



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 D1.1

Requirements for Next Generation intra-Data Centre Networks Design

Due date of deliverable: June 30th, 2014

Actual submission date: July 21st, 2014

Start date of project: January 1, 2014

Duration: 36 months

Lead contractor for this deliverable:
IBM, Katherine Barabash

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

Executive Summary

This document surveys and analyses the recent and the anticipated Data Centre (DC) trends and derives the requirements these trends bring towards the internal Data Centre Network (DCN).

The trends surveyed include the business purpose of the data centre, i.e. public or private clouds, enterprise, campus, service and network providers; the scale and growth patterns, the usage patterns, e.g. Infrastructure as a Service (IaaS), Platform as a Service (PaaS), or Software as a Service (SaaS). In addition, we survey the workload types that must be supported by the Data Centres and the services that are expected from the Data Centre Network.

The analysis is applied to tie the Data Centre level use-cases to the concrete requirements towards the Data Centre Network, including both the service-level and the infrastructural requirements.

The requirements derived include the Network as a Service (NaaS) level requirements, i.e. what applications need from the virtual network services they consume, as well as the Infrastructure level requirements, i.e. what Data Centre operators need from the network as a resource that can be flexibly allocated, monitored, and managed.

The purpose of this document is to provide the input to the subsequent stages of the project responsible to come up with the architectural design for the future Data Centre Network. That is why in this document we deliberately avoid surveying the existing and the emerging tools and technologies, leaving it to another task being executed in parallel, namely task T1.2.

Document Information

Status and Version:	Draft 1.6	
Date of Issue:	21/07/2014	
Dissemination level:	Public	
Author(s):	Name	Partner
	Katherine Barabash	IBM
	Jose Soler	DTU
	Sarah Ruepp	DTU
	Michael Berger	DTU
	Matteo Biancani	IRT
	Alessandro Predieri	IRT
	Giada Landi	NXW
	Nicola Ciulli	NXW
	Eduard Escalona	I2CAT
	Jose Ignacio Aznar	I2CAT
	Bingli Guo	UNIVBRIS
	Shuping Peng	UNIVBRIS
	Reza Nejabati	UNIVBRIS
	Georgios Zervas	UNIVBRIS
	Dimitra Simeonidou	UNIVBRIS
	Salvatore Spadaro	UPC
Edited by:	Katherine Barabash	IBM
Checked by :	Sarah Ruepp	DTU

Table of Contents

Executive Summary	2
Table of Contents	4
1 Introduction.....	5
1.1 Reference Material	5
1.1.1 Reference Documents	5
1.1.2 Acronyms and Abbreviations	5
1.2 Document History	6
2 Data Centre Trends and Usage Patterns	7
2.1 Data Centre Growth Trends.....	7
2.2 Data Centre Usage Patterns and Business Models	10
2.2.1 Classification of cloud services	10
2.2.2 Data Centre Usage by Business Sectors	17
2.3 Prevalent Data Centre Applications	18
2.3.1 Traditional Data Centre Applications	18
2.3.2 Migrating Traditional DC Application to the Cloud.....	21
2.3.3 Cloud-born Applications	23
3 Analysis of the Data Centre Network Use Cases.....	25
3.1 Virtual Data Centre.....	25
3.2 Multi-tenant Software Cloud	28
3.3 Data Centre Operations, Management, and Orchestration	30
4 Limitations and Challenges of the Current Data Centre Network Design.....	33
5 Data Centre Networks Requirements	37
5.1 Business Requirements.....	37
5.2 Service Level Requirements	39
5.3 Infrastructure Level Requirements	40
5.3.1 Data plane	40
5.3.2 Control Plane	43
5.3.3 Management and Orchestration	45
6 Classification and Prioritization of Requirements.....	47
7 Conclusions and Architectural Discussion	51
7.1 Major Components of the COSIGN Solution.....	51
7.1.1 Software Defined Networking	51
7.1.2 Network Virtualization	52
7.1.3 Orchestration	52
7.1.4 Optical DCN Technologies.....	53
7.2 COSIGN Architectural Blueprint	53

1 Introduction

1.1 Reference Material

1.1.1 Reference Documents

[1]	COSIGN FP7 Collaborative Project Grant Agreement Annex I - "Description of Work"
-----	--

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.

API	Application Programming Interface
BGP	Border Gateway Protocol
BSS	Business Support Services
CAGR	Compound Annual Growth Rate
CE	Customer Edge
CP	Control Plane
DC	Data Centre
DCN	Data Centre Network
DCIM	Data Centre Infrastructure Management
DMZ	Demilitarized Zone
IaaS	Infrastructure as a Service
IETF	Internet Engineering Task Force
ISO	International Organization for Standardization
IPS	Intrusion Prevention System
MOOC	Massive Open Online Course
MPLS	Multi-Protocol Label Switching
NaaS	Network as a Service
NIST	National Institute of Standards and Technology
OSS	Operational Support Services
PaaS	Platform as a Service
PCI DSS	Payment Card Industry's Data Security Standard
PE	Provider Edge
PoD	Point of Delivery
RAS	Reliability, Availability, Serviceability
RFC	Request For Comments
SaaS	Software as a Service
SDE	Software Defined Environments
SDN	Software Defined Networking
SLA	Service Level Agreement
TCAM	Ternary Content-Addressable Memory
TCO	Total Cost of Ownership
ToR	Top of the Rack
VDC	Virtual Data Centre
VLAN	Virtual LAN
VPN	Virtual Private Network
VRF	Virtual Routing/Forwarding

1.2 Document History

Version	Date	Authors	Comment
00	13/03/2014	See the list of authors	TOC first draft
01	24/03/2014		TOC updated after partners review
02	06/04/2014		TOC restructured to suit the official deliverable template and agreed with the WP leader
03	24/04/2014		TOC and initial contents imported to the official deliverable template
03_IRT_VDC	28/04/2014	IRT	VDC section from IRT
03_IRT_new	30/04/2014	IRT	Added sections 2.2 and 5.1
03_IRT_NXW	03/05/2014	NXW	CP requirements
03_Bristol_and_DTU	05/05/2014	Bristol	DP requirements
04	06/05/2014	See the list of authors	Integrated v03 contributions by IBM (2.1 and 2.3), IRT (2.2, 3.2, and 5.1), NXW (5.3.2), DTU and Bristol (5.3.1)
04_i2CAT	12/05/2014	i2CAT	Contributed section 5.2
04_IBM	19/05/2014	IBM	Completed section 2; added initial section 5.3.3 text
04_IRT	20/05/2014	IRT	Added section 4. Updated section 5.1
05	21/05/2014	See the list of authors	Integrated all the 04 contributions
05_IBM	26/05/2014	IBM	Editing of 5.3.3
05_UPC	30/05/2014	UPC	New contribution for section 5.3 and editing of 5.3.1
06	09/06/2014	See the list of authors	Integrated previous changes
1	15/06/2014	IBM	Integrated all the contribution as v1
1.1	19/06/2014	IBM	Homogenized section 5.3
1.2	20/06/2014	i2CAT	Add requirements prioritization and architectural guidelines
1.3	27/06/2014	IBM	Update Sections 6 and 7
1.4	04/07/2014	I2CAT	Update Section 6
1.5	21/07/2014	I2CAT	Final integration

2 Data Centre Trends and Usage Patterns

Recent years have been a very exciting period when utility computing started turning from vision to reality.

The vision of utility computing stems from realization that the ICT services are required for almost every organization, person, or business on the planet. Moreover, these services are very often interleaved so that information acquired and stored by a process in an organization is often consumed, transformed, and further stored by another process, potentially in another organization, further feeding additional processes. The need for the ubiquitous access to large volumes of information is one of the major driving forces for consolidating the ICT storage resources in high volume Data Centres (DCs). In order to ensure the timely, secure, and efficient access to the ever growing information sources, large computational capabilities are required. In addition, communication networks must evolve to deliver huge information volumes between the data sources and the data consumers all over the services chain.

An additional driver behind the recent Data Centre trends is the economy of scale whereby collocation of large amounts of ICT resources together and sharing them between multiple consumers helps reducing the operational expenses, e.g. the power and cooling expenses, the floor space, the wiring, the service personnel salaries etc., through sharing these expenses among the consumers. Moreover, consolidating large amounts of resources together aids achieving equipment and technology homogeneity, further driving down the acquisition and the maintenance costs through equipment simplification and commoditization.

In order to fully realize the economy of scale potential, all Data Centre resources, namely compute, storage, and network, must be brought to as full utilization as possible at all times. Server virtualization is one of the technological enablers helping to achieve this goal through consolidating more workloads on a single computational platform. For virtualized Data Centres, storage and network virtualization are required to keep up with the consolidation ratios, workload dynamicity, and multi-tenancy.

Large-scale consolidation and utility computing gave raise to the concept of the Cloud – new ICT delivery and consumption model where shared consolidated infrastructure resources are provisioned automatically on demand and are made available in a pay-as-you-go manner to individuals and businesses.

In the subsequent sections we present these and other Data Centre trends and discusses them in detail. In Section 2.1 we present the figures showing the achieved and projected DC growth. In Section 2.2 we discuss the business models and usage patterns that have already emerged or are expected to emerge due to ICT turning into a utility. In Section 2.3 we review the application base currently being deployed in scale-out data centres and available for ubiquitous access, as well as the evolution the applications are going through due to the changes in ICT resource provisioning and consumption models.

2.1 Data Centre Growth Trends

Several global trends fuel increasing demand in advanced ICT services. Over the past decade more and more individuals, businesses, and organizations have come to rely on ICT services to perform their day-to-day tasks, store their vital data, and govern their processes. Some of the recent global trends are: **MOBILE** – the explosion in the amount of mobile interconnected devices; **BIG DATA** – accumulation of huge amounts of information that can be beneficially analysed and transformed to knowledge and other assets; **SOCIAL** – interconnection of individuals and groups for leisure, commerce, and business, sharing the information, knowledge, and processes. These trends have created a new type of ICT consumers, not knowledgeable or even interested in infrastructural aspects of the services they consume. What users care about are ICT services, their ubiquitous availability, privacy and security, and global reach.

In addition, there remain traditional consumers of ICT – scientific computations, commercial and business workflows, regulated information systems, etc. These sectors have been traditionally deploying and maintaining their own ICT infrastructure and were concerned with the technology they consume as well as with the cost, performance, and maintainability of the infrastructure they own or hire.

ICT solutions providers industry that first appeared to cater to the needs of the traditional ICT infrastructure outsourcing, have grown significantly to accommodate the demand created by the MOBILE, SOCIAL, and BIG DATA trends. Today's largest industry players, e.g. Amazon [22], IBM [23], Google [17], Microsoft [24], CISCO [21], etc., build scale-out DCs and compete in a growing and demanding market for providing ICT infrastructure and/or services. To be successful, it is mandatory to drive down the costs and drive up the efficiency and profitability. Resource consolidation and usage optimization are key to meeting these needs, turning the ICT provider business facilities, Data Centres, into heavy-duty industrial-scale factory-style instalments. Governments and regulatory requirements apply to limit the environmental impact of the DCs, further driving the trend of constructing green-field DCs in suitable geographical areas where, for example, cooling costs can be saved by leveraging the natural resources without harming the environment. Construction and equipment costs of these endeavours are huge making it more beneficial to create larger instalments and leverage economy of scale.

Consolidation trends in DC construction started to happen in the early 2000s, mainly in North America and Central Europe. Over the years, the trend was continuing and spreading into more areas, so that today in the most developed areas new DC creation has slowed down while the size of newly created DCs continues to grow (according to a collocation/outsourcing report by DCD Intelligence, almost a quarter of all DCs in North America is now outsourced [18], while TechNavio's analysts forecast the Collocation and Managed Hosting Services Market in North America will further grow at a CAGR of 13.59 percent over the period 2013-2018 and the Data Centre Collocation market in the UK to grow at a CAGR of 14.43 percent over the same period [19]). In Asia, Far East, Latin America, and less connected areas of Europe, new DCs construction continues at fast speed bringing in huge investments.

In what follows we cite analyst research and vendor reports to back up the above information and to illuminate the current and the projected DC growth with numbers that can be used for deducing the requirements towards the COSIGN DCN architecture.

- **Construction costs.** Construction costs depend on the facility site and the location's economic conditions. For example, Google [17] has vastly invested in its now established worldwide consolidated infrastructure in between the 2007 and now. Total long-term DC investment by Google range from \$150 million USD planned to be invested in its new, to be put into the operation this year, data centre in the municipality of Quilicura, near Santiago, Chile, to \$1.2 billion invested into both the Caldwell County, NC and the Dallas facilities. TechNavio's analysts forecast the Global Data Centre Construction market to grow at a CAGR of 21.99 percent over the period 2013-2018.
- **Operating expenses.** In addition to the construction costs, there are maintenance and operating expenses, e.g. cost of power required to fuel the equipment, the cost of the cooling and UPC equipment, and labour cost. These costs depend on the reliability, redundancy, and the performance of the infrastructure, as well as on the equipment lifetime. There are many considerations for bigger up-front investments that allow for more capable, resilient and expandable infrastructures, versus right-sized initial instalments requiring more investment over time but being harder to scale with the demand, see for example network equipment TCO comparison in [25].
- **Floor space.** In 1999 the Data Centre was considered "Large" if it was located in 5,000 square feet facility. In 2004, this number grown to be 50,000 square feet, in 2009 – 500,000 square feet. In 2011, IBM has started construction of the IBM/Range Technology Data Centre in China (near Beijing) in a 624,000 square feet facility. **Figure 2.2** presents the total floor space taken by DCs in North America.




Figure 2-1: DCD report on DC floor space in North America [18].

- Energy consumption.** In 2001, several of the largest supercomputer centres in the world needed 10 MW of power, in 2011 there existed dozens of 10 MW data centres worldwide and the US Government was planning construction of ~60 MW data centres. Today, the North American DC market power consumption stands at 11.55GW, an increase of 6.8% over the previous year. In UK, total DC power consumption has grown from 2.85GW in 2012 to 3.10GW in 2013 and is predicted to grow at an average annual rate of 6.4%, to reach 3.68GW in 2016.
- Global Data Centre traffic.** According to *CISCO Global Cloud Index: Forecast and Methodology, 2012–2017* [20], by the end of 2017 annual global data centre IP traffic will reach 7.7 zettabytes, while global data centre IP traffic will reach 644 exabytes per month (up from 214 exabytes per month in 2012). The report classifies the traffic and traffic growth rates along several axes: by segment (business, individual consumer), by type (cloud, traditional), by geographical area. For COSIGN, one of the most interesting aspects is traffic classification by reach – within DC, between DCs, and between the client and the DC, presented in Figure 2.1.

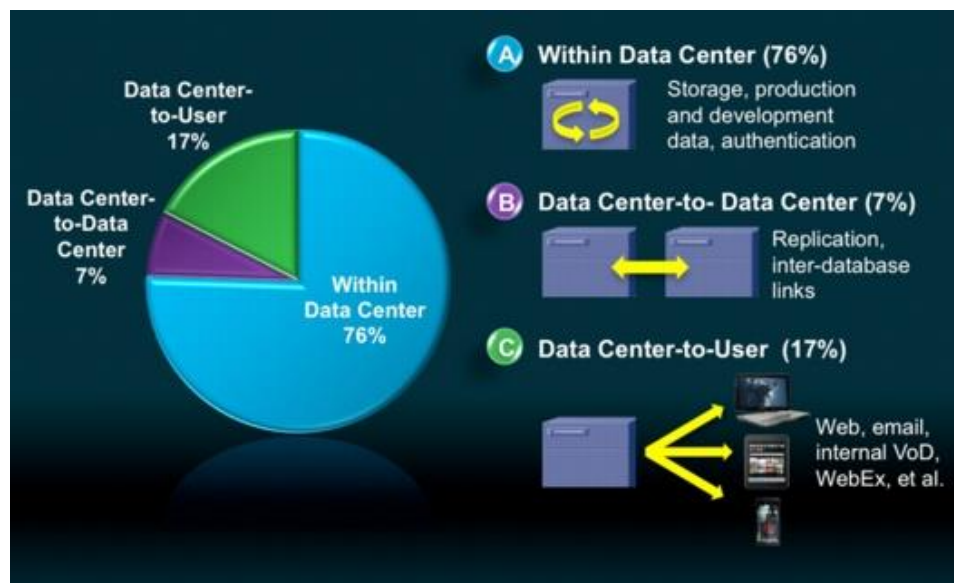


Figure 2-2: Global Data Center Traffic by Destination [20].

- **Data Centre Servers.** Increasing global capacity demand and the cost efficiency requirements push the DC server market towards increased density, higher utilization, and commoditization. The world's largest DC operators, e.g. Google, Amazon, Facebook, tend to custom build their servers in order to optimize the price performance trade-off for their specific workloads. Server manufacturers, e.g. IBM, HP, Dell, offer highly integrated DC-in-a-box systems that are either common purpose or optimized towards specific application – business management, web serving, and data processing – and include the compute, the storage, and the network resources packaged together in a single Point of Delivery (PoD).

We summarize the above by noticing that as the demand for the ICT services as well as their diversity and customer base grow, the whole range of technological advances are taking place. These advances span all the areas of the ICT – DC construction, servers, storage, virtualization, interconnect, management, sometimes pulling in opposite directions. For example, with today's servers it can be anticipated that the amount of servers in the typical data centre as well as its floor space will grow tremendously to catch up with the demand (Microsoft has reported they have 1,000,000,000 servers in their DCs in July 2013 [26]). But server technology is far from stagnant so with the servers of tomorrow, less floor space and fewer servers might be required to satisfy the same demand. This uncertainty makes it hard to quantify the requirements towards the data centre interconnect in terms of the amount of interconnected entities, ports, cables, etc. As a result, in COSIGN we are focused mostly on demand projections, namely latency, bandwidth, and throughput, as well as on qualitative requirements, namely flexibility, manageability, cost and power efficiency.

2.2 Data Centre Usage Patterns and Business Models

This section starts describing some possible classifications of cloud services widely adopted in the literature and that will be considered as reference for the rest of the work. These categories take into account the deployment, the type of entity (or entities) that owns and operates the cloud infrastructure or the type of service delivered to the final end-user (see paragraph 2.2.1). For each category, we analyse the characteristics and the business benefits of the cloud services and the main types of customer. Some examples of cloud service usage in different sectors are provided in paragraph 2.2.2, with main reference to relevant case studies from Interoute, the cloud provider in the COSIGN consortium.

2.2.1 Classification of cloud services

2.2.1.1 Private, public and hybrid clouds

The first classification of cloud infrastructures is based on the deployment models identified by NIST [3] and can be summarized as follows:

- **Private cloud:** cloud infrastructure provisioned for exclusive use by a single legal organization comprising multiple consumers (e.g. business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises. The main benefits of private cloud deployments are the enhanced security, the higher availability (resources are fully dedicated to a given group of users), and flexibility (services are usually tailored to specific user requirements).
- **Community cloud:** cloud infrastructure provisioned for exclusive use by a specific community of consumers from organizations with shared concerns (e.g. a common mission, specific security or compliance requirements, a given set of policies). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.
- **Public cloud:** cloud infrastructure provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.
- **Hybrid cloud:** cloud infrastructure given by a composition of two or more distinct cloud infrastructures (private, community, or public). These infrastructures remain unique entities,

but are bound together through standardized or proprietary technologies that enable data and application portability (e.g. cloud bursting for load balancing between clouds).

The most popular deployment model for many cloud consumers is the **public cloud model**, where cloud services are provided in a virtualized environment built with pooled, shared physical resources. A public cloud service is accessible over a public network, like Internet, to multiple customers and makes use of a single shared infrastructure. The most popular examples of cloud computing belong to the public cloud category since they are “publicly available”. They mainly offer Software as a Service (SaaS) applications (see section 2.2.1.2 for a formal description of SaaS, PaaS and IaaS), like communication or office applications. However, the public cloud model can be adopted also for IaaS and PaaS services, such as cloud based web hosting and development environments.

The main difference between public and private cloud is that the latter reserves the pool of physical resources and builds a separated cloud platform for each single organization, so it inherently guarantees higher levels of security compared to public cloud. Consequently, the main targets for public cloud services are private individuals since they do not usually require the same high level of infrastructure security as in case of enterprise customers. However, public clouds can be convenient also for enterprise customers whenever they need to improve the efficiency in the management and storage of non-sensitive materials, e.g. through the usage of online document collaboration tools.

As explained above, in the **private cloud model** a distinct and secure cloud based environment is created for each single organization, and this approach results in greater control and privacy. Due to this reason, private cloud is an ideal solution for any organisation, including enterprises that need to deal with private data or carry out sensitive tasks. An example of a potential target is a financial company that wants to make use of cloud computing services, but must store sensitive data internally by regulation.

It should be noted that, from a technical perspective, cloud operators may adopt a large variety of mechanisms to provide cloud services categorized as belonging to the private cloud model. This is the reason why the classification is usually preferred in terms of features offered to the customers, e.g. ring fencing of a cloud for the sole access and usage of a single customer and higher levels of network security, which can be provided through private leased lines or secure encrypted connections via public networks.

Finally, the **hybrid cloud model** combines private and public clouds to perform distinct functions for the same organization, mixing the peculiar benefits of each model for specific tasks. For example, public cloud services are often more cost efficient and scalable, so they can be preferred for all non-sensitive operations, while private clouds can be adopted for all the tasks where higher levels of control and privacy is required. A typical example of customer interested in hybrid cloud platforms is an enterprise with an e-commerce website hosted in a private cloud providing high security guarantees and a brochure site, with less stringent security requirements, hosted in a less expensive public cloud.

Hybrid cloud services can be entirely provided by single cloud providers or by separate cloud providers, each of them specialized in the private or public model, that join to provide the integrated service. An alternative solution is the composition of a private cloud managed directly by an organization together with a public cloud service rented by the same organization which integrates it into its own infrastructure.

The following Table 2.1 provides a summary of the main features and characteristics for the different types of cloud deployment models.

Public cloud	Private cloud	Hybrid cloud
High scalability Cloud resources are available on demand from the public clouds' large pools of resources, so that the applications running on them can respond seamlessly	Cloud bursting Some providers may offer the cloud bursting option in a private cloud scenario in order to better support spikes in demand. The provider moves	Scalability Private clouds offer a level of scalability that depends on their configurations (e.g. if hosted internally or externally); public cloud scalability has fewer

to load fluctuations.	some non-sensitive functions to a public cloud to free up more space in the private cloud for sensitive functions.	boundaries because resources are pulled from a larger cloud infrastructure. Thus, the hybrid cloud scalability can be improved moving non-sensitive functions to the public cloud, since it reduces the demands on the private cloud.
<p>Cost efficiency</p> <p>Public clouds bring together greater levels of resource thus benefitting from the largest economies of scale. Operation and management of the physical resources are shared across all cloud services and they require less bespoke configuration.</p> <p>The pay-as-you-go pricing model often adopted in public cloud services allows users to access the required resources when needed, and then only pay for what they use, therefore avoiding wasted capacity.</p> <p>Some mass market propositions can even be free to the client, relying on advertising for their revenue.</p>	<p>Cost and energy efficiency</p> <p>Private clouds can improve the allocation of resources within an organisation, since resource availability is tailored to the demands of single business functions or departments. This minimizes the investments in unused capacity and can also reduce an organisation's carbon footprint.</p> <p>However, private cloud is not as cost effective as public cloud services, due to smaller economies of scale and increased management costs.</p>	<p>Cost efficiency</p> <p>Public clouds are likely to offer more significant economies of scale, and hence greater cost efficiency, than private clouds.</p> <p>Hybrid clouds therefore allow organisations to access these savings for as many business functions as possible (the non-sensitive ones), while still keeping sensitive operations secure.</p>
<p>Lower security</p> <p>The public cloud model only guarantees a lower level of security if compared to the other models, so it is not suitable for storing or managing sensitive data.</p>	<p>Higher security and privacy</p> <p>Private clouds offer the highest security and privacy. They usually adopt techniques based on distinct pools of resources with access restricted to connections originated behind firewalls, dedicated leased lines and/or on-site internal hosting.</p>	<p>Security</p> <p>The private cloud element of the hybrid cloud model not only provides the security where it is needed for sensitive operations but can also satisfy regulatory requirements for data handling and storage where it is applicable.</p>
<p>Flexibility</p> <p>A wide variety of IaaS, PaaS and SaaS services based on the public cloud model are available on the market, for both private and enterprise clients.</p> <p>Businesses can even integrate public cloud services with their own private clouds, when the latter is needed for sensitive business functions.</p>	<p>More control</p> <p>Since a private cloud is only accessible by a single organisation, it can be managed and configured to create a tailored network solution. However, this level of control is obtained giving up the economies of scale generated in public clouds, where the hardware management is centralized and shared for multiple customers.</p>	<p>Flexibility</p> <p>The availability of both secure resource and scalable cost effective public resource can provide organisations with more opportunities to explore different operational avenues.</p>

<p>Reliability</p> <p>There is no single point of failure which would make a public cloud service vulnerable.</p> <p>The sheer number of resources involved in creating a public cloud and the redundancy configurations allow the cloud service to still run unaffected even in case of failures in a physical component.</p> <p>Moreover, in case of resources distributed among multiple data centres, cloud services could still remain running even if an entire data centre goes offline.</p>	<p>Reliability</p> <p>In case of cloud hosted with a third party, the level of reliability is comparable to the one of an infrastructure hosted within data centres.</p> <p>However, even if cloud resources are hosted internally, the creation of virtualised environments guarantees higher level of resilience to individual failures across the physical infrastructure. In fact, virtual partitions can use resources from the unaffected servers.</p>	<p>Reliability</p> <p>Reliability is improved through the virtualization mechanisms adopted in the private cloud element, mixed with the resilience guarantees offered by the public cloud element through the distribution and redundancy of critical functions.</p>
<p>Location independence</p> <p>Public cloud services are available through an Internet connection, wherever the client is located. This provides invaluable opportunities to enterprises such as remote access to IT infrastructure or online document collaboration from multiple locations.</p>	--	--

Table 2.1- Summary of public, private and hybrid cloud features

Depending on the entities that provide the cloud service and their relationships, we can also identify the following transversal models:

- **Single Cloud Model:** a model typically adopted by big cloud operators that deploy several data centres in different geographical areas, interconnected through Tier-1 and Tier-2 domains. Connectivity among different data centres is typically provided through the cloud operator's network infrastructure when both cloud and network are owned by the same entity (e.g. in the Interoute's case), or are based on economic agreements with ISPs. However, no agreements with other cloud operators are established in the single cloud model.
- **Federated Cloud Model:** a model typically adopted by smaller cloud operators (with connectivity on Tier-2 and Tier-3 domains). They join together to form a federation and achieve business advantages through the possibility to serve more end-users and offer a wider range of services. In this model, small cloud operators can use resources located in data centres that belong to other cloud operators who have joined the federation, while this procedure is completely transparent from the user's perspective. Some open issues for this model are related to: (i) protocols enabling the federation, (ii) mechanisms for user identity, security and privacy, and (iii) location of user's data.
- **Interconnected Cloud Model:** similar to the previous model, but without the need to establish any federation. Each cloud operator maintains its administrative role and establishes economic agreements with other partners to achieve service mobility and offload its computing and hosting capacity, with the objective to guarantee a better QoS/QoE and service reliability guarantees to its own customers. As for the federation case, the technical procedures and the agreements established among cloud operators are transparent for the end-user, who has an SLA with a single provider.

2.2.1.2 IaaS, PaaS and SaaS

NIST [3] defines three service models according to the capabilities of the service delivered to the end-users, as detailed in Table 2.2.

Service model	NIST definition
IaaS Infrastructure as a Service	<i>“The provision of processing power, storage, networks, and other fundamental computing resources, where the consumer is able to deploy and run arbitrary software, which can include 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; and possibly limited control of selecting networking components (e.g., host firewalls or machines with specific characteristics)” [3]</i>
PaaS Platform as a Service	<i>“The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the cloud provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment” [3]</i>
SaaS Software as a Service	<i>“The capability provided to the consumer is to use the cloud provider’s applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings” [3]</i>

Table 2.2- Cloud service models – NIST definitions

In general, a cloud service provides access to computing resource in a virtualised environment across a public connection, usually Internet.

In **Infrastructure as a Service (IaaS)**, the virtual resource provided to the user is a virtualised hardware, i.e. a computing infrastructure which can include virtual server space, network connections, bandwidth, IP addresses, load balancers, etc. The pool of hardware resources is pulled from a multitude of servers and networks distributed across data centres under the responsibility of a cloud provider. Customers are given access to the virtualised component, so that they can create their own cost effective and easily scalable IT platforms, but the complexity and the cost of managing the underlying hardware are outsourced to the cloud provider. The main benefits for IaaS services are the following:

- **Scalability.** IaaS resources are available as needed, without delays in expanding the capacity of the virtual infrastructure or wasting money for unused resources.
- **No investment in hardware.** The cloud provider is responsible to maintain the underlying physical infrastructure, with benefits in terms of time and cost saving for the customer.
- **Utility style costing (pay-as-you-go).** IaaS can be accessed on demand and customers pay only for the used resources.
- **Location independence.** IaaS can usually be securely accessed from any location through an Internet connection.

- **Physical security of data centre locations.** Services available through a public cloud, or a private cloud hosted externally with the cloud provider, benefit from the physical security afforded to the servers hosted within a data centre.
- **No single point of failure.** IaaS benefits from the resilience characteristics of cloud services: in case of server or network failures, the broader service is not affected due to the redundancy configuration, the local or remote backup and the service recovery options.

The following list provides some examples of IaaS usage:

- **Enterprise infrastructure.** Internal business networks (e.g. private clouds and virtual local area networks) can make use of server and networking resources to store their data and run the applications to be operated day-to-day. Expanding businesses can scale their infrastructure in accordance with their growth and exploits private clouds (accessible only by the business itself) to protect storage and transfer of sensitive data.
- **Cloud hosting.** Websites can be hosted on virtual servers based on pooled resources from underlying physical servers. The main benefits are the redundancy, guaranteed by a wide set of physical servers, and the scalability, in order to support unexpected peaks of demands.
- **Virtual Data Centres (VDC).** A virtualised network of interconnected virtual servers which can be used to offer enhanced cloud hosting capabilities, enterprise IT infrastructure or to integrate all of these operations within either a private or public cloud implementation. An analysis of the Interoute VDC solution is provided in section 3.2.

In **Platform as a Service (PaaS)**, the virtual resource provided to the user is a platform or an environment for developers, where building, testing and deploying applications and services. PaaS services are hosted in the cloud and accessed via web browser.

Cloud providers offer development tools to create software applications and they may include additional features, like per-user customization and add-ons, automatic updates and/or tool upgrades, and customer support. PaaS packages can widely vary, from simple offers of point-and-click pre-configured frameworks to more advanced development environments where customers can select their own infrastructure options. However, some common features and components usually included in PaaS packages include the operating system, a server-side scripting environment, a database management system, tools for design, development and testing, as well as storage and network resources.

The PaaS pricing model is usually on a subscription basis, where customers pay for what they use; the sharing of the underlying physical infrastructure results in cost reduction. The main targets for PaaS products are software and web developers. Moreover, businesses developing their own internal software can also utilise PaaS to create distinct ring-fenced development and testing environments.

The main benefits for PaaS services are the following:

- **No investment in physical infrastructure.** Software development companies do not need to purchase or manage hardware, so that they can entirely focus on their own business. Moreover, customers will only need to rent the resources they need, without large fixed investments.
- **Makes development possible for ‘non-experts’.** Some PaaS packages are designed to simplify application development, so that anyone can develop simple applications through the web-browser using one-click functionalities.
- **Flexibility and adaptability.** PaaS packages are characterized by different levels of control and customization: customers can choose pre-configured solutions or prefer to maintain the control over the tools to be installed. This last option allows developers to pick and choose the desired features in order to create a personalized platform, and even change it as required.
- **Easy cooperation of distributed development teams.** Teams in various locations can work together on the same application just using an Internet connection and a web browser.
- **Security.** PaaS guarantees security including data security, backup and recovery.

In **Software as a Service (SaaS)**, cloud consumers are able to access software applications, which are hosted in the “cloud”, through an Internet connection. Common examples of SaaS are Google, Twitter, Facebook and Flickr. Targets for SaaS services are both individuals and organizations; typical SaaS applications for enterprise users include accounting and invoicing, tracking sales, planning, performance monitoring and communications, like webmail and instant messaging.

SaaS customers rent the software, while in traditional models their purchase the software as a package to be installed in their own computer, with licences that may limit the number of users or devices where the software can be installed. SaaS products are usually purchased through subscription on a monthly basis and the applications are accessed and used on line with files that are maintained in the cloud.

The main benefits for SaaS services are the following:

- **No additional hardware or setup costs.** Processing power to run the applications is supplied by the cloud provider and the applications can be used as soon as the user subscribes.
- **Pay-as-you-go.** Software that is needed only for limited periods can be paid for just for those periods, with time limited subscriptions.
- **Scalable usage.** More storage or additional services can be added on-demand without needing to install new software or hardware.
- **Automated updates.** Updated software versions are automatically deployed and managed by the cloud provider, and made available online to the existing customers, often free of charge.
- **Cross device compatibility.** Access to SaaS applications is compatible across all internet enabled devices, from computers to smartphones and tablets, and can be accessed from anywhere.
- **Customization.** SaaS applications can be customised to suit the needs and branding of a particular customer.

Projections for the growth in cloud service spending for IaaS, PaaS and SaaS, as reported by IDC's forecasts [4], are shown in Figure 2.3.




Figure 2-3: IDC's forecast for growth in cloud service spending [4].

2.2.2 Data Centre Usage by Business Sectors

This section presents some case studies from Interoute, showing how cloud services can be tailored to efficiently meet the demands of a variety of customers. This survey will enable defining the main characteristics and requirements to be supported in cloud services, perceived as key factors by different categories of customers.

The European Space Agency (ESA) has designed and developed its SuperSites Exploitation Platform (SSEP) in collaboration with Interoute. The SSEP enables the on-demand processing of satellite data, with initial access to 13TB of data, encompassing 50,000 radar scenes from ESA. The SSEP runs on Interoute's Virtual Data Centre, innovating the way satellite data is distributed and processed in the scientific community. SSEP users have immediate and controlled on-line access to a wide set of contents, but also to remote facilities, instruments and collaboration tools allowing scientists across Europe to select the desired data and perform direct analysis using the tools hosted in the VDC. The SSEP model promotes the setup of virtual research communities with a collaborative creation of larger data pool available for scientists, enabled through the sharing of ICT tools and resources across different disciplines and technology domains [5].

The Interoute VDC has also become the infrastructure platform for the UEFA's development team, where testing and validation of new features and version upgrades are performed. The VDC solution provides an environment that is able to quickly scale-up and where virtual machines can be added on-demand through self-service procedures; moreover, it guarantees high levels of security that reflect the characteristics of a production environment. In particular, on the network side VLAN segregation, DMZ, firewall protection, LAN, WAN and Internet integration can be provided in real-time through the VDC portal. In the longer term, UEFA's strategy will be to migrate its production environment from its private cloud to the Interoute VDC [6].

Cloud services can also impact the e-learning area as exemplified by the appearance and the popularity of MOOCs. Marconi University, the first Italian Open University, offers virtual learning courses that complement the traditional classroom teaching methods. The IT infrastructure of the university, where the on-line Virtual Campus e-Learning platform runs, was previously hosted in a co-location facility in Rome, with servers and storage managed by a third party. The entire platform is now moved to the Interoute VDC and the integration with Interoute Virtual Private Network allows the connection of the headquarters and the regional offices with their own private cloud. The main requirements are in terms of service reliability and availability, combined with high performance, quality of service and network resilience to provide the students with a continuous access to the e-learning platform anytime and anywhere. The Interoute solution is based on a platform hosted in the London VDC facility, with backup of training and media materials at the Amsterdam VDC facility, interconnected through a secure VPN. The benefits in terms of cost reductions have been estimated in the order of 23%, and the design of auto-scaling mechanisms, to match the fluctuating usage patterns throughout the academic year, are in progress to fully exploit the pay-as-you-go pricing model [7].

An example where the integration between cloud and network is fundamental is given by the GC Europe case study [8]. GC Europe, a leading European manufacturer of dental care products, has a constant flow of heavy data traffic of large images and multimedia contents across its many locations. The usage of a network-integrated VDC allows the company to store and easily share huge files with employees and partners across the world. Moreover, the managed virtual private network, with the embedded VDC, connects the GC Europe dispersed offices, manufacturing, sales and research sites across Europe guaranteeing a fast, reliable and secure exchange of data. The results show a relevant increase of performances mixed with costs reduction: the transfer of 2GB of multimedia file across Interoute VDC is 260% faster and the TCO is estimated 20% to 25% lower than a previously chosen managed hosting service.

These case studies highlight the relevance of some cloud service characteristics, in particular:

- 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
- On-demand and self-service provisioning and scalability, possibly enabled through open APIs

- High-level of security, including the integration of firewall or IPS services
- Service reliability and availability, combined with network resilience
- Automated mechanisms for intra-DC and inter-DC content replications and fast service recovery
- Automation of scale-up and scale-down procedures
- Support of multiple and flexible pricing models, including commit, pay-as-you-go or mixed models
- User-friendly administration and control interfaces for resource consumption monitoring, server administration and management of network and data physical location

2.3 Prevalent Data Centre Applications

Since computers were first invented, their use has been expanding and now reaches into more and more areas. Moreover, as the technology develops, new application types are enabled, giving in turn a new boost to the technological development through more demanding requirements and use cases.

Back in 1940ies, when the computers were very expensive, rare, and isolated from each other, applications were monolithic and dedicated to the specific organization's needs. When first computer networks appeared in 1970ies, simple peer to peer and client server architectures have been conceived and implemented. Over time, application architectures become more and more distributed and dynamic, following the ICT trends on the one side and the use cases evolution on the other.

In Section 2.3.1 we cover the ICT applications running in today's established and emerging Data Centres, in Section 2.3.2 we briefly explain how traditional applications are adopted and migrated to the novel consolidated ICT platforms, while in Section 2.3.3 we discuss the novel applications that became possible due to the recent ICT shift towards cloud Data Centres and utility computing.

2.3.1 Traditional Data Centre Applications

2.3.1.1 Engineering Applications

Many manufacturing, construction, and design processes employ ICT simulation that allows engineers to design and verify products in a virtual, risk-free environment and minimize the need for physical prototypes and tests. Engineering simulations often involve resource savvy computations and often demand hardware acceleration of high-end graphics like 3D graphics applications, e. g. for rendering large models of space crafts, automobiles, or submarines.

In older environments, many engineers have workstations for 3D design, in addition to the equipment they use to run enterprise applications and collaboration tools such as email and instant messaging. This multiple workstation model was inefficient and costly and did not lend itself well for the type of collaboration necessary to enable real-time review of component designs.

In today's DC environments, engineering software, e.g. computer-aided design (CAD) or computer-aided manufacturing (CAM), computer-aided engineering (CAE), engineers work remotely, accessing specialized applications deployed in centralized product development centres from their private client equipment, as shown in Figure 2.4. In such product development environments, private or hosted, dedicated computing resources are typically deployed in support of a single workload, project, or organization. HPC clusters are often required to ensure adequate supply of resources to the engineering task at hand. These resources have to be managed and maintained, limiting the application portability, scale, and availability to geographically remote parts of the engineering team.

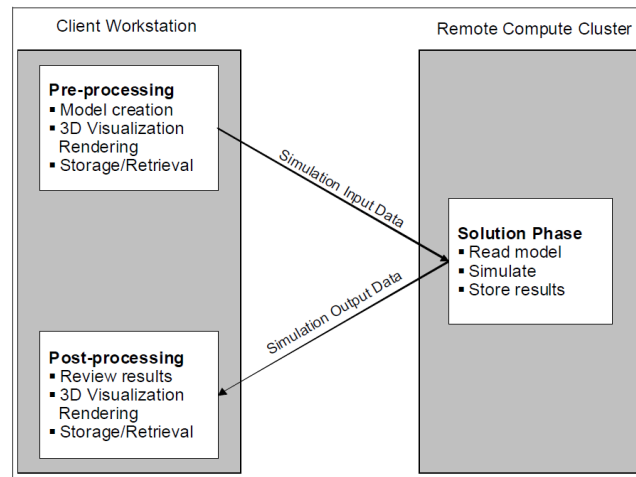


Figure 2-4: Engineering workflow with centralized computational resources

2.3.1.2 Scientific Applications

Scientific workloads typically involve huge quantities of data and very complex computations required to analyse the data and turn it into the knowledge according to the Computational Science Education Reference Desk (CSERD), see Figure 2.5. Moreover, the knowledge must be further organized, classified, and stored, to be accessed by researchers and practitioners, and to serve as a basis to more computations. Good examples are computational chemistry and bioinformatics where computations of human genome can easily require thousands of computing hours when sequentially computed on a modern computer, as well as thousands of Gigabytes of storage.




Figure 2-5: What is Computational Science [15]

These workloads traditionally require highly parallel infrastructure designs and low latency access to the vastly scalable storage resources. To achieve this, data integration and sharing solutions, such as data warehousing, allow researchers to tap into multiple heterogeneous data sources. Creation and maintenance of data warehouses with hundreds of data sources, as well as the technical effort and costs of the associated software are affordable to a limited number of very large organizations.

2.3.1.3 Financial Applications

Modern financial workloads combine the complexity of the computer and data intensive scientific workloads (mathematical, statistical computations, modelling, simulations, and forecasting) with the real-time transactional processing of highly sensitive data. In this industry, complex analytical computations are based both on past collected and stored data and on data arriving at high speeds in real time. According to regulatory requirements, financial institutions must be able to analyse new data as quickly as it is generated, and also store the data to further utilize it later on.

Financial applications can be classified according to two different categories: real-time versus long-running (batch) and data intensive versus computing intensive, as presented in Figure 2.6.

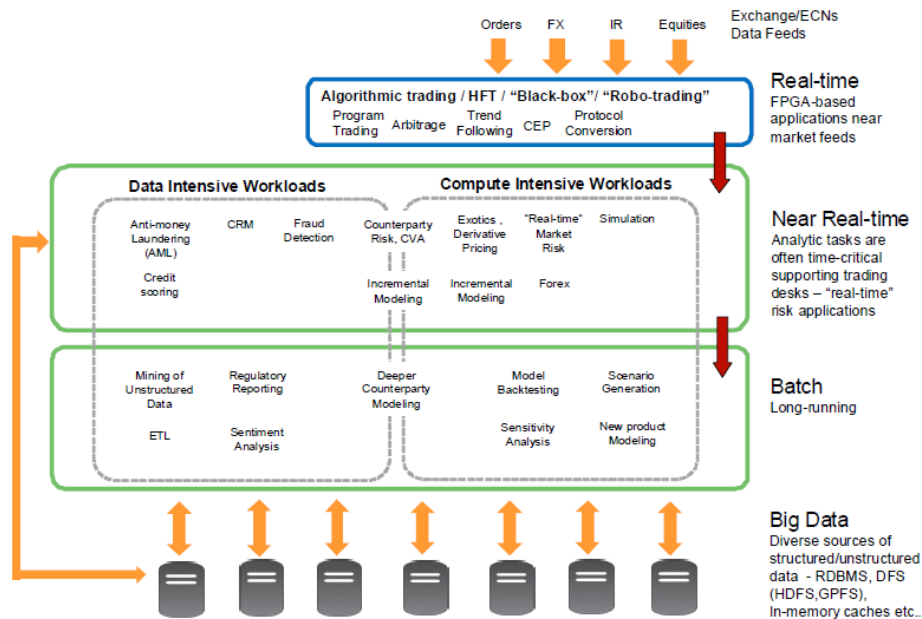


Figure 2-6: Diagram classifying some of the tasks that are run by a bank

For data intensive financial applications, the time to transfer data among computing nodes for processing can exceed calculation times and the network can get saturated as the number of data that need to be analysed grows rapidly.

2.3.1.4 Business Applications

Business applications are used by business users to increase productivity, to measure productivity and to perform business functions accurately. Depending on the size and the complexity of the organization, different classes and types of business software is employed. This software is often off-the-shelf vendor software developed for typical business needs, e.g. business reporting, document processing, employee scheduling, customer relationship tracking, inventory management...

Traditionally, each business or organization had to acquire, deploy, and manage the ICT resources and services required by its operations. As a result, the ICT management became an additional burden of each business. Over the years, business applications became sufficiently standardized to lend themselves to outsourcing and hosting deployment models. Therefore such applications are a very natural fit to every DC, including modern and emerging scale-out consolidated ICT environments, such as private, public, and hybrid clouds.

2.3.1.5 Exploration and Production Applications

Exploration and production industry is concerned with the discovery and harvesting of earth natural resources, such as oil and gas. To facilitate the process, organizations must acquire, store, and analyse large amounts of highly specialized unstructured data, simulate natural processes, and govern production processes as shown in Figure 2.7. As a result, this industry employs a set of ICT applications with mixed characteristics, some typical to engineering tasks like rendering 3D models, some typical to scientific tasks like simulations and data analysis. In addition, specialized business process management tasks are required to follow up on the execution and the production of the harvested natural resources.

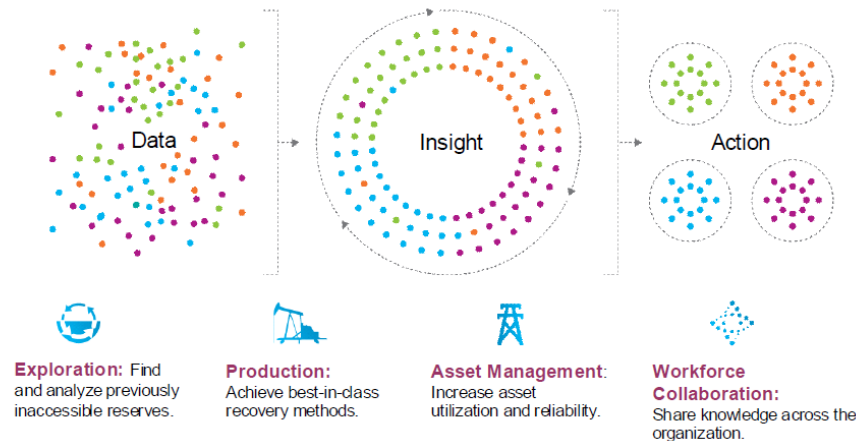


Figure 2-7: Processes in the exploration and production industry

2.3.1.6 Web Applications

Web applications appeared in the early 90^{ies} and were very simple client server applications where static content was accessed by users over a network through a thin client, i.e. a browser. Later on, web applications have evolved to become more interactive and dynamic and the server side architectures advanced to take care of the resulting complexity. Thus, Web 2.0 applications typically follow the multi-tier server side design with separate components realizing the presentation, the application logic, the data access logic, the data storage, etc. Web services architectures allow different Web 2.0 applications to interact and invoke service of one another. The emerging Web 3.0 paradigm adds intelligence to the services web so that data is continuously collected, analysed, and synthesised to create useful and valuable insights.

2.3.1.7 ICT Management Applications

ICT resources require management and maintenance like any other business assets. Organizations typically keep warehouse records for their hardware and software equipment, e.g. computers, networking, and storage systems listings, software versioning, and licence management. In addition, ongoing performance monitoring is a must for keeping track of the performance versus the demand, as well as the allocation of ICT resources to departments, tasks, and individuals inside the organizations. Traditionally, ICT resource allocation has been fairly static so that standard warehousing software, e.g. somewhat specialized ERM, was sufficient to take care of the ICT department needs. Modern DC workloads, however, are highly dynamic and require speedy resource provisioning and release following the demand cycle. As a result, modern ICT management is a complex dynamic task governed by enterprise level DCIM tools, e.g. IBM Tivoli, or cloud management stacks, e.g. OpenStack.

2.3.2 Migrating Traditional DC Application to the Cloud

Many of the traditional DC applications described above can be migrated into the consolidated Data Centre environments, some with more ease than the others.

On the one hand, cloud DCs pose difficulties for including workloads with real time response time demands, harsh regulated security requirements, or predictable resiliency contracts. On the other hand, modern consolidated DC environments facilitate elasticity, dynamicity, and cost efficiency thus enabling workloads that are impossible in the traditional DCs. Figure 2.8 illustrates how several common IT workloads fare on gain versus pain measures, based on IBM's research and experience. In general, the workloads that appear in the upper right quadrant have proven to be the best fit for cloud computing.

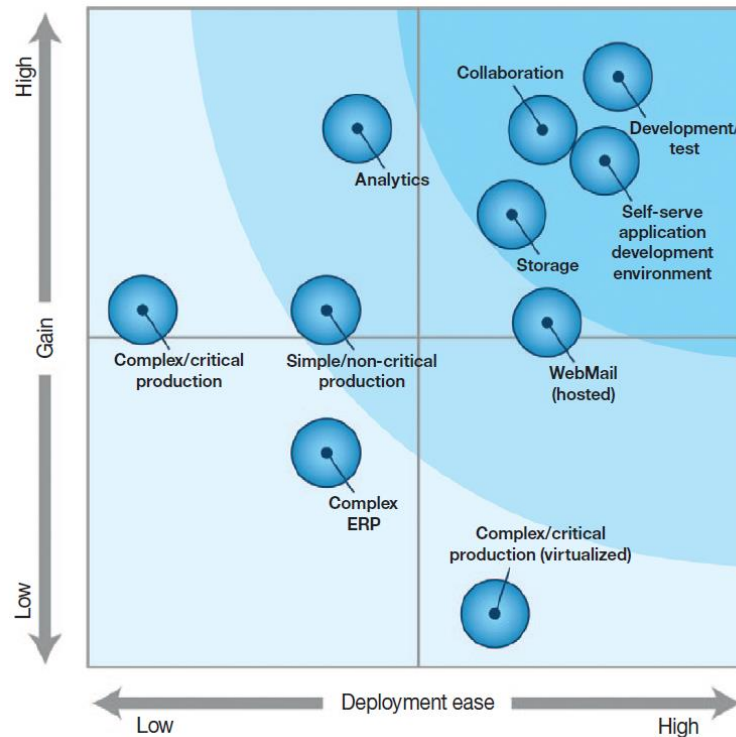


Figure 2-8: Workload affinity for cloud computing based on IBM experience [16]

It can be foreseen that DC technologies will evolve in the coming several years to overcome the challenges (security and privacy, trustable performance SLAs and performance isolation, management flexibility and smart orchestration) and to enhance the benefits so that a vast majority of imaginable computational workloads will be sustainable over the shared and federated utility computing environments

Here are, for example, DC requirements for some of the traditional workloads discussed in Section 2.3.1:

- **Engineering Applications** require specialized software, are computationally intensive and often long running. These application's computational demands are bursty – high demand during computation, no demand in between the runs when engineers evaluate results and refine their models. Network bandwidth is required in and out of the data center -- to transfer the model's data from the engineer's workstation and to deliver the results back to the engineer. Network latency also plays a significant role -- if the latency is greater than 40 ms between the engineer and the engineering cloud, the graphics will not display smoothly causing loss of productivity and complaints. Engineering workloads can benefit from elasticity in hardware resources allocation, e.g. compute clusters allocated to a task can grow and shrink, and software resources, e.g. flexible allocation of shared licenses to running instances of specialized application.
- **Scientific Applications** are both compute and data intensive, depend on large volume data warehouses allowing parallel low latency access, and consist of large amount of different components and system. One specific advantage of consolidated DC environments stems from abstracting the commonalities in software functionality of different components, thus allowing development, deployment, and even sharing of common middleware among different types of scientific tasks, e.g. MapReduce services, messaging services, parallel files systems and object storage solutions, cluster schedulers ...
- **Financial Applications** exhibit heavy resource usage, include a mix of long running tasks and short transactions, generate a mix of compute intensive and data intensive tasks, require high bandwidth, low latency networks. Consolidated environments can provide benefits to these workloads by offering flexible, automatic and rapid resource management and optimized middleware. The major challenge is high sensitivity of at least some of the processed and

stored data which prevents from deploying enterprise-level financial applications on public clouds. Private DCs, owned, managed, or hosted, can be invaluable in relieving financial organizations of the burden of ICT management.

2.3.3 Cloud-born Applications

Recent ICT trends have enabled totally new class applications, e.g. services mash-ups, scale-out social networks, ubiquitous mobile services, etc. In what follows we briefly cover the major trends. One must remember that the list below is in no way exhaustive as we witness new applications, use cases, and business models are being continuously conceived and implemented.

2.3.3.1 Big Data

As more and more devices are being connected to the global Internet, more and more data is generated, uploaded, and stored. Mining huge amounts of data for useful insights and finding creative ways to put the data and the insights to use is the subject of the Big Data field. Big Data applications are typically looking through vast amounts of data that cannot be stored in a single DB instance and thus require accessing lots of data storage locations.

2.3.3.2 Storage Clouds

Storage Clouds, e.g. Amazon EC3, are applications allowing users to consume data storage services in a pay-as-you-go manner. Instead of purchasing, maintaining, and operating their own storage devices and services, organizations and individuals alike enjoy cloud storage services to upload their data, to access it when needed, and to share it as they see suit.




Figure 2-9: Storage Cloud [27]

2.3.3.3 IaaS Management Stacks

Providing ICT infrastructure as a service requires high level of consolidation and resource utilization in order to be competitive and profitable. To govern ICT resources provisioning and management, IaaS management software is being used, e.g. OpenStack. Running and operating this software efficiently requires dedicated ICT resources, namely, compute nodes, storage to hold the IaaS inventory, and network to provide the management tools connectivity throughout the stack.

2.3.3.4 PaaS Offerings

PaaS paradigm requires, in addition to the ICT management stack needed by the IaaS layer, various bits and pieces that developers use to facilitate the software creation process, e.g. version control tools, collaboration software, development and debugging tools. In addition, commonly used application components start to be provided as a service by PaaS offerings. For example, message queue component, DB layer, REST layer, are expected “value add” for developers deciding to go with a certain PaaS provider, e.g. IBM BlueMix, Google App Engine, Salesforce.

2.3.3.5 New types of SaaS

Given the ubiquitous availability of ICT resources through SaaS and ease of application development through PaaS, the rapid pace of new applications appearance is not surprising. Novel mobile and cloud applications are being created and demonstrated in a matter of hours in various hackaton events. Of course, developing the product level service takes much more time but still, it is not closely comparable with the software development cycles the desktop or server applications have witnessed in the pre-cloud era. New types of applications provided as a service over the internet include social and e-commerce applications, location aware mobile services, smart city and smart government application, educational and entertainment services, online games, and more.

3 Analysis of the Data Centre Network Use Cases

In this section we present and analyse three representative DCN use case that are chosen to be handled as part of the COSIGN project. The use cases are: advanced IaaS provisioning (VDC) described in Section 3.1, advanced PaaS provisioning (multi-tenant software cloud) described in Section 3.2, and advanced infrastructure orchestration and management use case described in Section 3.3.

3.1 Virtual Data Centre

Interoute Virtual Data Centre (VDC) [9] is a scalable, fully automated Infrastructure as a Service (IaaS) solution that provides on demand computing, storage and applications, which can be easily integrated with the customer's IT infrastructure. The Interoute VDC solution is characterized by the simplicity and convenience of a public cloud, combined with the security of a private cloud on the same platform.

Interoute's cloud service platform relies on several Data Centres (DC) across Europe, directly connected and interconnected through the Interoute's pan-European network, which spans more than 60.000 km of lit fibre (Figure 3.1). Integrating computing and network virtualization, the Interoute VDC delivers virtual IT infrastructures and connects across Europe using the Interoute's virtualised MPLS fibre optic network. Customers can choose between VDC delivered as a private cloud service via the corporate VPN or as a cloud service via the public Internet, or both.



Figure 3-1: Interoute Data Centres across its network infrastructure in Europe.

Interoute VDC is a flexible solution that offers the same level of control as in physical DCs, but without the related costs. The virtual infrastructure is automatically deployed and provisioned in real-time and customized configurations can be created on-demand specifying RAM, CPU, storage, network and adding appliances (e.g. firewalls or IPSs) or options for scheduled local or remote backup and automated disaster recovery through the VDC Control Centre graphical interface. VDC components can be bound to an existing Interoute VPN network, a private VDC network or an Internet routable network.




Figure 3-2: Interconnection between VDC, corporate VPN and public Internet.

The process for requesting, creating and deploying a VDC is fully automatized and customizable, as shown in Figure 3.3. Customers can select the VDC zones to host their data and virtual resources from a variety of sites in Europe (e.g. London, Geneva, Amsterdam, Rome ...) and outside (Hong Kong and New York) to comply with local data legislation.

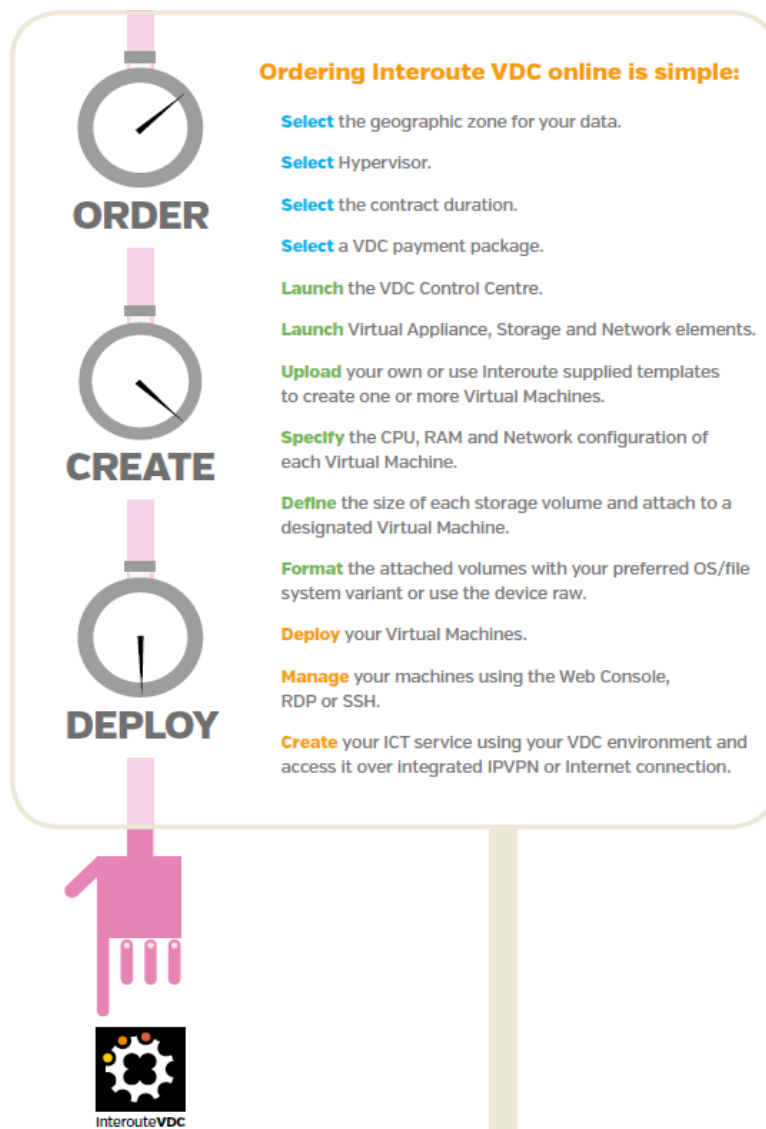


Figure 3-3: Process to deploy an Interoute VDC.

The Interoute VDC can expand and shrink to match the amount of provisioned machines, storage, and network to the business needs of the customers. APIs are available to automatically control the VDC infrastructure directly from the customer's applications, while monitoring reports allow tracking the usage of VDC services and controlling the resource consumption. Two price models are available: the utility pay-as-you-go pricing, where resources are priced per hour of allocation, and the commit pricing, with a fixed rate month for a package of resources in small, medium or large size.

The Interoute Unified ICT set of services, including private networking, computing and communications, are strongly focused on the customers' demands for information security, in data confidentiality, integrity and availability [10]. These requirements are addressed through the principles of logical separation and high availability, implemented across all of the technology domains. At the network layer, the traffic from each customer's domain is logically isolated, preventing data leakage and interference by external entities.

In the metro-area network, Interoute operate IEEE 802.3 standard Ethernet switching networks and implements the logical separation through the use of IEEE 802.1Q VLAN tagging of Ethernet frames and VLAN-aware TCAMs for Ethernet switching. On the other hand, high availability is guaranteed through the use of IEEE 802.1s (Multiple) Spanning Tree Protocol, constrained/partitioned VLAN-based MAC tables ("Port Security") and traffic rate policing.

In wide area networks, Interoute operates an IP/MPLS network to forward customer traffic to common wide-area destinations. The logical separation is implemented through the use of MPLS encapsulation (IETF RFC 3031 [11]), BGP-based MPLS VPNs (IETF RFC 4364 [12]) and Virtual Routing/Forwarding (VRF) tables. Customer networks are separated into groups of sites that are associated with a VRF table retaining information specific to a single customer, and a single Provider Edge (PE) router may accommodate a large number of VRFs (Figure 3.4). Beyond the customer-service provider links, the logical separation of the customer traffic on the service provider's internal link is preserved through MPLS encapsulation.




Figure 3-4: Logical separation of customers' traffic in the wide area network.

Interoute VDC allocates unique VLANs to a customer. These VLANs extend down to the partitioned resources within the VDC physical infrastructure, called "pod" (a pod is a switch infrastructure connecting compute and storage hardware, see Figure 3.5). Another separate range of VLANs can be acquired on-demand by the customers to build their own network topology.

Customers manage their virtual appliances/machines through VLANs, most commonly an MPLS VPN. Moreover, they can add other VLANs through the abstraction of adding "VDC networks", in order to provide access to the Internet or just a "private" network accessible only by appliances and machines in the customer's VDC.

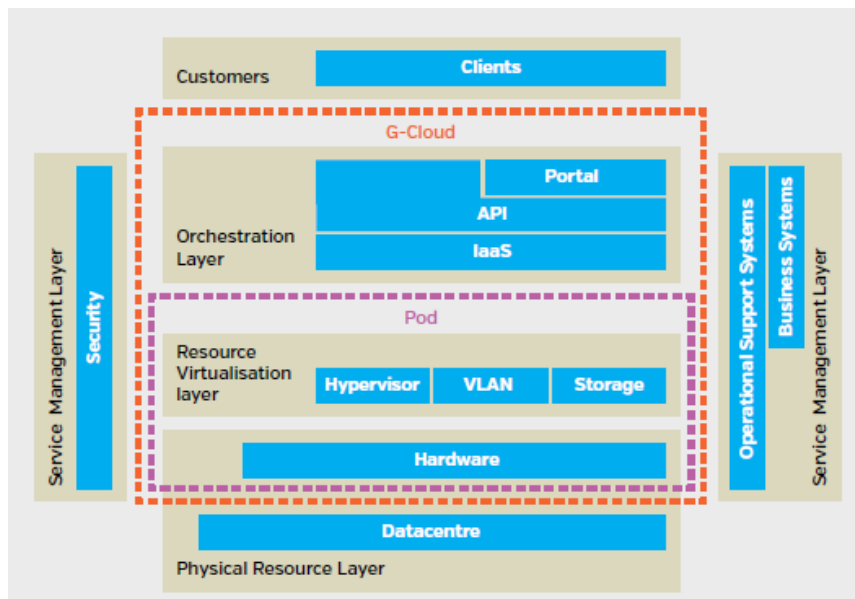


Figure 3-5: Architecture layers.

Analysing the Interoute VDC, we can derive some key use-cases and features to be supported by DCs and DC Networks (DCN), as follows:

- Automated and on-demand deployment and provisioning of customizable virtual infrastructures, integrating virtual appliances along with compute, storage and network elements;
- Elasticity features, monitoring service and programmable API to enable VDC control from customer's applications;
- Monitoring and accounting services supporting utility pay-as-you-go and commit pricing models;
- On-demand and scheduled backup in local or remote VDC sites;
- Disaster recovery and fast backup/restore procedures;
- Logical separation and high-availability;
- Integrated SLA monitoring and verification.

3.2 Multi-tenant Software Cloud

As can be seen from the previous use case description (Section 3.1), Virtual Data Centre is concerned with providing an advanced IaaS solution to multiple users over the single shared infrastructure. Use case described here is concerned with catering to another type of users – users unconcerned with the infrastructural aspects but interested in either the software development, deployment, and operations (PaaS), or in the services delivery or consumption (SaaS). From the networking perspective, the PaaS and the SaaS use cases are similar in that they do not require providing the users with trustful emulation of the physical network aspects. Instead, low level networking aspects need to be abstracted away from the application developers, operations, and consumers as much as possible, exposing only the important high level operational and functional aspects.

IBM develops cloud management software following the Software Defined Environment (SDE) principles, employing the power and the flexibility of software to the complex task of automation and specialization of the environment to the exact task at hand. Software cloud consumer is not only given the illusion of a dedicated environment like in a VDC case; this dedicated environment is also sufficiently abstracted so the consumer can interact with it without understanding the underlying infrastructural details and having to directly configure them. For example, in IBM Technical Compute Clouds, the technical computing users are presented with software Grid environments they are used to

work with, while the underlying hidden layers of infrastructure and middleware management take care of providing this illusion as presented in Figure 3-6.

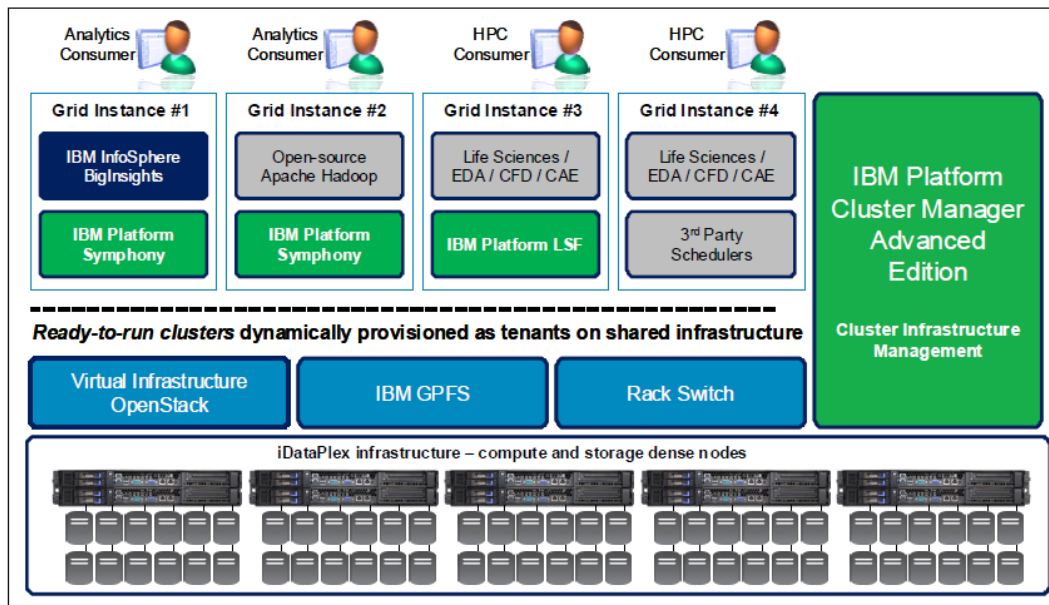


Figure 3-6: IBM Technical Computing Cloud architecture [14]

Another example is IBM Storage Cloud solutions [27] where the user consumes storage abstractions without configuring the details of the how the data is organized on the underlying media and how the underlying interconnect makes the data available for the usage, as soon as the data is sufficiently available, accessible, and secure. Such and other cloud solutions provide service developers and service providers with pattern-based experience of creating and managing their services. This paradigm allows the user to specify all the aspects important to its service functionality; specifying the infrastructural aspects is handled by the orchestration and management layers without burdening the user unnecessarily. In the networking domain, this pattern-based experience is realized, for example, as part of the IBM SDN VE, a software defined, overlay-based network virtualization solution [28]. In SDN VE, users define network connectivity blueprints that express the connectivity aspects of the application. SDN VE network blueprints can be used as patterns for deploying multiple application instances; the blueprint specifies the intended aspects of the network connectivity and the network services and not the operational or the configuration aspects of the specific networking equipment that happens to be used to realize the blueprint [29].

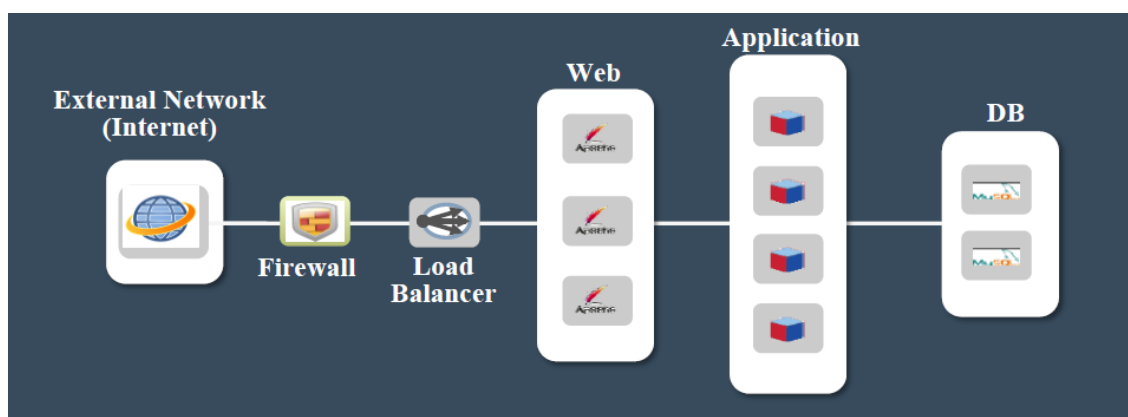


Figure 3-7: 3-tier Web 2.0 application architecture

Let's consider an example of Web 2.0 application depicted in Figure 3-7. This is a typical 3-tier application consisting of three sets of components – the set of Web Servers, the set of Application Servers, and a set of Database Servers. Web Servers need to communicate with the external clients and with the Application Servers but not with the Database Servers and not between themselves.

Application Servers need to communicate with the Web Servers and with the Database Servers but not with the external clients and not between themselves. Database Servers need to communicate between themselves and with the Application Servers but not with the Web Servers and not with the external clients. Additional requirements must be specified regarding the quality of the allowed connections, in terms of bandwidth, latency, packet loss, etc. Security and other service requirements, e.g. monitoring, traffic shaping, caching, can be specified by stating which of the allowed connections should be served by which types of service appliances. In traditional DCs, deploying the application described in Figure 2-1 usually requires interaction with the network administrator who will create network design, pre-provision the required subnets and services and let the application owner know where to deploy the application components and how to configure them over the network. Some of this work can be automated but the application owner is still exposed to network related aspects such as subnets, ports, gateways, IP addresses, etc. In VDC use case, the users are given an isolated slice of the networking underlying infrastructure where they can create and configure diverse networking designs suitable to the applications at hand. In multi-tenant software cloud use case, however, the user is the application owner and prefers to be isolated from the networking design aspects, concentrating on the application architecture and functionality. For a very large subset of typical cloud applications, e.g. multi-tier web applications, MapReduce applications, application clusters, networked storage application, etc., the patterns can be pre-provisioned. Moreover, both the cloud management stack and the underlying infrastructural resources can be optimized to better suit the specific patterns when they are deployed.

In this use case, the user is the application owner (who can be either the developer or the operator). The application owner is knowledgeable about the application architecture, its componentization and the connectivity requirements between the components. Based on the application architecture, the application owner specifies the network blueprint. The Data Center orchestration/management layer must be capable of realizing the networking configuration described by the blueprint through configuring the network devices, both at the deployment time and dynamically throughout the application life cycle.

This use case must be supported both over the bare metal Data Centers where multiple isolated applications belonging to different independent tenants are deployed over the shared DC infrastructure (simple IaaS) and over the Virtual Data Centers where these applications are deployed in the virtual slices exposed as part of the VDC use case (advanced IaaS).

3.3 Data Centre Operations, Management, and Orchestration

In order to efficiently operate the complex scale-out ICT infrastructure, the infrastructure provider must have tools for the DC management. These tools typically involve the equipment inventory services, with or without the automatic discovery, the equipment failure detection and health monitoring services, the software inventory with update, replicate, and health monitoring services, etc. Although no one is building the data centre exclusively for the ICT operations purposes, this class of DC applications pose a separate use case that must be serviced simultaneously with the main purpose of the infrastructure in case. In this section we discuss the network related aspects of this use case.

First, DCN operational use cases involve acquisition of the underlying network devices and network links state, monitoring their health and status, and management capabilities. At large scale, it is imperative that the management access for configuration or state check must be centralized as opposed to having to access the devices individually.

Second, network resource utilization visibility is very important for assessing the suitability of the infrastructure to the load, planned extension and evolution, as well as imperative for providing predictable services to the deployed use cases.

Third, 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.

According to the IBM Cloud Computing Reference Architecture, submitted to the Open Group Cloud Architecture Project and presented in **Figure 3-8**, the cloud management platform is an important part of the cloud management stack. Shown by pink boxes in **Figure 3-8**, the Common Cloud Management Platform comprises the Operational Support Services (OSS) and the Business Support Services (BSS).

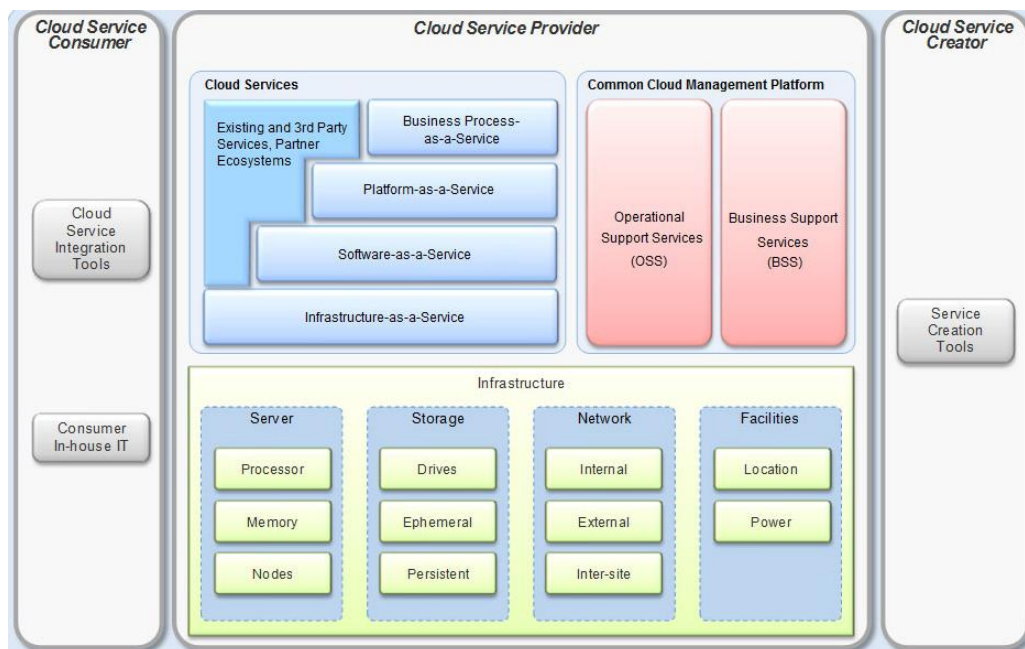


Figure 3-8: IBM Cloud Computing Reference Architecture

DCN infrastructure project like COSIGN is less concerned with the Business support Services. We narrow the Data Centre Operations, Management, and Orchestration use case to the Operational Support Services, specifically those related to the networking part of the infrastructure, presented in more detail in **Figure 3-9**.

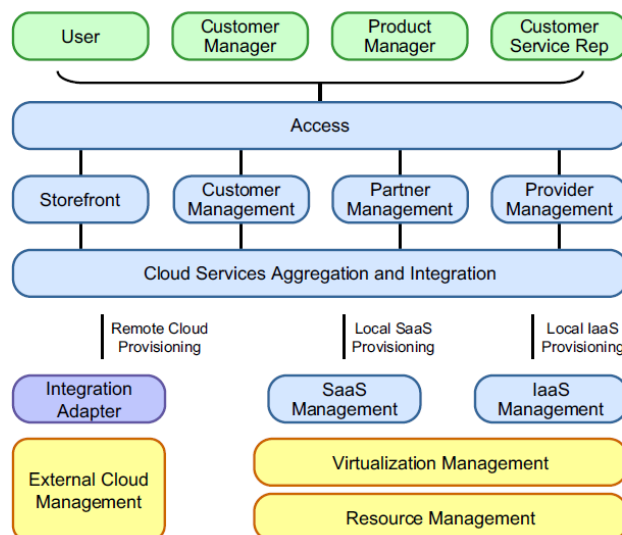


Figure 3-9: Operational support services [30]

From a networking perspective, there is a need to provide the functionality required by the typical users of operational support services, represented in the top part of the figure. This functionality is, starting from the bottom part upwards, comprises the network Resource Management, the network Virtualization Management, and the networking aspects of the IaaS and the SaaS management. To implement this use case, the COSIGN architecture will provide these functionalities and integrate them into the upper layers of the cloud management stack, e.g. OpenStack, so that users can access them through the interfaces and the application the cloud management stack provides. Network

Combining Optics and SDN In next Generation data centre Networks

provisioning and automation management will be thus provided on both the physical and the virtual levels, and will include management and discovery of network resources, dynamic provisioning and allocation, and monitoring across heterogeneous devices.

4 Limitations and Challenges of the Current Data Centre Network Design

The DC trends analysed in section 2.1 show a continuous traffic growth. Forecasts from CISCO [20] indicate a 25% CAGR, whereby the amount of DC traffic will reach 7.7 ZB per year in 2017 from 2.6 ZB per year in 2012 (Figure 4.1).



Figure 4-1: Global DC IP traffic growth – 2012-2017 [20].

Moreover, by 2017, over two-thirds of all data centre traffic will be based in the cloud, with cloud traffic representing 69% of total DC traffic and characterized by 35% CAGR in the period 2012-2017 [20], compared to the 12% CAGR of traditional DC traffic (Figure 4.2 and Figure 4.3). These trends are motivated by the rapid migration to cloud architectures and the capability of cloud DC infrastructures to handle significantly higher traffic loads. Virtualization and automation in particular are key factors to increase performance and guarantee high capacity and throughput.




Figure 4-2: Total DC traffic growth: cloud and traditional DCs – 2012-2017 [20].



Figure 4-3: Cloud DC traffic growth – 2012-2017 [20].

In terms of market segment, the CISCO Global Cloud Index indicates forecasts for the 2012-2017 period with 36% CAGR for consumer cloud traffic growth and 31% CAGR on the business side (Table 4.1). This means 4.3 ZB consumer traffic and 1.0 ZB business traffic to be reached in 2017, where business traffic must comply with stronger security requirements and provide fast and flexible access to large data archives, possibly integrating advanced analytics features. On the other hand, the consumer traffic growth is expected to be mostly motivated by video and audio-streaming, but also newer services like personal content lockers, especially accessed from mobile devices.

Cloud IP Traffic, 2012-2017							
	2012	2013	2014	2015	2016	2017	CAGR 2012-2017
By Segment (EB per Year)							
Consumer	918	1,384	1,923	2,581	3,368	4,310	36%
Business	259	371	496	644	810	1,004	31%
Total (EB per Year)							
Total cloud traffic	1,177	1,755	2,419	3,224	4,178	5,313	35%

Table 4.1- Global Cloud Traffic per segment – 2012-2017 [20]

Another trend, beyond the pure DC traffic growth, that is particularly challenging for DC networks is the increasing global DC virtualization, also reported in [20]. Even in this case, the trend shows a continuous transition of workflows from traditional to cloud DCs, with two-thirds of the whole workloads processed in cloud DCs in 2017 (Figure 4.4). The increasing degree of virtualization in the cloud (Figure 4.5) allows to deploy the workflows dynamically, to better meet the dynamic demands of cloud services. The cloud economics (e.g. for the physical infrastructure costs, service resiliency and scalability) suggest to distribute and migrate the workloads among a variety of servers, mainly inside DCs, but also across DCs and even in different geographical areas, e.g. for backup and service recovery strategies. It is clear that the highly distributed approach to accommodate the dynamic workloads generates huge traffic flows within and across DCs. The role of the DC network is fundamental to support this type of traffic and enable the dynamic migration of workflows across

servers. These reasons motivate the need to implement more automated procedures and reduce manual operator interventions for the configuration of the DCN, which needs to become an integrated part of the cloud service workflow.




Figure 4-4: DC workload distribution – 2012-2017 [20].




Figure 4-5: Trend of cloud virtualization – 2012-2017 [20].

Summarizing the analysis above, it is clear that the future DCNs must support huge amount of traffic, with unpredictable dynamicity and peaks. This traffic is characterized by flows in both vertical (i.e. traffic from the servers to the external network, for DC to DC and user to DC connections) and horizontal directions (i.e. traffic across servers located within the same DCs, belonging to the same or other racks). Moreover, most of the expected cloud applications are characterized by real-time and time-sensitive constraints, with high-bandwidth data exchange for both business and consumer segments, in support of big unstructured data analysis in the former case and audio/video streaming in the latter. Finally, intelligent management and control frameworks to automate the DCN allocation planning and re-configuration is fundamental to reduce the manual activities of the DC administrators and reach the best compromise between the cloud service performance and workload balancing (intra- and inter-DC) required to optimize the cloud infrastructure usage and operation.

These trends impose a transition from the hierarchical and over-provisioned DCN architectures with rigid configurations that are mostly adopted today to flat DCN architectures more suitable to (i)

support the east-west flows generated by distributed applications and workloads mobility and to (ii) scale better with the increasing DC size and traffic amount. In fact, the traditional DCN architectures have been mainly designed for user-to-DC traffic and are optimized to accommodate flows based on patterns that follow the client-server paradigm of vertical interaction. An example of this type of architecture is shown in **Figure 4.6**, with an access tier composed on Top of the Rack (ToR) switches connected to a layer of aggregation switches and finally to core switches for the Internet access.




Figure 4-6: Traditional DCN architecture multi-tier design.

This hierarchical organization does not scale well with a growing number of servers and ToRs, since it may require additional intermediate tiers at the aggregation level, composed of expensive equipment, increasing the overall costs. Moreover, this design does not optimize the performance of high-bandwidth and low-latency server-to-server traffic within and among racks that, as analysed above, constitutes an important component of the overall intra-DC traffic. In order to overcome these limitations, COSIGN will explore the feasibility of flat DCN solutions, introducing the high-bandwidth and low-latency capabilities of optical technologies at the data plane level as well as dynamicity and automation at the control and management level, through a DCN programmed by intelligent SDN-based control frameworks integrated with the overall cloud service management.

5 Data Centre Networks Requirements

In this section we specify the requirements towards the next generation data centre networks, deducing them from the use cases presented in Section 3 and from the limitations surveyed in Section 4. We are presenting the broad view on the requirements from the business perspective in Section 5.1, from the workload or applications perspective in Section 5.2, and from the infrastructure perspective in Section 5.3. Prioritization of the requirements and listing the ones most important for a successful next generation DCN architecture are left to Section 6, where architectural insights and guidelines for the rest of the COSIGN project are given.

5.1 Business Requirements

The analysis performed in sections 2.2.2 and 3.1 has allowed identifying a set of requirements for cloud services, as summarized in Table 5.1:

Requirement ID and Name	Requirement Description
R-BUS-01 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.
R-BUS-02 Elasticity	Self-service on-demand and/or automated scale-up and scale-down elasticity features, enabled through open and programmable APIs.
R-BUS-03 Monitoring tools	Integration of configurable monitoring tools to allow DC operators to keep trace of resource usage and service performance.
R-BUS-04 Virtual infrastructure control APIs	Secure open and programmable APIs to enable advanced control of a running VDC service from customer's applications. Examples of functions to be exposed by these APIs are monitoring, enforcement of automated elasticity rules or dynamic modification.
R-BUS-05 Accountability and SLA management	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.
R-BUS-06 Privacy and security	High-level of privacy and security, mechanisms for logical separation and high-availability.
R-BUS-07 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.
R-BUS-08	Service reliability and availability, combined with network resilience. Support of mechanisms for automated and fast disaster

RAS	recovery, integrated with on-demand or scheduled backups and replications of data in local or remote DC sites.
R-BUS-09 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.

Table 5.1- Cloud service requirements

The impact of these general requirements on the DC network is quite clear: since the network is a fundamental part of an integrated cloud service, all the features related to on-demand and automated deployment, provisioning, up-scaling/down-scaling, monitoring, accounting, security, logical separation and programmability of cloud resources are also valid for the network component. Moreover, further characteristics should be considered in support of specific management services. An example is the automated configuration of high-capacity and secure channels to allow the transfer of huge amounts of data within the same data centre or across different data centres (but still with impact on the internal networks of the involved data centres) to support on-demand or scheduled content replications without degrading the performance of the other running services.

Beyond the requirements reported in Table 5.1, it is interesting to note that a cloud report from Gigaom Research [13] has identified six key factors, weighted in terms of their relative importance, to evaluate cloud infrastructure providers in the European market for 2014. These factors reflect six areas where well established and emergent providers should compete to seek advantage.

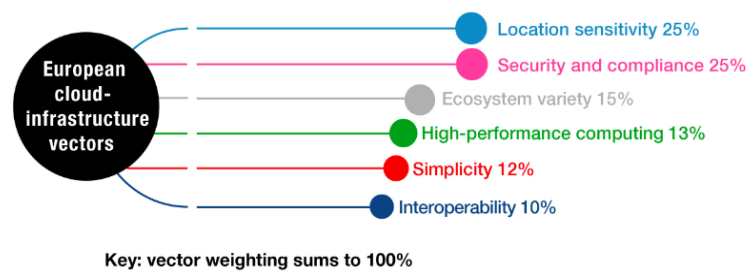


Figure 5-1: Disruption vectors for assessment of cloud providers in 2014 European market. Source: Gigaom Research [13].

These factors, represented with their weights in Figure 5.1, are the following:

- **Location sensitivity:** there are many reasons to consider the physical location of the data placed in the cloud as key for businesses. The first reason is related to the performance experience by the users to access the data: physical distance or network congestion between data centre and users increase the risk of poor performance that may become relevant, e.g. for high-speed businesses like algorithmic financial trading. Moreover, the current legislation in many European countries is very strict about the locations where information about European citizens are stored and processed.
- **Interoperability:** interoperability is taken in consideration from many customers when selecting a cloud provider, and this is especially true in this first phase when the market is still in evolution. A major concern is the interoperability and easy integration with the pre-existing IT infrastructure deployed on-premise, but customers may take also into account the risk of becoming locked in by proprietary features offered by their cloud providers. This explains the growing interest in open source cloud software and multi-cloud management platforms.
- **Security and compliance:** the possibility to demonstrate the compliance to standards for security and data protection (e.g. ISO 27001), personally identifiable information storage and processing, and domain-specific standards such as Payment Card Industry's Data Security Standard (PCI DSS) is a preliminary requirement to be accepted in some markets.

- **Simplicity:** the customer management interface must be intuitive and accessible, offering a clear product proposition. Moreover, it is necessary to reach an efficient trade-off between rich configuration options, especially suited for experienced users, and the simplicity of the offer for prospects and new customers.
- **High-performance computing:** the cloud offer should cover a wide range of computing instances, adapted and optimized for different levels of high-performance workloads to be applicable in a broad set of customers' use-cases. Moreover, the network performance within and across data centres should enable efficient distributed workloads.
- **Ecosystem variety:** the support of programmable and open APIs, possibly widely used or standard, is fundamental to create an active ecosystem of third-party users and system integrators around the cloud services offered by a cloud provider. This could result in the development of tools that increase the capabilities of the original cloud service, even though the integration of additional and external applications, with the potentiality to create added values for the offers initially proposed from a cloud provider.

5.2 Service Level Requirements

The main purpose of data centers is to provide storage and computing capacity to application software, to process and store data and content, and to support IT operations. Important and stringent requirements arise from the applications deployed in the data center, being as extensive as the wide range and variety of applications can be. Applications and service providers highly rely on the stability and reliability of cloud resources that could impair or even interrupt companies' regular operations. To provide a proper service level, an efficient design and management of application capacity, availability, continuity and security is fundamental. Service level requirements have to be considered both at the operation phase and at the design and planning phase.

Table 5.1 lists the requirements a service or an application can expect from the cloud.

Requirement ID and Name	Requirement Description
R-SERV-01 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
R-SERV-02 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
R-SERV-03 Service level monitoring	The data centre should facilitate application level monitoring required to decide whether service levels agreements are satisfied.
R-SERV-04 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.
R-SERV-05 Flexibility	The data centre needs to be flexible enough to accommodate different types and sources of data without sacrificing performance or latency and foresee any requirement for dynamically increasing or decreasing capacity (upscaling, downscaling) according to predefined policies

R-SERV-06 QoS provisioning	The data centre should provide mechanisms to satisfy application's QoS requirements (e.g., BW, delay, resiliency) and policies set to ensure the agreed service level
R-SERV-07 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.
R-SERV-08 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
R-SERV-09 Secure data management	Insurance of the end-to-end, secure use of data within applications as well as secured access to remote data sources
R-SERV-10 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)
R-SERV-11 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.
R-SERV-12 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.

Table 5.2 – DCN Service Level Requirements

5.3 Infrastructure Level Requirements

As it was presented in Section 2.1, DCs are required to provide increasingly more powerful IT capabilities: i.e. higher intra-DC bandwidth, lower latency, more efficient energy usage, and greater ease of service deployment and operation. Moreover, DCs must be sustainable at increasingly large scales, exhibit high levels of resiliency and availability in face of component failure or disaster, and be manageable and profitable.

All the above raises the importance of the infrastructure – its cost, operational complexity, ability to be repurposed on demand, etc. In this section, the DCN infrastructural requirements are discussed traditionally subdivided into the data, the control, and the management plane concerns.

5.3.1 Data plane

Data Centres are undergoing rapid changes. This includes the sizes, the business and the operational models, the computational and the storage equipment evolution, etc. Throughout these changes, networking infrastructure is required to efficiently interconnect all the communicating DC components. Networking data plane is the layer most affected by the changes in the rest of the DC infrastructure.

Before presenting the data plane requirements, we first briefly discuss the most important infrastructure and workload parameters that influence the data plane choices, both qualitatively and quantitatively.

5.3.1.1 Server Port Density and Capacity

The trend to multi-core and multi-virtual machine (VM) servers is accelerating the need to transition server connectivity from GbE/10 GbE, to 40 GbE/100 GbE. Also, more than 10GbE connectivity will also provide support for unified storage networking based on NAS, iSCSI, and FCoE. Figure 5-2 depicts the forecast for the server data-rates inside the data centres by Intel and Broadcom. It is estimated that by 2017 the majority of the Ethernet transceivers will be based on 40G modules.




Figure 5-2 Server Datarate Forecast by Ethernet connection type Source: Intel and Broadcom [31]

In addition, **Figure 5-3**, provided by the Ethernet alliance, shows Ethernet port shipment in millions for data centre applications.




Figure 5-3 Port shipments in Millions [32]

IEEE P802.3bs 400 Gb/s Ethernet Task Force was just established in March 2014, to provide Standard for Ethernet Amendment: Media Access Control Parameters, Physical Layers and Management Parameters for 400 Gb/s Operation. 400 GbE is expected to have a certain share in the horizon of COSIGN.

5.3.1.2 Traffic Matrix

Both the data centre environment and the applications deployed in the data centre have significant impact on the network design [33][34]. Studies analysing traffic in today's data centres show that 80% of the traffic stays within the rack, while for the intra-rack traffic the matrix is quite sparse, i.e., even the hot ToRs end up exchanging much of their data only with few other ToRs. While highly relevant for designing today's DCNs, this data does not necessarily project on the traffic generated by future data centre applications. Moreover, most of the data available today is collected for traditional tree-based DCN topologies and can lose its relevance for novel DCN designs. Therefore, generating new

data for the proposed topologies and benchmarks is an important part of the work to be done in COSIGN.

5.3.1.3 Resiliency and High Availability

Surveys show that high availability considerations top the service importance ratings for industrial data centres. According to [35], DC downtime, often caused by network outages, can cost anywhere from \$50K to over \$6 million per hour. Therefore, DCN data plane resiliency is of outmost importance to the next generation data centre architecture to be built in COSIGN.

Traditionally, availability requirements are quantified with X-nines notation, e.g. 99.999 for X=5 states that the system is allowed downtime of 5.26 minutes per year, 25.9 seconds per months, and 6.05 seconds per week. However, in a data center, as each failure event may trigger a set of follow-up actions that can further impact data-centre performance, it might be worth considering, whether one failure of 1 hour duration is preferable over 12 failures of 5 minutes duration each spread over multiple days.

5.3.1.4 The requirements

Summarizing the above considerations, Table 5.3 presents future DCN data plane.

Requirement ID and Name	Requirement Description
R-DP-01 Capacity	This requirement specifies both the aggregated and the link level DCN capacity. At link level DCN must support 10G to server to day 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.
R-DP-02 Latency	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.
R-DP-03 Reconfigurability/Flexibility	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.
R-DP-04 Resiliency and HA	DCN data plane should provide at least 5 nines availability and minimize the amount of systemic downtime events.
R-DP-05 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.
R-DP-06 Scalability and Extensibility	DCN data plane should support large-scale data centers and allow the existing data center to grow organically, both in number of servers, in number of the deployed workloads, supported amount of traffic, etc.

Table 5.3 – DCN Data Plane Requirements

5.3.2 Control Plane

This section presents the requirements for the COSIGN DCN control plane, collected through the analysis of the use-cases proposed in section 3, the current challenges of DCN architectures and Control Plane (CP) solutions and the new capabilities provided by the innovative optical technologies adopted in the COSIGN DCN data plane. The CP requirements are summarized in Table 5.1.

Requirement ID and Name	Requirement Description
R-CP-01 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.).
R-CP-02 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.
R-CP-03 Support for multiple 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.
R-CP-04 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.
R-CP-05 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.
R-CP-06 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.
R-CP-07 Network monitoring	The COSIGN CP must be able to provide monitoring functionalities for network resource usage, network service performance and faults detections.

R-CP-08 Programmable APIs	<p>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.</p> <p>These APIs should allow to expose some (limited) functionalities directly to the users and enable an easy integration with the overall DC control and management platform.</p>
R-CP-09 Support of network service monitoring and accounting	<p>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).</p>
R-CP-10 Support of scheduled network connectivity.	<p>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.</p>
R-CP-11 Network service resilience	<p>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.</p>
R-CP-12 Multi-tenant isolation	<p>The COSIGN CP must enforce suitable virtualization mechanisms to enable the sharing of a common physical network infrastructure among fully isolated connectivity services.</p>
R-CP-13 Support of SLA enforcement, runtime monitoring and assessment.	<p>The COSIGN CP must be able to support an SLA-driven provisioning and runtime management of intra-DC cloud network services. Moreover, it should provide automated mechanisms to monitor and evaluate the performance of the running services according to the metrics and KPIs specified in the enforced SLA and, when possible and required, automatically react to predicted or detected SLA breaches. Synchronization and cooperation procedures with the upper layer (e.g. the orchestration level) for SLA management should also be provided.</p>
R-CP-14 Easy interoperability with existing cloud management platforms for unified DC control.	<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>

R-CP-15 Integration with external connectivity services (inter-DC)	<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>
R-CP-16 Dynamic DCN reconfiguration for optimization strategies	<p>The COSIGN CP should support the automated re-planning of already established network services, in order 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.</p>
R-CP-17 CP architecture in support of scalable DCNs and network traffic	<p>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.</p>
R-CP-18 CP scalability	<p>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.</p>

Table 5.4 – CP Requirements Overview

5.3.3 Management and Orchestration

In order to support efficient operations of the DC workloads, DCN must be manageable as a standalone resource and as a resource integrated with the overall ICT infrastructure, namely with the compute and the storage resources. In addition, workload aware aspects are required to be provisioned at the management and orchestration level in order to deliver a coherent business environment as specified in Table 5.3 below.

Requirement ID and Name	Requirement Description
R-MGMT-01 Element Management	Management and inventory of network equipment, equipment management, maintenance, administration, configuration, and performance monitoring.
R-MGMT-02 Network Management	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.
R-MGMT-03	Provisioning and maintenance the of the network services required by the business workloads; forcing the traffic through the service

Network Service Management	enforcement points (appliances), service enforcement points configuration, elasticity of the network services tier and load balancing between the service instances.
R-MGMT-04 Self-service multi-tenancy	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.
R-MGMT-05 Workload awareness	Orchestrated DC management must be capable of receiving the workload hints or requests specifying network-related behaviours (connectivity, latency, path redundancy...) and or enforcing these requirements through the DCN controls and management tools. Moreover, the specified behaviours must be enforced throughout the dynamic changes allowed by the application life-cycle – component start-up and shut-down, migration, and failover.

Table 5.5 – Management and Orchestration Requirements

6 Classification and Prioritization of Requirements

Previous sections have described the requirements for COSIGN from different perspectives, not only technological but also for the applicability to the different stakeholders visions. In a data centre architecture, applications are located on the topmost layer, therefore, their provisioning and operation affect the underlying layers, driving the requirements in the COSIGN context. In this sense we can observe that some of the requirements for each of the identified levels (Business, Service, Management, Control, Data plane) are related because of the recursive way that functions have to be supported along the different layers.

In the following tables we classify these requirements regarding novelty, business value, knowledge and exploitability, trying to prioritize them in terms of impact. “Novelty” and “business value” are general requirements meanwhile “knowledge” and “exploitability” requirements have been retrieved based on the expertise of consortium partners, especially industrial ones.




Figure 6-1: Requirements – Novelty & Business Value

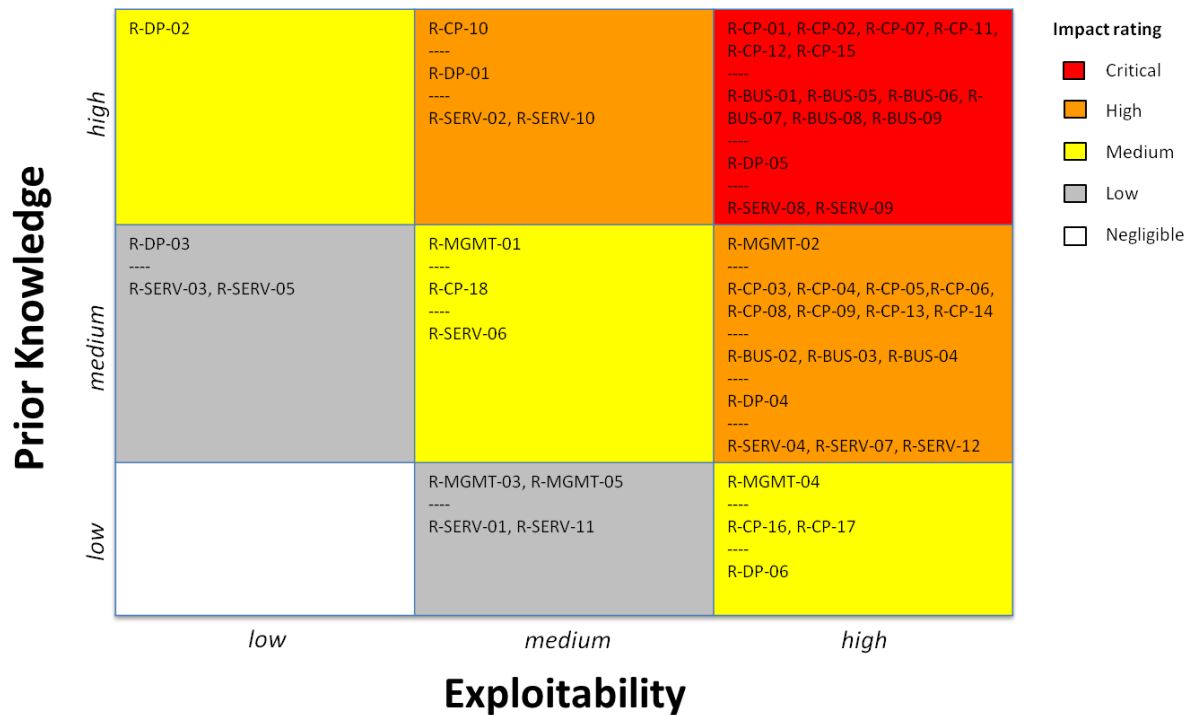


Figure 6-2: Requirements – Prior Knowledge & Exploitability

After having evaluated each requirement regarding novelty, business value, prior knowledge and exploitability, we can now assign a priority level to it by counting the number of dimensions in which it has been rated as "high". This provides us with an easy (but coarse-grained) way to group requirements in priority categories. The definitions of these categories are provided in the following table.

Category Name	Category Description
Very High Priority	Requirements that have been rated as "high" in all four dimensions
High Priority	Requirements that have been rated as "high" in three out of four dimensions
Medium Priority	Requirements that have been rated as "high" in two out of four dimensions
Low Priority	Requirements that have been rated as "high" in one out of four dimensions
Very Low Priority	Requirements that have been rated as "high" in zero dimensions

Table 6.1 – Priority Categories

In the following table we present in a summarized way the COSIGN requirements that belong to each of the above categories.

Requirement ID	Requirement Name
R-BUS-05	Accountability and SLA management
R-BUS-06	Privacy and security
R-SERV-09	Secure data management
R-BUS-01	Resource and Services Provisioning

R-BUS-04	Virtual infrastructure control APIs
R-BUS-07	Virtual infrastructure administration and management interface
R-BUS-08	RAS
R-BUS-09	Network Management Integration
R-SERV-07	Real Time control
R-SERV-08	Big data support
R-CP-01	Automated provisioning of intra-DC connectivity
R-CP-02	Support for customizable network services
R-CP-03	Support for multiple connectivity paradigms
R-CP-04	Multi-layer operation of COSIGN data plane optical technologies
R-CP-05	DCN resource usage optimization
R-CP-06	Elastic intra-DC connectivity
R-CP-07	Network monitoring
R-CP-08	Programmable APIs
R-CP-11	Network service resilience
R-CP-12	Multi-tenant isolation
R-CP-13	Support of SLA enforcement, runtime monitoring and assessment
R-CP-14	Easy interoperability with existing cloud management platforms for unified DC control
R-CP-15	Integration with external connectivity services (inter-DC)
R-CP-16	Dynamic DCN reconfiguration for optimization strategies
R-MGMT-04	Self-service multi-tenancy
R-BUS-02	Elasticity
R-BUS-03	Monitoring tools
R-SERV-02	Advanced Network Services Provisioning
R-SERV-10	Security management
R-SERV-12	Multi-tenant isolation
R-DP-01	Capacity

R-DP-03	Reconfigurability/Flexibility
R-DP-05	Scalability and Extensibility
R-DP-06	Scalability and Extensibility
R-CP-09	Support of network service monitoring and accounting
R-CP-17	CP architecture in support of scalable DCNs and network traffic
R-MGMT-02	Network Management
R-MGMT-03	Network Service Management
R-MGMT-05	Workload awareness
R-SERV-04	Adapting to application/service dynamicity
R-SERV-06	QoS provisioning
R-SERV-11	Self-service management interfaces
R-DP-02	Latency
R-DP-04	Resiliency and HA
R-CP-10	Support of scheduled network connectivity.
R-MGMT-01	Element Management
R-SERV-01	Seamless Network Resources Provisioning
R-SERV-03	Service level monitoring
R-SERV-05	Flexibility
R-CP-18	CP scalability

Table 6.2 – Requirements prioritized

7 Conclusions and Architectural Discussion

This deliverable has presented the set of requirements for the COSIGN architecture. Some of these are well known, some are emerging and some are innovative. The key requirements for each level of the COSIGN architecture have been identified, presented and discussed considering business, service level and infrastructure aspects, taking into account, in the latter, data plane, control plane and management and orchestration requirements, essential for the holistic end-to-end provisioning of data centre services.

A prioritization exercise of the requirements has been performed focusing on novelty, business value; prior knowledge and exploitability to allow COSIGN to concentrate on the most relevant requirements for its later support within the system features and functionalities designed in the architecture, optical infrastructure and software system.

7.1 Major Components of the COSIGN Solution

A first analysis of the requirements reveals the importance of four main areas that must be addressed by the architecture design which are Software Defined Networks (SDN), network virtualization, orchestration and the support of optical technologies.

7.1.1 Software Defined Networking

Software Defined Networking (SDN) is a novel networking architecture paradigm conceived to overcome some of the most acute shortcomings of today's DCN architectures, namely – the static nature, the management complexity, and the vendor lock in. Conceived in the academia for opening up the networking field to innovation and renewal, SDN has received lots of attention from the innovative industry segments, e.g. big data centre and cloud providers. Today, SDN concepts – decoupling the data plane from the control plane, logically centralizing the control plane, and providing well defined interfaces, both between the controller and the network elements and between the controller and the applications – are being actively incorporated into DCN solutions, by DCN owners and by networking vendors. In June 2104 ONF has published the SDN architecture document [36] where the basic principles are outlined and further elaborated to include the management plane and to develop advanced ideas, e.g. interaction between the controllers, both for multi-domain support and for building control hierarchies. **Figure 7-1** presents the concise representation of ONF's SDN architecture that includes all the major components – the controller, the network elements, the management, the applications, and the interfaces between them.

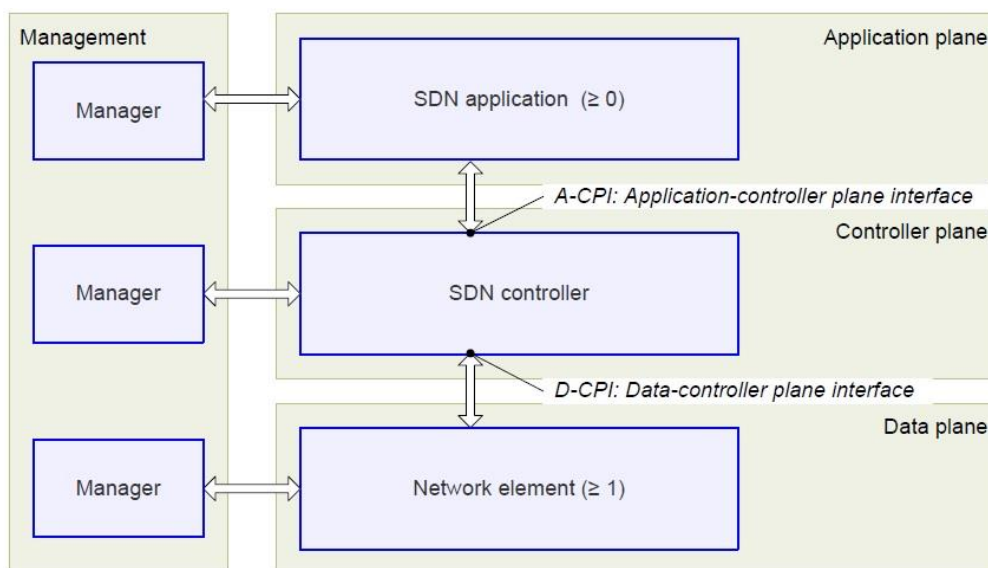


Figure 7-1: SDN overview, with physical data plane

We foresee, therefore, that basing COSIGN architecture on SDN principles will ensure the flexibility and the agility of a data centre control and management system and enable the support of many of the requirements associated with performance and adaptability.

7.1.2 Network Virtualization

Of all ICT resources, the network is the least sufficiently virtualized. See for example the comparison between the state of the compute virtualization with the state of the network virtualization in today's traditional DCN presented by James Hamilton of Amazon in his famous 2010 talk [37]. Although there were some positive changes in this area from the 2010, adequate virtualization of network resources – on the one hand capable of providing the applications with everything they might need from the networking layer while on the other hand harnessing all the advanced and/or specialized features of the underlying communication layer – does not yet exist.

In COSIGN, network virtualization will be approached with two goals in mind, goals matching the use cases presented in Sections 3.1 and 3.2. The first goal, matching the use case of the multi-tenant cloud, is creation of customizable application-level overlays, based on high level description of the application's connectivity requirements, including applying advanced services to the application traffic [29]. The second goal, matching the VDC use case, is creation of infrastructure-level network slices for multi-tenancy and simultaneous provisioning of networked cloud services. Both types of COSIGN network virtualization will be realized with the help of SDN principles, therefore will be programmable and composable allowing, for example, deployment of the application-level overlays over the virtualized network slices.

7.1.3 Orchestration

In order to fully realize the potential benefits of the novel DCN architecture, it must be integrated into the data centre orchestration layers, to build complex tailored cloud services dynamically and on demand. In COSIGN, interfaces and controls will be defined for integration with available orchestration mechanisms, exemplified by the integration with cloud management stack, e.g. OpenStack. As part of the COSIGN work, specific flows will be defined and realized demonstrating the advantage of orchestrated management of the networking resources as part of the overall data centre resource management. One example of such flow is presented in Figure 7-1.

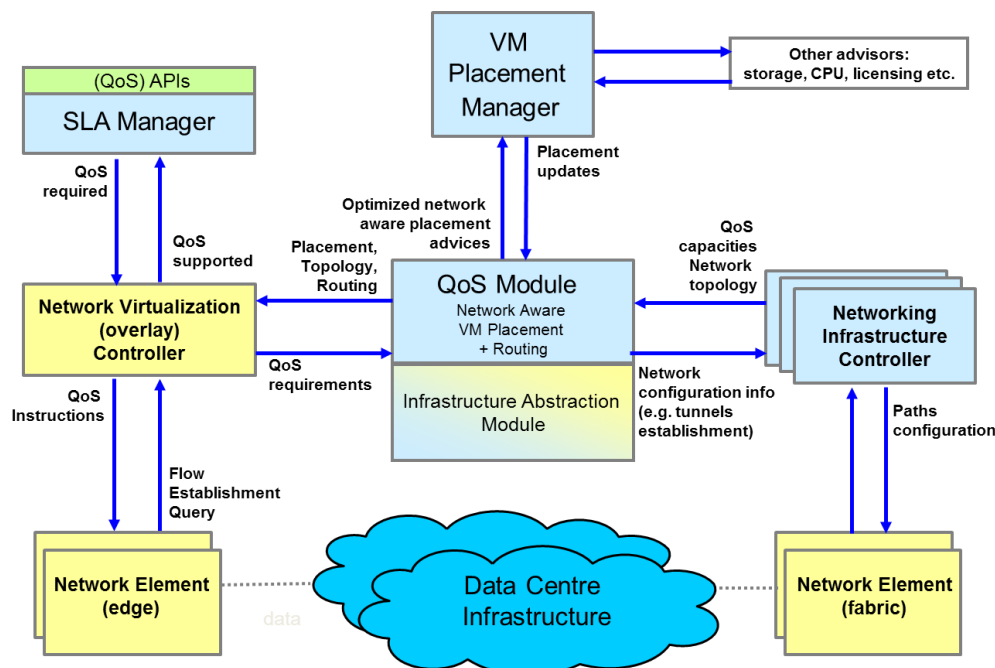


Figure 7-2: Orchestrated Example: Network-aware Placement with Network QoS Support

7.1.4 Optical DCN Technologies

As was shown in Section 2.1, data bandwidth requirements have grown very fast and this tendency continues, calling for fundamentally different underlying communication speeds. Copper links cannot satisfy the growing bandwidth, latency, resiliency, and density demands, making it mandatory to adopt alternative kinds of interconnect technologies – optics and photonics. As optical links replace copper in DCN environments, wasteful conversions from optics to electronics and from electronics to optics become a main source of overhead, both in added latency and in additional power dissipation. Advancements in optical and photonics research and manufacturing will eventually enable all optical DCN with ultra-high densities, low latencies and high flexible bandwidth. Taking into account the inherent limitation of optical technology – tendency to favour circuit switching over the packet switching, relatively slow reconfiguration time and lack of reliable buffering and per-packet processing, it becomes clear that SDN's programmability is not simply desirable, as it is for large scale switched Ethernet, but that it is mandatory.

In COSIGN, we will develop SDN control plane capabilities harnessing the benefits and making up for the shortcomings of optical interconnect technologies, so that the complexity of a flat architecture is controlled by a highly dynamic programmable software system.

7.2 COSIGN Architectural Blueprint

The COSIGN architectural blueprint is outlined in Figure 7-3. In the bottom part of the figure, there are physical architectural elements – network forwarding devices, both physical and virtual, compute hosts, links, network interface cards, etc. The left side of the figure, shown in grey colour, is collapsed and is present only to signify that upper layers of the management stack have access to computational resources for deployment, monitoring, management and operations. The right side of the figure, shown in blue, is the networking domain where most of the COSIGN innovation will reside.

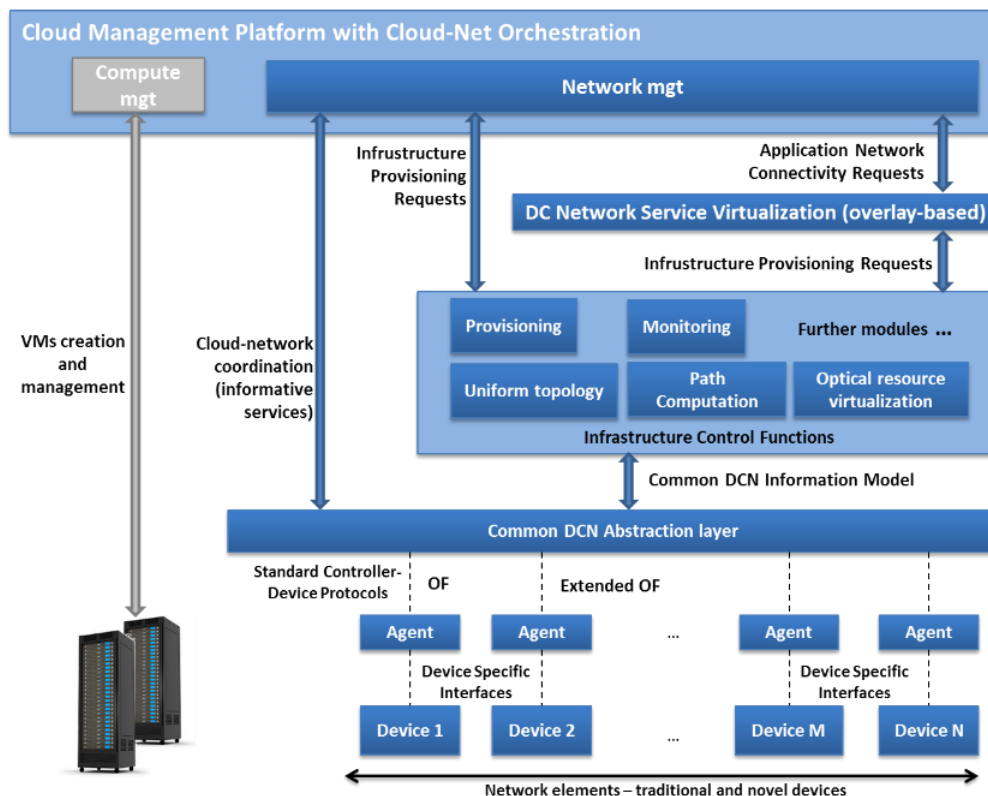


Figure 7-3: COSIGN Architectural Blueprint

In the networking part of the figure, we see the network elements, both traditional ones and the new ones that will be developed as part of COSIGN in WP2. These elements will be interconnected between themselves in a topology that will enable as much flexibility as possible, for any given type of constituent devices. As the technology is in constant flux, only specific types of hardware devices,

both fibre links and switches, will be available to the COSIGN team, according to the technologies provided by the HW members of the consortium. This makes it more important to create software architecture capable of harnessing a wide range of underlying technologies – their interconnect type, topology, non-functional aspects, etc. Such an architecture, exemplified in the COSIGN test bed, will remain relevant as the technology evolves, either as is or through organic extensions of the common data model layer or the infrastructure control protocols and applications presented in the figure. To demonstrate this aspect of the proposed COSIGN architecture; we plan to create several different underlying physical network designs capable of presenting the same logical/virtual view towards the data centre management and orchestration tools. In all cases where there is no need to expose additional underlying aspects to the integrated layers, all the technology differences will be accommodated in the control agent layer or the control protocol layer.

REFERENCES

- [2] "The NIST Definition of Cloud Computing", National Institute of Standards and Technology, US Department of Commerce, Special Publication 800-145, P. Mell, T. Grance, September 2011
- [3] "IDC Forecasts Worldwide Public IT Cloud Services Spending to Reach Nearly \$108 Billion by 2017 as Focus Shifts from Savings to Innovation", September 2013, <http://www.idc.com/getdoc.jsp?containerId=prUS24298013>
- [4] "European Space Agency delivers its SuperSite Exploitation Platform on Interoute Virtual Data Centre", <http://cloudstore.interoute.com/main/knowledge-centre/library/case-studies/european-space-agency-delivers-its-supersites-exploitation>
- [5] "UEFA & Interoute Virtual Data Centre", <http://cloudstore.interoute.com/main/knowledge-centre/library/case-studies/uefa-interoute-virtual-data-centre>
- [6] "Marconi University cuts cost of delivering e-Learning from the cloud by 23% using Interoute Virtual Data Centre", <http://cloudstore.interoute.com/main/knowledge-centre/library/case-studies/marconi-university-cuts-cost-delivering-e-learning-cloud-23>
- [7] "GC Europe Sees Data Transfer Speeds Increase 260 per cent", January 2014, <http://www.netcommseurope.com/index.php?section=1074&simple=go&type=newsEvent&id=4529>
- [8] "Interoute Virtual Data Centre. Hands on cloud control.", http://www.interoute.com/sites/default/files/product-instance/file-attachments/VDC_Brochure_UK_190214_ONLINE.pdf
- [9] "Interoute Virtual Data Centre and Security" (whitepaper) http://www.interoute.com/sites/default/files/product-instance/file-attachments/WhitePaper_Cloud_ComputingSecurity_ONLINE.pdf
- [10] E. Rosen, A. Viswanathan, R. Callon, "Multiprotocol Label Switching Architecture", IETF RFC 3031, January 2001
- [11] E. Rosen, Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", IETF RFC 4364, February 2006
- [12] P. Miller, "Sector RoadMap: the European cloud infrastructure market", Gigaom Research, March 2014, <http://cloudstore.interoute.com/main/sites/default/files/white-papers/Whitepaper-sector-roadmap-the-european-cloud-infrastructure-market.pdf>
- [13] IBM Technical Computing Clouds *An IBM Redbooks publication* SG24-8144-00 <http://www.redbooks.ibm.com/abstracts/sg248144.html>
- [14] The Computational Science Education Reference Desk (CSERD) <http://www.shodor.org/refdesk/>
- [15] Success in the cloud: Why workload matters. *Observations from IBM's own cloud transformation* IBM Office of the CIO, July 2013.
- [16] Google Data Center Locations. <http://www.google.com/about/datacenters/inside/locations/index.html>
- [17] DCD Intelligence. <http://www.dcd-intelligence.com/>
- [18] TechNavio Data Center Market Reports <http://www.technavio.com/data-center>
- [19] CISCO Global Cloud Index: Forecast and Methodology, 2012–2017 http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.html
- [20] Cisco and Partners to Build World's Largest Global Intercloud. Press release. <http://newsroom.cisco.com/press-release-content?type=webcontent&articleId=1373639>
- [21] Amazon Web Services Cloud Computing Platform Now Available from Datacenters in Australia. Press release. <http://phx.corporate-ir.net/phoenix.zhtml?c=176060&p=irol-newsArticle&ID=1757361>
- [22] IBM Commits \$1.2 Billion to Expand Global Cloud Footprint. Press release. <http://www-03.ibm.com/press/us/en/pressrelease/42956.wss>
- [23] Microsoft Datacenters website. <http://www.globalfoundationservices.com/>
- [24] A Comparison of Total Costs of Ownership of 10 Gigabit Network Deployments in the Data Center. Commscope whitepaper; November 2009. https://www.anixter.com/content/dam/Suppliers/CommScope/Documents/10GbE_TCO_whitepaper.pdf
- [25] Steve Ballmer: Worldwide Partner Conference 2013 Keynote. <http://www.microsoft.com/en-us/news/speeches/2013/07-08wpcballmer.aspx>
- [26] IBM Private, Public, and Hybrid Cloud Storage Solutions. *An IBM Redbooks publication* REDP-4873-01. <http://www.redbooks.ibm.com/redpieces/pdfs/redp4873.pdf>
- [27] Implementing IBM Software Defined Network for Virtual Environments. *An IBM Redbooks publication* SG24-8203-00. <http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg248203.html?Open>
- [28] An intent-based approach for network virtualization. Rami Cohen, Katherine Barabash, Benny Rochwerger, Liran Schour, Daniel Crisan, Robert Birke, Cyriel Minkenberg, Mitchell Gusat, Renato Recio, Vinit Jain *Proc. 2013 IFIP/IEEE International Symposium on Integrated Network Management (IM 2013)*

- [29] IBM SmartCloud: Becoming a Cloud Service Provider. *An IBM Redbooks publication* REDP-4912-00.
<http://www.redbooks.ibm.com/abstracts/redp4912.html?Open>
- [30] Christoforos Kachris, Keren Bergman, Ioannis Tomkos, **Optical Interconnects for Future Data Center Networks**, 2012, Springer
- [31] Fainman Y & Porter, G. "**Directing Data Center Traffic**". *SCIENCE* — October 2013, Volume 342, Issue 6155, pp. 202-203
- [32] Srikanth kandula, jitu padhye, and victor bahl, "**Flyways to De-Congest Data Center Networks**", in 8th ACM Workshop on Hot Topics in Networks, 2009.
- [33] Theophilus Benson, Ashok Anand, Aditya Akella, and Ming Zhang. 2010. "**Understanding data center traffic characteristics**", *SIGCOMM Comput. Commun. Rev.* 40, 1 (January 2010), 92-99.
- [34] www.visionsolutions.com, white paper: **Assessing the Financial Impact of Downtime**, 2010
- [35] SDN architecture. ONF Document Type: TR (Technical Reference), non-normative, type 2. ONF Document Name: SDN ARCH 1.0 06062014. Issue 1.0, June, 2014
- [36] Datacenter Networks are in my Way. James Hamilton's Blog RSS 2.0.
<http://perspectives.mvdirona.com/2010/10/31/DatacenterNetworksAreInMyWay.aspx>