Privacy and security	High							
R-BUS-07 - Virtual infrastructure administration and management interface	High							X
R-BUS-08 - RAS	High	X						
R-BUS-09 - Network management Integration	High			X				X
SERVICE REQUIREMENTS		Automation and optimization of management and control procedures	Increase network performance	Guarantee high capacity and throughput	Fast real time and time-sensitive service provisioning	Low latency	enabling innovation by means of efficient network virtualization	Gain Service control and management
R-SERV-1 - Seamless Network Resources Provisioning	High	X			X			
R-SERV-2 - Advanced Network Services Provisioning	High						X	
R-SERV-3 - Service level monitoring	High		X	X	X			
R-SERV-4 - Adapting to application/service	High							X

Combining Optics and SDN In next Generation data centre Networks

dynamicity								
R-SERV-5 - QoS provisioning	High		X	X	X	X		
R-SERV-6 - Real Time Control	High				X			
R-SERV-7 - Big data support	High		X					
R-SERV-8 - Secure data management	Very High							
R-SERV-9 – Security management	Very High							
R-SERV-10 - Self-service management interfaces	Medium	X						X
R-SERV-11 - Multi-tenant isolation	High		X				X	
R-SERV-12 - Service Traffic Profile	Medium					X		
R-SERV-13 - Generalized information model	Medium						X	
R-SERV-14 - Service connectivity disjunction	High							
R-SERV-15 - VDC optical BW control	High			X			X	X
R-SERV-16 - Service Calibration	High		X					X
R-SERV-17 - Intelligent/selective physical failure awareness	Medium							
R-SERV-18 - equipment inventory services tool	Medium							
MANAGEMENT REQUIREMENTS		Automation and optimization of management and control procedures	Increase network performance	Guarantee high capacity and throughput	Fast real time and time-sensitive service provisioning	Low latency	enabling innovation by means of efficient network virtualization	Gain Service control and management
R-MGMT-01 - Element Management	High							X
R-MGMT-02 - Operator	High		X			X		X

Network Management mechanisms								
R-MGMT-03 - Self-Service multi-tenancy	High						X	X
R-MGMT-04 - Automatic resources discovery	Medium	X						X
CONTROL PLANE REQUIREMENTS		Automation and optimization of management and control procedures	Increase network performance	Guarantee high capacity and throughput	Fast real time and time-sensitive service provisioning	Low latency	enabling innovation by means of efficient network virtualization	Gain Service control and management
R-CP-01 - Automated provisioning of intra-DC connectivity	Medium	X						
R-CP-02 - Support for customizable network services	Medium				X	X	X	
R-CP-03 - Dynamic Support for multiple automated connectivity paradigms	Medium							
R-CP-04 - Multi-layer operation of COSIGN data plane optical technologies	High		X				X	
R-CP-05 - Resource usage optimization	Medium		X					
R-CP-06 - DCN Elastic intra-DC connectivity	Medium			X				
R-CP-07 - Network monitoring	Medium		X	X				X
R-CP-08 - Programmable APIs	High							X
R-CP-09 - Support of network service monitoring and accounting	Medium							
R-CP-10 - Support of scheduled network connectivity	high	X	X					X

Combining Optics and SDN In next Generation data centre Networks

R-CP-11 - Network service resilience	Medium			X				
R-CP-12 - Multi-tenant isolation	High		X					
R-CP-13 - Interoperability with existing cloud management platforms for unified DC control.	Medium							X
R-CP-14 - Integration with external connectivity services (inter-DC)	High							X
R-CP-15 - Dynamic DCN reconfiguration for optimization strategies	Medium		X					
R-CP-16 - CP architecture in support of scalable DCNs and network traffic	High			X				X
R-CP-17 - scalability	High			X				
DATA PLANE REQUIREMENTS		Automation and optimization of management and control procedures	Increase network performance	Guarantee high capacity and throughput	Fast real time and time-sensitive service provisioning	Low latency	enabling innovation by means of efficient network virtualization	Gain Service control and management
R-DP-01 - Capacity	High			X				
R-DP-02 - Latency	High					X		
R-DP-03 - Reconfigurability/Flexibility	Medium		X					
R-DP-04 - Resiliency and HA	High		X	X				
R-DP-05 - Traffic isolation	High		X					
R-DP-06 - Scalability and extensibility	High			X				

Table 2: Matching of the requirements against the overall COSIGN ambitions and objectives.

5 COSIGN Layered Architecture and Functional Specification

This section provides the network architecture design of the reference COSIGN architecture. The architecture is discussed by identifying, describing and analysing the different layers, functional blocks and interfaces as well as the features and innovations that each layer brings to COSIGN infrastructure level. The physical layer (a.k.a, Data Plane layer) presents the scenarios and technologies that the consortium aims to deploy. The physical level analyses the 4 possible scenarios that have been proposed from a short, medium and long-term point of view. The Control Plane architecture is based on SDN technology and consists of multiple layers, with functional control blocks defined from a high level point of view with the relationship between them. The orchestration layer has been described together with its responsibilities and the main functionalities to coordinate the underlying COSIGN layers with the application/server layer while computing data transfer paths and configuring data plane devices to enforce the computed paths. Figure 2 shows the overall COSIGN layered architecture.

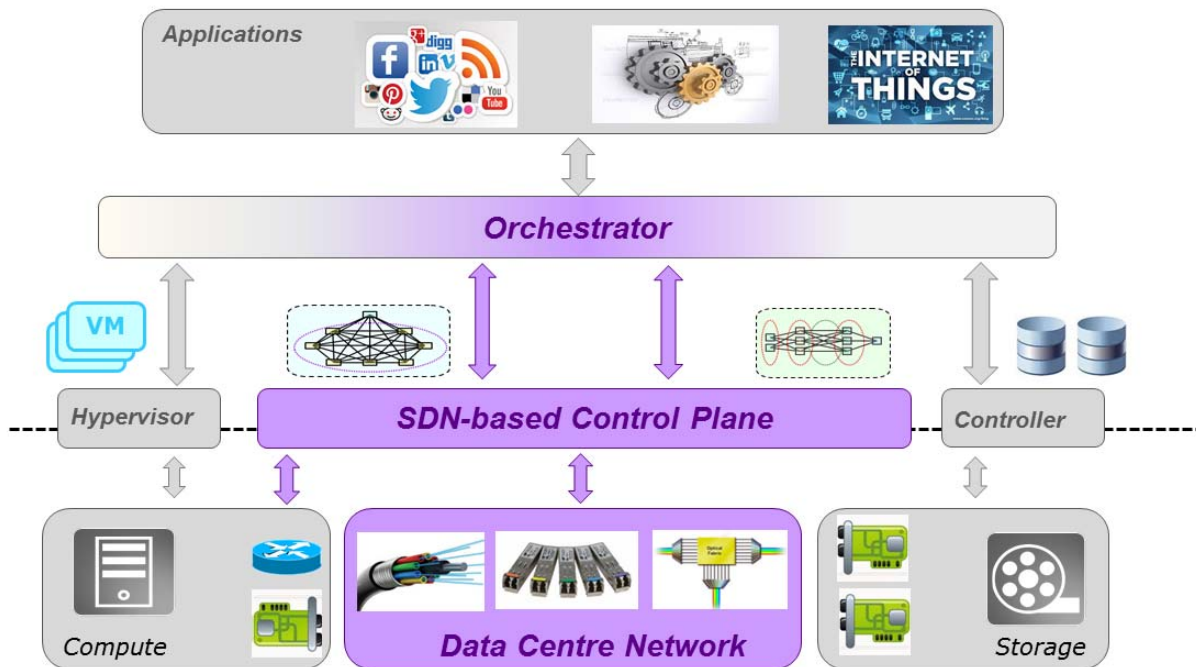


Figure 2 Overall DCN COSIGN architecture.

5.1 Physical layer – Data Centre Network

COSIGN novel optical components/devices are conceived to bring a number of added value features to DCN cloud environments aiming to meet the data plane requirements elaborated in D1.2 [11], develop the three short, medium and long terms physical scenarios and demonstrate them. These features include:

- Large and scalable top-of-the-rack (ToR) switches.
- Scalable, ultra-high capacity and ultra-low latency interconnections.
- High spatial dimensioning to support large port density as well as scalability of the architecture utilizing new data-com fibre technologies such as multicore fibres.

The novel features employed in COSIGN, bearing in mind the data plane requirements elaborated in D1.2 [11], will be introduced in this section. Three DCN physical architectures integrated with different technologies are proposed. In particular, short-term scenario consists of an electronic Ethernet switch with mid-board optics following the traditional layered architecture; its performance will be benchmarked to justify the benefit of involving the proposed optical technologies. The mid-term scenario proposes a flatten structure with large scale switch combined with the electronic ToR switch. Last but not least, the long-term scenario aims to provide all optical server-to-server interconnection by introducing optical NIC and optical ToR switch.

5.1.1 Taxonomy of COSIGN physical resources features

The following technologies are conceived to build a flexible, scalable and cost efficient DCN architecture in COSIGN:

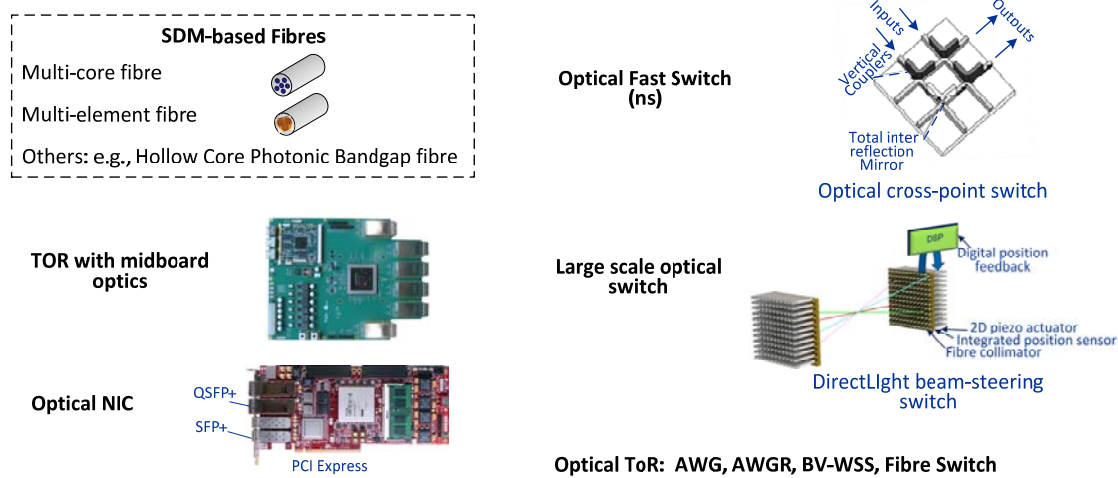


Figure 3 Devices to be employed in COSIGN DCN architecture

- High Radix optically enabled Ethernet switches using mid-board optics to allow reduced cost, power consumption and increased front panel bandwidth density.
- High spatial dimensioning and low latency optical fibre: the use of spatial division multiplexing fibres (e.g., multi-core fibre (MCF)) will reduce cabling complexity and allow dynamic build of logical networks on top of the installed physical infrastructure. In addition, improved end-to-end data latency can be achieved by replacing standard single mode fibres (SMFs) with hollow core photonic band-gap fibres (PBGFs).
- Highly-scalable core switching nodes: A flattened software defined meshed network topology is preferred for interconnecting ToR switches based on large port count and low loss free space 3D-beam steering optical switches. In this way, cluster switches can be eliminated by reducing east-west communication delays to reduce physical reconfiguration requirements. Optical circuit switch will also support port and matrix partitioning (hardware virtualization) to assist SDN based virtualization and slice isolation.
- Fast optical switch (OXS): fast optical switches with highly modular scalability will be investigated for optical packet and TDM switch in a synchronous intra-rack mesh network. This will double or quadruple the current scale.
- Optical NIC: NIC communicates with server through PCI-E interface (read/write data from/to RAM and disk) and could send out the data through optical interface (e.g., SFP+ and QSFP+) in various ways (e.g., TDM and Ethernet).

- Optical ToR: Optical switch sits on top of the rack and integrates/switches optical connection from server to core switch, which is implemented based on AWG, WSS or fibre switch.

All the features described above and properly combined will mark very significant innovations in optical switching technologies.

5.1.2 Foreseen Data Centre physical scenarios/architectures

The different scenarios are designed based on availability of different technologies. Specifically, in the short-term scenario, the TU/e ToR electronic switch with on-board optics is employed to replace the traditional ToR switch but still follow the layered architecture. Therefore, the short-term solution is designed to align with the current DCN architecture and aims to increase the port count of a single ToR switch hence reducing the total number of switches in the entire network.

In the mid-term scenario, combined with TU/e ToR electronic switch, a more flattened architecture is designed with large scale optical cluster switch. Comparing with the layered structure, an optical-enabled flattened DCN architecture could avoid the bandwidth bottleneck caused in the aggregation/core switch by increasing the connectivity of ToR switch and further reduce the traffic latency. In this scenario, the scalability of the optical switches is critical to meet production network requirements and control/management scheme also needs to be investigated.

In the long-term scenario, data intensive application will overwhelm the DCN and more bandwidth will be required to reduce cost and power, coming up with efficient VM-to-VM communication in cloud service. Instead, the power consumption of electronic switch with high processing capability would be an important limitation. Therefore, in the long-term scenario, besides large scale optical cluster switch and optical fast switch, optical NIC in server and optical ToR switch are employed to enable high capacity server-to-server optical connections, as well as much lower latency.

A. Short-term Scenario

The short-term scenario is built in a layered architecture (e.g., fat-tree) with electronic ToR switches. The design, fabrication and testing of a ToR switch will be carried out in this scenario. The hardware configuration of the ToR switch will support a dynamic reconfiguration of any number of channels. The on-board optical interfaces developed will be coupled to the ToR front panel through MPO (Multipath Push On) multi-fibre connectors providing maximum bandwidth density on the ToR front panel. The ToR design, having optical interfaces placed nearby the switch fabric IC, delivers a simplified and cheaper printed circuit board.

The performance of this scenario will be investigated in a benchmarking exercise to show the potential advantage of optical switch and SDN technology.

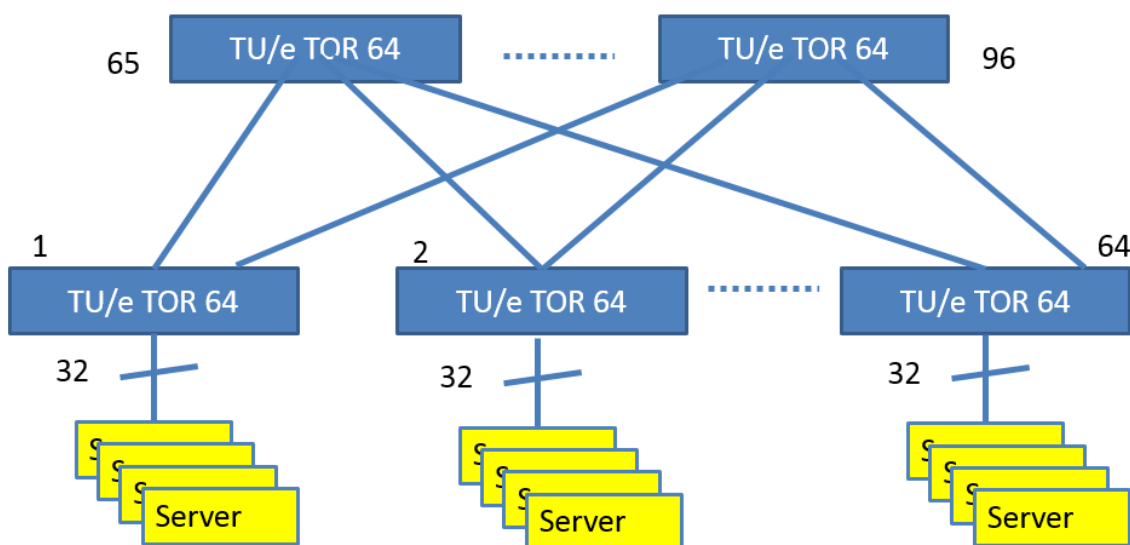


Figure 4 Short-term DCN architecture

The short-term architecture scenario already entails a number of benefits driven from the data plane towards the upper layers. The TU/e ToR electronic switch with on-board optics is employed to replace the traditional ToR switch but still follow the layered architecture. The switch is still essentially an electronic switch but several major changes have been made to it on both HW and SW aspects in order to enable great savings in power consumption, to save costs by reducing the number of optical modules, to reduce the switch size and to be able to run OpenFlow, which means they can be controller using an SDN controller. This will lower the power consumption of the switching ASIC (SW change).

Figure 5 shows the different benefits that the short-term scenario will bring to the different layers of the architecture thanks to the adoption of the TU/e ToR electronic switch with on-board optics.

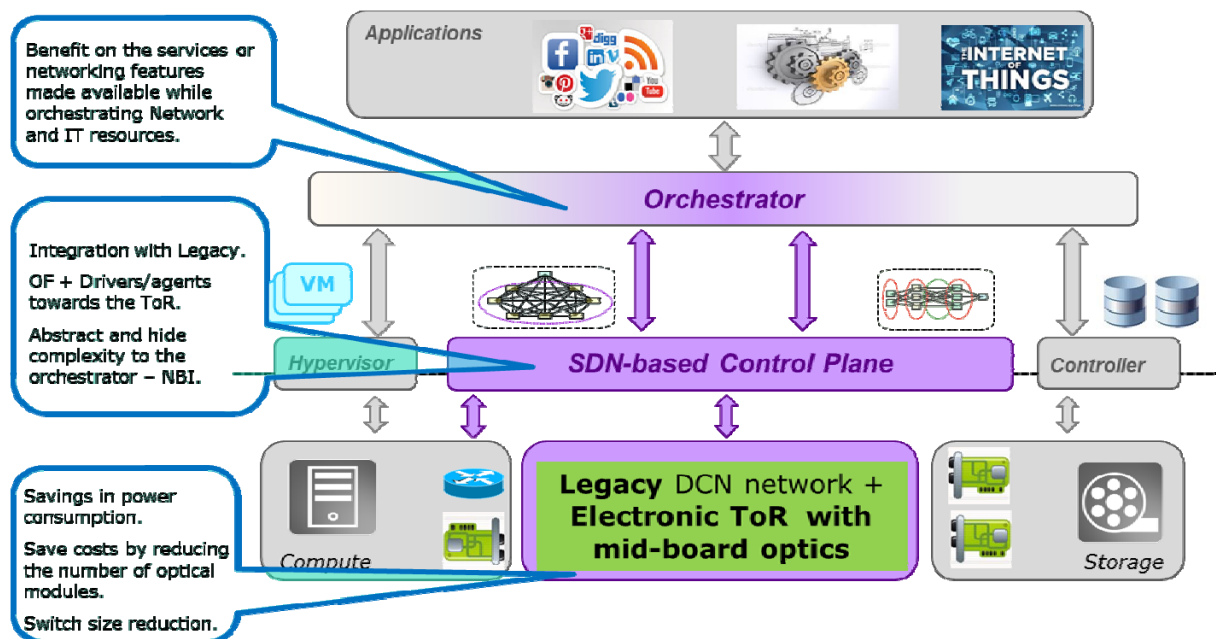


Figure 5 per layer expected short-term DCN architecture benefits.

B. Medium-term Scenario

Medium-term scenario shown in Figure 6 is designed as reconfigurable flat DC network based on high radix optical circuit switches in combination with electronic ToRs. High radix optical switch is meant to improve the flexibility of ToR-to-ToR interconnection.

In COSIGN, we aim to scale the DirectLight optical matrix [15] switch size by at least 2x beyond the current SoTA with a target of 500x500 single mode fibre ports and also advance matrix resilience, energy efficiency and reduced cost per port. The scalability of the high performance OXS switches will be increased to 8x8 monolithically, potentially 16x16. Additionally, the ToR switches and network switches are interconnected through spatial division multiplexing (SDM) based or low latency optical fibres, as well as the Mux/Demux components (as shown in Figure 7).

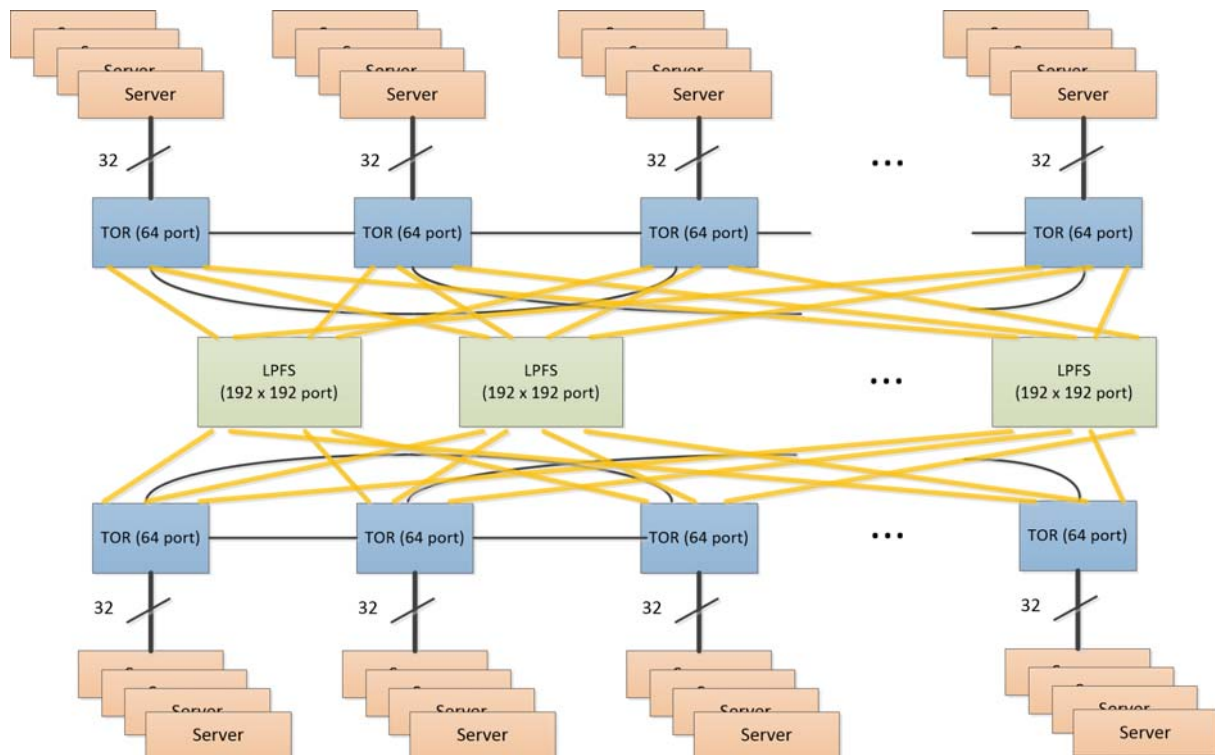


Figure 6 Medium-term DCN architecture (TOR – Top Of Rack Optical Ethernet Switch, LPFS – Large Port Count Fibre Switch).




Figure 7 SDM-based fibres and Mux/Demux

This architecture supports a meshed ToR-to-ToR circuit connection through a Large Port Count Fibre switch (LPFS) and also there are some direct connection between ToRs (which improves the adaptability of various traffic features, especially for the short-lived and mice flows). The LPFS and ToR switch will be configured on a per-flow basis. Also, servers with built in optical NICs could forward traffic in a time-slot way predefined by the control plane and OXS could switch optical TDM packets in ns scale. In this way, the traffic between servers within the same rack will be offloaded from the ToR switch whose load will be reduced. So, comparing with the short-term scenario, we could have benefits in terms of traffic latency and network throughput.

Thus, the mid-term scenario includes all of the advantages of short-term scenario but offers an additional optical overlay network which can be configured in a dynamic manner. Since the optical overlay switch is also SDN controlled, one can orchestrate the TORs and overlay switches together for specific use cases. Comparing with the layered structure, an optical-enabled flattened DCN architecture could avoid the bandwidth bottleneck caused in the aggregation/core switch by increasing the connectivity of ToR switch. Figure 8 shows the identified benefits that the mid-term scenario, by means of the optical overlay network, is expected to bring to future DCN network architectures

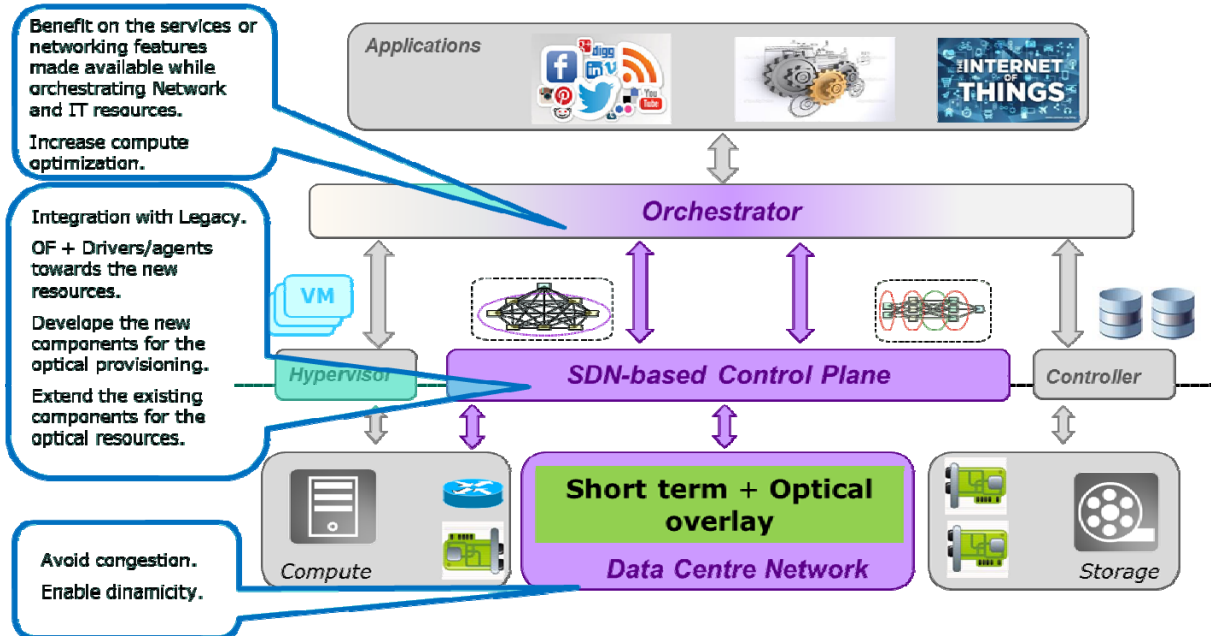


Figure 8: Per layer expected mid-term DCN architecture benefits.

C. Long-term Scenario A

In the long-term scenario A, optical NIC and optical ToR are utilised. With the above mentioned technologies, COSIGN designs a scalable, ultra-high capacity and low latency optical interconnection for intra-DCs communication. More specifically, the proposed intra-DC data plane architecture provides three different kind of server-to-server optical connection to adapt different service/communication requirements: optical TDM (OTDM) connection, optical packet/burst switching (OPS/OBS) connection and optical circuit connection (OCS).

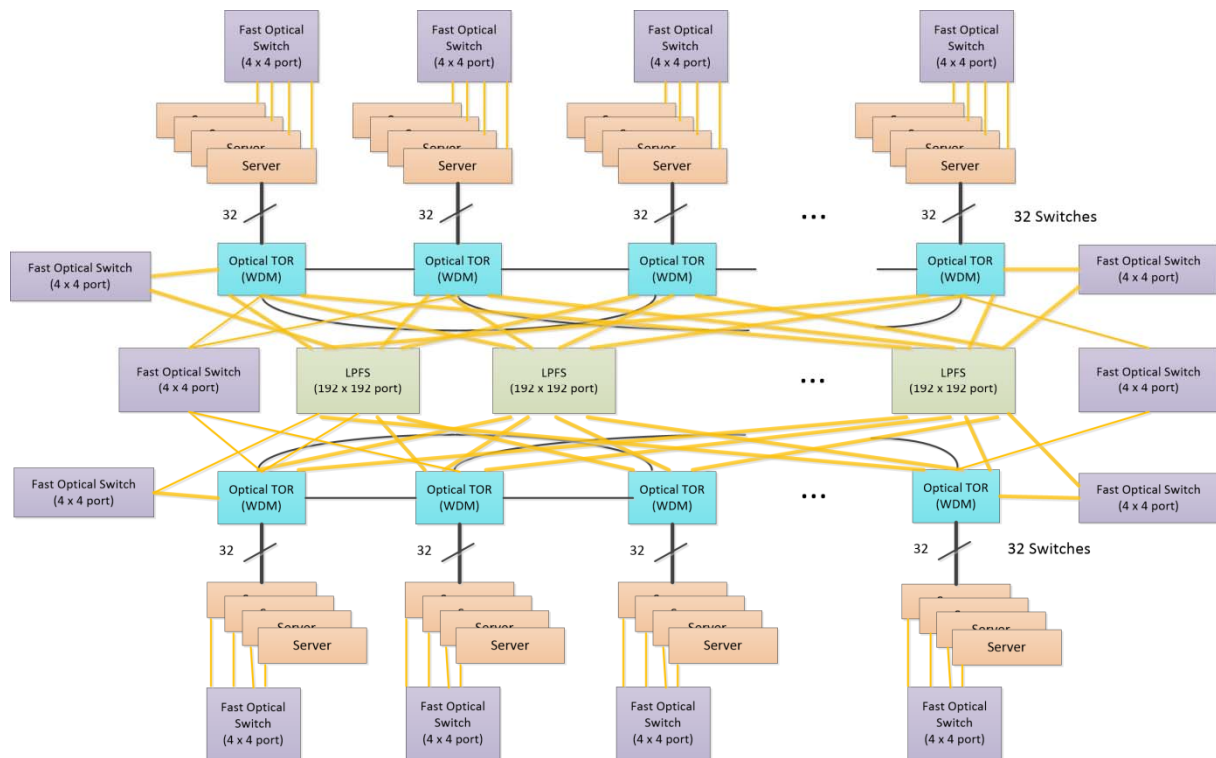


Figure 9 - Long-term DCN architecture (LPFS – Large Port Count Fibre Switch, WDM – Wavelength Division Multiplexed)

Figure 9 depicts NIC with optical interface (Optical NIC) that will be widely deployed in future DC servers, and optical ToR switch are being used to provide circuit connection for intra- and inter-rack communication demands. In this case, a WMD optical ToR switch can be implemented with different devices depending on the networking scenarios, i.e., AWG(R) or Bandwidth Variable-WSS (BV-WSS). Complementary, fast optical switch can provide OTDM or OPS/OBS based switching to efficiently accommodate the short-term and bursty traffic between servers. Compared with an OPS/OBS based connection, the OTDM connection provides contention free transmission with guaranteed QoS. Furthermore, the Optical NICs need to be fully-programmable with hybrid function, i.e., supporting packet assembly under different working modes - circuit and OTDM/OPS/OBS. Figure 9 shows inter-rack connection that is implemented by a direct rack-to-rack circuit connection through a large scale fibre switch or OTDM/OPS/OBS connection with a fast optical switch. Furthermore, structure with cascaded fast and large scale switch also could be considered with improved connectivity re-configurability. In this way, according to the traffic flow characteristics, flexible server-to-server optical connection could be provided by a flattened intra-DC network. Also, the introduction of parallel fibre connectivity overcomes great challenges of the cabling engineering, and COSIGN introduces spatial division fibres (which have higher port/bandwidth density and cost efficiency) to resolve this problem, interfacing via a spatial mux/demux component. At the same time, new alternative fibres with attractive characteristics (e.g., lower propagation latency) are also investigated.

D. Long-term scenario B

In this final scenario, we are moving beyond the fat-tree architecture and aim at evaluating different network configuration options. A ring network structure, like the one shown in Figure 10, is preferable to achieve robust connectivity. Switching at various networking layers is necessary to address the demand in future DCs for high bandwidth, high flexibility, scalability and robustness. In order for a DC to be maintainable, extensible, scalable and fault tolerant, it is required that physical infrastructure is architected in a well-defined modular way.

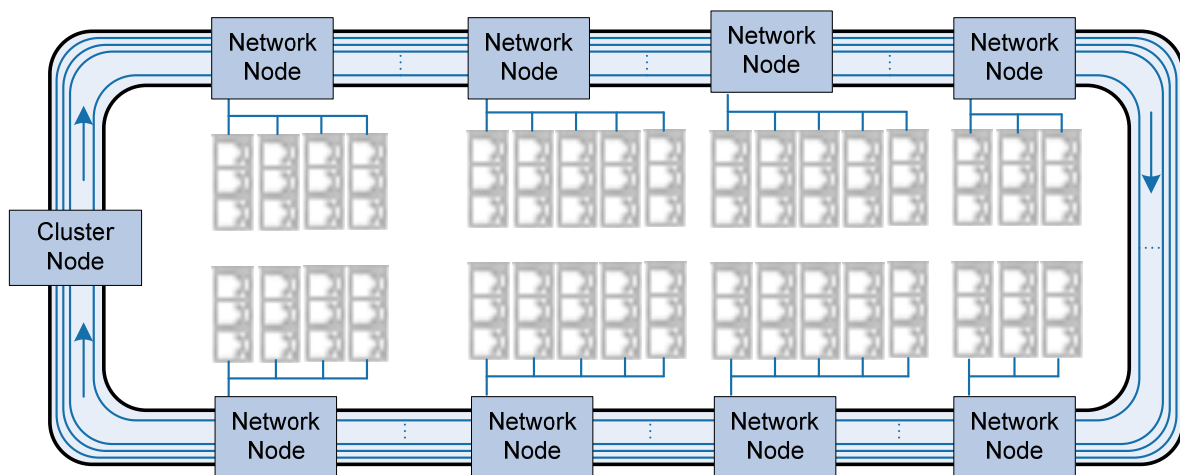


Figure 10 Generic data centre ring network

Each network node serves a number of racks, where each rack contains a number of servers (e.g. running Virtual Machines) and a ToR switch. ToR switches are connected to network nodes through fibres or multicore fibres, when required. The ring network interconnects network nodes through multicore fibres, thus forming a cluster. Clusters can be interconnected through the Cluster node, often duplicated for resiliency and typically connected in a ring, tree, star or mesh topology. An example of a self-healing ring cluster interconnect is illustrated in Figure 11.




Figure 11 Cluster nodes interconnected by self-healing ring

Figure 12 provides a more detailed view of a Network node, where all switching layers are represented. At the lowest layer, switching at multi-core granularity is performed. On the next level,

fibre switching is executed, which allows switching between cores in the multicore fibres, as well as add/drop of fibre cores to the wavelength switch. The next level up performs switching at wavelength level and may allow wavelength conversion. The higher top layer provides TDM and/or packet switching functionality, allowing sub-wavelength granularity by using circuit or packet based paradigms. The reasoning behind TDM capabilities requirement is to allow all-optical top-to-bottom transmission using optical timeslots. However, packet switching could also provide that type of connectivity and accommodate traffic burstiness on smaller time-scales.

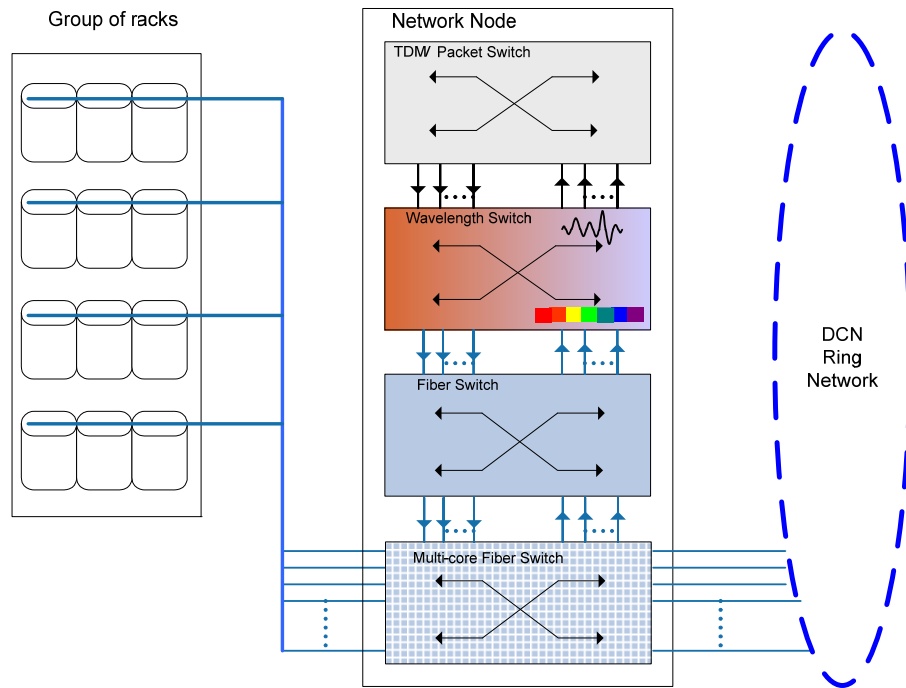


Figure 12 Network node

Each network node in the ring is equipped with different layers of switching granularity, such that switching capabilities do not need to be available at all nodes. Figure 13 illustrates end to end connectivity for a TDM connection which traverses several switch nodes. To optimise energy consumption and cost, it is desirable to switch at the lowest possible layer and only perform switching at higher granularity when and where necessary.

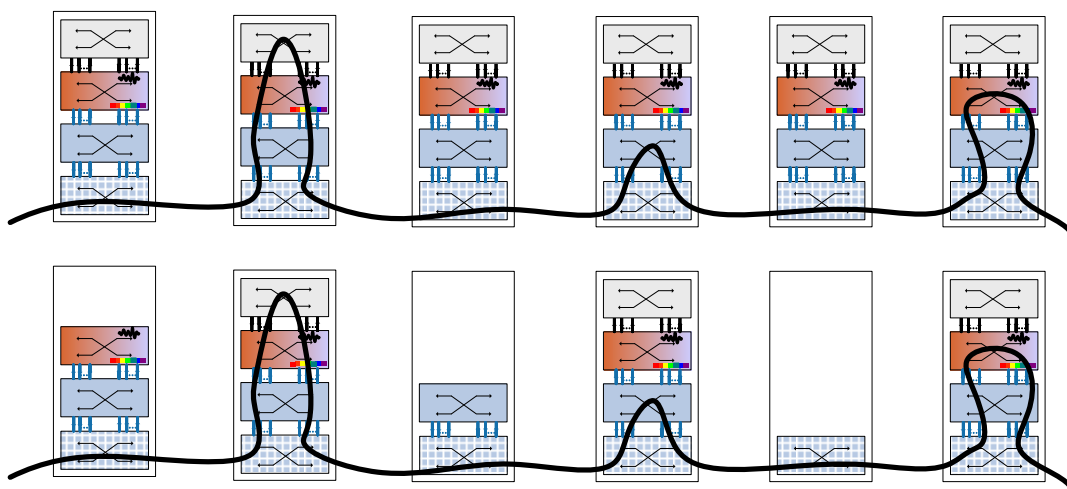


Figure 13 Example connection crossing several layers

It is not feasible to provide a figure explaining the benefits for the various layers for the long term scenarios. The ring-scenario which is the one I know best is focusing at realizing an all-optical data plane. It is the target to show high throughput, low latency and a simple cabling scheme. The demonstrations will mainly focus on data plane advances.

As a summary, Table 3 presents a roadmap of our data plane infrastructure demonstration.

Scenarios	2015	2016
Short-term scenario	Finish the first prototype of electronic ToR switch and system demonstration.	Improve the function of the device.
Medium-term scenario	Besides the ToR switch, finish the first prototype of novel fibres, fast and large-scale cluster switch.	Improve the function of each device/component and demonstrate the integrated scenario.
Long-term scenario A	Besides the devices/components in medium-term scenario.	Improve the function of each device and demonstrate the integrated scenario.
Long-term scenario B	Develop the hierarchical optical network node and demonstrate its functionality.	Demonstrate ring topology comprised of hierarchical optical network nodes.

Table 3 Roadmap of the scenario demonstration

5.1.3 Orchestration, Management and Control aspects of physical resources

To enable the SDN-based control and management of the physical resources, novel features of the physical infrastructure should be abstracted and exposed to the control plane in a generic way. The features to expose include:

- different features of optical switches
 - large scale cluster switch: expose as a port based switch
 - fast optical switch: expose as a time slot-based or optical label-based switch
- optical channels: optical TDM or OCS channel
 - channel abstraction/definition
 - TDM: path + {start_time_slot and granularity}
 - OCS: path + bandwidth
- monitoring information
 - hardware failure event, e.g., optical port
 - performance of optical channel/path/link
 - OpenFlow software interface.

5.2 Control layer

The COSIGN control layer is based on the Software Defined Networking (SDN) principles, which allow to cope with the dynamic and highly variable profiles of the DC traffic through a fast and coordinated reconfigurations of the DCN data plane. The DCN programmability is key to achieve the effective convergence between the DC cloud environment and the network which interconnects the servers, through automated procedures which integrate the deployment of the VMs with the instantiation of the associated virtual network topologies.

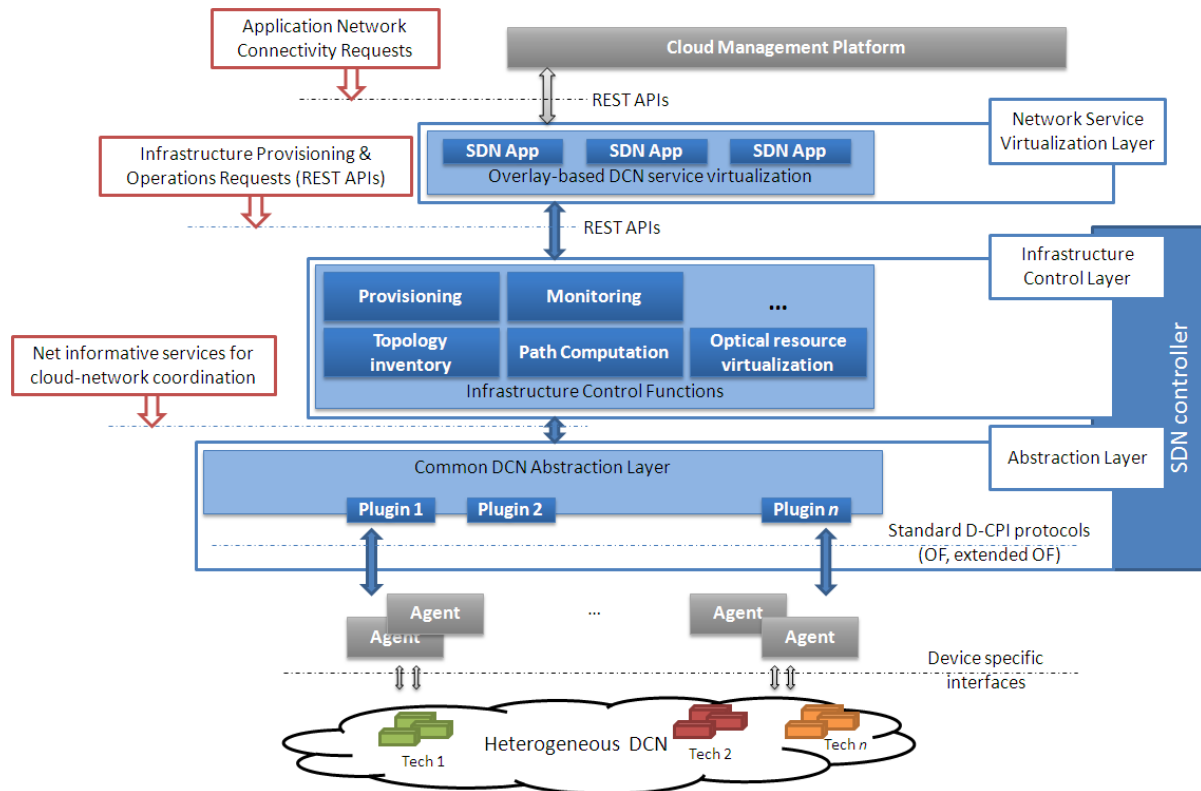


Figure 14 – COSIGN control plane architecture and services

The COSIGN control plane architecture, detailed in deliverable D3.1 [2], is shown in Figure 14, with its three main layers:

- **Abstraction layer:** provides commands and notifications to interact with the COSIGN optical data-plane through a unified DCN information model. In particular, the abstraction layer provides common mechanisms to retrieve the capabilities of the optical devices deployed in the data plane, to configure them and collect their alarms or statistics. The abstraction layer is structured in several protocol-specific drivers, implementing the OpenFlow extensions defined for each COSIGN data plane technology.
- **Infrastructure control layer:** provides the basic network functions for virtualization of optical resources and dynamic provisioning of optical connectivity. These functions are implemented through internal modules in the SDN controller, for example a topology manager, a statistics manager, a connection provisioning manager, etc. which interact to setup dynamic intra-DC paths and build virtual slices coexisting over the DCN. The high-level services implemented in the SDN controller are exposed through the A-CPI using REST API and they can be consumed by further SDN applications at the control layer, but also by the cloud orchestrator.
- **Network service virtualization layer:** provides the overlay-based network virtualization, with traffic encapsulated at the end points and tunnelled across the DCN. It is implemented through

an SDN application which configures directly the network edges but makes usage of the SDN controller internal functions to establish dynamically the required optical connectivity inside the network boundaries. This interaction allows the cooperation with the infrastructure control layer, enabling the exchange of information between the two layers and their mutual adaptation, overcoming the limitations of the simple overlay-based virtualization.

5.2.1 Preliminary functional blocks of the control layer

Deliverable D3.1 has specified the functional components of the COSIGN control layer. Figure 15 provides an overview of the main components, their positioning in the control plane and their interaction.

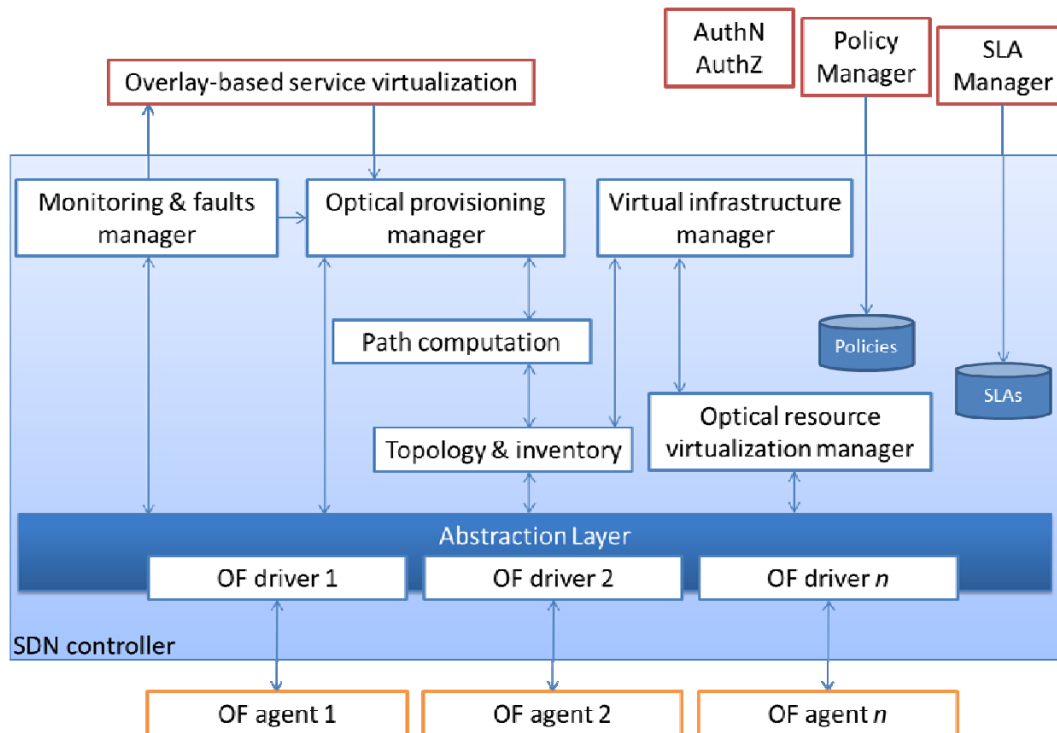


Figure 15 – COSIGN control plane functional blocks

The entities which implement the main control plane services within the SDN controller are the following:

- **Optical Provisioning Manager:** provides the mechanisms to setup and delete intra-DC optical connections between servers with QoS guarantees. It invokes the **Path Computation** block to request the computation of optical paths with a set of constraints (defined by the operator policies and the customers' SLAs). These paths are computed over the DCN topology exported by the **Topology & Inventory** component, able to describe the optical characteristics of COSIGN data plane. For each path, the Optical Provisioning Manager configures the required cross-connections on each node of the physical DCN, using unified commands exposed by the **Abstraction Layer** to interact with the data plane. The Optical Provisioning Manager is also responsible to coordinate the recovery of the services in case of failures, which are notified through the monitoring component.
- **Virtual Infrastructure Manager:** provides a service to compose and deliver virtual optical slices with specific capabilities and topologies. The physical topology of the DCN, represented through the Topology & Inventory component, is partitioned in isolated logical network infrastructures composed of a set of interconnected virtual nodes. The management of the single virtual resources is handled through the **Optical Resource Virtualization Manager**, which manages and abstracts the specific constraints of the optical technology, mapping the

commands for the logical resources to the corresponding configuration of the physical network.

- **Monitoring & Faults Manager:** Collects stores and distributes information about status, statistics and alarms for the different physical nodes of the DCN. This service is consumed by other SDN controller components (e.g. the Optical Provisioning Manager) or by external SDN applications via REST APIs.

All the functional blocks which need to interact with the data plane use a unified interface exposed by the **Abstraction Layer**, with primitives to send commands, retrieve capabilities information from the DCN nodes and subscribe for notifications about alarms or change of status. The translation between these primitives and the corresponding OpenFlow messages exchanged on the south-bound interface between physical nodes and SDN controller is handled by several **OpenFlow drivers**, dealing with the specific protocol extensions defined for each category of optical devices. On top of each node at the data plane, an **OpenFlow Agent** acts as mediator between the SDN controller and the specific device, implementing the vendor specific protocols used to interact with the device itself.

The overlay-based network virtualization is provided through a dedicated SDN application which runs over the SDN controller. It interacts with the Optical Provisioning Manager to request the on-demand creation of optical connections, as required to support the underlying connectivity for the overlay networks built above. On the other hand, it may receive information from the Monitoring & Faults Manager to take decisions which take into account the status of the physical DCN.

Other SDN applications, mainly with a management scope, are also placed on top of the SDN controller. One of these applications (AuthN/AuthZ) manages authentication and authorization to permit a secure access to the different actions available on the physical DCN or on its logical partitions. Other applications allow the DCN administrator to configure the operator policies and the SLAs for the different customers in the SDN controller.

5.2.2 Functional blocks features and innovations

The following table summarizes the innovative features that will be implemented in COSIGN.

Functionality: Multi-technology optical resource virtualization
Functionality description
Heterogeneous physical resource from COSIGN DCN is virtualized to provide elementary and abstract virtual optical resources which can be combined to create complex infrastructures.
Involved SD-CP functional blocks
<ul style="list-style-type: none"> - Optical Resource Virtualization Manager - Abstraction Layer - OpenFlow drivers and agents#
Input
<p>Creation of virtual resources: capabilities and characteristics of physical optical resources.</p> <p>Configuration of virtual resources: configuration commands for the virtual resource.</p>
Output
Creation of virtual resources: capabilities and characteristics of the corresponding virtual optical resources.

Configuration of virtual resources: configuration commands for the corresponding physical resource.

Table 4 Innovative functional blocks in the control layer (1)

Functionality: Provisioning and management of virtual optical infrastructures
Functionality description
Composition and delivery of on-demand and isolated virtual network slices.
Involved SDN-CP functional blocks
<ul style="list-style-type: none"> - Virtual Infrastructure Manager - Topology & Inventory - Monitoring & Faults Manager - Optical Resource Virtualization Manager - Abstraction Layer - OpenFlow drivers and agents
Input
<p>Specification of the desired virtual slice, e.g.:</p> <ul style="list-style-type: none"> - topology - bandwidth capacity - expected traffic profiles - resiliency - elasticity rules and thresholds - programmability options
Output
Result of the virtual slice provisioning action, with details about the delivered virtual infrastructure and how to access it.

Table 5 Innovative functional blocks in the control layer (2)

Functionality: Provisioning and management of intra-DC optical connectivity
Functionality description
Computation and setup of intra-DC optical paths to provide underlying connectivity between DC servers over the physical DCN.

Involved SDN-CP functional blocks
<ul style="list-style-type: none"> - Optical Provisioning Manager - Path Computation - Topology & Inventory - Monitoring & Faults Manager - Abstraction Layer - OpenFlow drivers and agents
Input
<p>Specification of the desired connectivity:</p> <ul style="list-style-type: none"> - connection type - end-points - QoS parameters - service resiliency (optional) - monitoring information (optional) - time constraints (optional)
Output
<p>Result of the connection provisioning.</p>

Table 6 Innovative functional blocks in the control layer (3)

5.3 Orchestration layer

As previously described in section 5.1, the physical layer of the DCN is responsible for transferring data between network clients and control layer is responsible for computing data transfer paths and configuring data plane devices to enforce the computed paths. The orchestration layer is responsible for feeding the control plane with high level application information required to optimize the computation data transfer paths taking into account also the IT resources. This way, the overall DC goals are achieved, be it global policies, workload or tenant specific performance requirements, etc.

While the physical layer described in Section 5.1 and control layer described in Section 5.2 mainly look after the networking domain, the orchestration layer also comprises the DC management stack and takes decisions which involve both the DC resources (compute, storage, etc.) and the network.

Thanks to its position in the COSIGN architecture stack, the orchestrator owns a clear perspective on the infrastructure layers of the architecture. This fact enables the orchestrator to match the application layer requirements with the DC infrastructure resources.

In typical data centres today, e.g. in public cloud environments, the application layer is very dynamic and is composed of many multi-component workloads with varying communications requirements. There is no single application pattern that can be used for static or semi-static decisions regarding, for example, bandwidth allocation. Therefore, the orchestrator must dynamically monitor the infrastructure-level and the application-level KPIs and adjust the resource allocation decisions to meet them optimally.

The orchestrator requires from the full visibility and control into the entire infrastructure silos – the compute, the storage and the network; meaning to be aware of availability and utilization level of the compute, compute-storage network proximity and availability of the network links. This is a must, for instance, in order to guarantee that the orchestrator takes the optimal decisions regarding the placement of the application components and data. In order to perform the decisions, the orchestrator must be capable of commanding the infrastructure controllers through well-defined technology independent interfaces. Generic functional layers of the orchestrator layer are as follows:

- **The internal control and optimization engine.** This is the heart of the orchestrator which owns the full underlying infrastructure view. The orchestrator engine comprises the functionalities that enable matching the application level resource requirements to the infrastructure level resource capabilities.
- **The application/client (north-bound) interfaces.** Northbound interfaces are responsible for accepting the application/client layer requests and hold the service and business logics. The type of expected requests highly depends on the type of application. For instance, considering the proposed COSIGN use cases, such request may range from a parametrized application template for application provisioning request to a full VDC description for a case of virtual infrastructure provisioning.
- **The infrastructure (south-bound) interfaces.** There are two generic types of interaction with the infrastructure layer, namely monitoring (south-to-north interaction) and control (north to south interaction). The former retrieves the infrastructure resources' monitored information and feeds it to the orchestration layer to facilitate the decision making process and to realize some business logic scenarios, e.g. billing, SLA controls, etc. The later type of interaction required, the control, flows from the opposite direction: the orchestrator pushes the configuration rules towards the infrastructure control layer to enforce the required behaviour of all the controllable infrastructure aspects to match application layer expectances.

The interfaces must be flexible enough to serve a wide range of applications and use cases, while capable of harnessing the unique capabilities of the COSIGN physical fabric based on optical and photonics technology.

5.3.1 Preliminary layers and functional blocks of the COSIGN orchestration layer

Preliminary architecture of the COSIGN orchestrator is specified as part of D4.1 and outlined in Figure 16. There are three major architectural layers – the layer interfacing the DC workloads/clients, the layer interfacing the infrastructure devices/controllers and the layer implementing the orchestrator functionality for continuous monitoring the infrastructure and workload KPIs, as well as continuous matching of the application level resource requirement to the infrastructure level resource capabilities.

- **Client Facing Interfaces** layer of COSIGN orchestrator is designed to hold components interfacing the different types of clients – the **Application Blueprints** component to interface with application provisioning requests, the **VDC Requests** components to interface with the virtual infrastructure provisioning requests and the **KPIs** component to interface with the DC administration and operation requests.
- **Data Models, Algorithms, and Decision Making** layer of COSIGN orchestrator is responsible for main orchestrator functionality and includes the **Admission Control** component that makes decisions regarding embedding new workloads and jobs into the infrastructure, the **Provisioning** component that makes placement decisions and resource allocations for admitted workloads and the **Periodic Adjustments** component that monitors the KPIs of both the DC and the workloads and modifies the resource allocation and/or configuration decisions when required.
- **Infrastructure Control Clients** layer of COSIGN orchestrator is responsible for the interface with the infrastructure, across all the types of resources and includes the **Compute** component for the compute interface, the **Storage** component for the storage interface and the **Network** component for DCN interface.

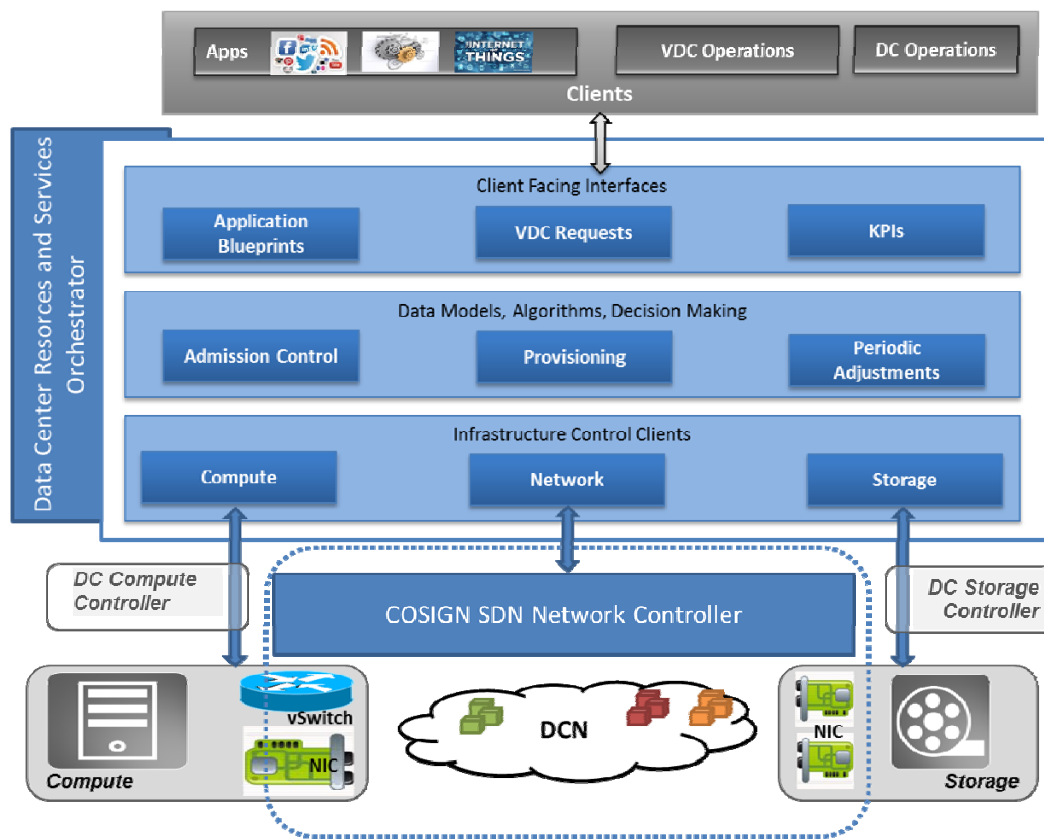


Figure 16 – COSIGN Orchestration layer – functional blocks and interfaces

5.3.2 Functional blocks features and innovations

Functionality: Workload-level DCN demand specification.
Functionality description
Allow for specifying workload and client requirements towards the network in high level, technology independent and application domain terms.
Involved functional blocks
<ul style="list-style-type: none"> - Application Blueprints - VDC Requests
Input
High level description of application or VDC components and their communication requirements – connectivity patterns, QoS, network services, etc.
Output
Workload placement decisions are made through considering networking aspects and the networking infrastructure is automatically reconfigured as part of workload provisioning

Table 7 Innovative functional blocks in the orchestration layer (1)

Functionality: Dynamic DCN re-provisioning.
Functionality description
Allow DCN re-provisioning triggered by monitoring data from the DCN and the application domains.
Involved functional blocks
- Periodic Adjustments
Input
Monitoring and KPI data collected from the infrastructure and, where possible, from the application hooks.
Output
DCN resource re-provisioning to optimize the usage, e.g. configuring fast optical circuit to eliminate congestion.

Table 8 Innovative functional blocks in the orchestration layer (2)

Functionality: Integrated ICT orchestration
Functionality description
Allow DCN re-provisioning triggered by automatic DC-wide features like application and service elasticity, storage transfers, etc.
Involved functional blocks
- Provisioning
Input
Explicit elasticity settings and policies, coupled with the application-level events like sudden load increase
Output
DCN resource re-provisioning to satisfy application-level policies, e.g. configuring fast optical circuit for upcoming storage transfer or expanding virtual network to grow application cluster.

Table 9 Innovative functional blocks in the orchestration layer (3)

6 Architecture inter-Layer Interfaces Preliminary Specification

Section 6 defines the interfaces between the different layers (e.g. Control plane - Data plane and Control plane - Orchestration plane) and the relationship between them from a high level point of view. The heterogeneous data plane requires a suitable southbound interface at the network control plane which allows discovering, manipulating, configuring and monitoring the new types of resources, in order to exploit their capabilities, connection granularity and dynamicity. On the other hand, the interaction between network controller and orchestrator needs to rely on a powerful interface to provide an access to the control plane services with abstract primitives that hide the data plane technology details but expose network programmability and monitoring with a sufficient level of details. A first high level specification of the interfaces as the means to request for services and configure and manage the DC network is next presented. Nevertheless, A full specification of such interfaces, including protocols, workflows specification will be provided in D3.2 (SDN framework north-bound and south-bound interfaces specification) and D4.2 (COSIGN orchestrator low level architecture and prototype design).

6.1 Control plane and Orchestration layer interface

The interface between the control plane and the orchestration layer scheme is represented in Figure 17. It includes all the mechanisms used by the Orchestration layer, through the Cloud Management Platform, to request services on the DC network or collect information about the network infrastructure which may be used for decisions at the whole DC level. The interface is typically based on REST APIs which support CRUD operations on network related resources.

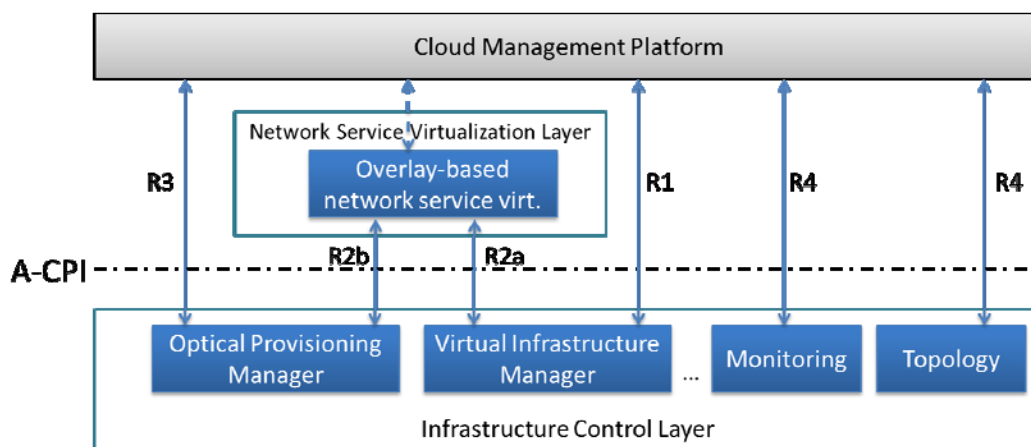


Figure 17 – COSIGN CP functional blocks. Interaction with orchestration layer.

We can identify the following main interactions:

- A. **Request for customized virtual networks.** This type of request is used to provide and manage the network part of an IaaS offer and it is typically invoked by the Neutron component at the Cloud Management Platform. The request can be directed to the SDN application for the provisioning of overlay-based virtual networks or to the Virtual Infrastructure Manager within the SDN controller, depending on the request type. A request which maps the basic concepts of Neutron resource (e.g. network, subnets, ports, routers) is typically served through overlay virtualization. A request for a complex virtual topology which specifies requirements in terms of optical capabilities must be handled by the Virtual Infrastructure Manager and mapped on a proper virtual network slice.

- B. Request for dynamic optical connectivity within the DC.** This request, when triggered by the orchestration layer, is usually associated to the setup of optical paths to support management actions at the DC level, e.g. the replication of VMs between different servers which requires the transfer of a large bulk of data. In this case it is typically invoked by internal management functions in different OpenStack components and, at the control plane level, it is handled by the Optical Provisioning Manager.
- C. Collection of information about the DC network.** This interface allows the orchestrator layer to retrieve details about the behaviour of the underlying network, which can be used in internal decisions to optimize the usage of the whole DC resources or to react to unexpected behaviour at the network level. DCN related information covers different aspects of the network, for example failures, status of the devices, level of congestion in some network segments, power consumption, but also topology details like nodes and links, even if highly abstracted. The control plane entities which expose this type of information are the Monitoring & Faults Manager and the Topology & Inventory components respectively. Beyond the REST mechanism to collect data polling a resource with a GET request, the interface should also support subscriptions and notifications (e.g. via web sockets) for asynchronous notifications about alarms and failures.

6.2 Control plane and data plane interface

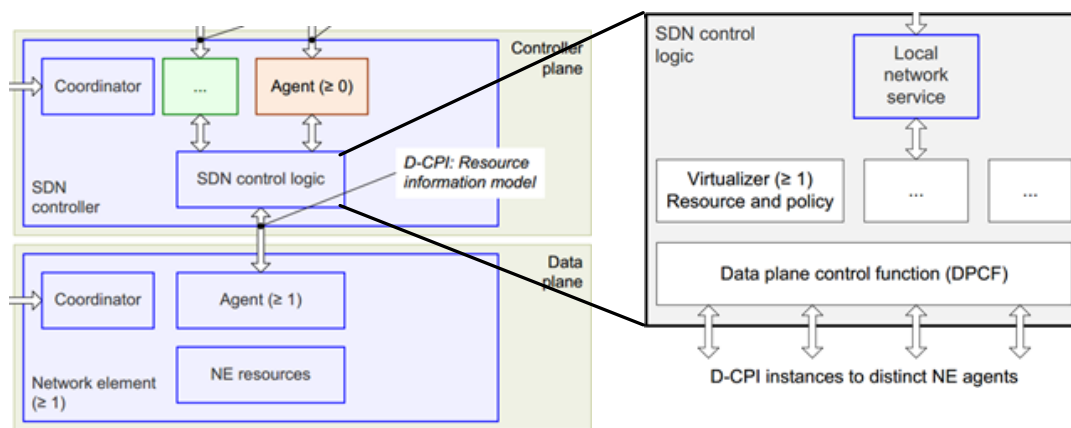


Figure 18 Control Plane and Data Plane interface defined in SDN architecture

According to the definition in [10], the interface between data and control planes (D-CPI) includes functions such as:

- Programmatic control of all functions exposed by the Resource Database (RDB)
- Capabilities advertisement
- Event notification

A controller-agent model is appropriate for the relation between a SDN controller and a controlled entity. The controlled entity is a designated agent, a functional component that represents the client's resources and capabilities. It is also responsible for carrying out the commands of the controller and notifying the controller of events in categories specified by the controller. The Network Element (NE) provides at least one logical data-controller plane interface (D-CPI) that allows its functions to be managed and controlled by the SDN controller, as shown in Figure 18.

The data plane agent is the entity that executes the SDN controller instructions in the data plane. At the lowest layer of recursion, data plane resources are physical entities (including soft switches). At higher levels of abstraction, however, data plane resources need not be physical (e.g., virtual NEs). Similarly to the other planes, the SDN architecture operates on an abstract model of the data plane and

as long as the functions advertised by the model are correctly executed, the architecture is blind to the difference. In the data plane, a resource data base (RDB) models the current information model instance and the necessary supporting capabilities of the NE. These data plane resources are represented as a set of virtual NEs (VNEs) through D-CPI. Different data plane agents represent data plane function abstraction in different abstraction level or different scale.

6.3 Orchestration layer and Service/application layer interface

Figure 19 shows how the orchestration and the services layers interact with the lower and upper layers. The Orchestration layer coordinates the provisioning and management of DC resources (i.e. compute, storage, and network) in order to provision cloud services to the customer. Additionally, the orchestrator (the entity comprising the orchestration layer) may perform management of the DC resources for optimization and efficiency purposes. The *service layer* includes functional entities that do not necessarily coordinate DC resources, but still provides required functionality such as informative service, billing or other management services. In order to fulfil its purpose, the orchestration layer interacts with the different resource controllers.

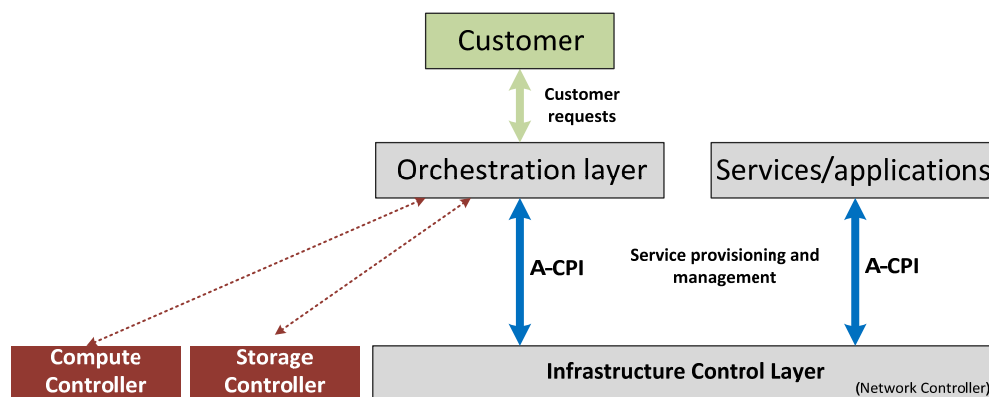


Figure 19 Interfaces for the orchestration and services layer

The network controller, which is generally referred as the *Infrastructure Control Layer*, enables programmatic control of abstract network resources through a set of APIs exposed at the A-CPI (Application-Controller Plane Interface). The programmatic control of network resources may comprise services for provisioning of network resources (e.g. optical connectivity, optical slices), services for policy and SLA configuration, informative services (monitoring data), etc. Using the A-CPI and the interfaces towards the other resource controllers (i.e. compute and storage) the orchestrator realizes the coordinated provisioning of cloud services to the customers. The message sequence chart in Figure 20 shows a high level interaction between various entities for the purpose of provisioning cloud services to customers. As it can be noticed in the figure, the orchestrator receives the customer's request and further provisions the required DC resources using the capabilities exposed by the resource controllers (compute and network).

Similarly, the services layer consumes the APIs exposed at the A-CPI in order to realize their functions. It is important to emphasize that the A-CPI provides capabilities for event notification. This would enable various services to receive notifications about network events and react to them accordingly.

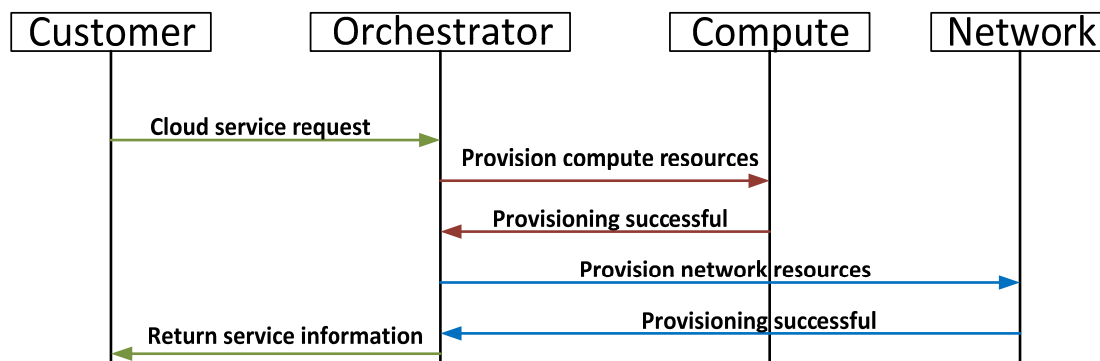


Figure 20 High level interaction with the orchestrator

7 COSIGN service delivery model

In this section it is analysed a list of involved actors and roles suitable to be adopted in the context of COSIGN and it is provided an overview on the overall services delivery model phases which enable to validate the proposed architecture framework.

7.1 COSIGN service model actors and roles

A role names the behaviour of an entity participating in a particular [13] context and generally is used to identify it [14]. A role groups a set of responsibilities and actions that are associated with a resource of a given service provisioning state. Roles are completely disjoint. An actor is a given materialization of at least one role, although it can cope with one or more roles simultaneously into a single entity. It may be the case that roles and actors match 1-to-1 and some other situations in which an actor adapts several roles depending on the type of service provided.

7.1.1 COSIGN roles:

SDN based technologies, virtualization and service abstraction implemented in COSIGN within its layered architecture enables to define a number of roles that COSIGN actors can adopt. The following roles have been identified:

- Data Centre Physical Infrastructure Provider (DC-PIP) owns the physical infrastructure (both IT and networking resources) and makes it available to others with the final aim to gain revenue from it. It owns and manages one or more datacentres to provide virtual infrastructure services (i.e. IaaS services) to its customers (e.g. DC virtual operators).
- Network operator: the provider of network services for connections between datacentres located in different sites. Depending on the scenario, the network operator and the datacentre operator roles may be covered by the same actor.
- Data Centre Virtual Infrastructure Provider (DC-VIP) acts as a virtual infrastructure resource broker, responsible for combining IT and Net virtual resources and provides Virtual Infrastructures as a Service (transfers the control rights over it).
- Cloud service provider (Cloud-SP): it provides a cloud services to its customers, relying on its own physical infrastructure or on a rented virtual infrastructure.
- Cloud service customer: represents the consumer of a generic cloud service and establishes an SLA with the cloud service provider. It may be an enterprise customer or a final end-user.

Roles combine between each other depending on the type of service to be provided. Each role provides a further step in the chain value towards providing the final service. Figure 21 shows the stack of potential COSIGN roles interrelation matching the proposed architecture and services described at each of the interfaces in section 6.

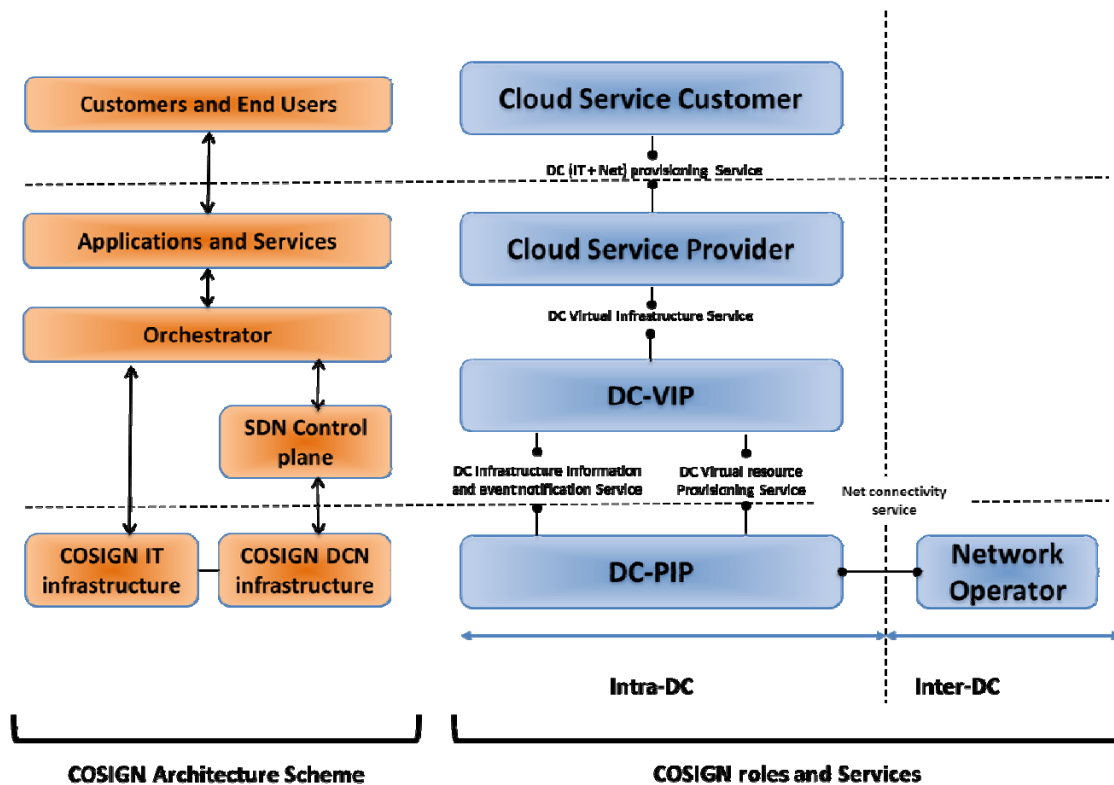


Figure 21 COSIGN architecture and roles relationship

7.1.2 COSIGN actors:

It is possible to define the actors involved in the context of future DCN environments based on the architecture specification and previous roles.

- **DC Operator:** Infrastructure provider relying on its own physical infrastructure and performing physical DCN and IT resources partitioning and virtualization. Integrated DCN services handled by the physical DC operator include:
 - Provisioning requests for low level DCN components like links, ports, FW rules, etc.
 - Monitoring requests for low level DCN stats like port level of flow level counters readings
 - Automated services like fencing off a specified part of the DC for the sake of maintenance or replacement. Such requests are higher level and will trigger a series of DCN operations like circuit configurations to migrate the existing workloads and data from the specified DC part, followed by circuits demolition and then by blocking off all the data connectivity of the fenced DC part leaving only the management port up, etc.
- **Virtual DC Operator** working on top of a rented virtual infrastructure; acts as a broker, managing virtual data centre infrastructures and providing them as a service for applications and service developers. As with DC operations, the Virtual DC operations include networking operations and will also have network administrators as part of the team. The difference from the DC Operator use cases is that virtualized infrastructure can be composed of ‘devices’ and ‘technologies’ different from those the underlying DCN is built with. Integrated DCN services offered by this type of actor may include:
 - Provisioning requests for virtualized DCN topology – switches, switch ports, links, etc.

- Provisioning requests for virtualized DCN services – connectivity, port-level of flow-level policies, FW rules, etc.
 - Monitoring requests for the amount of physical resources consumed by the VDC infrastructure, e.g. for business controls.
 - Monitoring requests for virtual DCN stats like port level of flow level counters readings
 - Automated services like VDC replication over remote DC locations or policy based right-sizing of the VDC infrastructure to match the dynamically changing VDC load.
- Vertically integrated DC provider: the actor owns the physical infrastructure but also creates virtual DC Network slices to be provided as a service. Thus, the DC provider covers both physical and virtual joint IT and DCN infrastructure as a service. The SDN based control framework together with the orchestration layer permits the complete DC network management.
- End user: End users act as DC services consumer. Such end users may act differently according to the service model of their own business:
 - IaaS consumers demanding a whole DC infrastructure (processing, storage computing and networking resources) on top of which the consumer is able to deploy and run arbitrary software, including operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage and deployed applications.
 - PaaS consumers usually request a platform (Net and IT resources) on top of which they are able to deploy cloud infrastructure applications created using programming languages, libraries, services and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, nor storage, but has control over the deployed applications and possibly configuration settings of the application-hosting environment.
 - SaaS consumers. In this case, the consumers only demand the usage of providers' applications running on top of a cloud infrastructure and are able to commit limited user-specific application configuration settings. The low level networking details are totally supercilious as the application owner is just concerned with application components connectivity service and its properties being the network logic transparent for them.

7.2 Service model phases

7.2.1 Service request

The service consumer defines the requirements for the desired service instance and sends the request to a service provider. This phase ends with the establishment of a final agreement between service consumer and service provider. The process can be fully automated (e.g. through a request to a web portal) or may require human intervention (e.g. for the delivery of a customized solution). The figure below shows, as an example, the web page to order an Interoute VDC service, where the customer can choose from a range of pre-configured options and is guided through a simple configuration. For more complex deployments, the web page shows the service provider contact points.

Interoute Virtual Data Centre

Build Your Cloud On Top of Our Global Network

The Interoute Virtual Data Centre (VDC) has sites in London, Paris, Amsterdam, Madrid, Berlin, Geneva, Milan, Hong Kong and New York. The package you select below covers all these zones and you can select where you want to place your infrastructure when you deploy the VM.

Hypervisor: Contract Duration:

	VMs	VMs	VMs	VMs
VMs	100	250	500	1000
VMs (GB)	100	250	500	1000
Storage (GB)	100	250	500	1000
Storage (GB)	100	250	500	1000
Storage (GB)	100	250	500	1000
Storage (GB)	100	250	500	1000

Connect up your VDC

You can connect up your VDC zones using a free MPLS VPN now... or later if you would prefer. If you already have an MPLS VPN with Interoute, this will be added to your existing network by default. If you don't already have one, maybe you would like to connect your VDC network up to a corporate network? If so contact sales.support@interoute.com.

Zone	IP Address	Subnet Mask	Default Gateway
<input checked="" type="checkbox"/> Berlin	10.0.3.0	24	254
<input checked="" type="checkbox"/> Paris	10.0.4.0	24	254
<input checked="" type="checkbox"/> Milan	10.0.7.0	24	254
<input checked="" type="checkbox"/> Amsterdam	10.0.8.0	24	254
<input type="checkbox"/> Hong Kong		24	
<input type="checkbox"/> New York		24	
<input type="checkbox"/> Slough		24	
<input checked="" type="checkbox"/> London	10.0.12.0	24	254
<input checked="" type="checkbox"/> Geneva	10.0.13.0	24	254
<input checked="" type="checkbox"/> Madrid	10.0.14.0	24	254

Figure 22 – Web pages to request and configure an Interoute VDC service

7.2.2 Service deployment

The cloud service specified in the previous phase is deployed by the service provider. This phase should be as much automated as possible.

A cloud service is delivered through two steps: a deployment phase, where the resources are allocated and instantiated and an operation phase, where the instantiated resources are configured and provisioned to the user. During the operation phase, the user has access to the service, which should be available according to the SLA defined at the service request phase.

7.2.3 Service Operation

7.2.3.1 *Service Monitoring*

The service provider continuously monitors the service to verify the compliance with the SLA. In some scenarios, per-service monitoring facilities should be provided to the service customer. Monitoring should have a limited impact on the running services.

7.2.3.2 *Service re-planning and restoration*

Service re-planning and restoration can be triggered to avoid SLA violations or to mitigate their impact, e.g. in case of service disruption. The datacentre operator may also use re-planning for maintenance or performance/usage optimization, but without disrupting the running services.

7.2.4 Service termination

The service is decommissioned and the resources are de-allocated, to become available for new services.

8 Matching the requirements to COSIGN architecture

The **requirements** proposed to satisfy the use cases and validate the COSIGN DCN architecture entail a number of functionalities that attain to different layers. Driven by the added value that the COSIGN technologies bring to the project, there have been identified the component modules that contribute to accomplish such requirements. COSIGN partners have conducted a thorough analysis to clearly identify the layers, components and interfaces that enable the requirements. The analysis (see Table 10) clearly shows the matching between the Service, management and control plane requirements to the specific modules and interfaces of the COSIGN architecture. The analysis aims to assess and validate the suitability of the technologies chosen to satisfy the requirements. Business requirements have been omitted in the analysis since they bring very high level functionalities that are already taken into consideration at a lower level with the Service, Management or control plane requirements. Also the Data plane requirements are not considered, since those requirements are not specifically covered at the data plane, but they may be satisfied by means of the cross-coordination among the upper layers of the COSIGN architecture.

The different component modules that the SDN controller manages at both control and orchestration layers help to satisfy the requirements mapping some of their functionalities directly to these components. Nevertheless, and as it has been pointed out, the critical point that enables to fully satisfy them is the cross-coordination and interfacing among the layers to serve the requirements and the use cases proposed; this is achieved by exposing and exerting management, control and configuration over the DCN resources of the COSIGN architecture.

REQUIREMENTS	MATCHING OF THE REQUIREMENTS TO LAYERS, COMPONENTS AND INTERFACES			
SERVICE REQUIREMENTS	FUNCTIONALITIES	LAYER(S)	FUNCTIONAL COMPONENTS	INTERFACES
R-SERV-01: Seamless Network Resources Provisioning	<ul style="list-style-type: none"> • <u>SDN controller</u>: HW resource controlled by SDN controller layer. • <u>Resource APIS</u>: Resources' features and capabilities available to be consulted by the service and application layers. 	Application/Service layer SDN controller layer – CPI, north-bound plugins. Orchestration Layer.	Orchestrator: <ul style="list-style-type: none"> • Heat • Neutron • Nova • Physical observer • Algorithms for joint computation of network and computing resources SDN controller: <ul style="list-style-type: none"> • Overlay Controller • Optical provisioning manager • Path computation 	Orchestrator: <ul style="list-style-type: none"> • Heat API SDN Controller: <ul style="list-style-type: none"> • A-CPI for connectivity requests • D-CPI for resources provisioning
R-SERV-02: Advanced Network Services Provisioning	<ul style="list-style-type: none"> • <u>HW Configuration tools</u> to delegate and make available at user level the possibility to handle the HW as it was of its ownership. • <u>Cross layer mapping</u>: Mechanisms to map such advanced services to the application layer available. 	Application / Service Layer Orchestration layer SDN controller layer Data plane layer	Orchestrator: <ul style="list-style-type: none"> • Neutron • Neutron Client • DCN orchestrator • DB Module SDN controller: <ul style="list-style-type: none"> • Optical provisioning manager • Path computation • Topology Inventory • Monitoring and fault manager 	Orchestrator: <ul style="list-style-type: none"> • Neutron Client component REST API SDN Controller: <ul style="list-style-type: none"> • A-CPI for connectivity requests • D-CPI for notification of data plane failures
R-SERV-03: Service level monitoring	<ul style="list-style-type: none"> • <u>Service monitoring tools</u> enabled, consuming network monitored parameters. 	Application/Service layer Orchestration Layer SDN controller Layer. CP	Orchestrator: <ul style="list-style-type: none"> • Ceilometer • Ceilometer client SDN Controller:	Orchestrator: <ul style="list-style-type: none"> • Ceilometer client component REST APIs SDN Controller: <ul style="list-style-type: none"> • D-CPI for notification of data plane failures

		Northbound and southbound interfaces. Data plane Layer	<ul style="list-style-type: none"> Monitoring and fault manager 	
R-SERV-04: Adapting to application/service dynamicity	<ul style="list-style-type: none"> <u>Service update methods</u>: Service auto-reconfiguration and re-provisioning tool available. Include <u>resources re-provision mechanisms</u> as part of the monitoring tool. 	Application/Service layer Orchestration layer SDN controller layer and Southbound CPI. Data plane layer	Orchestrator: <ul style="list-style-type: none"> Heat Physical Observer SDN Controller: <ul style="list-style-type: none"> Overlay Controller Optical provisioning manager Monitoring and fault manager Virtual Infrastructure manager 	Orchestrator: <ul style="list-style-type: none"> Heat Component REST API SDN Controller: <ul style="list-style-type: none"> A-CPI for configuration of operator policies A-CPI for modification requests D-CPI for collection of statistics or failure notifications
R-SERV-05: QoS provisioning	Enable a <u>monitoring tool</u> that maps service metrics into levels of performance.	Application/Service layer Orchestration layer SDN controller layer	Orchestrator: <ul style="list-style-type: none"> Heat Neutron Physical Observer Ceilometer Ceilometer client SDN Controller: <ul style="list-style-type: none"> Overlay Controller Optical provisioning manager Virtual Infrastructure Manager 	Orchestrator: <ul style="list-style-type: none"> Heat Component REST API Neutron component REST API Ceilometer client component APIs SDN Controller: <ul style="list-style-type: none"> A-CPI for connectivity requests with user-constraints specification
R-SERV-06: Real Time Control	<u>Automated re-provisioning mechanisms</u> : Linked to (or even embedded in) the monitoring tool, it there could be included a mechanism to automatically react on real-time whenever changes are required.	Application/Service Level Orchestration layer SDN controller layer Data plane layer	Orchestrator: <ul style="list-style-type: none"> Heat Neutron Physical Observer Ceilometer Ceilometer client SDN Controller: <ul style="list-style-type: none"> Monitoring and fault manager 	Orchestrator: <ul style="list-style-type: none"> Heat Component REST API Neutron component REST API Ceilometer client component REST APIs SDN Controller: <ul style="list-style-type: none"> A-CPI for connectivity requests

			<ul style="list-style-type: none"> Optical Provisioning Manager Overlay Controller 	<ul style="list-style-type: none"> D-CPI for notification of data plane failures
R-SERV-07: Big data support	<ul style="list-style-type: none"> <u>Algorithms manager</u>: Algorithm based on relevant metrics for Big Data support monitoring. <u>Energy efficiency mechanisms</u>: Moving data is bad for data. It really cost a lot of money in terms of energy. 	Application/Service layer Orchestration layer	Orchestrator: <ul style="list-style-type: none"> Algorithms manager 	<ul style="list-style-type: none"> Heat Component REST API Neutron component REST API
R-SERV-08: Secure data management	<ul style="list-style-type: none"> <u>Security NFVs</u> complementary to current HW security. Data and service <u>isolation mechanisms</u>. <u>Authorization, authentication and availability mechanisms</u> to ensure secure data management. 	Application/Service layer Orchestration layer SDN controller layer	Orchestrator: <ul style="list-style-type: none"> Algorithms for joint computation of network and computing resources guaranteeing service isolation SDN Controller: <ul style="list-style-type: none"> Virtual Infrastructure manager (for isolation of virtual infrastructures) Monitoring and fault manager AAA for authorization (OpenStack Keystone)	SDN Controller: <ul style="list-style-type: none"> A-CPI to expose monitoring and accounting information
R-SERV-9: Security management	<ul style="list-style-type: none"> <u>Security NFVs</u> complementary to current HW security. Data and service <u>isolation mechanisms</u>. Authorization, authentication and availability mechanisms to ensure secure data management. 	Application/Service layer Orchestration layer SDN controller layer	Orchestrator: <ul style="list-style-type: none"> Algorithms for joint computation of network and computing resources guaranteeing service isolation SDN Controller: <ul style="list-style-type: none"> Virtual Infrastructure manager (for isolation of virtual infrastructures) AAA for authorization (OpenStack Keystone)	SDN Controller: <ul style="list-style-type: none"> A-CPI to expose monitoring and accounting information
R-SERV-10: Self-service management	<u>Self-service management and configuration Interfaces</u>	Orchestration layer	Orchestrator: <ul style="list-style-type: none"> Heat 	Orchestrator: <ul style="list-style-type: none"> Heat Component REST API

interfaces			<ul style="list-style-type: none"> • DCN Orchestrator 	
R-SERV-11: Multi-tenant isolation	<ul style="list-style-type: none"> • Implement <u>partitioning/slicing techniques</u> at different levels: <ul style="list-style-type: none"> • NIC • L2 (VLAN) • L3 Network using L3 (VXLAN, GRE) • L3 Network using L3 (MPLS, GRE, IPSec) 	Application/Service layer SDN controller layer Data plane layer	SDN Controller: <ul style="list-style-type: none"> • Virtual Infrastructure manager • OpenStack – multi-tenancy functions 	
R-SERV-12: Service Traffic Profile	<ul style="list-style-type: none"> • <u>Interface for service specification</u> • <u>Service orchestration</u> with coordinated network configuration. 	Application/Service layer Orchestration layer	Orchestrator: <ul style="list-style-type: none"> • Heat • Neutron 	Orchestrator: <ul style="list-style-type: none"> • Heat Component REST API • Neutron component REST API
R-SERV-13: Generalized information model	<u>Interface for service specification</u>	Application/Service layer	Orchestrator: Heat	Orchestrator: <ul style="list-style-type: none"> • Heat Component REST API
R-SERV-14: Service connectivity disjunction	<u>Restoration mechanism</u> offering disjointed path provisioning.	Application/Service Level Data plane layer	SDN Controller: <ul style="list-style-type: none"> • Optical Provisioning Manager • Path computation manager 	
R-SERV-15: VDC optical BW control	<u>Cross layer mechanism</u> to enable this feature. This may involve the HW, control and management layers.	Application/Service Level Orchestration layer SDN controller layer Data plane layer	SDN Controller: <ul style="list-style-type: none"> • Optical provisioning manager • Virtual Infrastructure Manager Orchestrator: <ul style="list-style-type: none"> • VDC algorithms • Neutron 	Orchestrator: <ul style="list-style-type: none"> • Neutron component REST API
R-SERV-16: Service Calibration	<u>Application layer interface</u> to enable DC tenants to control physical HW network resources.	Application/Service Level Orchestration layer	Orchestrator:	SDN Controller: <ul style="list-style-type: none"> • A-CPI for configuration of operator policies

		SDN controller layer Data plane layer	SDN Controller: <ul style="list-style-type: none"> Optical provisioning manager Virtual Infrastructure Manager 	
R-SERV-17: Intelligent/selective physical failure awareness	<ul style="list-style-type: none"> Service functionality <u>filtering failure awareness to service consumer</u> according to specific metrics which measure the seriousness of the failure. Integration with monitoring tools. 	Application/Service Level Orchestration layer SDN controller layer Data plane layer	SDN Controller: <ul style="list-style-type: none"> Monitoring and fault manager SLA manager Monitoring Orchestrator: <ul style="list-style-type: none"> Ceilometer client Ceilometer 	SDN Controller: <ul style="list-style-type: none"> A-CPI for SLA specification A-CPI and D-CPI for collection of statistics or failure notifications.
R-SERV-18: equipment inventory services tool	<ul style="list-style-type: none"> <u>Service catalogue</u> visible at application layer 	Orchestration layer	Orchestrator: <ul style="list-style-type: none"> DB Module SDN Controller: <ul style="list-style-type: none"> Inventory and Topology manager 	Orchestrator: SDN Controller:
MANAGEMENT REQUIREMENTS	FUNCTIONALITY	LAYER	FUNCTIONAL COMPONENT	INTERFACE
R-MGMT-01: Element Management	<ul style="list-style-type: none"> <u>Management interface</u> to manage individual devices deployed in the DC. <u>CLI/GUI</u> 	Orchestration layer	Orchestrator: <ul style="list-style-type: none"> Horizon Heat Neutron client Nova client 	Orchestrator: <ul style="list-style-type: none"> Heat Component REST API Neutron component REST API Nova component REST API SDN Controller.
R-MGMT-02: Operator Network Management mechanisms	<ul style="list-style-type: none"> <u>Network topology manager</u> <u>Network services provisioning interfaces.</u> 	Orchestration layer SDN controller layer	Orchestrator: <ul style="list-style-type: none"> Horizon Heat Neutron client Nova client SDN Controller: <ul style="list-style-type: none"> Inventory and Topology manager 	Orchestrator: <ul style="list-style-type: none"> Heat Component REST API Neutron component REST API SDN Controller: <ul style="list-style-type: none"> A-CPI for configuration of the network.
R-MGMT-03: Self-Service multi-tenancy	<ul style="list-style-type: none"> <u>DCN slicing mechanisms</u> <u>Virtual network</u> 	Orchestration layer SDN controller layer	Orchestrator: <ul style="list-style-type: none"> Heat Neutron client 	Orchestrator: <ul style="list-style-type: none"> REST API for L2 Network modelling using VTN

	<u>management as service</u> (in order to delegate network slices management to tenants)		SDN Controller: <ul style="list-style-type: none"> Virtual Tenant Network 	SDN Controller: <ul style="list-style-type: none"> A-CPI for forwarding VTN related information.
R-MGMT-04: Automatic resources discovery	<ul style="list-style-type: none"> <u>Topology manager</u> <u>Topology inventory</u> 	Orchestration layer SDN controller layer Data plane	Orchestrator: <ul style="list-style-type: none"> Heat Neutron client Nova client SDN Controller: <ul style="list-style-type: none"> Inventory and Topology manager Data plane: <ul style="list-style-type: none"> Automatic discovery protocols 	Orchestrator: <ul style="list-style-type: none"> Heat Component REST API Neutron component REST API Nova component REST API SDN Controller: <ul style="list-style-type: none"> A-CPI for communication of additional resources. D-CPI for collecting information about new network resources.
CONTROL PLANE REQUIREMENTS	FUNCTIONALITY	LAYER	FUNCTIONAL COMPONENT	INTERFACE
R-CP-01 Automated provisioning of intra-DC connectivity	Network service provisioning	SDN controller layer	SDN Controller: <ul style="list-style-type: none"> Optical provisioning manager Path computation 	SDN Controller: <ul style="list-style-type: none"> A-CPI for connectivity requests
R-CP-02 Support for customizable network services	Additional services inventory services computation algorithms	SDN controller layer	SDN Controller <ul style="list-style-type: none"> Optical provisioning manager with scheduling and recovery support Path computation for scheduled or disjoint paths Topology Inventory with calendar extensions Monitoring and fault manager 	SDN Controller <ul style="list-style-type: none"> A-CPI for connectivity requests with user-constraints specification D-CPI for notification of data plane failures
R-CP-03 Multiple automated connectivity paradigms	Computation and provisioning of point-to-point, point-to-multipoint or anycast connectivity	SDN controller layer	SDN Controller <ul style="list-style-type: none"> Optical provisioning manager Path computation with support for P2MP and anycast paths 	SDN Controller <ul style="list-style-type: none"> A-CPI for P2MP and anycast requests
R-CP-04 Multi-layer DCN operation	Operation of multiple optical technologies with different granularities	SDN controller layer	SDN Controller <ul style="list-style-type: none"> Agents and south-bound plugins for each data plane technology Service Abstraction Layer for optical resource abstraction 	SDN Controller <ul style="list-style-type: none"> D-CPI for data plane configuration

			<ul style="list-style-type: none"> • Optical resource virtualization service • Optical provisioning manager with multi-layer capabilities, including re-adaptation of the virtual server layer • Path computation with multi-layer capabilities 	
R-CP-05 DCN resource usage optimization	Computation algorithms policy management	SDN controller layer	SDN Controller <ul style="list-style-type: none"> • Virtual Infrastructure manager • Optical provisioning manager with re-planning features • Path computation with Global Concurrent Optimization algorithms • Policy manager 	SDN Controller <ul style="list-style-type: none"> • A-CPI for configuration of operator policies • D-CPI for collection of statistics about network performance
R-CP-06 Elastic intra-DC connectivity	Update methods to enable the dynamic modification of established connectivity or Virtual Infrastructures on-demand or based on automatic elasticity rules	SDN controller layer	SDN Controller <ul style="list-style-type: none"> • Virtual Infrastructure manager with modification capabilities • Optical provisioning manager with modification capabilities • Monitoring and fault manager 	SDN Controller <ul style="list-style-type: none"> • A-CPI for modification requests • A-CPI for specification of elasticity rules • D-CPI for collection of statistics or failure notifications
R-CP-07 Network monitoring	Monitoring services and statistics on the DCN resources usage	SDN controller layer	SDN Controller <ul style="list-style-type: none"> • Monitoring and fault manager 	SDN Controller <ul style="list-style-type: none"> • D-CPI for collection of statistics or failure notifications • A-CPI to expose monitoring information
R-CP-08 Programmable APIs	Programmable APIs to request provisioning of network connectivity, even over virtual slices.	SDN controller layer	SDN Controller <ul style="list-style-type: none"> • Policy manager • AAA for authorization 	SDN Controller <ul style="list-style-type: none"> • A-CPI
R-CP-09 Network service monitoring and accounting	Monitoring services and statistics on the DCN resources usage	SDN controller layer	SDN Controller <ul style="list-style-type: none"> • Monitoring and fault manager • AAA for accounting 	SDN Controller <ul style="list-style-type: none"> • A-CPI to expose monitoring and accounting information • D-CPI for collection of statistics or failure notifications

R-CP-10 Scheduled network connectivity	Computation and provisioning of scheduled connectivity or virtual slices.	SDN controller layer	SDN Controller <ul style="list-style-type: none"> Optical provisioning manager with scheduling support Path computation for scheduled paths Topology manager with calendar extensions 	SDN Controller <ul style="list-style-type: none"> A-CPI for requests with timing constraints
R-CP-11 Service resiliency	Failure detection and notification mechanisms: Detection and notification through the D-CPI and A-CPI of failure occurrence. Automated procedures for network service recovery	SDN controller layer	SDN Controller <ul style="list-style-type: none"> Optical provisioning manager with recovery support Path computation for disjoint paths Monitoring and fault manager 	SDN Controller <ul style="list-style-type: none"> A-CPI for failure notifications D-CPI for notification of data plane failures
R-CP-12 Multi-tenant isolation	Resource virtualization	SDN controller layer	SDN Controller <ul style="list-style-type: none"> Virtual Infrastructure manager 	--
R-CP-13 Interoperability with cloud platforms		SDN controller layer	SDN Controller <ul style="list-style-type: none"> Virtual Infrastructure manager 	SDN Controller <ul style="list-style-type: none"> A-CPI for the interaction with the cloud orchestrator
R-CP-14 Integration with external connectivity services	Support of traffic for inter-DC communications	SDN controller layer	SDN Controller <ul style="list-style-type: none"> Optical provisioning manager 	SDN Controller <ul style="list-style-type: none"> A-CPI in case on multi-domain interactions coordinated at the orchestrator level
R-CP-15 DCN reconfiguration for optimization strategies	Automated service re-provisioning tool. Algorithms manager: automated re-optimization engine based on different (strategic) criteria.	SDN controller layer	SDN Controller <ul style="list-style-type: none"> Virtual Infrastructure manager with re-planning features Optical provisioning manager with re-planning features Path computation with Global Concurrent Optimization (GCO) algorithms Monitoring 	SDN Controller <ul style="list-style-type: none"> A-CPI for configuration of operator policies D-CPI for collection of statistics about network performance
R-CP-16 CP architecture in support of scalable DCNs and network traffic	SLAs manager.	SDN controller layer	SDN Controller <ul style="list-style-type: none"> SLA manager Monitoring 	SDN Controller <ul style="list-style-type: none"> A-CPI for SLA specification D-CPI for collection of statistics or failure notifications

R-CP-17 Scalability	Per network resource optimizer functionality (leaning on virtualization). Per Network optimizer functionality (leaning on virtualization)	SDN controller layer	SDN Controller <ul style="list-style-type: none"> Virtual Infrastructure manager 	SDN Controller <ul style="list-style-type: none"> A-CPI for the interaction with the cloud orchestrator
----------------------------	--	----------------------	---	--

Table 10: Matching of the requirements to layers, functional components and interfaces.

9 Summary

This deliverable provides a comprehensive description of COSIGN innovations and scope by depicting the novel COSIGN layered architecture including the supported functionalities which address some of the future DCN requirements and challenges. Starting from a SoTA analysis in terms of DCN architectures and an extensive compilation of requirements, WP1 work team has identified the architecture building blocks, functionalities and interfaces. COSIGN vision is to move away from today's highly hierarchical, vendor specific, manually controlled DCs, which offer performance and escalation limited DC solutions towards flat, scalable, automatic and optimized DC infrastructures of tomorrow [1].

In terms of data plane, the different scenarios described in 5.1 are defined based on their prospective horizon for deployment/implementation. They all include innovative concepts: The **short-term scenario** puts emphasis on flattening the network by scaling the number of high speed ports on a switch. The switch is still essentially an electronic switch but several major changes have been made to it on both HW and SW aspects. On the HW aspect the optical interconnects have been moved from the edge of the board where they are traditionally located to the perimeter of the switch ASIC. This leads to great savings in power consumption as well as cost. Power is saved by making signal re-timers and repeaters unnecessary as well as allowing to lower the signalling voltages (and hence power) of the high speed signals (between SERDES and optical module). Cost is saved by reducing the number of optical modules needed, removing the retime ASICs and allowing the use of lower cost PCB technology since transmission lines (between ASIC and optical modules) are down to a few centimetres). On the SW aspect, the switches will run open flow which means they can be controlled using an SDN controller. This will lower the power consumption of the switching ASIC. With regards to the **mid-term scenario**, it includes all of the advantages of short-term scenario but offers to create an additional overlay network which can be configured on a longer time scale to answer needs of resource/bandwidth reallocation for specific and scheduled tasks in a dynamic manner. The technology for the overlay switches allows to scale the switch size to >500x500 ports which means that a single overlay switch can service a large number of traditional TORs. Since the optical overlay switch is also SDN controlled, one can orchestrate the TORs and overlay switches together for specific use cases. Regarding the **long-term scenario "A"**, the optical overlay switch will be augmented by small scale high speed optical switches. In this way, additional speed and flexibility can be introduced into the overlay network. Thus one can use the optical overlay not only for scheduled large flow switching but also for packet transmission and TDM media access control. The fast switches can also be integrated close to the servers and support meshing of servers. In the **long-term scenario "B"**, we are abandoning some of the main constraints in conceptualizing the network, such as cost and complexity and try to envision the most flexible node architecture and network topology possible. The key advantages behind this architecture are that it supports both SDM, WDM & TDM all with maximum granularity. While the eventual network node is very complex and most likely very expensive in today's technology, the many advantages of this architecture are interesting and will be explored in part by building parts of such a node as well as in simulation. The different scenarios proposed, will be used by WP5 partners as basis for the demonstration and validation of COSIGN.

Referring to the upper software layers of the COSIGN architecture; the **SDN-based control layer** enables the exploitation and integration of the innovative optical data plane in data centre environments. Control plane internal features are designed to maximize the utilization of the optical infrastructure while delivering connection services and virtual optical slices with high degrees of programmability, flexibility and dynamicity. The interfaces exposed at the north-bound are the key to achieve the full cooperation with between network control and cloud **orchestration layer** and enable joint decisions and operations over the whole set of DC resources. The orchestration layer proposes a unified framework for the joint management of network and IT resources.

Regarding the **benchmarking** of COSIGN against some reference **SoTA** related projects, it shows up a progressive adoption of optical technologies as basis to deploy DC network infrastructures at the

data plane level. Nevertheless, COSIGN constitutes the forefront of technological research an innovation, by also bringing SDN technology ,virtualization and the orchestration of IT and Network resources to efficiently provide with advanced DC services in a flexible and scalable way.

Also, from the analysis of **the impact that the requirements exert on the ambitions and objectives of COSIGN architecture** (see section 4.2.2), some conclusions can be also retrieved (see Figure 23):

- The requirements that have been retrieved are consistent and aligned with the different topics that COSIGN ambitions to improve by proposing the design.
- “Gain service control and management” and “increase network performance” ambitions are represented by a larger number of requirements, an aspect which reveals the strong focus and importance that has been conferred to address the use case’s needs. “Gain service control and management” is consistent with the cornerstone idea of bringing an SDN based network control plane (WP3) that is orchestrated (WP4) with the IT resources. Integrating SDN technology enables a better automation and control of the DC network, whereas the orchestration with the IT part of the solution brings a better resource management while providing services.
- SDN is not just about separating the control and data plane, but more importantly enabling a fully virtualized model. SDN exploits automation, streamlined operations and an unprecedented openness to achieve a dramatic improvement in time-to-new features and services, network simplification, and significant cost reduction. All those aspects clearly contribute to “increase network performance” and can be utilized by providing APIs to the orchestration layer to lastly benefit services provisioning and performance.
- Finally, the influence of the requirements identified for the data plane exerts impact on the ambitions and objectives less related to control and management. Data plane technologies and the added value they bring to the development of future DCNs constitute the basis of the architecture and the fact that COSIGN data plane strength relies on the interfacing and coordination with the control and orchestration layers to make the most of the optical technologies developed in WP2.




Figure 23: Per-layer identified requirements against COSIGN ambitions and overall objectives.

Appendix: Complete list of COSIGN and Future DCN architectures requirements

To ensure a suitable full analysis of the requirements and the implications they may impose to the COSIGN architecture design, there has been provided a template (see Figure 1) aiming to cover as many aspects as possible to make the best of each of them. The proposed template is organized as follows:

1. Requirement description.
2. Identify the KPI(s) that will enable us to measure that we are accomplishing such a requirement. Both quantitative and qualitative KPIs have been identified. For the quantitative ones, some COSIGN partners have provided some reference SoTA and expected future numerical values (concrete number or ranges) when possible. It must be taken into account that the numbers for some KPIs are not available or constitute private information.
3. Impact: Which impact do these business requirements have? (HIGH, MEDIUM, LOW) for NG-DC design?
4. Novelty: Which is the novelty of the Business requirements? (HIGH, MEDIUM, LOW) for NG-DC design?
5. Required Functionalities: Provide (if possible) suggestions about possible functionalities required to accomplish each requirement and the architecture layer (Data, Control, Management, Application) in which such functionalities could be allocated.
6. Layer: Which layers are supposed to be affected and should take this requirement into account.
7. Related Data Plane (DP) requirements: WP2 needs to know the consequences (physical requirements) that upper layer requirements may impose on the Data plane. This is really an important field of the template that may contribute to overcome and minimize the gap between SW and HW layers while defining interfaces and workflows among them in WP2 and WP3 specially. Not only the requirements coming from the SW layers have a direct impact on the HW layer, thus, each of the requirements has been tagged with an specific colour depending on whether the requirement has a direct impact on the physical HW layer or not. The colour code is as follows:
 - Red: There are no direct HW implications.
 - Yellow: It is not clear if there are direct HW implications or not.
 - Blue: There are HW implications.

<Requirement code>	<Name> <Description>
KPIs	
Impact	
Novelty	
Functionalities	
Layer	
Related DP requirements	

BUSINESS REQUIREMENTS

R-BUS-01 Resource and Services provisioning

R-BUS-01 Description	Resource and Services provisioning Self-service on-demand and/or automated deployment and provisioning of customizable virtual infrastructures, integrating virtual appliances, storage and network elements. Customers must be able to specify a virtual infrastructure beyond the pure IaaS model, including virtual appliances like firewalls, IPSs, VPN services, load balancers, network monitoring and analysis services in the service request.		
KPIs	<u>Quantitative KPIs</u> Service provisioning time	<u>SoTA value:</u> A self-service Interoute VDC can be enabled with 3 clicks. A virtual machine can be up and running and connected to the internet or a VPN in less than 5 minutes . Complex configurations may require the manual intervention of an operator.	<u>COSIGN target values</u> SoTA values are enough to ensure a good level of service. However, the same values must be valid for more complex VDC requests (e.g. including optical nodes)
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Virtual appliances that can be specified in the service request Enable certain level of customization of the virtual infrastructure service 		
Impact	High. For DC operators, self-service provisioning and automation are fundamental to limit the manual management of the DC infrastructure and optimize the usage of the available resources. Moreover, customers are often interested in complex infrastructures offering integrated advanced services, but easy to manage and configure.		
Novelty	Low in terms of self-service provisioning. Current virtual DC services allow the customer to specify the desired infrastructure and most of the procedures to setup VMs and simple network services are automated. However, more complex network configurations and the deployment or provisioning of additional virtual network appliances is not yet integrated and automated.		
Functionalities	<ul style="list-style-type: none"> <u>Service orchestration</u>: Orchestration of service components in the IaaS area, combined with virtualized network functions. <u>Automated provisioning</u> of cloud resources and multiple virtual network functions and applications. 		
Layer	Orchestration layer SDN controller layer – D-CPI and south-bound plugins.		
Related DP requirements	From the infrastructure perspective it means flexible reconfiguration capabilities that can be implemented at the controller layer or at the HW layer, or both. While defining the SDN framework (WP3) it will be needed to define which are the exact implications that this requirement exerts on the HW layer. In principle, this requirement has no direct HW implications.		

R-BUS-02 Elasticity

R-BUS-02 Description	Elasticity Self-service on-demand and/or automated scale-up and scale-down elasticity features, enabled through open and programmable APIs.		
KPIs	<u>Quantitative KPIs</u>	<u>SoTA value:</u> Dynamic scaling of a VM depends on the	<u>COSIGN target values</u>

	Time required to upgrade or downgrade the virtual infrastructure and re-configure the running services	hypervisor, the operating system (and the particular version), and the operating system source template used to deploy the VM. In some cases, only one of CPU or RAM scaling may be available. Alternative names for dynamic scaling include 'hot add' of RAM and 'hot plug' of CPU. In the current VDC, dynamic scaling can only be upwards for CPU-only, RAM-only, or CPU & RAM. Time depends on various conditions and is limited upwards to 5 minutes (time to provision a VM from scratch). See more info at the end of the appendix.	The target value should be in the same range of the SoTA values, but considering service elasticity related not only to VMs but also to network reconfiguration.
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Impact on the existing services; • Efficiency of scale-up and scale-down strategies in terms of resource utilization • Availability of real-time monitoring information through the APIs to determine the load and the performance of the service (for on-demand scale-up and scale-down) 		
Impact	High. Customers with service requests that are highly variable with the time period are interested in using infrastructures customized to their actual business needs, without paying fixed fees for an over-provisioned infrastructure. Moreover, some of them are interested in controlling and manipulating themselves the size and characteristics of the rented infrastructure, using their own algorithms and applications to define the best environments for their current service load. In this case, the system must provide interoperability with external applications through APIs that allow collecting monitoring information and dynamically regulating the dimension of the virtual infrastructure.		
Novelty	Medium. Existing cloud platforms already implement elasticity mechanisms; however the programmability of these operations through open APIs is still limited.		
Functionalities	<ul style="list-style-type: none"> • <u>Monitoring and detection</u> of service performance degradation. • <u>Update methods</u>: Dynamic modification of running virtual infrastructures possibly associated to re-optimization of resource placement. • <u>Open APIs</u> to get monitoring information (e.g. service performance or real-time availability of unused resource in the virtual infrastructure) and request infrastructure updates from external applications. 		
Layer	SDN controller – D-CPI and south-bound plugins.		
Related DP requirements	From the infrastructure perspective it means flexible reconfiguration capabilities that can be implemented at the controller layer or at the HW layer, or both. While defining the SDN framework (WP3) it will be needed to define which are the exact implications that this requirement exerts on the HW layer. In principle, this requirement has no direct HW implications.		

R-BUS-03 Monitoring Tools

R-BUS-03 Description	Monitoring tools Integration of configurable monitoring tools to allow DC operators to keep trace of resource usage and service performance.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> • Number, frequency and type of available monitoring measurements • MTR (Time required to detect failures, the 	<u>SoTA value:</u> <ul style="list-style-type: none"> • Every 5 minutes for monitoring measurements. • In the basic Interoute VDC service no specific SLA is publicly available for MTR on specific resources. Resource MTR KPIs may be negotiated 	<u>COSIGN target value:</u> <ul style="list-style-type: none"> • On demand granularity for monitoring info, instead of default frequency (still in the scale of

	location and recovery)	on a customer basis only for managed VDC solutions where additional professional services can be added to the VDC offer. However, currently, failure resolution requires the manual intervention of the operator. Depending on the level of service and on the severity level, the time for the initial response can vary from 1 hour to 8 hours	minutes). <ul style="list-style-type: none"> The MTR should be reduced through the automation of the recovery procedures. So the target is reducing the manual intervention. Acceptable values are in the range of few hours.
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Configurability of monitoring mechanisms 		
Impact	High. Monitoring is fundamental to evaluate the performance of the services, detect misbehaviours or failures, keep trace of the resource usage and verify the SLAs established with the customers (or federated partners if applicable).		
Novelty	Low. Monitoring mechanisms are already integrated in most of the cloud platforms. However, there are still limitations in the integration of additional monitoring tools oriented to the customers.		
Functionalities	<ul style="list-style-type: none"> <u>Monitoring</u> mechanisms. Integration of additional and customized monitoring tools, e.g. through automated deployment of dedicated virtual network applications. 		
Layer	SDN controller layer – D-CPI, south-bound plugins Data plane layer		
Related DP requirements	From the infrastructure perspective this requirement entails the need to monitor physical resources information and forward towards the SDN controller layer to be consumed by Orchestration and management layer. More specifically, it must be specified the parameters which can be monitored at the physical layer and also the interface to forward them towards the upper layers.		

R_BUS_04 Virtual infrastructure control APIs

R-BUS-04 Description	Virtual infrastructure control APIs Secure open and programmable APIs to enable advanced control of a running VDC service from customer's applications. An example of functions to be exposed by these APIs is monitoring, enforcement of automated elasticity rules or dynamic modification.
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Number, type and configurability of the functions to manipulate the VDC service, as exposed through the open APIs: The existing APIs of the Interoute VDC service (just used as example) is available at the following link: https://cloudstore.interoute.com/main/knowledge-centre/library/api-command-reference
Impact	High. Customers obtain more control on their own VDC and are able to adopt their own proprietary applications. These applications need to be easily modified to make use of the open and programmable APIs, to fully customize the VDC behaviour and characteristics.
Novelty	High. The programmability of the VDC can be extended with a richer set of functions and configuration options. The security of the APIs, with different levels of access, must be guaranteed.
Functionalities	<ul style="list-style-type: none"> <u>Open APIs</u> to allow external applications to programme the different components of a VDC instance.
Layer	SDN controller layer
Related DP requirements	From the infrastructure perspective it means flexible reconfiguration capabilities that can be implemented at the controller layer or at the HW layer, or both. While defining the SDN framework (WP3) it will be needed to define which are the exact

	implications that this requirement exerts on the HW layer. In principle, this requirement has no direct HW implications.
--	--

R_BUS_05 Accountability, billing and elastic SLAs management (UC1)

R-BUS-05 Description	<p>Accountability, billing and SLA management</p> <p>Monitoring and accounting mechanisms to support of multiple and flexible pricing models, including commit, pay-as-you-go or mixed models. Integrated mechanisms for end-to-end (network + cloud) SLA monitoring and verification.</p> <p>While building the DC network infrastructure for future DCs and providing new added value services, it is important to consider new more flexible and attractive accounting and billing mechanisms.</p> <p>Related to requirements (1 and 2), it would be interesting to enable the possibility of different types of users accessing to the service with different level of access on the infrastructure provided (premium, gold, silver, etc.) From a technical perspective, the solution could for instance provide with full network programmability and then hide or show depending on the user type.</p> <p>With regards the SLAs elasticity, it refers to Service Re-shaping (i.e. computing/storage capabilities) SLAs in an automated way. Two layers would be here implied: the virtual infrastructure layer while adding VMs and the physical layer while adding machines. The billing model that could be applied: “Pay as you go”.</p>		
KPIs	<p><u>Quantitative KPIs:</u></p> <p>Capability and time required to detect or predict SLA violations and time of reaction. SLA violations awareness time.</p>	<p><u>SoTA value:</u></p> <p>In the basic Interoute VDC service no specific SLA is publicly available for MTR on specific resources. Resource MTR KPIs may be negotiated on a customer basis only for managed VDC solutions where additional professional services can be added to the VDC offer. However, currently, failure resolution requires the manual intervention of the operator. Depending on the severity level, the time for the initial response can vary from 1 hour to 8 hours.</p>	<p><u>COSIGN target values:</u></p> <p>SoTA values should be reduced introducing more automation in service recovery and limiting the manual intervention. Acceptable values are in the range of few hours depending on the severity level</p>
	<p><u>Qualitative KPIs</u></p> <ul style="list-style-type: none"> Supported accounting models Flexibility of SLA models for complex VDC services SLA verification 		
Impact	<p>High. The flexibility of the pricing model is key to attract and maintain different categories of customers, each of them with widely different requirements. Moreover, the VDC performance and the compliance of the service with the guarantees negotiated in the SLAs are fundamental to enable customers to move their service platforms on the VDC.</p>		
Novelty	<p>High. Mechanisms for cloud SLA management need to be improved, especially in case of complex services with a variety of interacting components. They should integrate procedures to detect SLA failures in advance, to allow the system to properly react in time.</p>		
Functionalities	<ul style="list-style-type: none"> <u>Service monitoring</u>: Integrated network and cloud service monitoring. <u>Integrated SLA validation</u> for composed services. <u>Accounting mechanisms</u> for commit, pay-as-you-go and mixed models. 		

Layer	Orchestration layer. SDN Control plane layer.
Related DP requirements	From the infrastructure perspective it means flexible reconfiguration capabilities that can be implemented at the controller layer or at the HW layer, or both. While defining the SDN framework (WP3) it will be needed to define which are the exact implications that this requirement exerts on the HW layer. In principle, this requirement has no direct HW implications.

R_BUS_06 Privacy and security

R-BUS-06 Description	Privacy and security High-level of privacy and security, mechanisms for logical separation and high-availability.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Maximum number of coexistent and isolated VDC instances. Percentage of availability of the service. 	<u>SoTA values</u> <ul style="list-style-type: none"> No information available for VDC service Reference value: 99,7% 	<u>COSIGN target values</u> <ul style="list-style-type: none"> No information available for VDC service SoTA value is OK.
Impact	High. The privacy, isolation and high availability of the VDC service are a fundamental requirement for the acceptance of the service.		
Novelty	Medium. Mechanisms for virtualization and logical separation of service instances are already available. Efficiency of scheduled backups and service restoration in case of failure can be improved.		
Functionalities	<ul style="list-style-type: none"> <u>Isolation</u>: Logical separation of each customer's domain. At the network level examples of technologies to provide logical separation are VLAN tagging, MLPS encapsulation, BGP-based MPLS VPNs and Virtual Routing/Forwarding (VRF) tables. <u>Service replication and fast restoration</u> 		
Layer	Orchestration layer. SDN controller layer		
Related DP requirements	From the infrastructure perspective it means flexible reconfiguration capabilities that can be implemented at the controller layer or at the HW layer, or both. While defining the SDN framework (WP3) it will be needed to define which are the exact implications that this requirement exerts on the HW layer. In principle, this requirement has no direct HW implications.		

R_BUS_07 Virtual infrastructure administration and management interface

R-BUS-07 Description	Virtual infrastructure administration and management interface User-friendly interfaces for VDC administration and control with differentiated levels of (i) user authorization and (ii) resource configurability. The interface should support mechanisms for resource consumption monitoring, server administration and management of network and data physical location.		
KPIs	<u>Quantitative Values</u> <ul style="list-style-type: none"> Time to Transfer a VM (Or Data: e.g. 500TB). 	<u>SOTA Values</u> <ul style="list-style-type: none"> 500TB/10Gig + boot = ~5 days 	<u>COSIGN Target Values</u> <ul style="list-style-type: none"> ~10 Seconds
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Functions and configuration details supported by the administration and management interface. Ease of use of the administration and management interface. Granularity and configurability of the access levels. 		
Impact	High. Customers need to have an easy, secure and complete access to the main administrative and management functions of their own VDC instance.		

	Differentiated levels of authorization are required to avoid undesired and possibly disruptive configuration actions and to regulate the access to restricted and reserved information.
Novelty	Medium. Most of the cloud platform and VDC-like services offer a graphical user interface to enable the main management actions and visualize historical and real-time usage data.
Functionalities	<ul style="list-style-type: none"> • <u>Administration and management interfaces</u> with authentication and policy-based authorization procedures.
Layer	Management layer. SDN controller layer
Related DP requirements	From the infrastructure perspective it means flexible reconfiguration capabilities that can be implemented at the controller layer or at the HW layer, or both. While defining the SDN framework (WP3) it will be needed to define which are the exact implications that this requirement exerts on the HW layer. In principle, this requirement has no direct HW implications.

R_BUS_08 RAS

R-BUS-08 Description	<p>RAS</p> <p>Service reliability and availability, combined with network resilience. Support of mechanisms for automated and fast disaster recovery, integrated with on-demand or scheduled backups and replications of data in local or remote DC sites.</p> <p>RAS is one of the acronyms expanded in the beginning of the document. Means “Reliability, Availability, Serviceability”. From the networking HW perspective means failure detection and reporting capabilities and enough redundancy to enable the software to overcome the failures.</p>		
KPIs	<p><u>Quantitative KPIs</u></p> <p>Service Levels for the VDC Managed service by Interoute:</p> <ul style="list-style-type: none"> • Backup Recovery Point Objective (RPO: point in time of the last backup (or snapshot) that reflects the maximum amount of data loss due to the time intervals between backups): • Backup Recovery Time Objective (RTO: time required to recover to the most recent recovery point backup or snapshot): • EBS (External Block Storage) Snapshot recovery: 	<p><u>SoTA Values</u></p> <ul style="list-style-type: none"> • RPO: 24 hours • RTO: 2 hours +1 hour per 50GB of data recovered. • EBS: 4 hours 	<p><u>COSIGN target Values</u></p> <p>Customizable RPO (in a reasonable range of a few hours). The variable portion of the RTO could be reduced exploiting the capacity of the optical technologies (e.g. in the range of seconds).</p>
Impact	High. Service reliability and availability are features that need to be guaranteed for VDC services. Customers should also be able to request customized strategies for backups, in terms of frequency, number of replications or data locations, with differentiated prices.		
Novelty	Medium. Recovery mechanisms for network services are already available. However, the main novelty is in their integration with the overall service recovery and the quick reconfiguration of the network infrastructure to make use of the data replicated in other servers within the same data centre or in remote data centres. Moreover, integrated network configuration mechanisms should allow for an efficient utilization of the resources during scheduled backups, when huge amounts of data are transferred among servers or entire DCs in pre-planned time intervals.		
Functionalities	<ul style="list-style-type: none"> • Network or service <u>failure detection</u> • Automated network <u>service recovery</u> • Automated <u>network reconfiguration</u> in support of cloud service recovery, spanning across the local data centre or to external data centres, in case of remote data back-ups 		
Layer	Orchestration layer.		

	SDN controller layer
Related DP requirements	From the networking HW perspective means failure detection and reporting capabilities and enough redundancy to enable the software to overcome the failures.

R_BUS_09 Network Management Integration

R-BUS-09 Description	Network Management Integration Integration between cloud and network, to allow a fast transfer of huge amounts of data within and among data centres, as well as to users distributed around the world. Global connectivity is out of the scope; also, it is not clear how this requirement relates to the integration between the network and the rest of the cloud. May be it is a result of co-editing and needs fixing. On the other hand, We need to cover the local provisioning for the external connectivity and this includes sending data outside the DC and receiving it from the outside, catering both for user requests and data transfers. On HW level this means we need to provision bandwidth for the sake of external connectivity.
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Performance of network connectivity within a data centre and between different data centres during cloud management procedures (e.g. to transfer VMs or user contexts/data into different locations) • Performance of network connectivity when users access data on their own VDC from remote locations • Level of automation of network management within cloud infrastructure workflows
Impact	High. Users need to have a guaranteed and efficient access to their own data and services located or running on the VDC. On the other hand, management actions that involve transfer of large data need guarantees about the network performance. From a data centre operator perspective, the management of the network connectivity associated to these services should be fully integrated, to limit the manual intervention.
Novelty	Medium. Basic coordination of cloud and network services has been studied in previous projects. However COSIGN can analyse the adoption of SDN technologies in network management and a potential more advanced cooperation between cloud and network control, exploiting the network programmability offered by the SDN architectures.
Functionalities	<ul style="list-style-type: none"> • <u>Orchestration</u>: coordination of cloud and network services.
Layer	Orchestration layer SDN controller layer
Related DP requirements	There is a need to cover the local provisioning for the external connectivity and this includes sending data outside the DC and receiving it from the outside, catering both for user requests and data transfers. On HW level this means we need to provision bandwidth for the sake of external connectivity.

SERVICE REQUIREMENTS

R-SERV-01: Seamless Network Resources Provisioning

R-SERV-01 Description	Seamless Network Resources Provisioning Network resources should be provisioned to applications flexibly and transparently; underlying IT infrastructure should be configured automatically, to achieve fast (minutes) application/service deployment. Configuration time is a very relevant
------------------------------	--

	requirement for the data plane. It is most likely that this process will be dominated by the SW and not the HW delays.
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • E2E connectivity service established within the DC. • E2E service matches expected SLAs.
Impact	High. A suitable resources provisioning mechanism is key to ensure a proper service performance. It is important to ensure that the resources provided match services SLAs during the service lifecycle.
Novelty	Low. Resources provisioning is a well-known requirement.
Functionalities	<ul style="list-style-type: none"> • <u>SDN controller</u>: HW resource controlled by SDN controller layer. • <u>Resource APIS</u>: Resources' features and capabilities available to be consulted by the service and application layers.
Layer	Application/Service layer SDN controller layer – CPI, north-bound plugins.
Related DP requirements	Not related HW requirements.

R-SERV-02: Advanced Network Services Provisioning

R-SERV-02 Description	Advanced Network Services Provisioning The data centre should be able to provision advanced network services on demand, allowing dynamic provisioning and configuration, and including service chaining and orchestration for complex services. This requirement is mostly SW related and belongs to management layer. Some services can be implemented by specialized HW (like offload, acceleration, filtering). Building such HW is out of the scope of the project. Nevertheless, it is a requirement to be considered at the HW (independently of its feasibility to be implemented).
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • E2E connectivity service established within the DC. • E2E service matches expected SLAs.
Impact	High. A suitable resources provisioning mechanism is key to ensure a proper service performance. It is important to ensure that the resources provided match services SLAs during the service lifecycle.
Novelty	High. Making available specific on-demand and configurable network services making the most of optical HW at DCs is definitely new. Leaning on the SDN control and the management layers, it would be possible to provide with new features to network services in DC.
Functionalities	<ul style="list-style-type: none"> • <u>HW Configuration tools</u> to delegate and make available at user level the possibility to handle the HW as it was of its ownership. • <u>Cross layer mapping</u>: Mechanisms to map such advanced services to the application layer available.
Layer	Application / Service Layer Orchestration layer SDN controller layer Data plane layer
Related DP requirements	Some services can be implemented by specialized HW (like offload, acceleration, filtering).

R-SERV-03: Service level monitoring

R-SERV-03 Description	Service level monitoring In order to provide service layer monitoring, the monitoring mechanisms must be deployed at the physical infrastructure layer and forwarded to the upper service layer. The number of monitoring alarms forwarded to the service layer will be lower than the number of total alarms managed at the resource level and probably the
------------------------------	--

	alarms' information will be abstracted in the service layer. The data centre should facilitate application level monitoring required to decide whether service levels agreements are satisfied. There are also some requirements regarding the data collection in networking HW can be derived (through the controller APIs).
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Service performance real-time monitoring. • Service level matches expected SLAs.
Impact	High. Service monitoring is key to show service performance and enable service reconfiguration in case service degradation, for instance.
Novelty	Low.
Functionalities	<ul style="list-style-type: none"> • <u>Service monitoring</u> tools enabled, consuming network monitored parameters.
Layer	Application/Service layer Orchestration Layer SDN controller Layer. CP Northbound and southbound interfaces. Data plane Layer
Related DP requirements	Pure physical data monitoring, compilation and forwarding of networking HW.

R-SERV-04: Adapting to application/service dynamicity

R-SERV-04 Description	Adapting to application/service dynamicity Resources allocated for a service or an application should be dynamically re-provisioned according to the changes in demand and/or in requirements of the application/service.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> • Service reconfiguration as a consequence of modifying allocated resources. 	<u>SoTA Values</u> Dynamic service reconfiguration and scaling depends the VMs running the service, and VMs depends on the hypervisor, the operating system (and the particular version), and the operating system source template used to deploy the VM. In some cases, only one of CPU or RAM scaling may be available. Alternative names for dynamic scaling include 'hot add' of RAM and 'hot plug' of CPU. In the current VDC, dynamic scaling can only be upwards for CPU-only, RAM-only, or CPU & RAM. Time depends on various conditions and is limited upwards to 5 minutes (time to provision a VM from scratch).	<u>COSIGN target values</u> The target value should be in the same range of the SoTA values, but considering service elasticity related not only to VMs but also to network reconfiguration.
Impact	Medium. Automation of resources reallocation is highly desirable. Nevertheless manual mechanisms are already available, so that this feature would add value to the service, but is not a fundamental one.		
Novelty	High. Automated re-provisioning systems are not typical of operation environments. NOCs still feel reluctant to lose control over specific network provision mechanisms and due to the fact that the number of re-provisioning events is low, they prefer to keep it in a manual way.		
Functionalities	<ul style="list-style-type: none"> • <u>Service update methods</u>: Service auto-reconfiguration and re-provisioning tool available. • Include <u>resources re-provision mechanisms</u> as part of the monitoring tool. 		
Layer	Application/Service layer Orchestration layer SDN controller layer and Southbound CPI.		

	Data plane layer
Related DP requirements	The ultimate consequence of enabling re-provisioning mechanisms at the service layer, would consists on the effective re-provisioning at the HW layer. Therefore the same CPIs that handle HW layer resources should be enabled for dynamical reconfiguration.

R-SERV-05: QoS provisioning


R-SERV-05 Description	QoS provisioning The data centre should provide mechanisms to satisfy application's QoS requirements (e.g., BW, delay, and resiliency) and policies set to ensure the agreed service level. Again, the global SLAs must be translated into network QoS by the SW layers.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Achievable bandwidth with optical technologies Latency 	<u>SoTA value:</u> <ul style="list-style-type: none"> No information available No information available 	<u>COSIGN target value:</u> <ul style="list-style-type: none"> No information available No information available
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Service performance matching SLAs. 		
Impact	High. Service SLAs are fundamental in terms of service provisioning. Therefore it is important to keep service QoS above the minimum acceptable thresholds.		
Novelty	Low. QoS provisioning is a well-known addressed topic.		
Functionalities	<ul style="list-style-type: none"> Enable a <u>monitoring tool</u> that maps service metrics into levels of performance. 		
Layer	Application/Service layer Orchestration layer SDN controller layer		
Related DP requirements	Not direct related HW requirements.		

R-SERV-06: Real Time Control

R-SERV-06 Description	Real Time Control Control of network resources should allow real time detection and reaction to changes in the state of the infrastructure resources to cope with performance constraints. This will help reducing the impact of failures and unexpected behaviours. HW in the data plane should report faults to higher level– timely failure detection and reliable failure reporting.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Fast service performance recovery. 	<u>SoTA Values</u> Service Levels for the VDC Managed service by Interoute: <ul style="list-style-type: none"> RPO: 24 hours RTO: 2 +1 hours EBS: 4 hours 	<u>COSIGN target Values</u> <ul style="list-style-type: none"> Customizable RPO (in a reasonable range of a few hours). The variable portion of the RTO could be reduced exploiting the capacity of the optical technologies (e.g. in the range of seconds).
	<ul style="list-style-type: none"> <u>Qualitative KPIs</u>: Fault reporting. 		
Impact	High. Real time fault detection is fundamental to detect network service degradation.		
Novelty	Medium.		
Functionalities	<ul style="list-style-type: none"> <u>Automated re-provisioning mechanisms</u>: Linked to (or even embedded in) the monitoring tool, it could be included a mechanism to automatically react on real- 		

	time whenever changes are required.
Layer	Application/Service Level SDN controller layer Data plane layer
Related DP requirements	HW in the data plane should report faults to higher level– timely failure detection and reliable failure reporting. Additionally, decisions taken on the SDN control plane should take an effect at the HW layer to react upon service degradation.

R-SERV-07: Big data support

R-SERV-07 Description	Big data support Applications making use of big data should be supported by future data centres, having an impact on capacity, latency, access, security, cost and flexibility and requiring Big data analytics. This requirement has direct HW implications as it is fundamental to provide enough capacity to deal with big-data traffic.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Network energy consumption reduction while moving Data. Proper flow rates while moving large amount of Data. 	<u>SoTA Values</u> Information on energy consumption measurements and targets from Interoute are available only at the DC level. Data centre energy consumption model for Interoute data centres is based on the combination of the following parameters: <ul style="list-style-type: none"> IT space usage: 40% DC efficiency (server, network, cooling usage): 47% IT cooling usage: 39% IT power usage: 26% DC infrastructure availability (server, network and cooling uptime): 100% <p>The following graph shows the final KPI (red line), the minimum target (blue line) and the measurements (green line) for each parameter in the Interoute Berlin DC.</p>  <p>The DC operation (IT, network and cooling facilities) is responsible for nearly ¾ of the overall power consumption (73%). Within this segment, the network devices (including intra-DC switches and PoP devices to interconnect the DC to the Interoute pan-European network) are responsible of 4% of energy consumption and 11% of energy cost per year. For further details see</p>	<u>COSIGN target values</u> <ul style="list-style-type: none"> The target values are reported in the picture (blue and red lines). <ul style="list-style-type: none"> IT space usage: 80% DC efficiency (server, network, cooling usage): 80% IT cooling usage: 80% IT power usage: 80% DC infrastructure availability (server, network and cooling uptime): 100% <p>With regards the flows while moving Data, it highly depends on the underlying infrastructure.</p>

		LIGHTNESS D2.4 “Analysis of the application infrastructure management using the proposed DCN network architecture”.	
		With regards the flows while moving Data, it highly depends on the underlying infrastructure.	
Impact		High. When it comes to Big Data, network infrastructure is in fact one of the top priorities. No matter how you define Big Data, it's all about large volumes of data that need to move around a network. Bandwidth, latency, and the importance of network infrastructure are key relevant to Big Data initiatives.	
Novelty		High. Big Data is a relative new concept and the movement of large volumes of data within and across DCs is becoming more important.	
Functionalities		<ul style="list-style-type: none"> • <u>Algorithms manager</u>: Algorithm based on relevant metrics for Big Data support monitoring. • <u>Energy efficiency mechanisms</u>: Moving data is bad for data. It really cost a lot of money in terms of energy. 	
Layer		Application/Service layer Orchestration layer SDN controller layer	
Related DP requirements		Not related HW requirements, others than providing monitoring data to the upper layers.	

R-SERV-08: Secure data management

R-SERV-08 Description	Secure data management Insurance of the end-to-end, secure use of data within applications as well as secured access to remote data sources.		
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • DC security checklist applied, ranging from room security access to brute force and sneak attacks. • Preventive controls • Detective controls • Corrective controls. 		
Impact	Very High. Service security in DCs (at any level, not only at service level) is fundamental for Cloud related business. As simple, as if services are not secure, the client will move to the competence.		
Novelty	Low.		
Functionalities	<ul style="list-style-type: none"> • <u>Security NFVs</u> complementary to current HW security. • Data and service <u>isolation mechanisms</u>. • <u>Authorization, authentication and availability mechanisms</u> to ensure secure data management. 		
Layer	Application/Service layer Orchestration layer SDN controller layer		
Related DP requirements	Typical HW later security requirements which may impact to the service layer.		

R-SERV-9: Security management

R-SERV-9 Description	Security management Provide a way of managing security mechanisms for a coordinated physical, network, data and user security for different stakeholders (user, service provider, cloud provider).		
-----------------------------	--	--	--

KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • DC security checklist applied, ranging from room security access to brute force and sneak attacks. • Preventive controls • Detective controls • Corrective controls.
Impact	Very High. Service security in DCs (at any level, not only at service level) is fundamental for Cloud related business. As simple, as if services are not secure, the client will move to the competence.
Novelty	Low.
Functionalities	<ul style="list-style-type: none"> • Security NFVs complementary to current HW security. • Data and service <u>isolation mechanisms</u>. • Authorization, authentication and availability mechanisms to ensure secure data management.
Layer	Application/Service layer Orchestration layer SDN controller layer
Related DP requirements	Typical HW later security requirements which may impact to the service layer.

R-SERV-10: Self-service management interfaces

R-SERV-10 Description	Self-service management interfaces Several interfaces are required for the management of the services and applications deployed on the data centre. The interfaces should allow the configuration of self-service policies, notifications, information retrieval and monitoring and the configuration of the elements of the service, independently of other services running in the system.		
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Full orchestrated services management based on previous described tools and functionalities. • Automation of management procedures. 		
Impact	Medium.		
Novelty	High. Management operational tasks are commonly controlled by NOCs and processes automation and features are not so common.		
Functionalities	<ul style="list-style-type: none"> • Self-service <u>management and configuration Interfaces</u> 		
Layer	Orchestration layer		
Related DP requirements	Not related HW requirements.		

R-SERV-11: Multi-tenant isolation

R-SERV-11 Description	Multi-tenant Isolation Each application or service deployed in the cloud should be sufficiently isolated from the rest of the workloads, in terms of data isolation, management isolation, and performance isolation. DC networks isolation techniques strongly rely on virtualization mechanisms and abstraction tools.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> • Maximum number of coexistent and isolated VDC instances. 	<u>SoTA values</u> <ul style="list-style-type: none"> • This value depends on the DC facilities. The objective is to maximize the DC 	<u>COSIGN target values</u> <ul style="list-style-type: none"> • This value depends on the DC

		utilization.	facilities. The objective is to maximize the DC utilization.
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Implementation of per-tenant service isolation mechanisms. 		
Impact	High. Service isolation is important to ensure service integrity and make clients trust on the provided service security.		
Novelty	Medium. Service isolation is not a new topic; Nevertheless the large variety of isolation technologies based on virtualization which are arising are quite novel and add value to service isolation at different architectural layers.		
Functionalities	<ul style="list-style-type: none"> • Implement <u>partitioning/slicing techniques</u> at different levels: <ul style="list-style-type: none"> • NIC • L2 (VLAN) • L3 Network using L3 (VXLAN, GRE) • L3 Network using L3 (MPLS, GRE, IPSec) 		
Layer	Application/Service layer SDN controller layer Data plane layer		
Related DP requirements	Multi-tenant isolation can be done at different layers. Performing such mechanisms at physical layers is a possibility.		

R-SERV-12: Service Traffic Profile (From UC2)

R-SERV-12 Description	Service Traffic Profile Each service or application to be deployed in the Data Centre should provide information about its expected traffic profile and connection requirements, so that the DC network can be configured to offer the required connectivity at service runtime. More specifically, traffic profile information and requirements may include the following parameters: max packet loss, minimum guaranteed bandwidth, peak rate, jitter and latency, type of connectivity (point-to-multi-point, point-to-point, asymmetric/symmetric connectivity), recovery strategies, and redundancy. Since VMs may be instantiated on the same DC or distributed among different DCs, the service traffic profile may impact in the intra-DC connectivity or also the inter-DC connectivity.		
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Type of traffic profile parameters supported in the service description 		
Impact	Medium. Traffic profile metrics must be included in the service description, handled at the orchestrator level and mapped into specific parameters in the requests towards the SDN controller for network configuration. At the control level, this requirement is related to R-CP-02, with high impact on CP components (A-CPI and internal mechanisms for customized network provisioning).		
Novelty	Medium. Current architectures offers pre-defined cloud service templates with the possibility to specify a limited number of network-related parameters.		
Functionalities	<ul style="list-style-type: none"> • <u>Interface for service specification</u> • <u>Service orchestration</u> with coordinated network configuration. 		
Layer	Application/Service layer		
Related DP requirements	Not direct related HW requirements.		

R-SERV-13: Generalized information model (From UC2)

R-SERV-13	Generalized information model The specification of cloud applications and services should be based on a		
------------------	---	--	--

Description	generalized information model, able to describe a wide variety of applications in a powerful, but user-friendly and abstracted manner, completely technology-independent. The information model must be able to support complex topologies and interactions between the logical entities composing the service. Moreover, it should support the description of the traffic profiles specified in R-SERV-13.
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Level of service abstraction enabled by the information model. • Complexity of the scenarios that can be specified through the information model.
Impact	Medium. Some existing languages for cloud service specification and abstract modelling already exist. They may be properly extended to include the additional characteristics and metrics of novel application environments.
Novelty	Medium. Generalized information models for cloud service specification are already widely used. However, the main innovation is in the flexibility of the model so that it can be adapted to describe a wide variety of application environments, with different levels of complexity, abstraction and details about traffic profiles and network topology.
Functionalities	<ul style="list-style-type: none"> • <u>Interface for service specification</u>
Layer	Application/Service layer
Related DP requirements	Not direct related HW requirements.

R-SERV-14: Service connectivity disjunction (From UC2)

R-SERV-14 Description	Service connectivity disjunction It would be nice to be able to specify two fully disjoint connections in order to guarantee redundancy. This requirement at the service level may impact also the infrastructure layer. It was discussed about the convenience of this requirement due to the implications it imposes at the infrastructure layer. This is a “nice to have in the future” requirement that probably we will not be able to incorporate in the implementation stage.
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Service redundancy based on disjointed network connectivity.
Impact	High. On one hand, this constitutes a nice way to guaranty service redundancy. In the other hand, the costs derived for maintaining such redundancy with regards the HW layer, makes difficult to justify this requirement.
Novelty	Low. Redundant connections are well-known protection mechanisms. Nowadays, restoring mechanisms seem to be more effective and reasonable.
Functionalities	<ul style="list-style-type: none"> • <u>Restoration mechanism</u> offering disjointed path provisioning.
Layer	Application/Service Level Data plane layer
Related DP requirements	Redundant connections at physical layer.

R-SERV-15: VDC optical BW control (or control of physical network aspects of the IaaS service provided) (From UC1)

R-SERV-15 Description	VDC optical BW control (or control of physical network aspects of the IaaS service provided) VDC operator should be able to control the optical BW (e.g. enable access to the optical wavelength or slot). In current VDC architectures, VIOs (Virtual Infrastructure Operators) are able to provide with the virtual connection and that's it. One step beyond is also to enable them with access to the BW of the connections they are providing so that VIOs are also enable to sell Capacity. VIO is the acronym for Virtual Infrastructure Operator.
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • A VIO has full control on the connectivity service.

	<ul style="list-style-type: none"> VIO monitors, controls and sells the BW of the network connectivity service.
Impact	High: Ability for VIOs to sell capacity on-demand and modify the features of the connections they are provisioning. Most probably this cannot be done today, but it would entail better control on specific resources from the upper layers' perspective.
Novelty	High. Enable to VIOs with the opportunity to fine tune specific physical layer parameters, is not currently an available feature.
Functionalities	<ul style="list-style-type: none"> <u>Cross layer mechanism</u> to enable this feature. This may involve the HW, control and management layers.
Layer	Application/Service Level Orchestration layer SDN controller layer Data plane layer
Related DP requirements	Enable with the a functionality at the service level that enables VIOs to own specific rights to control and tune certain optical features of the connections they are providing. "Same level of control as if you were working on the physical devices". From WP2 it would be expected to know the features at the infrastructure layer that can be exposed to the upper layers, and which ones would be interesting to have and which ones should the control not be delegated to under any circumstance to avoid problems.

R-SERV-16: Service Calibration (From UC1)

R-SERV-16 Description	Service Calibration This requirement may enable users to own the resources that he EXACTLY wants to. This is especially interesting from the application layer perspective.		
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> DC End-user is capable to control the physical HW provided by the DC network infrastructure provider. 		
Impact	High. Enabling to end users with the control of a network HW equipment which is shared with other end users is a huge challenge, but a desirable long-term feature		
Novelty	High. To the best of our knowledge, multi-tenancy mechanisms do not yet enable such feature.		
Functionalities	<ul style="list-style-type: none"> <u>Application layer interface</u> to enable DC tenants to control physical HW network resources. 		
Layer	Application/Service Level Orchestration layer SDN controller layer Data plane layer		
Related DP requirements	At HW layer, there should be exposed the interfaces that enable to control the physical resource to the upper (in this case application) layers. Besides, this should be combined with all the virtualization and abstraction mechanisms that enable current multi-tenancy and isolation features.		

R-SERV-17: Intelligent/selective physical failure awareness (From UC1)

R-SERV-17 Description	Intelligent/selective physical failure awareness Depending of the gravity and seriousness of the failures that may take place at the physical layer, this should be hidden or noticed to the IaaS service consumer. Two main service behaviours: <ul style="list-style-type: none"> The physical failure does not impact users' SLAs → the user must not be aware of the failure. The physical failure impacts users' virtual infrastructure or the system cannot immediately recover → It must be noticed. 		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Resources MTR (Time required to detect failures, the location and 	<u>SoTA Values</u> <ul style="list-style-type: none"> Depending on the level of service and on the severity level, the value can vary from 	<u>COSIGN target values</u> <ul style="list-style-type: none"> SoTA values could be reduced by

	recovery)	1 hour to 8 hours	an half
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Forwarding of physical layer alarms/notifications towards the service/user layer. 		
Impact	Medium: It is an added value feature that adds value to VIOs (Virtual Infrastructure Operators) to manage and control the services performance.		
Novelty	High: To the Best of this consortium knowledge, this feature is not enabled for VIOs while providing services. VIO is the acronym for Virtual Infrastructure Operator.		
Functionalities	<ul style="list-style-type: none"> Service functionality <u>filtering failure awareness to service consumer</u> according to specific metrics which measure the seriousness of the failure. Integration with monitoring tools. 		
Layer	Application/Service Level Orchestration layer SDN controller layer Data plane layer		
Related DP requirements	HW monitoring information must be available to be forwarded to the upper service layer.		

R-SERV-18: equipment inventory services tool (From UC3)

R-SERV-18 Description	equipment inventory services tool It would be desirable to enable to service consumers with a catalogue of the whole suite of DC network services enabled, the list of options which each service includes and their functionalities. Service customer is this way aware of the whole network connectivity services that are possible to be managed in the DC which is making use of.		
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Service catalogue available for service consumer. 		
Impact	High. This is a nice feature to have from the “usability” perspective. It adds value to the overall service at the time it enables to service consumers with additional information on DC network services.		
Novelty	Medium. Usually service catalogues only provide a high level overview on the available services, without getting into too much detail. While providing to a deeper control and management		
Functionalities	<ul style="list-style-type: none"> <u>Service catalogue</u> visible at application layer 		
Layer	Orchestration layer		
Related DP requirements	Not related HW requirements.		

INFRASTRUCTURE LEVEL REQUIREMENTS **MANAGEMENT & ORCHESTRATION LAYER** **REQUIREMENTS**

R-MGMT-01: Element Management

R-MGMT-01 Description	Element Management Management and inventory of network equipment, equipment management, maintenance, administration, configuration, and performance monitoring. This means that all the devices must be able to lend themselves to this type of management – to have management interface, to be able to report its features,		
------------------------------	--	--	--

	negotiate its settings, etc., through some of the standard or extended protocols).
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Ask for the topology • Ask for routers appliances • Ask for routing algorithms (the consumer would be the customer consuming the VI (Virtual Infrastructure) operated by the VIO (Virtual Infrastructure Operator)). • Request for resources at specific locations: This requirement comes from the application layer. A certain application provider needs 5 co-allocated and 5 de-allocated machines to ensure the proper behaviour of the application. • VLAN control and map visualization.
Impact	High. This feature enables a more intuitive management of the elements/resources of the service.
Novelty	Medium. There are already some Cloud and Network management tools which include monitoring and infrastructure visualization tools based on system queries.
Functionalities	<ul style="list-style-type: none"> • <u>Service GUI/CLI</u>: It should be enabled a visualization functionality (GUI/CLI, etc.) to the service consumer, to access this information. • <u>Management interface</u>: to manage individual devices deployed in the Data Centre.
Layer	Orchestration layer SDN controller layer Data plane layer
Related DP requirements	HW resources must be able to lend themselves to this type of management – to have management interface, to be able to report its features, negotiate its settings, etc., through some of the standard or extended protocols).

R-MGMT-02: Operator Network Management mechanisms

R-MGMT-02 Description	Operator Network Management mechanisms Operator-scoped mechanisms for provisioning, monitoring, and control of network paths and circuits. Network paths must enable setting, monitoring, and enforcing the QoS and QoE attributes, like bandwidth, latency, amount of redundant paths between the source and the destination, etc.
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Network admin is able to request for DCN resources provisioning. • Network admin is able to operate (monitor, modify) the characteristics of an already existing provisioning service. • Network admin is able to finish a provisioned service.
Impact	High. The full control of DC networking resources is a MUST HAVE feature, very much appreciate to fully operate the DCN infrastructure.
Novelty	Medium. Despite network resources management and operation are not novel requirements, in DCN environments it is not so common network administrators experiencing the full management of the network resources while provisioning services.
Functionalities	<ul style="list-style-type: none"> • <u>Network topology manager</u> • <u>Network services provisioning interface.</u>
Layer	Orchestration layer
Related DP requirements	Not related HW requirements.

R-MGMT-03: Self-Service multi-tenancy

R-MGMT-03	Self-Service multi-tenancy
Description	Each tenant of the DC must be enabled to self-manage the networking on behalf of its workload, in a way totally independent of the physical network infrastructure management and of the management of the other tenants' networks. On the other hand, simple and expectable out-of-the-box experience should be provided for tenants unwilling or unable to take care of their own networking administration. This is high level requirement, with no direct link to HW.
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Several network slices provisioned to different tenants and managed without colliding. • The tenant fully manages the requested DCN slice by means of a GUI.
Impact	High. Added value in the service provided to DCN tenants. An additional management functionality that could be provided as a service.
Novelty	High. Slicing of virtual network mechanisms and the possibility to offer the management of such slices to the tenants constitutes a novel requirement. In practice, it's not foreseen to be incorporated in the short-term, since DCN network admins would not be very keen on this to happen.
Functionalities	<ul style="list-style-type: none"> • <u>Virtual tenant network</u> to provide network virtualization support. • <u>Virtual network management as service</u> (in order to delegate network slices management to tenants). • <u>Rest API for utilizing the VTN</u>
Layer	Orchestration layer SDN controller layer
Related DP requirements	This is high level requirement, with no direct link to HW.

R-MGMT-04: Automatic resources discovery

R-MGMT-04	Automatic resources discovery
Description	Resources maybe added or removed from the infrastructure and it is highly desirable from the management point of view to automatically identify when a resource is up or down in the system, so that it becomes part of the pool of resources available/no longer available to be utilized.
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Network topology is automatically updated when a new network resource is configured in the DCN. • IT facilities are updated when a new storage or computation facility is enabled in the DC.
Impact	Medium. The automation of this type of processes are nice to have features, but not fundamental for the proper behaviour of the overall management system.
Novelty	Low. This is an already well known requirement.
Functionalities	<ul style="list-style-type: none"> • <u>Topology manager</u> • <u>Topology inventory</u>
Layer	Orchestration layer SDN controller layer Data plane layer
Related DP requirements	This is high level requirement. Nevertheless, the new physical resource deployed in the DCN environment requires enabling some mechanisms (beacons, broadcast, and "hello" messages) to make the upper layers aware of their presence. Thus, it has an impact on the HW layer.

CONTROL PLANE LAYER REQUIREMENTS

R-CP-01 Automated provisioning of intra-DC connectivity

R-CP-01 Description	Automated provisioning of intra-DC connectivity The COSIGN CP must provide mechanisms for an automated DCN configuration to establish and destroy intra-DC connectivity services on-demand (triggered by external requests) or in support of internal procedures (e.g. re-optimization of DCN resource allocation, connectivity restoration, etc.).		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Connection setup time. Connection tear-down time. 	<u>SoTA Values</u> <ul style="list-style-type: none"> No specific measure available. It must be taken into account that the time it will take to establish/tear-down connectivity will be limited by the SW (control and/or orchestration) and not the physical layer. Here also it should be clear that optical technologies are not superior in time to establish a link but will deliver improved latency due to time-of-flight nature of the connection. No specific measure available. Same reason as the connection setup time. 	<u>COSIGN target values</u> <ul style="list-style-type: none"> No specific measure available No specific measure available
Impact	Medium. The granularity of connection provisioning must be handled considering the technologies available at the data plane.		
Novelty	Medium. The novelty relies in the SDN approach applied to the DCN, where the provisioning of the DCN resources is controlled through the SDN controller(s) and in the automation of the provisioning.		
Functionalities	<ul style="list-style-type: none"> <u>Network service provisioning</u>: Computation and provisioning of optical paths with user-constraints (QoS, service resiliency, scheduling) 		
Layer	SDN controller – network application and A-CPI.		
Related DP requirements	Mainly a control plane requirement. At the data plane level, R-DP-01 to R-DP-03 are enablers for this requirement.		

R-CP-02 Support for customizable network services

R-CP-02 Description	Support for customizable network services The COSIGN CP must provide mechanisms to provide intra-DC connectivity services compliant with a variety of user-level constraints, including QoS parameters (e.g. bandwidth), timing constraints, service resilience guarantees.		
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Supported service constraints. 		
Impact	High. Requires a service interface at the A-CPI powerful enough to specify the desired constraints and control plane mechanisms to support them at the provisioning level (e.g. mechanisms for QoS configuration, network service scheduling, recovery, etc.).		
Novelty	Medium. Current architectures are able to provide pre-defined and basic network services.		
Functionalities	Connection provisioning service; <ul style="list-style-type: none"> <u>Additional services</u> as required by the specific constraints (scheduling, recovery, etc.); Correlated services like extended <u>inventory services</u> or <u>computation algorithms</u> dedicated to QoS-enabled connections, able to operate on scheduled resources or 		

	to compute disjoint routes, etc.
Layer	SDN controller layer – mainly network applications and A-CPI, but also core services.
Related DP requirements	Mainly a control plane requirement. At the data plane level, R-DP-01 to R-DP-03 are enablers for this requirement.

R-CP-03 Dynamic Support for multiple automated connectivity paradigms

R-CP-03 Description	Dynamic Support for multiple automated connectivity paradigms In order to support a variety of virtual infrastructure topologies, the COSIGN CP must provide mechanisms to establish point-to-point, point-to-multipoint or anycast connectivity, in unidirectional or bi-directional, symmetric or asymmetric mode.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Connection setup time for P2P, P2MP and anycast connections. Connection tear-down time for P2P, P2MP and anycast connections. 	<u>SoTA Values</u> <ul style="list-style-type: none"> No specific measure available. It must be taken into account that the time it will take to establish/tear-down connectivity will be limited by the SW (control and/or orchestration) and not the physical layer. Here also it should be clear that optical technologies are not superior in time to establish a link but will deliver improved latency due to time-of-flight nature of the connection. No specific measure available. Same reason as the connection setup time. 	<u>COSIGN target values</u> <ul style="list-style-type: none"> No specific measure available No specific measure available
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Efficiency of network resource utilization 		
Impact	Medium. Requires a service interface at the A-CPI able to specify multiple connectivity models and control plane mechanisms to support them at the provisioning level.		
Novelty	High. Current architectures are able to provide pre-defined and basic network services.		
Functionalities	<u>Network service provisioning</u> + correlated services (computation algorithms for P2MP or anycast services).		
Layer	SDN controller layer – network applications and A-CPI		
Related DP requirements	The control of the P2P, P2MP and anycast connectivity is beyond the scope of the data plane. However, the efficiency of the resource allocation may require suitable architectures of the network topology.		

R-CP-04 Multi-layer operation of COSIGN data plane optical technologies

R-CP-04 Description	Multi-layer operation of COSIGN data plane optical technologies The COSIGN CP must be able to efficiently operate the heterogeneous optical technologies adopted at the COSIGN DCN data plane. This means it must be able to handle their different constraints (even through resource virtualization), as well as
----------------------------	--

	the matching among the resource granularities offered for each specific technology. Multi-layer mechanisms must be adopted to coordinate the cross-technology resource allocation and maximize the efficiency for the whole DCN utilization, still in compliance with the requirements of the running services. Where possible, the usage of open, standard protocols for the interaction with the underlying physical technologies should be preferred, adopting dedicated extensions where required.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Granularity and dynamicity of end-to-end connectivity services. 	<u>SoTA Values</u> <ul style="list-style-type: none"> No specific measure available 	<u>COSIGN target values</u> <ul style="list-style-type: none"> No specific measure available
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Efficiency of resource allocation in multi-layer environments (e.g. through the sharing of a connection established at a server-layer among multiple connections at a client-layer with higher granularity). 		
Impact	High. Requires additional internal functionalities within the SDN controller in order to manage the provisioning and re-adaptation of the multi-layer connectivity. At the interface level, it requires the common extensions to manage optical resources at the D-CPI.		
Novelty	High. Current architectures usually provide static network services, without exploiting the multi-layer capabilities of the transport technologies.		
Functionalities	<ul style="list-style-type: none"> <u>Optical resource virtualization</u> + multi-layer connection service provisioning + dynamic re-allocation of the resources providing the connectivity at the server-layer. <u>Computation and provisioning of point-to-point, point-to-multipoint or anycast connectivity</u> 		
Layer	SDN controller – D-CPI, south-bound plugins and core services (for resource abstraction). SDN controller – resource virtualization service. SDN controller – network application for multi-layer service provisioning and dynamic re-allocation. Agents to interact with the optical devices.		
Related DP requirements	Multiple optical technologies with different granularities needs to be available at the data plane (linked to R-DP-01 and R-DP-03 requirements).		

R-CP-05 DCN Resource usage optimization

R-CP-05 Description	DCN resource usage optimization The COSIGN CP must provide mechanisms to provide intra-DC connectivity services compliant with a variety of operator-level constraints, including load-balancing strategies, energy efficiency, paths with minimum cost, etc.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Time required by the computation algorithms to select the most efficient set of resources with the given constraints. 	<u>SoTA Values</u> <ul style="list-style-type: none"> No specific measure available. Time to elaborate the network configuration for a new VM is included in the time required to setup the VM itself. With reference to the VDC service developed in the LIGHTNESS project and limited to the network aspects, the algorithm runs in the VDC composition application with an execution time of 35.12 ms. 	<u>COSIGN target values</u> <ul style="list-style-type: none"> The target values should be in the same range of the SoTA, also when dealing with the new COSIGN technologies.
Impact	Medium. It requires the implementation of suitable optimization algorithms within the SDN controller. The parameters to be used as input for these algorithms needs to be notified at the D-CPI level when related to the transport technologies (e.g.		

	power consumption parameters) or configured through management interfaces when related to operator policies (e.g. “cost” of a link between two network nodes).
Novelty	High. The optimization of the resource usage allows the operator to increase the overall utilization of its Data Centre (e.g. to accommodate more services or users) while reducing the operating costs.
Functionalities	<ul style="list-style-type: none"> • <u>Computation algorithms</u> and <u>policy management</u> within the SDN controller to operate the DCN with operator-level constraints
Layer	SDN controller layer – network application for resource selection. SDN controller layer – network application for policy management. SDN controller layer – A-CPI (for policy configuration). SDN controller layer – D-CPI (support of resource parameters to be used as input for resource selection algorithms).
Related DP requirements	The CP algorithms may need to operate on parameters that must be made available from the hardware (e.g. energy consumption parameters).

R-CP-06 DCN Elastic intra-DC connectivity

R-CP-06 Description	Elastic intra-DC connectivity The COSIGN CP must be able to dynamically scale-up and scale-down the capacity of the established connections, in support of the elastic features of the cloud services. The decisions about connectivity upgrade/downgrade may be taken internally at the CP layer or coordinated by upper-layer entities (e.g. at the orchestration level), but always in compliance with the established SLAs.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> • Time required to detect an inefficient condition in the resource allocation for a given service (at the CP layer) and initiate a modification procedure. 	<u>SoTA Values</u> <ul style="list-style-type: none"> • Dynamic service reconfiguration and scaling depends the VMs running the service, and VMs depends on the hypervisor, the operating system (and the particular version), and the operating system source template used to deploy the VM. In some cases, only one of CPU or RAM scaling may be available. Alternative names for dynamic scaling include 'hot add' of RAM and 'hot plug' of CPU. In the current VDC, dynamic scaling can only be upwards for CPU-only, RAM-only, or CPU & RAM. Time depends on various conditions and is limited upwards to 5 minutes (time to provision a VM from scratch). However, depending on the level of service and on the severity level, the value can vary from 1 hour to 8 hours 	<u>COSIGN target values</u> SoTA values could be reduced by an half
Impact	Medium. It requires the support of modification requests at the A-CPI level and internal function within the SDN controller to support the modification of existing connections.		
Novelty	High. This feature allows applying the same elasticity concepts commonly used for the IT resources in cloud environments to the network domain, with great benefits in terms of service dynamicity and automatic adaptation to the user requirements.		
Functionalities	<ul style="list-style-type: none"> • <u>Update methods</u>: Modification of connection services on-demand or based on internal triggers. Dynamic modification of established connectivity or Virtual Infrastructures on-demand or based on automatic elasticity rules. 		
Layer	SDN controller layer– A-CPI and network applications.		
Related DP requirements	These are mainly CP requirements, but it is associated to the dynamic reconfiguration capabilities of the data-plane (see R-DP-03).		

R-CP-07 Network monitoring

R-CP-07 Description	Network monitoring The COSIGN CP must be able to provide monitoring functionalities for network resource usage, network service performance and faults detections.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Time required to detect and solve failures. 	<u>SoTA value:</u> <ul style="list-style-type: none"> No public information available. In the basic VDC service no specific SLA is publicly available for MTR on specific resources. Resource MTR KPIs may be negotiated on a customer basis only for managed VDC solutions where additional professional services can be added to the VDC offer. However, depending on the level of service and on the severity level, the value can vary from 1 hour to 8 hours 	<u>COSIGN target values:</u> <ul style="list-style-type: none"> On demand granularity for monitoring info, instead of default frequency (still in the scale of minutes).
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Statistics on the DCN resources usage and performance. 		
Impact	Medium. It requires the support of monitoring information at the D-CPI level. It requires the support of monitoring information at the management interface level.		
Novelty	Low. This feature allows to promptly react to any failure that could happen in the DCN as well as it allows to apply policies (e.g. load balancing) to better adapt to the user requirements while maximizing the utilization of the network resources. The main novelty relies on the usage of the SDN approach.		
Functionalities	Monitoring services and statistics on the DCN resources usage.		
Layer	SDN controller layer – D-CPI, Management		
Related DP requirements	It is associated to the dynamic reconfiguration capabilities of the data-plane (see R-DP-03), for recovery purposes in this case.		

R-CP-08 Programmable APIs

R-CP-08 Description	Programmable APIs The main configuration options and functions (e.g. service provisioning and tear-down, modification, queries for monitoring data) offered by the COSIGN CP must be exported through programmable APIs (e.g. based on the REST paradigms), with different levels of capabilities depending on authorization profiles. These APIs should allow exposing some (limited) functionalities directly to the users and enable an easy integration with the overall DC control and management platform.
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Set of configuration options and functions exported supported by the programmable APIs.
Impact	High. It requires the support of programmable APIs to enable users to perform provisioning and modification actions. It requires the support of the requests triggered by the users at the A-CPI level
Novelty	High. Current architectures do not allow exporting network statistics and information to users.
Functionalities	<ul style="list-style-type: none"> APIs: Programmable APIs to request provisioning of network connectivity, even over virtual slices
Layer	SDN controller layer – A-CPI and network applications.

Related DP requirements	This is mainly a CP requirement. No HW implications.
--------------------------------	--

R-CP-09 Support of network service monitoring and accounting

R-CP-09 Description	Support of network service monitoring and accounting The COSIGN CP must produce and expose through a suitable management interface a set of monitoring information about DCN resource usage and allocation, as required to support cloud service accounting for various pricing models (e.g. pay-as-you-go or commit model).
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Co-existence of network services with different pricing models
Impact	Medium. It requires the support of management interfaces between SDN controller and OSS to gather DCN resources and allocation and usage statistics.
Novelty	Medium. This feature allows gathering information about the network usage to support the various pricing models.
Functionalities	<ul style="list-style-type: none"> <u>Monitoring</u>: Make available information about the usage and performance of DCN resources
Layer	Orchestration layer SDN controller layer – D-CPI
Related DP requirements	This is mainly a CP and management layer requirement.

R-CP-10 Support of scheduled network connectivity

R-CP-10 Description	Support of scheduled network connectivity The COSIGN CP should be able to support scheduled or periodical connectivity services, between intra-DC resources or in support of connectivity among resources located in different DCs. Suitable synchronization procedures, in cooperation with upper layer entities (e.g. at the orchestration level), must be provided to coordinate the enforcement of the overall scheduled actions across the different DC resources. However, in-advance resource reservations at the network level (i.e. without actual configuration) should be planned and automatically updated to guarantee resource availability and optimize the global DCN utilization among the shared physical resources.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Time required to automatically provide scheduled connectivity services. 	<u>SoTA Values</u> <ul style="list-style-type: none"> No specific measure available. It must be taken into account that the time it will take to establish/tear-down connectivity will be limited by the SW (control and/or orchestration) and not the physical layer. Here also it should be clear that optical technologies are not superior in time to establish a link but will deliver improved latency due to time-of-flight nature of the connection. 	<u>COSIGN target values</u> <ul style="list-style-type: none"> No specific measure available
Impact	High. It requires the support of automated configuration of the DCN resources for the scheduled events requiring scheduled connectivity services. It requires the support of cooperation through the A-CPI interface between the SDN controller and the orchestration level.		
Novelty	High. This feature allows achieving a global optimization of the shared DCN resources to provide both scheduled and on-demand connectivity services.		

Functionalities	<ul style="list-style-type: none"> • <u>Computation and provisioning</u> of scheduled connectivity or virtual slices.
Layer	SDN controller layer – A-CPI and orchestration level.
Related DP requirements	This is mainly a CP requirement. No HW implications

R-CP-11 Network service resilience

R-CP-11 Description	Network service resilience The COSIGN CP must be able to detect network service failures and, where possible, react in an automated manner through fast connection recovery procedures. Depending on the established SLAs and the original service requests, protection or restoration mechanisms can be applied. When network service restoration procedures are not applicable, asynchronous failure alerts must be produced and notified to the upper layer (e.g. to the orchestration level) to enable recovery escalation strategies. Simple restoration mechanisms can probably be implemented directly in the physical HW (or close to it). Restoration based on SLA is beyond the scope of data plane alone.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> • Time to automatically recover from a failure is support of the established SLA and original service characteristics. • Time to detect and notify to the orchestration level the failure occurrence. 	<u>SoTA value:</u> <ul style="list-style-type: none"> • No information available. In the basic VDC service no specific SLA is publicly available for MTR on specific resources. Resource MTR KPIs may be negotiated on a customer basis only for managed VDC solutions where additional professional services can be added to the VDC offer. • No information available. 	<u>COSIGN target values</u> <ul style="list-style-type: none"> • On demand granularity for monitoring info, instead of default frequency (still in the scale of minutes).
Impact	Medium. It requires the support of (fast) recovery actions at both the A-CPI and D-CPI level.		
Novelty	Medium. The novelty relies on the usage of the SDN approach and the overall orchestration of the DCN resources.		
Functionalities	<ul style="list-style-type: none"> • <u>Failure detection and notification</u> mechanisms: Detection and notification through the D-CPI and A-CPI of failure occurrence. • <u>Automated procedures for network service recovery</u> 		
Layer	D-CPI - SDN controller – A-CPI and network applications.		
Related DP requirements	This is mainly a CP requirement, but it is associated to the dynamic reconfiguration capabilities of the data-plane according to the defined recovery strategy (see R-DP-03).		

R-CP-12 Multi-tenant isolation

R-CP-12 Description	Multi-tenant isolation The COSIGN CP must enforce suitable virtualization mechanisms to enable the sharing of a common physical network infrastructure among fully isolated connectivity services.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> • Given the DCN resources, the number of fully isolated virtual slices (tenants). 	<u>SoTA values</u> <ul style="list-style-type: none"> • It depends on the amount of dedicated resources, and the specs of such resources. 	<u>COSIGN target values</u> <ul style="list-style-type: none"> • It depends on the amount of dedicated resources, and the specs of such resources.
	<u>Qualitative KPIs</u>		

	<ul style="list-style-type: none"> Usage and performance statistics of the DCN resources for each tenant (virtual slice).
Impact	High. It requires the support of virtualization techniques for the efficient allocation of the shared DCN resources.
Novelty	High. This feature allows maximizing the efficient allocation of the shared DCN resources.
Functionalities	<ul style="list-style-type: none"> <u>Resource virtualization</u>: Virtualization of DCN resources isolated virtual slices allocation.
Layer	SDN controller layer– A-CPI and network applications.
Related DP requirements	This is mainly a CP requirement.

R-CP-13 Easy interoperability with existing cloud management platforms for unified DC control.

R-CP-13 Description	<p>Easy interoperability with existing cloud management platforms for unified DC control.</p> <p>The COSIGN CP should provide powerful APIs, possibly based on the REST concept and HTTP protocol, to enable an easy integration with existing orchestration systems and cloud management platforms adopted to handle the overall DC management. Where relevant, the usage of open, standard interfaces should be preferred to guarantee an easy interoperability.</p>
KPIs	<p><u>Qualitative KPIs</u></p> <ul style="list-style-type: none"> Joint cloud and network management provisioning tool available.
Impact	Medium. Cloud and Network management services are separated items. Usually IT guys are not experts on Network service provisioning too. Therefore, it would be interesting to create a tool which enables to centralize both management systems
Novelty	High. With the new advancement of cloud and network management tools, it has appeared the possibility to achieve such interoperability by exposing cloud management platforms APIs.
Functionalities	<ul style="list-style-type: none"> Cloud and Network management interface. Integration with cloud management platforms for NaaS provisioning
Layer	Orchestration layer SDN controller layer.
Related DP requirements	It's not be directly linked to HW

R-CP-14 Integration with external connectivity services (inter-DC)

R-CP-14 Description	<p>Integration with external connectivity services (inter-DC)</p> <p>The COSIGN CP must be able to configure the intra-DC network to efficiently support also the traffic generated among computing resources located in different data centres. This traffic could have various characteristics and requirements, since it could belong to running cloud applications distributed across different sites or it could be generated by management procedures, e.g. for inter-DC content replication.</p> <p>Mechanisms and interfaces for integrated management of intra-DC and inter-DC connectivity should be supported, with reference to different deployment and business models (e.g. inter-DC connectivity offered by an external provider or by the same administrative entity that manages the DC infrastructure).</p> <p>Open standard interfaces and protocols should be preferred where applicable to enable and simplify the (multi-domain) interoperability.</p>
----------------------------	--

KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> E2E network service orchestration tool available. E2E inter-domain Network service delivery with guaranteed QoS.
Impact	High. This feature may enable to control the intra-DC network connectivity and provisioning service but also to coordinate with other connectivity services managed by other NOCs, and agree on specific SLAs to guarantee the support of multi-domain services.
Novelty	Medium. Not directly, but Operator agreements and SLAs are carried out to support the traffic from each other's clients. In this case, it's not proposed to only propose to support network connectivity, but performing such support according to the SLAs of the DC network services.
Functionalities	<ul style="list-style-type: none"> Support of traffic for inter-DC communications
Layer	Orchestration layer SDN controller layer
Related DP requirements	Not be directly linked to HW

R-CP-15 Dynamic DCN reconfiguration for optimization strategies

R-CP-15 Description	Dynamic DCN reconfiguration for optimization strategies The COSIGN CP should support the automated re-planning of already established network services. The automated re-planning will bring the possibility to autonomously re-adapt the resource allocation to the dynamicity of the real-time DC loads, according to global re-optimization criteria. However, DCN re-configuration procedures must avoid any disruption of the existing services.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Reduction of the total number of SLAs violations indicator. 	<u>SoTA values</u> <ul style="list-style-type: none"> No information available. 	<u>COSIGN target values</u> <ul style="list-style-type: none"> No information available.
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Queue congestion reduction at network resources. Better DC network services overall performance. 		
Impact	Medium. Network resources and services re-planning usually impacts network services performance while dealing with the changes. DC networks resources are commonly over-provisioned, so that re-planning strategies are not very popular among DC network managers, especially to avoid services disruption.		
Novelty	Medium. Network services re-planning are already there and usually involve Load balancing techniques or overprovisioning of network resources.		
Functionalities	<ul style="list-style-type: none"> <u>Automated service re-provisioning tool.</u> <u>Algorithms manager:</u> Automated re-optimization engine based on different (strategic) criteria. 		
Layer	SDN controller layer– D-CPI. Orchestration layer		
Related DP requirements	Not be directly linked to HW.		

R-CP-16 CP architecture in support of scalable DCNs and network traffic

R-CP-16 Description	CP architecture in support of scalable DCNs and network traffic The COSIGN CP must be designed to efficiently operate DCNs of different sizes, scaling well with an increasing number of servers and network devices and an increasing amount of intra-DC and inter-DC cloud application and management
----------------------------	---

	traffic with flows of different dimension. Scalability is a major concern in the data plane considerations. However, this requirement is here more connected to applications and management traffic. It is more about management than about ports and links. Scalability can depend on the amount of tenants, deployed applications, users, etc.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> • Total number of SLAs violations indicator. 	<u>SoTA values</u> <ul style="list-style-type: none"> • No information available 	<u>COSIGN target values</u> <ul style="list-style-type: none"> • No information available
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Capacity plan available • Resources utilization: Percentage of network resources in use. 		
Impact	High. Properly scaling network service in DCs is fundamental to ensure basic DC services operation in both intra and inter DC environments. It is as simple as if the network do not scale, DC services cannot be provided within the SLA agreed levels, and consequently, “you are out of business”		
Novelty	High. Traditionally, the scalability of network services relies on the HW layer. A proper monitoring of network resources together with the intelligence applied at SDN control and management layers may also enable to enhance the scalability of network services.		
Functionalities	<ul style="list-style-type: none"> • SLAs manager. 		
Layer	SDN controller layer – D-CPI. Orchestration layer.		
Related DP requirements	Not be directly linked to HW.		

R-CP-17 Scalability

R-CP-17 Description	Scalability The scalability of the CP is a key requirement for the proper operation of the data centre. Therefore, the size (in terms of servers and optical devices to be managed) as well as the expected huge number of traffic flows among servers should not affect the properly working of the CP operation.
KPIs	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> • Capacity plan available • Availability of information regarding resources utilization: Percentage of network resources in use.
Impact	High. Properly scaling network service in DCs is fundamental to ensure basic DC services operation in both intra and inter DC environments. It is as simple as if the network do not scale, DC services cannot be provided within the SLA agreed levels, and consequently, “you are out of business”.
Novelty	Medium. Network service scalability is a MUST from the very beginning of a DC design. Nevertheless, novel virtualization techniques and the improvement of CP mechanisms enable with new approaches and reduce DC complexity through fewer physical connections. Through virtualization, you get better resource utilization which at the end turns into a better scalability solution.
Functionalities	<ul style="list-style-type: none"> • <u>Per network resource optimizer functionality</u> (leaning on virtualization). • <u>Per Network optimizer functionality</u> (leaning on virtualization)
Layer	SDN controller layer - D-CPI. Orchestration layer. Data plane layer.
Related DP requirements	The HW layer must combined with CP techniques must provide to the upper layers with the information to determine the scalability of DC network resources to be able to manage the DC network services.

DATA PLANE/HW LAYER REQUIREMENTS

R-DP-01 Capacity

R-DP-01 Description	Capacity This requirement specifies both the aggregated and the link level DCN capacity. At link level DCN must support 10G to server today and in the near future. In COSIGN horizon, 40G to server links must be considered as well. On an aggregated level, we have to consider the amount of server ports that have to be supported in typical DCs. Few hundred thousands of servers are typical within the world's largest DCs, bringing the aggregated capacity requirement to be considered in COSIGN to millions of Gb/s.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> data rate supported by server NIC card, switch interface and fibre/cable switch dimension and port density 	<u>SoTA values</u> IBM values <ul style="list-style-type: none"> Switch dimesion:2-4 Units per rack Switch port density: 24 per Unit Server NIC card rate: 10G Tue values: <ul style="list-style-type: none"> Switch dimension: TOR switches we have shrunk a typical 1RU, 19" wide switch to the size of 20x20cm board. This represents a port density increase for a single 1RU box of about x6. Polatis values <ul style="list-style-type: none"> Switch port density: 192 duplex ports. 	<u>COSIGN target values</u> IBM Values <ul style="list-style-type: none"> Switch dimesion:1/2 units per rack Switch port density: 196 per Unit Server NIC card rate: 100G <u>Tue Values</u> <ul style="list-style-type: none"> Switch dimension: Future work on compact transceivers carried out in WP2 can lead up to a density improvement of x10.
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Section bandwidth 		
Impact	High. The DCN capacity planning would have direct impact on amount of services that the DCN could accommodate (which reflects the cost efficiency) and their performance (e.g., implementation/response time, end-to-end delay).		
Novelty	High. In previous DCN design, EPSs with different dimension and capability are structured in a hierarchical way, which causes the fundamental capacity constraints of DCN. By employing fibre with high bandwidth/spectrum efficiency and optical switch with high data forwarding/transmission capability in DCN, these constraints would be tackled in a cost-efficient and power-efficient way.		
Functionalities	<ul style="list-style-type: none"> <u>Resource monitoring</u>: Network resource utilization monitoring functionality 		
Layer	Data plane layer		
Related DP requirements	The HW component and topology design must ensure the scalability of DCN.		

R-DP-02 Latency

R-DP-02 Description	<p>Latency</p> <p>Depending on the application, very low (microsecond) latencies can be required to some types of traffic, while some other types can thrive with longer (tens of milliseconds) response times. It is therefore of the outmost importance to be able: 1) to provide the lowest possible latencies for the chosen flows and 2) to be able to distinguish the flows requiring the low-latency paths.</p> <p>“Always on” low latency connectivity is crucial. The reasoning for that is that from a DC perspective, it would be valuable to always have “a little bit” of bandwidth available between the nodes, but the bandwidth should also be highly flexible so that it can be adjusted (cranked up or down) with very low latency.</p>		
KPIs	<p><u>Quantitative KPIs</u></p> <ul style="list-style-type: none"> Round Trip Time (RTT) and Jitter Average hop count of server-to-server path Switching delay (go through the switch) 	<p><u>SoTA values</u></p> <ul style="list-style-type: none"> Round Trip Time (RTT) and Jitter: No information available. Switching delay: It must be taken into account that the time it will take to establish/tear-down connectivity will be limited by the SW (control and/or orchestration) and not the physical layer. The average hop count is also a function of the size of the DC. As we have been able to shrink down the size and power consumption of the 128 port electronic switch (TOR) we can imagine trying to package 6 of them into a single 1RU box creating effectively a 256 switch. While this switch will still have poorer latency than a true 256 silicon ASIC, it will allow to build larger networks with less layers/hops improving the eventual number of hops KPI. Also including the OCS from Polatis can drastically lower the number of hops if we use it to bypass some of the aggregation layer switches. Even if used only on the top layer of the tree the hop count should at least improve with one hop. 	<p><u>COSIGN target values</u></p> <ul style="list-style-type: none"> Round Trip Time (RTT) and Jitter: No information available. It depends on the specs of the resources available. Switching delay: No information available. Average hop count: No information available.
Impact	High. Latency is a critical QoS guarantee for cloud service/application		
Novelty	Medium. Packets suffer from unpredictable delay caused by queuing and congestion in current EPS based DCN design, and reconfiguration time of device and network (e.g., path) Also, topology design flatten way		
Functionalities	<ul style="list-style-type: none"> <u>Per flow based service monitoring</u> functionality 		
Layer	Data plane layer		

Related DP requirements	The equipment used at the HW layer should ensure guaranteed low latency.
--------------------------------	--

R-DP-03 Reconfigurability/Flexibility

R-DP-03 Description	Reconfigurability/Flexibility <p>Traffic flow characteristics will have a significant impact on the network performance. Most flows are small <10 KB and last only a few 100 of milliseconds, requiring the network to be re-provisioned at a very high rate.</p> <p>Resource usage optimization is required for profitability. Resource optimization from the infrastructure owner perspective can come in conflict with the optimization goals of the deployed services. For example, workload optimizers tend to increase the amount of instances when the service experiences a peak in demand; for that it might be required to power on standby servers. Taking into account the wear and tear of frequent power on and power down operations is typically not part of the consideration of the workload manager, although it can be of outmost importance to the infrastructure operator.</p>		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Switch reconfiguration time Service/traffic blocking rate 	<u>SoTA values</u> <ul style="list-style-type: none"> Switch reconfiguration time: It must be taken into account that the time it will take to establish/tear-down connectivity will be limited by the SW (control and/or orchestration) and not the physical layer. Here also it should be clear that optical technologies are not superior in time to establish a Service/traffic blocking rate: The blocking rate is dependent on the cross section bandwidth in the network and the choice of the architecture be it: blocking, re-arrange able non block or strict non-blocking. The addition of an optical switching layer or optical overlay can on one hand remove some of the blocking/congestion but will take away ports from the electronic switches reducing the cross section bandwidth. The study of eventual blocking rate is probably application specific and architecture specific. It is not clear whether the hybrid architectures will provide a clear improvement in blocking rate for traffic. 	<u>COSIGN target values</u> <ul style="list-style-type: none"> Switch reconfiguration time: No information available. Service/traffic blocking rate: No information available.
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Path (e.g., bandwidth) provisioning, flexibility and time efficiency 		
Impact	Medium. The reconfigurability and flexibility of DCN would significantly influence its service/traffic accommodation capability, since the traffic in DC (between ToRs and servers) varies quickly in time and space domain.		
Novelty	Medium. Through employing large scale optical switch, fast optical switch and spatial division multiplexing (SDM) based fibres (e.g., MEF, MCF), the dimension of network resource and end-to-end path provision will be increased significantly (e.g., hybrid SDM, TDM/OPS connection).		
Functionalities	<ul style="list-style-type: none"> Underlay (physical layer) resource utilization/performance <u>monitoring</u> Flexible <u>end-to-end path provision functionality</u> 		

Layer	Data plane layer
Related DP requirements	It is required to have flexible configurable paths.

R-DP-04 Resiliency and HA

R-DP-04 Description	Resiliency and HA DCN data plane need provide high service availability and minimize the amount of systemic downtime events, i.e. it is required to have n fully redundant network paths between each pair of endpoints.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Amount of disjoint path between end points (ToRs or servers) 	<u>SoTA values</u> <ul style="list-style-type: none"> 1 or 2 	<u>COSIGN target values</u> <ul style="list-style-type: none"> N (Number of Spines)
Impact	High. An interruption of just seconds to normal data access can result in enormous cost to the business and if lengthy, may impact it to such a degree that it cannot recover. By building resilience into the DCN infrastructure, the risk of service interruption could be reduced.		
Novelty	Medium. In DCN structure design, the infrastructure resiliency should be considered properly to recover from hardware layer failure. And a lot of previous research on survivable optical network and DCN design could be refereed.		
Functionalities	<ul style="list-style-type: none"> <u>HW Failure recovery</u> functionality 		
Layer	Data plane layer.		
Related DP requirements	It is required to have fully redundant network paths between each pair of endpoints. This is more related with the topology.		

R-DP-05 Traffic isolation

R-DP-05 Description	Traffic isolation DCN data plane should be capable of isolating the traffic on the prescribed granularity – workload owner, application, application transaction, application tier, etc. In addition to the physical isolation, the management and the performance isolation must be provided. Isolation here read more like a logical request (virtual) then a physical request. But there is there also a requirement for physical isolation of parts of the data plane in the DC. For instance, physical isolation can be required in some cases (e.g. where we do not use overlays). This can be ensured by fully isolated paths, by using multiple cores in fibres, different wavelengths, etc.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Granularity of resource allocation 	<u>SoTA values</u> <ul style="list-style-type: none"> No information available 	<u>COSIGN target values</u> <ul style="list-style-type: none"> No information available
Impact	High. Physical network isolation (e.g., path isolation) could provide a more trustable traffic isolation to support multi-tenant cloud service.		
Novelty	High. SDM technology and optical switches with different features could provide different scale physical network isolation. Comparing with overlay based traffic isolation, it could provide more privacy, security, and robustness to isolation failure to client.		
Functionalities	<ul style="list-style-type: none"> <u>Overlay mechanisms</u>: Combined overlay-based and underlay (physical)-based network virtualization 		
Layer	SDN controller layer Data plane layer		

Related DP requirements	Physical isolation might be required in some cases (e.g. where we do not use overlays). This can be ensured by fully isolated paths, by using multiple cores in fibres, different wavelengths, etc.
--------------------------------	---

R-DP-06 Scalability and Extensibility

R-DP-06 Description	Scalability and Extensibility DCN data plane should support large-scale DC and allow the existing DC to grow organically, both in number of servers, in number of the deployed workloads, supported amount of traffic, etc.		
KPIs	<u>Quantitative KPIs</u> <ul style="list-style-type: none"> Performance/cost linearity to involve more servers/switches 	<u>SoTA values</u> <ul style="list-style-type: none"> No information available 	<u>COSIGN target values</u> <ul style="list-style-type: none"> No information available
	<u>Qualitative KPIs</u> <ul style="list-style-type: none"> Topology/structure stability for involve more servers/switches 		
Impact	High. The DCN scalability implicates the DCN performance stability (e.g., latency) when extension happens, as well as the cost efficiency.		
Novelty	High. The DCN design with optical switches should be investigated to better utilize the scalability of optical switch dimension (e.g., spatial, spectrum) and the capability of large scale optical switch.		
Functionalities	<ul style="list-style-type: none"> N/A 		
Layer	Data plane layer		
Related DP requirements	The scalable structure design should allow an easy way to extend with modular subsystem, as well guarantee its flexibility.		

Additional information on specific “R-BUS-02 Elasticity” requirement KPIs

(Time required upgrading or downgrading the virtual infrastructure and re-configuring the running services)

Dynamic scaling of a VM depends on the hypervisor, the operating system (and the particular version), and the operating system source template used to deploy the VM. In some cases, only one of CPU or RAM scaling may be available. Alternative names for dynamic scaling include 'hot add' of RAM and 'hot plug' of CPU. In the current VDC, dynamic scaling can only be upwards for CPU-only, RAM-only, or CPU & RAM. Time depends on various conditions and is limited upwards to 5 minutes (time to provision a VM from scratch).

The following table shows OS templates that have been tested for correct performance in dynamic scaling. However there are some specific limitations to be aware of:

- There is a RAM limitation for Linux 64-bit machines: a running VM with initial RAM 512 MB or 1 GB can be scaled upwards only as far as 2 GB. To go further it is necessary to stop the VM and change the RAM to 4GB.
- For VMs with more than 8 CPUs on the ESXi hypervisor, the VM may not be able to access in full the additional performance of an added unit of CPU, and memory performance can also be affected.

Test results for vertical scaling							
Hypervisor	OS	Template name	Static scaling (UP/DOWN)	Dynamic scaling (UP)	Dynamic scaling (DOWN)	Notes	
ESXi	Windows 2008 R2	Windows Server 2008 R2 (64-bit)	CPU and RAM	CPU and RAM	No	ESX limitation on dynamic scaling down	
ESXi	Windows 2012	Windows 2012	CPU and RAM	CPU and RAM	No	ESX limitation on dynamic scaling down	
ESXi	Windows 2012 R2	Windows 2012 R2	CPU and RAM	CPU and RAM	No	ESX limitation on dynamic scaling down	
ESXi	Windows 2012 R2 (non-Internet activation)	Windows 2012 R2 M	CPU and RAM	CPU and RAM	No	ESX limitation on dynamic scaling down	
ESXi	CentOS 6.5	IRT-CENTOS-6.5 (04/04/2014)	CPU and RAM	Not working	No	Scaling fails with this template	
ESXi	CentOS 6.4	Centos 6.4 (64-bit) [found in Community tab]	CPU and RAM	CPU and RAM	No	ESX limitation on dynamic scaling down	
ESXi	Debian 7.4	IRT-DEBIAN-7.4 (04/04/2014)	CPU and RAM	No	No	Not enabled for Interoute template**	
ESXi	Ubuntu 12.04	IRT-UBUNTU-12.04 (07/04/2014)	CPU and RAM	No	No	Not enabled for Interoute template**	
ESXi	Red Hat (RHEL) 6.5	IRT-REDHAT-6.5 (16/04/2014)	CPU and RAM	No	No	Not enabled for Interoute template**	
ESXi	pfSense 2.1.3	IRT-PFSENSE-2.1.3	CPU and RAM	No	No	Not enabled for Interoute template**	

** For Linux OS where dynamic scaling is not enabled, it is possible for the user to deploy from ISO image source to enable dynamic scaling.