Universitetet i Stavanger

FACULTY OF SCIENCE AND TECHNOLOGY

MASTER'S THESIS

Study programme/specialisation: Industrial Asset Managment	Spring semester, 2019 Open
Author: Vanessa Mejia Hove	Vanessa Hove (signature of author)
Programme coordinator: Jayantha P. Liyanage Supervisor(s): Jayantha P. Liyanage (UIS), Grete Breivik (ConocoPhillips)	
Title of master's thesis: Design of a Software-Defined Networking (SDN) Data Center to support further implementation of hybrid-cloud in ConocoPhillips	
Credits: 30	
Keywords: Software-Defined Networking (SDN), Application Centric Infrastructure (ACI), Cloud Computing, hybrid cloud, digital transformation	Number of pages: 94 + supplemental material/other: 24 Stavanger, 15. June/2019 date/year

Design of a Software-Defined Networking (SDN) Data Center to support further implementation of hybrid-cloud in ConocoPhillips

by

Vanessa Hove

Thesis submitted in partial fulfillment
of the requirements for the degree of
INDUSTRIAL ASSET MANAGEMENT
(Msc)

Faculty of Science and Technology

2019

Abstract

The industry is facing a digital revolution that is disrupting traditional business models; companies need to develop strategies to adopt emergent technologies that help to improve their value chain. Cloud integration is part of ConocoPhillips' digital strategy, as it provides access to many emergent technologies in an agile and cost-effective manner. This master thesis aims to design an SDN data center network solution that facilitates the migration of ConocoPhillips' computing assets to public clouds. The research methods consist of a comprehensive literature review of the digital revolution, cloud computing, and data center network technologies, followed by fieldwork based on qualitative action research. The findings from this project underline the need for an SDN data center network that provides high capacity, scalability, programmability, automation, multisite support, and cloud integration.

Keywords: *Software-Defined Networking (SDN), Application Centric Infrastructure (ACI), Cloud Computing, hybrid cloud, digital transformation.*

Acknowledgments

Being a full-time master student, full-time employee, and full-time mother might eventually be challenging. However, I am happy to look back and see how much I have achieved during the last years, and how this journey back to the academy has helped me to grow professionally.

None of this would have been possible without the support of my family, my mentors, and my job. I would first like to thank my thesis supervisor, Jayantha Prasanna Liyanage, for his dedication, time, and guidance during this project. He acknowledged the value of my thesis project and was always available to provide his advice. I would also like to thank ConocoPhillips for giving me the flexibility and opportunities that allowed me to enroll in this adventure. I am exceptionally indebted to my colleague, co-researcher, and friend Bjoern Fosse. Without his contribution, the fieldwork, business needs, and technical assessments would have never been accomplished.

I would also like to acknowledge my ConocoPhillips' advisor Grete Breivik. I am particularly grateful to her for her invaluable comments on this thesis.

I also wish to thank my mentor Anne Karin Borlaug; she is a self-effacing leader that knows how to encourage people to chase their dreams. I was lucky to meet her and to report to her one time. She inspired me to pursue my career purposes.

I would also like to thank my family for their love, support, and encouragement during these years of study. In particular, I want to thank my parents and brother for believing in me, and my son Martin for his patience, for giving me the time to do this, and for his unconditional love.

Finally, I want to thank my husband Kristian; he did not only took care of our child, dog, and home while I was writing, but he also was the second reader of this work and provided valuable input on this thesis.

Contents

List of Figures

List of Tables

Acronyms

ACI Application Centric Infrastructure. i, v, 39, 40, 56, 57, 66, 67, 70

AI Artificial Intelligence. 4, 9, 16, 17

AM Additive Manufacturing. 16

API Application Programming Interface. 67

APIC Application Policy Infrastructure Controller. 56, 64–67, 70

ASP Application Service Provider. 2, 20

AWS Amazon Web Services. 3, 22, 23, 25, 28

BA Business Analyst. 12

BD Bridge Domain. 67

BGP Border Gateway Protocol. 32, 69

BGP EVPN Border Gateway Protocol - Ethernet VPN. 39–41

BU Business Unit. 8, 44, 51, 53, 54, 69

CapEx Capital Expenditures. 28, 58

CDN Content Delivery Network. 30

CDPI Control Data Plane Interface. 42

CFO Chief Financial Officer. 1

CLI Command Line Interface. 3, 4, 29, 32

CMS Content Management System. 23

COP ConocoPhillips. v, 46, 47, 50

COPNO ConocoPhillips Norway. iv, 5, 68–70

CPU Central Processing Unit. 21

CRM Customer Relationship Management. 24

CSC Cloud Service Customer. 19, 20, 22, 29

CSP Cloud Service Provider. 19, 22, 25, 28–31, 45

DDoS Distributed Denial of Service. 23

DR Disaster Recovery. 56

EC2 Elastic Compute Cloud. 22

EIGRP Enhanced Interior Gateway Routing Protocol. 48, 64

EoL End of Life. 46, 50, 55, 58, 70

EoR End of Row. 48, 64

EPG Endpoint Group. 66, 67

F& L Flood and Learn. 39, 40

FEX Fabric Extender. v, 47, 48, 51

FO Fiber Optics. 44, 48, 53, 56, 64, 65

FTS Follow the Sun. 11, 54

Gbps Gigabits per second. 45, 47, 48

GRE Generic Routing Encapsulation. 38

GUI Graphical User Interface. 32, 43, 57, 65

HA High Availability. 33, 45, 46

HTTPS Hypertext Transfer Protocol Secure. 19

I/O Input/Output. 45, 53

IaaS Infrastructure as a Service. 22, 23, 26, 45

IoT Internet of Things. 4, 7, 9, 16, 17

IP Internet Protocol. 21, 33, 34, 37–41, 64–66, 70

IT Information Technology. 24–29, 70

L2 Layer 2. 55, 66

L3 Layer Three. 45, 55, 67

L4-L7 Layer 4 to Layer 7. 43

LCC Life Cycle Cost. v, 56–58

MAC Media Access Control. 21, 32–34, 36, 38–41

MC-LAG Multi-Chassis Link Aggregation. 35

MOC Management of Change. 67, 71

MPLS Multiprotocol Label Switching. 38

MTU Maximum Transmission Unit. 39

N2K Nexus 2000. 47, 51

NBI North-Bound Interface. 43

NGDC Next Generation Data Center. 63

NIC Network Interface Card. 21

NIST National Institute of Standards and Technology. 2, 19, 20, 24, 25

NSX Network Virtualization and Security Software. 39, 55, 70

NVE Network Virtualization Edge. 37, 39

ONF Open Networking Foundation. 41

OOB Out-of-Band. 65, 66

OpEx Operational Expenditures. 28, 59

OS Operating System. 21

OSPF Open Shortest Path First. 32, 39

OVN Open Virtual Network. 39

PaaS Platform as a Service. 22, 23

PHP Hypertext Preprocessor. 23

PIM Protocol Independent Multicast. 39

PoC Proof of Concept. 11

PPDIOO Prepare, Plan, Design, Implement, Operate, and Optimize. v, 50

QoS Quality of Service. 32, 54

RAM Random Access Memory. 21

REST Representational State Transfer. 67

S3 Simple Storage Servers. 23

SaaS Software as a Service. 2, 22, 24, 25, 30

SDN Software-Defined Networking. i, iii, 1, 3, 4, 26, 37, 41–43, 55, 56, 59, 69–72

SME Small and Medium-sized Enterprises. 28

SNMP Simple Network Management Protocol. 32, 65

SotA State-of-the-Art. 4, 53

SPAN Switched Port Analyzer. 54

SSH Secure Shell. 32

SSL Secure Sockets Layer. 23

STP Spanning Tree Protocol. 32–35

TB Terabyte. 23

TCP Transmision Control Protocol. 19

ToR Top of Rack. 48

TTL Time to Live. 35

UDP User Datagram Protocol. 19, 39, 41

VLAN Virtual LAN. 32–36, 39, 66

VM Virtual Machine. 21, 22

VMM Virtual Machines Monitor. 21

VNI Virtual Network Identifier. 39, 41

vPC Virtual Port Channel. v, 35, 36, 46, 48

VPN Virtual Private Network. 25, 38

VRF Virtual Routing and Forwarding. 66

VTEP VXLAN Tunnel Endpoint. 37, 39–41

VXLAN Virtual Extensible LAN. iii, v, 37–41, 55, 69, 70

WAN Wide Area Network. 32

XML Extensible Markup Language. 32

1 Introduction

The term *"digital transformation"* is often referred to during the last couple of years and professionals in all areas are getting occupied with how to create solutions that allow their companies to be part of the technological revolution.

The digital transformation goes hand in hand with the exponential increase of computing power, storage, and transmission capacity, that enables innovation and new business models to challenge the status quo, by disrupting the traditional value chains [37]. Among the disruptive technologies with significant impact on industry it is worth to mention: additive manufacturing or 3D printing, Robotics, Artificial Intelligence, Block Chain, Internet of Things, Big Data and Cloud Computing. The last one being one of the most important, as it provides the computational resources and characteristics that allow other technologies to exist *ibid.*

Cloud computing provides features such as ubiquitousness, automation, self-provisioning, and on-demand capacity, creating an ideal environment for new technologies to emerge, and allowing organizations to benefit from those technologies [52].

However, for enterprises with technological awareness wanting to benefit from cloud services, a strategy that adresses the requirements, pitfalls and weaknesses is essential as well as the engineering of a physical infrastructure that facilitates the integration between existing computing resources and public clouds. According to Forbes, *"74% of the Chief Financial Officers (CFOs) say cloud computing will have the most measurable impact in their business in 2017"*. However, even though most of the companies understand that they should have a migration strategy to the cloud, they do not know when and how to do it, [24].

The development of a cloud strategy is a comprehensive effort that requires collaboration between highly qualified professionals in the different IT fields such as Networks, Servers, Storage and applications. However, according to a survey conducted by LogicMonitor in November 2017, one of the most significant gaps in the interviewed organizations is the level of cloud expertise among their IT staff, [33]. Organizations need to focus on strengthening their internal cloud competencies and create collaboration spaces that allow their technical professionals to rethink and re-engineer the physical infrastructure to achieve a smooth and secure migration of their mission-critical assets to the cloud.

The Network is one of the most critical and complex components of the physical infrastructure in an Enterprise and plays a crucial role in the "cloud readiness" journey. Traditional

Data Center Networks need to change to provide support based on applications needs, such as dynamic locations, high bandwidth, traffic priority, network security, continuous monitoring, and automation capabilities, [39]. The demands of cloud networking have a substantial impact on the complexity of the networks, making the conventional configuration and maintenance methods outmoded and driving an evolution from distributed to central administration. [39]

There is not a standard network architecture that fulfills the Data Center requirements of any organization wanting to move assets to the cloud. Based on the nature of their business, regulations and internal policies, companies should assess their goals, technological maturity and constraints to identify and engineer the data center solution that adapts best to their needs, [48].

1.1 Background

Even if the Cloud Computing term is relatively new, the concept of outsourcing IT operations has its origins back in the 1980s with the upsurge of data center colocation services as the outgrowth of the Internet evolution [17] and in late 1990s with the emergence of Application Service Providers (ASPs) [29]. In the colocation model, the provider is responsible for everything related to physical infrastructure such as rack-space, physical security, temperature control, and electricity, while the customers own the physical equipment, their administration and support. ASPs on the other hand, not only provide the physical infrastructure, but also own the physical devices and support operating systems and customers applications, which in cloud terms is called Software as a Service (SaaS) *(ibid)*. Therefore, it is reasonable to assume that cloud computing is not an emerging technology, but a transformation of the IT operations outsourcing model enabled by the evolution of IT technologies, *viz.*: virtualization; networking; servers; storage; and transmission.

The National Institute of Standards and Technology (NIST), defines cloud computing in the following terms:

> *"Model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction [35]."*

According to this definition, cloud computing offers highly scalable resources and applications, with simple provisioning and accessible from anywhere. The definition is brief and indeed gives a general description of the cloud, but it does not clarify the intended public and the level of technological maturity required from organizations to take advantage of the offered services.

For most of the companies, a fully public cloud strategy, i.e., moving all their assets to public clouds such as Amazon Web Services (AWS) or Azure, is not viable. Hence they are adopting hybrid cloud deployments, that allow them to interconnect their in-house Data Centers or private clouds to one or more public clouds [30]. However, private and hybrid cloud deployments are technically demanding and require a correctly designed and well-dimensioned data center network infrastructure [23] that aligns with the cloud computing definition, *viz*: ubiquity; simple provisioning; and easy configuration [35].

The characteristics and topology of traditional Data Center Networks, have become a challenge in the deployment of hybrid clouds [23]. These networks are very static, and even if they have the technological capability to be extended across geographical locations, their complexity will increase to the point where they would become unmanageable. Moreover, their operation is distributed, meaning that the initial setup and further changes and provisioning, require network administrators to log in on each of the devices using a Command Line Interface (CLI) to insert the configuration code [27]. It is worth noting that the traditional data center networks might become a bottleneck for the inter-cloud connectivity, as they lack of fundamental attributes of cloud computing resources, *viz*: pervasiveness, self-provisioning; and automation.

The purpose of Next-generation data centers built upon Software-Defined Networking (SDN) is to provide a full integration between network resources, services and applications independent of their geographical location, manufacturer or virtualization features [26]. To achieve this level of elasticity, and simplify management and operations, SDN data centers are policy based and control all network equipment from a central device, offering programmability and orchestration capabilities *(ibid)*.

The transition from a traditional to an SDN data center could be challenging because of the broad scope, the initial investment, required expertise, change management and training of network administrators to support and maintain the new infrastructure. SDN data centers require new hardware arranged in a different topology, and do not offer the possibility for full integration between legacy equipment and new SDN devices [26], i.e. in an SDN implementation most of the existing network devices are not supported within the SDN fabric and need to be replaced, increasing the scope and cost of the project.

Further, an SDN deployment represents an organizational change and consequently a potential resistance or discontent from the technical staff supporting the network operations, and the customers that may be impacted by inconveniences that such implementation might cause, e.g., outages, redesign or less availability. In fact, this kind of upgrade to data center networks, requires a paradigm shift in the way network administrators manage and configure network devices, evolving from a distributed model based on CLI to centralized management through an infrastructure controller that serves as the primary interface to the fabric [27]. Consequently, some of the companies that have recently replaced their data center network equipment have adopted the traditional model, instead of the SDN approach, because their engineers did not feel comfortable with the new operational model and considered the required training as a burden *(ibid).* It would seem that engineers might have a resistance to change as a result of their apprehension of losing what Peacock defines in [43] as *"expert power"*, meaning they are afraid they will not be able to be on top of the required knowledge to operate and support the new system.

As it will be discussed, an SDN Data Center implementation is part of a long term strategy for companies to advance in their cloud readiness maturity and requires a high investment in equipment, planning, training, and expertise. However, according to some customers that are already operating SDN data centers, this technology also brings several advantages that might generate significant value for organizations, for instance, cost savings, enhanced agility, performance, security, reliability, scalability, compatibility, and integration capabilities [22].

By building on heuristics, research and work experience, this study will dive into the alternatives, advantages, and disadvantages of traditional versus software-defined data centers, and their role in a Cloud Strategy.

1.2 Overall study aim and individual objectives

The overall aim of this research is to design a data center network solution for ConocoPhillips Norway, that aligns with its global cloud strategy, positioning the company for the adoption and integration with emerging technologies, such as cloud computing, Internet of Things (IoT), Artificial Intelligence (AI), and Analytics.

However, before advancing on the implementation of the new network infrastructure, it is necessary to gain in-depth knowledge of the current status and the desired outcome, as well as the State-of-the-Art (SotA) and the available technologies. This research will be carried out through a detailed analysis of existing documentation, and comprehensive study of empirical data.

More specifically, the foundation of this study resides on the following individual objectives:

1. Understand next generation Data Center Network's concepts, services, value and its role in the inter-cloud journey.
2. Assess COPNO requirements and specifications for the data center network
3. Evaluate the technologies critically and architect a solution for the replacement of the data center network in COPNO.
4. Describe the Data Center Network implementation process, achievements, challenges, and topics for future research.

The first objective -*Data Center Network concepts and cloud integration*- comprises a thorough analysis of legacy and next-generation Data Center technologies and their role in the cloud and digital revolution context.

The second objective -*organization's requirements*-, will be approached by describing the current environment, the assets hosted by the local data center and the projected changes that might have an impact on the Data Center Network scope. This study will be carried out by the scrutiny of the existing Data Center Network and the collection of empirical data.

The third objective -*Architect the solution*- will evaluate the available technologies and possible solutions based on the identified business requirements and technical specifications. It will also provide a detailed network design based on the chosen solution.

The fourth objective -*Network Implementation*- describes the initial setup of the Data Center Network, the interim topology, and the migration process. It will provide the timeline of the project and recommendations for network optimization and future research.

The topics listed below are not part of the scope of this work:

- *Detailed configuration procedures and guidelines:* Configuration is not included in the scope, because guidelines, procedures, and best practices are well documented and possible to find on the internet. However, references to such documentation will be done throughout the report.
- *Network security analysis and considerations:* Network security is taken into consideration during the design phase, but it will not be included in the scope of this work, because I believe the subject itself is a topic worthy of dedicated research; and another department in the ConocoPhillips organization is responsible this matter.

However, the company's security policies and procedures are followed during the design and implementation phases.

- *Servers' and applications' analysis and Integration with public clouds:* The principal focus of this work is to design and implement a data center network that facilitates the integration with public clouds. The analysis of ConocoPhillips' applications, adoption of new cloud technologies and the migration or integration with public clouds is a topic of future research.

1.3 Structure of the Thesis

This thesis is divided into six sections and is structured as follows:

Introduction

This section introduces the relation between industrial assets management and the digital transformation, narrowing the scope of the thesis to what is the focus topic: The Data Center Network and integration with the cloud. It also presents the background and frame of the project, finalizing with the introduction of the overall aim and individual objectives.

Research Methodology

This section presents the research strategy, including the motivation and background for choosing such a model. It also describes the methodology, tools, and systems used to collect and analyze empirical data.

Literature Review

The purpose of this section is to provide a theoretical foundation for the main topics of this thesis, including the fourth industrial revolution or digital transformation, cloud computing concepts, and Data Center Networking.

State of the Art

This section explores the state of the art of Data Center Networking, analyzing the evolution from legacy networks and identifying the challenges or requirements that next-generation technologies aim to address.

ConocoPhillips Solution

This section discusses the framework and processes utilized to design and implement the ConocoPhillips' Data Center Network. It first presents a short introduction to ConocoPhillips Norway and their current data center implementation. Then it studies the company's business requirements to translate them into technical specifications. Further, it presents three different alternatives and the final decision based on a life-cycle cost and GAP analysis. Finally, it goes through the technical design and practical implementation of the solution.

Discussion and Concluding remarks

This section summarizes the findings related to individual objectives; Then it reflects over the learning and challenges faced during the project execution, finalizing with general recommendations and suggestions for future implementations.

2 Research Methodology

The purpose of this project is to look into the Data Center Network technologies and architect a solution that fulfills ConocoPhillips' business requirements and technical specifications. My reasons for diving into this topic are partially personal. I have worked for ConocoPhillips Norway during the last eight years: The first four years as a Telecommunication Engineer for major facility projects and the last four years as a Network and security engineer in the IT department.

During my time in the company, I have developed a particular interest in emerging technologies such as cloud computing and IoT. I also have been able to observe how networks are architected, managed and configured, and I believe the support model needs to evolve and become more agile in order to catch up with the changes we are facing in the digital context.

My motivation for studying the Data Center Networks in the cloud and digitalization context are also practical. ConocoPhillips Data Center Network equipment was starting to reach the end of support and needed to be replaced. This was an opportunity to evaluate the current solutions, emerging technologies, technical specifications, and business requirements.

Two persons from the Network Team were assigned to the project; a colleague and myself. My colleague had more than 15 years of experience in network technologies and broad exposure in different industries in the public and private sector, including government, telecom, education, transport, culture-art, banking, manufacture, and energy.

We worked together during all phases of the project. He provided insight into the existing technologies, protocols, standards, software, and hardware, while I focused more on the theoretical and practical part of the new implementation.

The first phase of the project was mainly to describe the theoretical framework for the technical solution. The objective was to study and understand the available Data Center Network technologies in the digitalization context we are currently facing.

The available literature about the topics I wanted to study was extensive; *"everybody was talking about digitalization and cloud computing."* For this reason and with the purpose of selecting reliable and constructive sources, I decided to carry out a systematic literature review and subsequently a systematic data collection and analysis.

The second phase of the project was the analysis of ConocoPhillips' requirements and the design of a solution that positioned the company for future public cloud integration. In this stage, I realized the need for an empirical study to collect, store, process and analyze information and data that could help to answer questions related to:

- *The current data center network*: everything from the architecture and technical details to the operation, performance, and challenges.
- *ConocoPhillips' IT strategy*: what were the goals of the company, the taskforces dedicated to achieving those goals, and the efforts done by Norway and other Business Units (BUs) to align with the IT strategy?
- *Support model:* What were the pros and cons of the current support model? What were the administration and maintenance requirements and how could we strengthen the competencies in-house to operate the new deployment? Did the network and server team feel comfortable operating the current solution? How did they react to the change?
- *Experts advise and best practices:* What were others doing in the Data Center Network field? What did the experts recommend and how could we proceed to implementation?

Given the nature of the questions and the population involved in the study, I determined that a quantitative research approach did not embrace the whole scope of the project. Moreover, I was interested in understanding the personal experiences of the users, operators, and engineers maintaining the network; This according to [36] relates to qualitative research design.

This chapter will go through the selected research design explaining in detail the strategy to collect and analyze the empirical data.

2.1 Research Strategy

Merriam and Tisdell, define in [36] four characteristics of qualitative research: *"the focus is on process, understanding and meaning; the researcher is the primary instrument of data collection and analysis; the process is inductive; and the product is richly descriptive."*

Rather than testing a hypothesis, I was interested in understanding the ConocoPhillips Data Center environment based on the needs and experiences of the company, network administrators and users, *viz: the focus is on process[36].* I was part of the team selected to do the key observations (the infrastructure team), and I conducted the interviews and analyzed the data, *viz: the researcher is the primary instrument of data collection* *and analysis[36].* After doing a systematic literature review, I understood that there was not a documented procedure about what technology companies should choose and how it should be implemented. In other words, I claim that it was necessary to part from existing quantitative and qualitative theories to find the solution that could best address the requirements of this specific case, calling for an *inductive process[36].* Finally, this work provides a comprehensive analysis and description of existing literature, the ConocoPhillips' case, requirements, alternatives, chosen solution, implementation, and lessons learned in the process, *viz: the product is richly descriptive[36].*

In conclusion, I chose to conduct a qualitative research study from a constructive/interpretive philosophical perspective: I departed from a conceptual point of view and built knowledge based on observations and interviews done on a selected population[36].

Qualitative research is a broad field that comprises several designs depending on the methodology and the research questions. It was important to narrow the focus and to adopt a design that could best contribute to the structure, data collection, and analysis in a systematic manner. After reviewing the key concepts and forms of qualitative research, I decided to do a thorough evaluation of three of the qualitative research designs: Grounded Theory, Case Study and Action Research.

Grounded theory was discarded in an early stage, as the objectives were clearly defined and I was not seeking to build up a theory, but to apply the existing concepts to solve a particular problem.

Deciding whether to choose a case study or action research took further investigation. Even if I was studying a particular case in a company, I was not documenting the work someone else had done. My purpose was to understand the current implementation of the Data Center Network, and also to produce and implement a new solution that contributed to positioning the company for the adoption of emerging technologies such as Cloud Computing, AI, big data, IoT, and analytics.

Moreover, I was part of the team involved in the research; the IT infrastructure team is divided into several groups: Network Services (my group), network security, servers, storage, onsite support, remote login, and service desk. This, according to Cunningham [25], cited by Biggam in [20] corresponds to action research;

> *"Action research is where the researcher starts with a particular problem that he wants to solve or understand better, usually within the environment where he is working."*

Furthermore, this study aligns with the principles of action research defined in [36]:

1. *"Focuses on a problematic situation in practice":* This work was focused on improving the performance of the Data Center Network, simplify its operation, and plan for future public cloud integration.
2. *"The design of the study is emergent,.., oriented toward some action or cycle of actions in which researchers and participants engage to improve practice":* This work was based on a plan and actions implemented principally in the network and server team, followed by observations and reflections about the results.
3. *"Researcher engage participants as co-investigators":* The research was done together with one of the team members as a co-researcher. The rest of the team provided input during the whole process and received periodic updates about the progress.
4. *"The researchers and co-investigators collect and analyze multiple forms of data in a systematic way":* four types of data were systematically collected and analyzed. The theoretical framework through a literature review; Input from other team members through interviews and observations; Experts' opinions through external interviews; Study of company documentation and existing equipment.

In brief, this project has as a purpose to generate value by producing two main outcomes: From the technical perspective, to enhance the performance and simplify the operation of the data center network. From the societal point of view, to improve the user experience and facilitate access to new technologies. Based on the characteristics, desired outcomes, focus, and stakeholders of this study, I decided to adopt a *Technical Action Research* approach.

2.2 Data Collection: Site and Sample Selection

The site selected for this project is the Data Center in the main office of ConocoPhillips Norway, located in Tananger. Rather than conducting a comprehensive analysis of all features, services, and applications hosted by the Data Center, this study focuses on network technologies whose primary purpose is to provide connectivity.

The empirical part of this work seeks to address the second and the third individual objectives: (2) Assess the organizations' requirements and specifications for the data center network; (3) Evaluate the technologies critically and architect a solution for the replacement of the data center network.

For the second objective -*Assess the organizations' requirements and specifications for the* *data center network*- the primary data collection methods were interviews and observations of the sample population; the secondary data collection was based on ConocoPhillips documentation review.

The interviews and observations were focused on the infrastructure, business partners and application teams in ConocoPhillips Norway; Specifically the technical staff responsible for the servers, network, and applications hosted by the data center. This sample population is accountable for most of the data center's assets and they have insight into the advantages, challenges, needs, and future projects. Focusing on this group allowed me to narrow the scope, providing reliable technical information and feedback. Figure 1 illustrates the sample population.

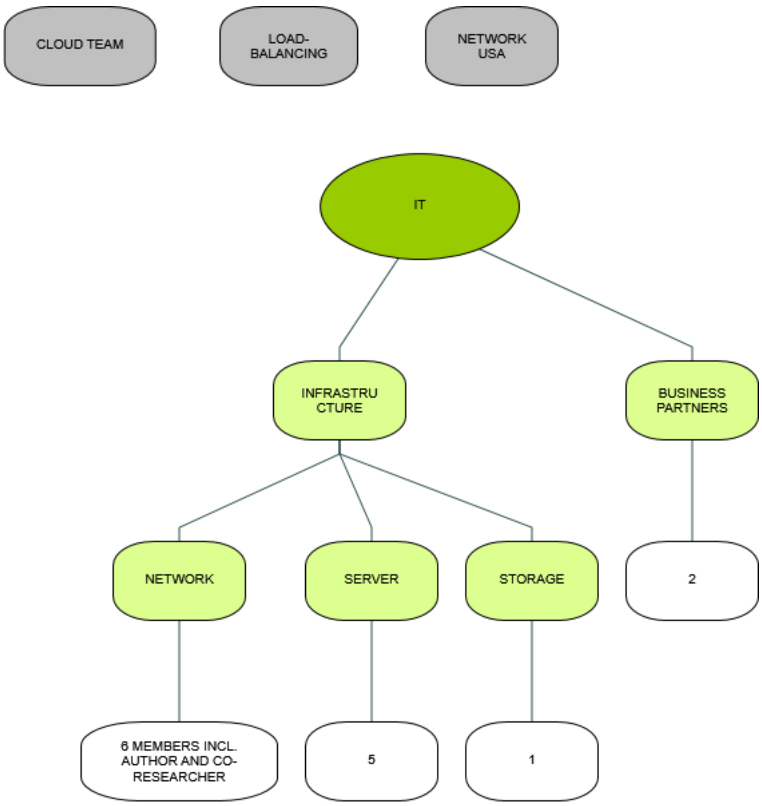

Figure 1: Sample population for the empirical study

Other groups interviewed were the global cloud, load-balancing, and network teams. They shared their experience with similar deployments, provided information about the global data center network philosophy, Follow the Sun (FTS) support model, cloud strategy, and results of performed Proof of Concepts (PoCs).

The interviews were conducted in a semistructured way: A list of questions was prepared followed by a meeting with the focus person or groups. This technique opened for brain-storming and allowed for good discussions with the experts in the different disciplines. The results provided a good foundation to understand the existing deployment and the requirements for the new solution, while the rest of the team was empowered to contribute to the design optimization by sharing their knowledge, covets, ambitions, and concerns.

The secondary data was collected through a company documentation review. The principal sources were:

- ConocoPhillips global IT department goals and cloud strategy.
- Information e-mails from the global IT management team.
- Documentation of the current Data Center Network.
- Network Topology Diagrams.
- Configuration files of the existing network equipment.
- Topology diagram of the cloud integration solution in ConocoPhillips USA.
- Sharepoint site of the global hybrid cloud team.

For confidentiality reasons, the transcription of the interviews will not be attached to this report. However, the analysis and results of the collected data will be provided in section 5. The collection of questions used for the interviews are documented in the following appendices:

- Appendix A: Collection of questions for server staff interviews
- Appendix B: Collection of questions for network staff interviews
- Appendix C: Collection of questions for USA network interviews
- Appendix D: Collection of questions for interviews with Business Analysts (BAs)
- Appendix E: Collection of questions for external experts

For the third objective -*Evaluate the technologies critically and architect a solution for the replacement of the data center network*-, the primary data was based on the review of existing literature and state of the art; this is documented in sections 3 and 4. The secondary data was necessary to validate the design through interviews with experts in Data Center Networking, and attendance of international technical conferences.

The principal sources of this secondary data include:

- Cisco Live Barcelona 2018

- Five days of consultancy with a Data Center expert from Atea
- Meeting with a Cisco expert
- Interviews with the network team in USA

2.3 Data Analysis

The empirical data and existing literature were structured and analyzed using the software NVIVO v12. After performing the interviews, the data was classified in different nodes that helped to address the second and third individual objectives: *(2) Assess the organizations' requirements and specifications for the data center network; (3) Evaluate the technologies critically and architect a solution for the replacement of the data center network.* Figure 2 illustrate the NVIVO nodes used to sort and analyze the collected empirical data.

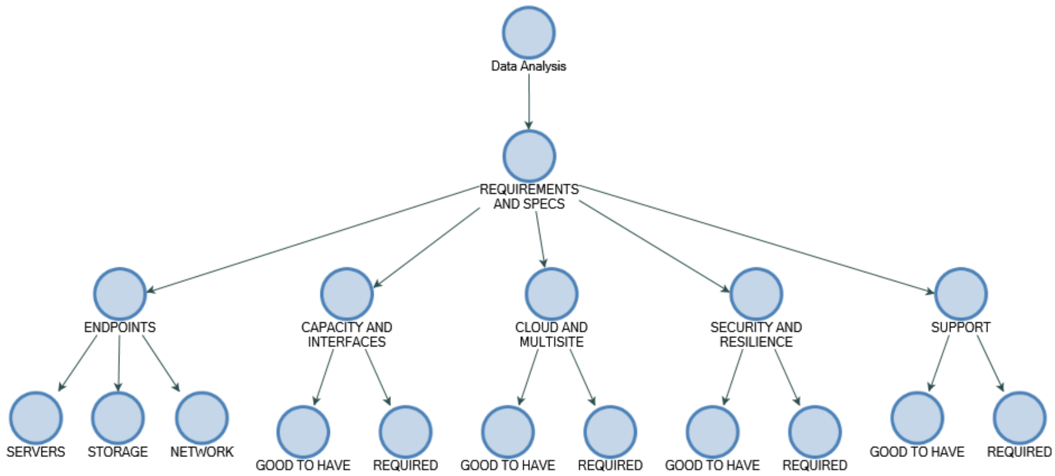

Figure 2: NVIVO Nodes for Data Analyzis

The interviews were performed using meetings and discussing the topics with the sample population; It was done in this way to make the interviewed persons feel comfortable and included, rather than questioned about their duties. The data was collected taking notes and was validated with the participants at the end of each meeting.

After analyzing the data and correlating with the literature review, it was sometimes necessary to interview some of the persons again to clarify concepts, issues or to gather more information, i.e., the empirical part of this work was a joint effort with the involved parties and required a high level of collaboration to understand the specifications and optimize the design.

The results of the data analysis are presented in section 5.

3 Literature Review

This chapter is linked with the first individual objective of this master thesis: *Understand cloud computing concepts, services and value to organizations, and the Data Center Network role in the inter-cloud integration.* The focus of this thesis is to dive into the data center solutions for companies that cannot adopt a full Public Cloud approach, where hybrid cloud represents an option that simplifies the private to public cloud integration.

The first section gives an introduction to the digital transformation and how it is disrupting the way everything operates from the industrial point of view.

The second section, briefly explains the main concepts of Cloud Computing, including the definition, advantages, types of deployments and some of the leading cloud computing solutions. The purpose of this section is to get an overall understanding of Cloud Computing without analyzing the model in detail.

The third section explores the Data Center Network technologies and features, focusing on those that are relevant for the integration of on-premises assets with public clouds.

3.1 The Digital Transformation

The industries, societies, and governments are in the middle of a revolution triggered by the emergence of innovative technologies that are disrupting the way everything operates. Schwab in [49] cited by Trailhead in [55] defines an industrial revolution as:

> *"the appearance of new technologies and novel ways of perceiving the world [that] trigger a profound change in economic and social structures."*

According to Schwab, the digital transformation we are currently experiencing is leading to a *"fourth industrial revolution."* Indeed building on its predecessor, *"the computerization of Industry"* and thanks to the advances in the computing power, storage capacity and transmission rates, the digital transformation is acting as an enabler for the development of new technologies that are changing the traditional business models, the industry and the societies[34].

Traditional business models are based on a *"vertically integrated value chain,"* where companies are responsible for most of the processes and activities which are highly integrated and dependent on each other[37]. Technology is disrupting the vertical value chain, forcing an evolution towards a "stack-based structure," that allows organizations to benefit from emerging technologies applicable in small segments of processes (ibid). For instance, a company could use IoT to gather data from sensors in a facility, while other activities such as data processing and monitoring are performed as before. This solution targets only a small portion of the operations section of a value chain illustrated in figure 3.

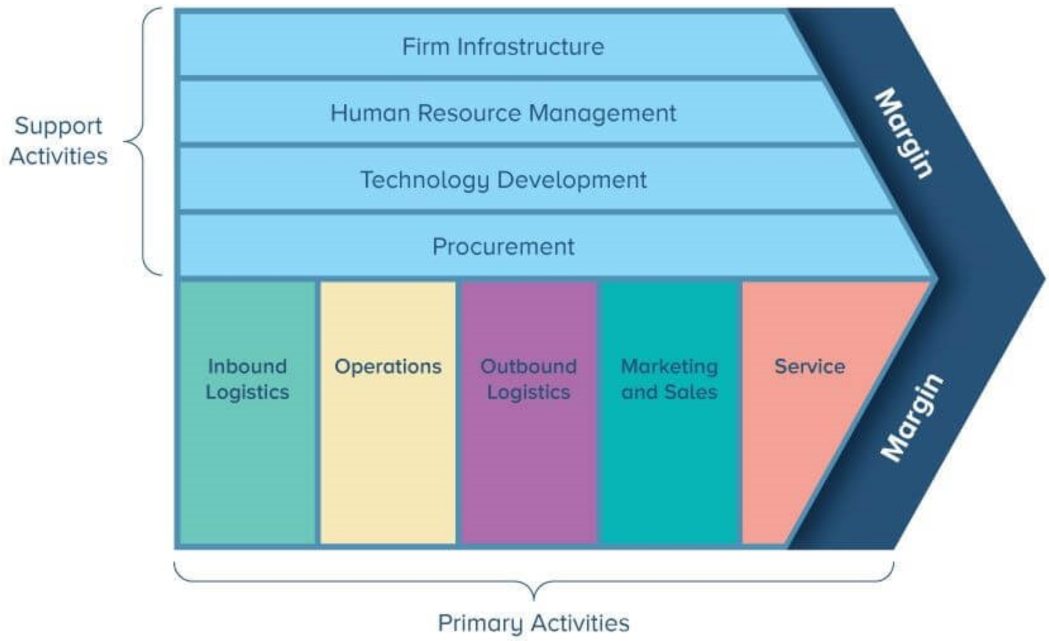

Figure 3: Porter's Value Chain Model, taken from [53]

Although next-generation information technology is disruptive for traditional value chains, it also generates opportunities for organizations. It provides ubiquitous access to data and computational resources, creates room for innovation, opens for a broader competition giving small actors entrance to the market, creates additional value, and allows customers to evaluate and choose from a broad portfolio of solutions that best address their needs[37].

Some of the technologies that are making the digital transformation possible by disrupting the structure of organizations include:

- Additive Manufacturing (AM) is also referred to as 3D printing and it is a technique that builds tridimensional objects using different materials such as plastic, ceramics, concrete or metal. AM can be used to print small objects, prototypes, machine parts, and even human organs or houses.
- Robotics: Robots are changing the way humans do physical activities such as building cars, vacuum cleaning, cutting the grass or driving cars[55].
- AI Uses algorithms to collect massive amounts of data that is analyzed to identify patterns and predict possible outcomes[37]. It can help companies to make decisions,

forecast production, or target customers.

- Block Chain is a digital register that records data about transactions. It is open and distributed, meaning that no one owns or control the applications and anyone can access the information[55]. Some possible uses include elections, personal banking, the sale of goods and patents *(ibid)*.

- IoT refer to objects and people connected to the internet via wireless and mobile networks, sharing data that is analyzed to provide intelligence to processes such as driving, building, and manufacturing. E.g., sensor networks in the roads that interact with self-driving cars, intelligent buildings, or condition monitoring sensors in process facilities.

- Big Data is defined in [37] as:

  > *"Enormous amount of unstructured, fast-moving data." "it can be traced, connected, and analyzed to generate business value and even to transform whole business models."*

  For example, in the case of the oil and gas industry, companies collect massive amounts of data that needs to be processed and analyzed for different purposes, such as understanding reservoirs, correlating well behavior, or performing predictive maintenance.

- Cloud Computing is a virtualized pool of computing resources self-provisioned on demand that can be accessed from everywhere via devices with a network connection. Cloud Computing services emerged from 2006 to 2009 and have evolved since then, providing an extensive portfolio of services within Computing, Storage, Machine Learning, Artificial Intelligence, Analytics, Block Chain and many others.

  Cloud computing is an enabler of the digital transformation as it hosts numerous resources that can be accessed on demand by everyone, i.e., customers only pay for what they consume, they are not committed to fixed contracts and can upscale or downscale their consumed services whenever they want based on their needs. This model allows small and big companies to develop solutions and make them available without significant upfront investments, opening the market for innovation and ideas, many of which are disrupting the traditional business models.

  As mentioned in the introduction, the focus of this research is to design a data center network infrastructure that aligns with the "all in cloud" strategy of ConocoPhillips Norway. Cloud computing definition, concepts, service models, deployments, advantages, and limitations are explained in detail in the next section.

3.2 Cloud computing concepts and evolution

There are many perceptions about the definition of Cloud Computing, among which, the one from NIST in their publication 800-145[35], is the most officially recognized:

> *"a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."*

This definition was published in 2009 as an effort to define a road-map towards cloud computing. Even if it describes a general understanding of the concept, it might appear to be an oversimplification of it. In [35], NIST also introduces five *Essential Characteristics* of cloud computing services:

- *On-demand self-service[35]:* Cloud Service Customers (CSCs) can request cloud computing services through a computer interface and without human interaction with the Cloud Service Provider (CSP)[52].
- *Broad Network access[35]:* Services can be accessed and provisioned from any area as long as the CSC has an endpoint device with internet or private network connection to the CSP. The connection must use standard protocols and ports such as Transmision Control Protocol (TCP), User Datagram Protocol (UDP), or Hypertext Transfer Protocol Secure (HTTPS)[52].
- *Resource pooling[35]:* The CSP, offers computing as a pool of resources using a logical separation between CSCs denominated multi-tenancy, i.e., two or more CSCs share virtual computing resources such as storage, computing capacity, and bandwidth[52].
- *Rapid elasticity[35]:* CSCs have the possibility to self-provision or cancel cloud computing services in real time or within an acceptable deferment[52].
- *Measured service[35]:* CSPs measure the CSC's utilization of computing resources for billing or analytic purposes, providing the consumers with a full overview of their consumption[52].

Pena Lopez, in[44], states that services in cloud computing not necessary have to comply with all the above characteristics. This statement can be argued according to the further clarification from NIST in[52], that defines *"essential"* as a requirement CSPs need to supply for their services to qualify as cloud computing solutions, giving CSCs the option to choose within the features, and utilize those that best satisfy their demands.

Building forward on NIST definition, Weinman[58] analyzes cloud computing from an economic perspective and extracts the essence of its characteristics using a simple and easy to remember mnemonic: C.L.O.U.D. *"Common, Location-independent, Online, Utility, on-Demand service[58]."*

Further, Weinman[58] uses the five cloud attributes summarized in the mnemonic C.L.O.U.D. to present a compelling argument against some common misinterpretations of the Cloud Computing model that, being partially valid, void the real context and value of Cloud Computing services. He argues that Cloud Computing might be based on, but is not equivalent to older technologies, concepts and IT business models, such as virtualization, data center colocation, ASP, Internet, and others, i.e., these older concepts in their own do not satisfy the five features of Cloud Computing.

The definitions from NIST in[35, 52] and Weinman in[58] seem to be clear and objective, and develop a framework that could provide a guide to CSCs, potential consumers and IT professionals to perceive the advantages of Cloud Computing and identify potential use cases. Some examples of misconception about the meaning of cloud computing are the situations presented below, which are based on own experiences.

The first one was a discussion between IT professionals; Some of them claimed that cloud computing has existed for a long time, just under other names such as ASP or IT outsourcing. Doing a brief evaluation of these IT service models, it would seem that they comply with two of the five cloud features: they could probably be accessed from anywhere, and the provider could use virtualization or multi-tenancy. On the other hand, on-demand self-service, rapid elasticity, and measured service are not typical characteristics of ASP or IT outsourcing; these type of providers usually demand fixed contracts for several years, and the deployment time typically is from weeks to several months.

The other one, was an occasion when someone asked an IT infrastructure team to "think cloud," without a further explanation of the meaning or purpose of such assignment. "Think cloud" could be many things, such as connecting an enterprise to the cloud, securing assets in the cloud, integrating on-premises data centers with public clouds, and using cloud services to support the business. Subsequently, during the evaluation of this diffuse assignment, another perplexing question was asked: "Are we cloud ready?", to which the answer was "sure we are, we have internet." It can be concluded that this was a brilliant answer to a question without context, and therefore supports the importance of a clear framework that helps to assess and understand the value that a powerful model such as Cloud Computing could provide to an organization.

3.2.1 Virtualization

Before moving on to the pure cloud computing concepts, it is essential to briefly introduce the concepts of Virtualization, Virtual Machine (VM), and Hypervisors. As mentioned before, the purpose of Cloud is to provide computing as a pool of resources, most of them shared among several customers. Virtualization is the enabler of that purpose, allowing multiple users to run different instances or VMs in a single physical device[29].

Virtual Machines running on a server, are entirely isolated from each other, and they have separate Operating System (OS) and applications, but they do not control the physical resources. Something called *"the hypervisor"* manages the hardware. Hypervisors or Virtual Machines Monitor (VMM) are an extra layer between the Virtual Machines and the physical server, and have several functions:

1. Administrate the physical resources, virtualizing and assigning them to the Virtual Machines. i.e., each virtual machine gets an Internet Protocol (IP), a virtual Media Access Control (MAC) address, virtual Network Interface Cards (NICs), and a portion of Central Processing Unit (CPU), storage and Random Access Memory (RAM), similar to a physical instance.
2. Create or delete Virtual Machines.
3. Forwards internal and external traffic, i.e., Internal traffic between VMs and external from/to VMs to outside.
4. Isolate Virtual Machines from each other.

Salam in[46] summarizes the definition and functions of hypervisors as, *"Hypervisors are the software, firmware or hardware that manage the complete life cycle of a Virtual Machine, including creating, monitoring usage and deletion."*

Hypervisors and virtualization, provide several advantages including:

- Optimization of physical resources as they are shared among several Virtual Machines
- Secure logical segregation between VM
- Virtual Machines are like a set of data, meaning they are easy to create, back up, destroy, duplicate or migrate.

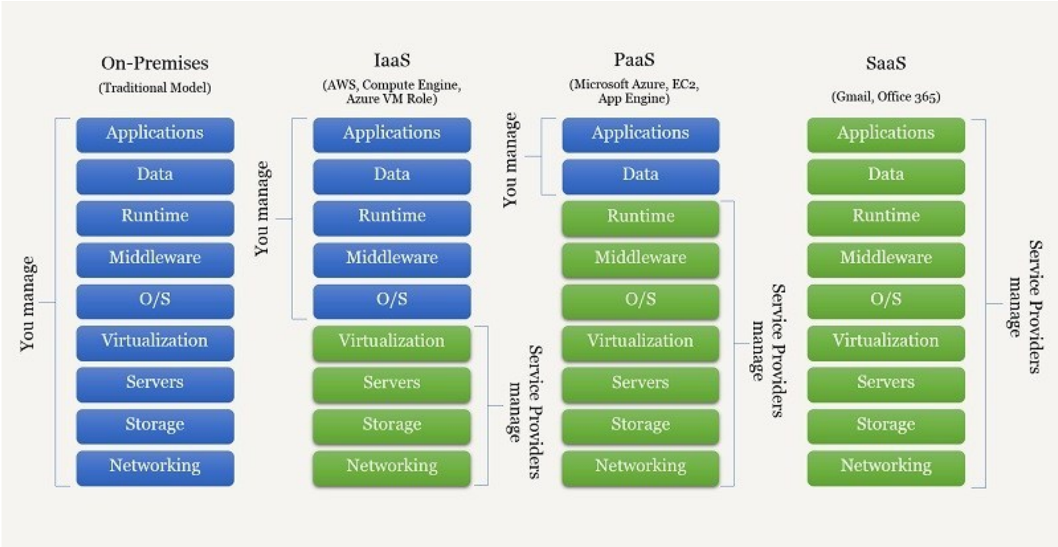

Figure 4: Cloud Computing Service Models taken from [11]

3.2.2 Cloud Computing Service Models

Depending on the level of control and responsibilities CSCs have over the computing resources, Cloud Computing services are classified into three groups: Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS)[29]. Figure 4 shows a comparison between the different Cloud Computing service models and the on-premises approach. In the traditional IT model, the user is also responsible for running and maintaining the Data Center's facilities.

3.2.2.1 Infrastructure as a Service (IaaS) In the IaaS service model, the CSP uses a hypervisor to provision VMs. The customer then requires dedicated operating system and can run any application on top of it. The service is billed as a utility, meaning it can be terminated or shut down on demand[47]. In IaaS customers manage the server and applications running on it while the CSP manage the underlying infrastructure as ilustrated in figure 4.

An example of IaaS is the Elastic Compute Cloud (EC2) from AWS. EC2 provides Virtual Machines in the Cloud which can be provisioned via AWS console within minutes[56]. Customers can choose between four pricing models:

- **On-demand:** Billed by seconds, the customer is only charged when the VM is up and running. Customers can turn off or delete the VM to stop the billing[10].
- **Reserved:** Customers sign a contract for one or three years to reserve certain

capacity[10].

- **Spot:** Customers bid a price and when the spot price hits the bid; the computational resources are assigned. This works for applications that are flexible or that do not require permanent uptime[10].
- **Dedicated Hosts:** Customers get a dedicated physical server billed by hours. This type of service is used for applications that do not support multi-tenancy virtualization, i.e., licenses that require a physical server that does not share resources with other virtual instances[10].

Another example of IaaS is the storage service, such as Simple Storage Servers (S3) from AWS. In S3, customers can upload an unlimited number of files that can be up to 5 Terabyte (TB). The files or *"objects"* are stored in folders called *"buckets"* where names must be unique as they can be accessed globally[56].

AWS provides six S3 classes depending on how often the data needs to be accessed and how fast the customers require the data to be available. Customers are billed depending on the class of storage, storage requests, data transfer, transfer acceleration, and cross-region replication *ibid.* Transfer acceleration is a service that allows the traffic to use AWS backbone instead of the internet for file transfer from the S3 bucket to the end user location *ibid.*

3.2.2.2 PaaS Platform as a service is an environment where customers can develop applications based on available programming languages and features. In these service models, customers only manage their applications and their data, while the service provider is responsible for the operation and maintenance of the underlying infrastructure, including the operating system.

An example of a platform as a service is `www.godaddy.com`. This platform is mainly a hosting provider that offers services from 25 Norwegian Kroner a month with plans for private users and enterprises[8]. Customers can design their websites using developing tools available in the platform such as Hypertext Preprocessor (PHP), Python, Cloud Linux, or they can use a Content Management System (CMS) program, for instance, Word Press and Drupal *ibid.* Billing is based on several factors such as the number of websites, computational resources, bandwidth, storage capacity, database requirement, and Secure Sockets Layer (SSL) certificates *ibid.*

Godaddy has availability of 99.9% and all their plans include 24/7 monitoring, Distributed Denial of Service (DDoS) prevention, more than 125 developing tools, public domains, database storage and scaling on demand; meaning customers only need to design their websites while the provider takes care of the operation and maintenance of the underlying services *ibid.*

3.2.2.3 SaaS In the software as a service class, the provider delivers an application that customers can use without having any responsibility for the underlying infrastructure. NIST[35] defines SaaS as:

> *"The capability that is provided to the consumer is to use the 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 (for instance, web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure. This infrastructure includes network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings."*

Some examples of software as a service are:

- Gmail is an Email service in the cloud
- SalesForces is a Customer Relationship Management (CRM) application in the cloud that provides employees overview and collaboration tools to handle customer relationships[2].
- Dropbox is data storage in the cloud and provides solutions for private customers and enterprises[3]
- Office 365 provides the traditional Microsoft office suite [1] in the cloud, allowing customers to access and edit documents from different end user devices. Additionally, customers have access to applications and collaboration tools such as Teams, Skype, Planner, Sway and many others[9].

3.2.3 Cloud Computing Deployments

Cloud Computing opens a world of possibilities when it comes to the way companies acquire, operate and maintain their Information Technology (IT) assets. For instance, Cloud could be an enabler for emerging businesses or greenfield deployments, as they can purchase computing on demand avoiding the start-up expenses of on-premises infrastructure and growing according to the business needs. However, for existing companies

[1] Word, Excel, Visio, Powerpoint, and others

with large data centers and owned computing assets, the journey to the cloud might not be that straightforward as it involves a fair amount of resources dedicated to research, design, planning, and development of strategies that allow a transition with minimum service disruptions.

Migrating an organization to the cloud might take several years depending on factors such as the size of the organization, the core business, the level of standardization, compliance and the motivation to change[29]. Despite these limitations, companies can benefit from cloud computing as they integrate their existing assets with public clouds *(ibid)*. This type of deployment is a hybrid environment and for many enterprises is the best path to cloud implementation. This section explores three types of cloud computing deployments defined by NIST in[35]: Public, hybrid and private clouds.

3.2.3.1 Public Cloud Public Cloud is the most common cloud computing deployment, and its main characteristic is that all resources are operated and maintained by the CSP[46]. Moreover, in a public cloud environment customers have the ability to self-provision their IT assets and scale the computational resources on demand, reducing their upfront investments and operational expenses *(ibid)*. However, customers do not have full visibility of the underlying infrastructure; this can be an issue for companies that require multilevel logs to perform IT forensics[58].

In order to comply with automation requirements and optimize the use of computing assets, CSPs rely heavily on multitenancy and virtualization, using hypervisors to administrate computing resources shared among multiple users. On the other hand, the use of shared environments requires a certain level of standardization[58], making public clouds less practical for highly customized applications.

The principal mean of connectivity to public clouds is the internet[29], and customers can choose whether to access their resources directly or to establish Virtual Private Networks (VPNs) that help them to assure confidentiality, origin, and integrity of their data. Additionally, CSPs provide support for dedicated connections through services such as direct connect by AWS or ExpressRoute by Azure[1, 5].

Some of the uses cases of public clouds include:

- Applications hosting, storage and backups[46].
- Disaster recovery sites [18].
- Application and web development using SaaS[46].

- Providers that use cloud computing to deliver IT services[46]

According to Weinman in[58], some of the services that might not be suitable or profitable for Cloud Computing are:

- Applications or services with a constant consumption of computational resources
- Highly customized solutions
- Applications with real-time or low latency requirements
- Transfer and storage of data that according to regulations or internal policies requires special handling
- Migration of large legacy code and services

3.2.3.2 Private Cloud Private cloud is a cloud computing environment that uses dedicated IT infrastructure for users in a single organization. Some organizations own, maintain and operate the assets themselves either on-premises or in co-location facilities, while others outsource the operation and maintenance to third parties[29].

The purpose of private clouds is to provide IT as a pool of resources for internal use, relying on automation, programmability and orchestration tools that help to optimize the way IT teams maintain and operate the underlying infrastructure[50]. In order to provide these capabilities, Data Centers are based on virtualization, hypervisors, and SDN networks: virtualization to allow resource sharing, hypervisors to administrate and manage the virtual resources assigned to virtual instances and SDN networks as a mean of communication and orchestration enabler[46].

In some cases, companies that require IaaS to host their applications and code; public clouds are not necessarily the most cost-effective alternative, especially for companies with existing Data Centers and large IT infrastructure[58].

Some typical use cases of private clouds include:

- Banking companies with existing data centers that require full control over the IT infrastructure for security, compliance, and forensics[46]. However, public clouds provide a high level of security, flexibility, and elasticity, making them suitable for hosting bank services, e.g., Capital One Bank cloud strategy seeks to migrate all their assets to the public cloud, including those that are mission critical[19].
- Health institutions with regulatory restrictions that prevent them from moving records to assets shared with other organizations[46].

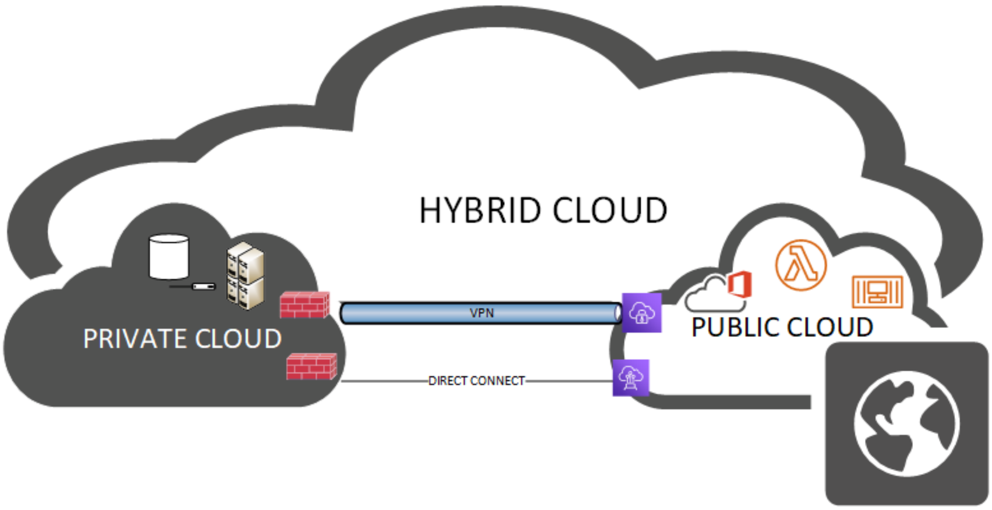

Figure 5: Hybrid Cloud Model

- Military organizations *(ibid)*.
- *"Government Institutions that require massive pools of dedicated computing (ibid)."*
- Development of applications and web services hosted in dedicated servers *(ibid)*. e.g., applications with low latency requirements or customized solutions[58].
- Services with flat computational resources consumption[58]

3.2.3.3 Hybrid Cloud A hybrid cloud is an approach that benefits from both public and private clouds. In this model, on-premises infrastructure is integrated with public clouds to provide scalability on demand or to access services and applications hosted in the public cloud[50]. The figure 5 illustrates the intercloud connectivity that the hybrid model seeks to achieve.

The Hybrid deployment is not only an alternative for a less disruptive migration to public clouds, but could also be the preferred solution for organizations with policy or regulatory restrictions that prevent them for moving all their assets to public clouds.

According to the enterprise cloud index, performed by Nutanix in 2018[40], 91% of the companies surveyed agreed on the Hybrid Cloud model deployment being the current ideal model. The principal reason to choose this type of implementation is probably the feasibility and simplicity hybrid clouds provide to move IT resources between clouds without significant business impact, and locating applications and services in the best suitable environment according to internal criteria, e.g., policy, regulations, cost optimization, functionality and others[40].

According to Salam in[46], hybrid cloud is the ideal model for the following types of organizations:

- *"Small and Medium-sized Enterprises (SME)s having multiple IT resource requirements and using a public cloud to balance traffic/computational load from private to public cloud."*
- *"IT service providers using a combination of private and public cloud models to service their clients."*
- *"Application/software developers using a public cloud for testing, integration, and deployment."*

However, I believe the hybrid cloud is an ideal deployment -at least as an interim solution- for almost every organization that own its IT infrastructure, as it provides the flexibility and orchestration capabilities to benefit from both private and public clouds, allowing a soft migration to public cloud or a multi-cloud interoperability that is transparent to the end users.

3.2.4 Advantages of Cloud Computing and their value to organizations

As discussed above, cloud computing is based on two solid business models: IT outsourcing and utility services, and relies on advanced technologies such as virtualization, programmability, high bandwidth networks, cryptology, and analytics, to provide customers multiple advantages and possibilities. Some of these advantages will be discussed in this section.

AWS in[56] extracts six of the most important advantages of cloud computing:

"Trade capital expense for variable expense" (ibid): This could be an advantage depending on the company. For example, a company that does not have enough cash flow might be interested in moving Capital Expenditures (CapEx) into Operational Expenditures (OpEx) to make projects viable as they reduce the required initial investment[59]. On the other hand, organizations generating enough cash flow from existing assets might prefer to have a higher CapEx and lower OpEx to make the business more attractive as they increase revenues[59].

"Benefit from massive economies of scale" (op. cit.): Cloud Service Providers need to have a vast infrastructure available to provide services that satisfy the requirements of cloud computing, which means CSPs purchase a large amount of equipment, having the opportunity to press down the prices. This is defined by Porter in[45] as *"buyers power. Powerful customers-the flip side of powerful suppliers- can capture more value by forcing down prices, demanding better quality or more service."* Conversely, companies whose primary business is not to provide IT services, do not have the same influence on the suppliers, increasing the procurement cost of their IT assets.

"Stop guessing capacity" (op. cit.): During the design phase of the IT physical infrastructure, engineers must dimension the data center capacity to host and provide services required by the organization during the life cycle of the assets. This engineering process, usually results in over- or under-dimension of the infrastructure, incurring unnecessary expenditures.

On the other side, using Cloud Computing, companies can increase or reduce capacity[2] *(op. cit.)* according to their needs, without paying for unused resources or running out of capacity. This advantage is especially beneficial in those occasions when organizations need to handle peaks of traffic during certain periods[59], e.g., during routine backups, or while executing irregular tasks such as analytics or launching new internet products.

"Increase speed and agility" (op. cit.): The automation and orchestration capabilities of Cloud Computing, allow CSCs to self-provision computing resources within minutes[29], on the contrary, it could take weeks, and even months to provision IT services in a traditional infrastructure *(ibid).* E.g., Applications can be developed and provisioned in the Cloud through simple steps within minutes, while the manual process of provisioning them on-premises could take significantly longer time as it requires highly human intervention including:

- Project leads to make the liaison between applications and infrastructure team.
- Network engineers to look at the application and place it somewhere in the network, configuring box by box using a CLI.
- Server experts to place the application on a server according to the computing requirements.
- Administrative personnel to grant access, create roles and define rights.

[2]Computing, storage and bandwidth

"Stop spending money running and maintaining Data Centers" (op. cit.): On-premises Data Centers, demand high costs associated with operations and maintenance. The equipment requires physical space, temperature control, power, physical security, software upgrades, and security patches. Depending on the business, Operation Centers might also be required to perform 24/7 monitoring, and network and server personnel should be available on call.

Moving assets to the cloud will transfer the maintenance and operations responsibilities and costs to the CSPs.

"Go global in minutes" (op. cit.): Cloud Computing infrastructure is spread around multiple geographical locations and is built on technologies that allow CSPs to offer services such as Content Delivery Networks (CDNs) or transfer acceleration, providing users low latency access to IT resources from any location on the world. This pervasiveness is very difficult to achieve using on-premises infrastructure.

In addition to the benefits highlighted by AWS and discussed above, Weinman in[58] identifies other characteristics of cloud computing that add value to organizations:

"Access to Competencies" (op. cit.): Through Cloud Computing, customers not only have the benefits from IT infrastructure expertise, but also from many other areas through the SaaS services, viz: Applications *op. cit.*, i.e. companies can access applications developed and maintained by highly qualified professionals, instead of dedicating own resources to in-house development and support.

Availability (op. cit.): Cloud Computing services provide availability up to 99.99% and durability of 99.999999999%[56]. For an enterprise, this availability translates in the possibility of using Cloud Computing as a disaster recovery site or having their IT assets replicated in several cloud locations[58].

Comparative Advantage and Core versus Context (op. cit.): Moving assets to the cloud, allow companies to focus on their core business, leaving the operations and maintenance of IT assets to those that have the best resources in that field. Moore[38] cited by Weinman in[58], holds the view that *"companies focus on core activities -those that strategically differentiate them from the competition- and leave context activities -even ones that are perfectly capable of doing- to others."* According to this statement, it is reasonable to assume that companies whose primary business is not related to IT, should consider acquiring IT services from external providers, e.g., an oil company do not need a large team of engineers to develop and maintain standard applications such as email servers, instead those can be purchased from CSPs as a service.

Customer and User Experience and Loyalty (op. cit.): : The capacity on demand and the quality of applications delivered through Cloud Computing, might improve the end user experience *(op. cit.)*, i.e., Cloud Computing services can scale on demand, providing higher capacity during the peak periods. Additionally, many applications delivered via Cloud, provide high quality, user-friendly and less problematic interface, as the developer companies have teams dedicated to operate and enhance those applications, or as it was described above: the application is their core business.

Employee Satisfaction (op. cit.): According to Weinman[58] *"Cloud services can enhance autonomy."* It can be deduced that employees with access to cloud computing services have the opportunity to create and explore different alternatives and applications, that would not be readily available on-premises. For example, companies using Office 365, not only have access to the traditional office packages such as Word and Excel, but also to an extensive suite of applications and collaboration tools, for instance, Teams, Planner, and Skype.

Community and Sustainability (op. cit.): As reported by Lavalle[32] *"data centers consume up to 3 percent of all global electricity production while producing 200 million metric tons of carbon dioxide".* Cloud Computing could help to reduce the carbon footprint by, first providing shared computing resources and second equipping their data centers with greener power and cooling alternatives[58].

Competitive Vitality and Survival (op. cit.): Cloud Computing services, can help customers to be on top of technology by giving them the freedom to expand without IT limitations and providing them with access to a broad portfolio of services and applications. Additionally, companies have the opportunity to focus on their core activities while CSPs deliver, operate and maintain their IT assets. These possibilities *"can help improve the chances of survival"*[58].

3.3 Data Center Network Technologies

Networking in Data Centers comprises all devices and software instances responsible for traffic forwarding using switching at layer 2[3] and routing at layer 3[4].

In addition to data forwarding, there are also other types of network devices that provide additional services at layer 4 to 7, including load balancers, firewalls, and Wide Area Network (WAN) optimizers [28].

Load Balancers: The principal function of load balancing is to distribute the traffic among multiple servers with the same characteristics. This technique improves the availability and scalability of applications and services [51].

Firewalls: are security devices that control the traffic that is allowed between hosts in a network. Firewalls can inspect the traffic and allow or deny protocol ports based on a security policy previously defined by the administrator. These devices can also examine applications data to determine if they perform as expected; dropping the traffic or activating alarms in case of suspicious behavior.

WAN Optimizers: perform Quality of Service (QoS), data compression and shaping to allocate WAN resources efficiently [28].

3.3.1 Functions of Network Devices

Figure 6 illustrates the principal building blocks of network devices: *"The management plane, the control plane, and the data plane[4]."*

The Management Plane is the interface network administrators use to manage and configure network devices. The most used interface in legacy networks is the CLI accessed via protocols such as Telnet and Secure Shell (SSH). However, there are network devices that also support configuration using Graphical User Interfaces (GUIs), Simple Network Management Protocol (SNMP), and Extensible Markup Language (XML) among others[4].

The control plane is the suite of layer two and layer three protocols that network devices use to communicate with each other. Some examples of control plane protocols are the Spanning Tree Protocol (STP) used to prevent loops in the network, the LLDP to discover neighbor devices, and routing protocols such as Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF)[4].

[3] Forwarding of traffic between endpoints in the same network or broadcast domain. Is done using MAC Addresses and Virtual LAN (VLAN)s [28]

[4] Forwarding of traffic between different broadcast domains or VLANs.

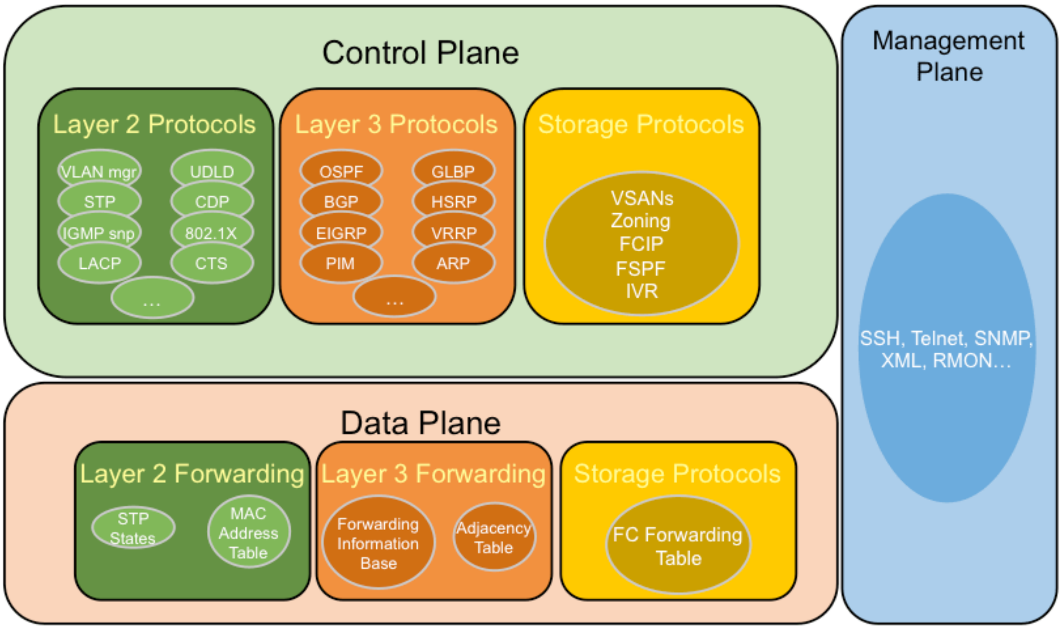

Figure 6: Functions of Network Devices, taken from [4]

The Data Plane is the way data is forwarded across the network. This process is usually done using IP and MAC addresses tables[4].

3.3.2 Legacy Data Center Networks

Legacy Data Center Networks, are deployed in a three-tier topology using the STP to provide loop prevention in High Availability (HA) configurations and VLANs to provide logical segregation.

The Spanning Tree Protocol, standarized in IEEE 802.1D-2004, has as a primary function to block links in redundant connections based on their path cost[5] to avoid broadcast storms[6]. STP then monitors the active links, and if one or more go down, it recalculates the path costs and brings up the relevant connections[12].

This type of Data Center deployment uses logical separation based on VLANs. VLANs are standardized in IEEE 802.1q, and are logical broadcast domains that allow segregation of end hosts, i.e., using VLANs, Network Administrators can group end hosts depending on different criteria such as function, location or security.

Figure 13 depicts a legacy data center topology. The connections marked with the red

[5]The path cost is calculated based on the bandwidth and the configured priority. The link with the highest path cost is blocked

[6]a broadcast storm is a loop that can reduce the throughput and even take down the whole network[31]

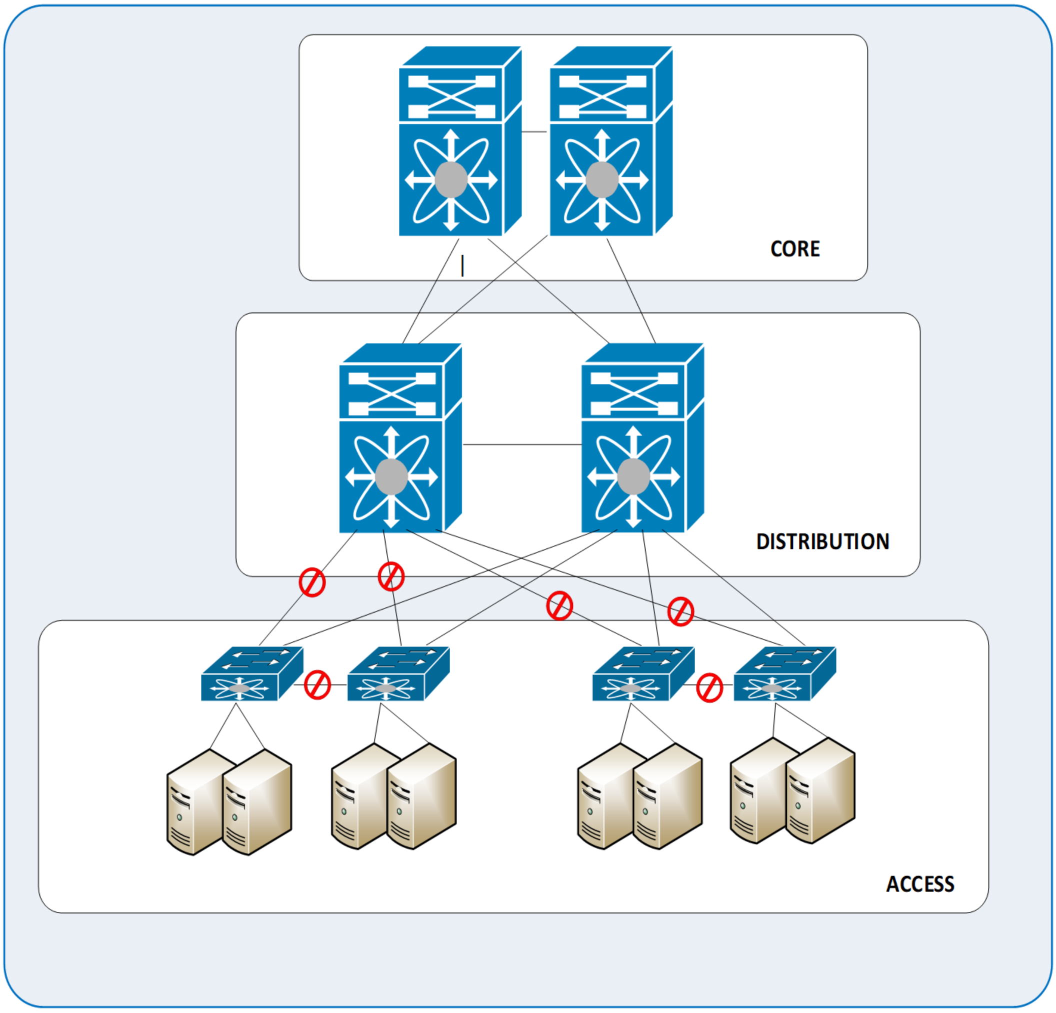

Figure 7: 3-tier Legacy Network Topology

signs illustrate the link blockage performed by STP. In this scenario, the **access tier** provides Layer 2 connectivity to physical servers and internal VLAN communication i.e., *"forwarding is based on a MAC address-based lookup[31]."* The **distribution tier** provides Layer 2 aggregation for the access switches and inter VLAN communication based on IP[7]. The **Core tier** is the perimeter device and provides the interface with external data centers, outside devices and users, i.e., the core is the interface to the outside world and is responsible for the user-server forwarding which is also called north-south or vertical traffic and in legacy networks is the predominant traffic pattern.

However, according to [31] the use of STP and VLAN segregation in large Data Center networks present several issues and limitations including:

- **Convergence:** If one or more links go down, the STP can take several seconds

[7] Layer 3 communication

to re-converge and recalculate the new path costs[12], resulting in significant traffic loss especially in high bandwidth links typical of Data Centers, *viz.*: 10G, 40G, and 100G.

- **Unused Links and Lack of dual-homing support:** STP blocks some of the links to avoid loops. This feature prevents the devices from doing multihoming[8] and results in underutilization of the interfaces.
- **Suboptimal Forwarding:** STP builds a tree that is always utilized to forward the traffic. As a result, some of the frames may not use the best path to reach their destination. The desired situation is a continuous calculation of the best route and the option of having more that one active path, as it happens in layer three networks.
- **Loop Risk:** STP needs to be configured very carefully because any failure in the configuration might lead to a broadcast storm that may reduce the performance of the network considerably or even take it down. Layer 3 networks do not present this issue, as the packets have a counter called Time to Live (TTL) that decreases by each hop dropping the packet when its value gets to zero.
- **Scalability Limitations:** The IEEE 802.1q header supports maximum 4096 VLANs. This limit can be easily exceeded in a Data Center Deployment.

To overcome the limitations and issues of STP, technologies that allow multihoming better known as Multi-Chassis Link Aggregation (MC-LAG) were introduced [31]. vPC is an MC-LAG technology that allows dual-homing, making a pair of upstream switches look like a single logical device for the rest of the network. I.e., access switches are attached to two different upstream devices with both links in an active state adding together their forwarding and computing capacity.

Figure 8 shows the previous scenario improved with vPC. See figure 13 for the original topology. In the vPC deployment, the devices in the access-tier see the distribution switches as one logical device having two active connections to it, which means the bandwidth of the up-links is increased.

The implementation of vPC in Data Center Networks mitigates most of the constraints of STP, but it does not address the scalability limitations and the requirement to extend layer 2 deployments through large networks[31]. Additionally the three-tier topology and the use of VLANs converts the access and distribution tier in a bottleneck as all the

[8] Multihoming is a feature that allows a device to connect to two or more switches.

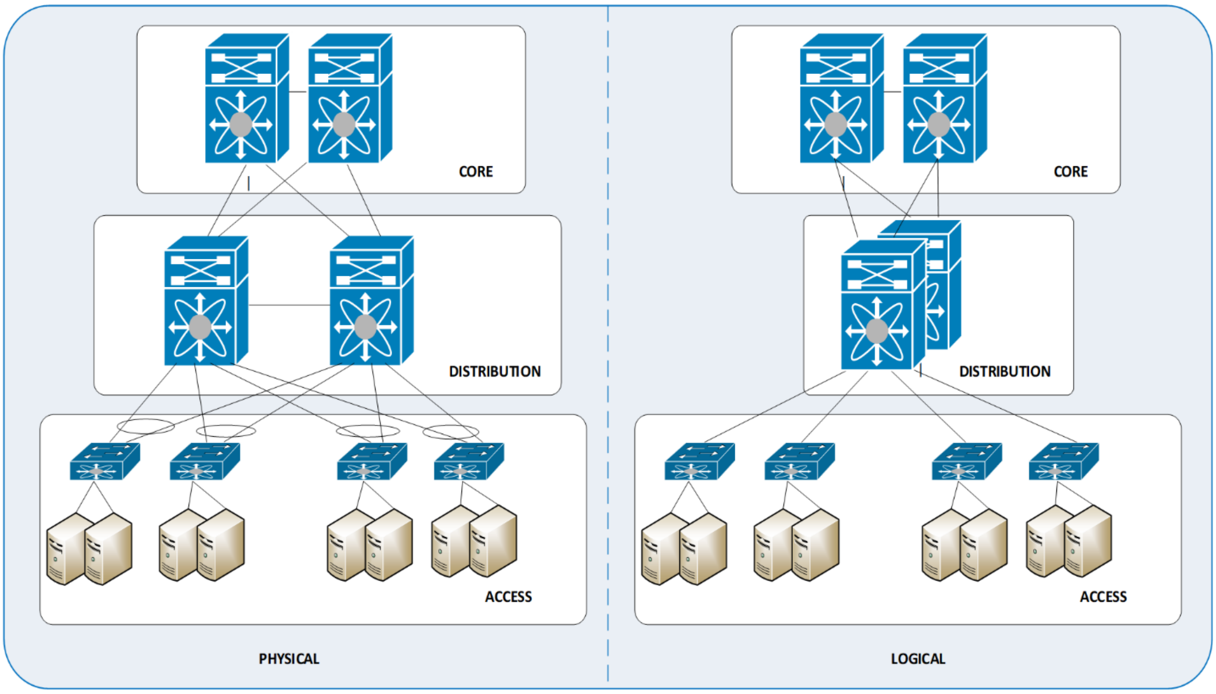

Figure 8: Legacy Data Center Topology improved with vPC

switches need to learn the MAC addresses of every endpoint connected to the network and there is a limitation of 4096 possible VLANs[23].

Moreover, with the evolution of virtualization and cloud computing, the horizontal traffic[9] has become predominant in the Data Center, and the philosophy of the network must adjust in order to support today's applications landscape.

Data Center Networks need to evolve to overcome the limitations of legacy data center technologies. Next-generation data center networking is an approach that seeks to improve the agility, scalability, and flexibility of the current model.

[9]Also called east-west traffic, is the communication between hosts in the Data Center

4 Next Generation Data Center Networks - State of the Art

With the increase of virtualization and the emergence of cloud computing, the traffic in Data Centers increased considerably, and the traffic patterns evolved from a being mainly vertical -from users to servers-, becoming mostly horizontal -between servers-. Data Center networks became then a limitation for the evolution of data center technologies.

In order to address data center requirements, a next-generation network is proposed. This model is based on a new topology and an overlay protocol on top of the traditional IP network. This section will go through the State of the Art of Data Center Networks describing their topology, overlay technologies, and Software-Defined Networking briefly.

4.1 Topology

The topology of legacy data center networks did not provide the scalability and elasticity requirements of today's data centers. The number of intermediate devices increased the diameter of the networks and the hops between endpoints, incrementing the latency and the power consumption [23].

To address these weaknesses, a two-tier topology in combination with VXLAN overlay networking is recommended; this approach is called Data Center Network Fabric. This topology, depicted in figure 9 defines two types of devices: the leaves or edge devices and the spines. The leaves are also called Network Virtualization Edge (NVE) devices because they host the VXLAN Tunnel Endpoints (VTEPs).

Layer 3, hosted in legacy networks by the distribution tier, is moved down to the leaves which are the access devices, and provides connectivity inside and outside the Data Center, while the spines' primary function is to forward the traffic between leaves[31]. This network topology not only improves the latency and reduces the power consumption, but also provides high scalability, resiliency, and efficiency *(ibid)*.

To increase the resiliency, the devices are configured in groups[10], and the recommendation is to connect each endpoint to a pair of leaves. If one of the devices goes down, perhaps a spine or a leaf, the bandwidth might be reduced, but the fabric will continue to operate and forward the traffic until the failure is remedied.

[10] usually pairs, but groups of more than two devices are also supported

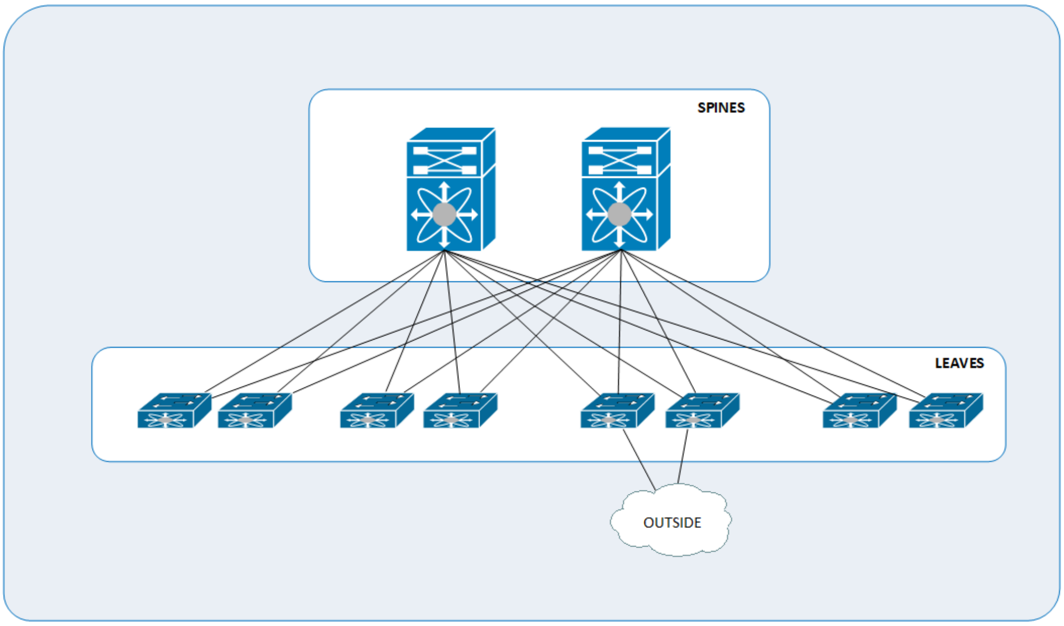

Figure 9: Next Generation Data Center Network Topology

Additional leaves or spines can be added to the network at any moment to upgrade the capacity of the fabric; this provides endless scalability that legacy networks did not support.

4.2 Network Virtualization Overlays

Overlay network technologies can be used to address some of the issues of legacy data center networks, *viz.: Scalability, Layer two connectivity, virtualization and growth of east-west traffic[31].*

Overlay Networking uses software to create virtual tunnels on top of a physical network, or to put it in another way; an overlay is for networking what the hypervisor is for servers. The data is encapsulated and tagged at the origin and then decapsulated at the destination point by the end network device or the host. Some of the overlay network technologies include VPN, Multiprotocol Label Switching (MPLS), Generic Routing Encapsulation (GRE) and VXLAN.

Overlay protocols have several features:

- *Identity and Location:* Refers to the host IP or MAC address and the edge device where the tunnel is terminated[31].
- *Overlay service:* The function of the overlay, type of origin and encapsulated traffic

ibid. E.g., IP/IP, MAC/IP, MAC/MAC.

- *Underlay transport Network:* The network used to transport the encapsulated traffic, e.g., IP network *ibid.*
- *Maximum Transmission Unit (MTU):* Overlay protocols add overhead to the packets or frames to include the identifications fields of the protocol. It is important to use the maximum MTU to allow the flow of overlay traffic without issues *ibid.*
- *End Host Mapping:* How the overlay protocol identifies the end host devices. This can be done using mechanisms such as Flood and Learn (F& L) or Border Gateway Protocol - Ethernet VPN (BGP EVPN) *ibid.*
- *Transport of Multidestination traffic:* How overlay protocols forward the traffic to multiple destinations using protocols such as IP multicast and ingress replication.

4.3 Virtual Extensible LAN (VXLAN)

VXLAN is an overlay protocol that encapsulates layer two frames over IP[11] packets and was introduced to address the requirements of large layer two network deployments, overcoming the 4096 possible VLANs' limitation with a network segment of 24-bits: *"16 million broadcast domains[31]."*

In Data Center Networks with spine-leaf topology, the leaves or NVE devices establish VXLAN tunnels to provide connectivity to the end-hosts. The spines are used as an IP underlay network to forward the traffic between leaves and standard routing protocols such as OSPF are used to manage the routing. The latency is significantly lower as there is only one hop between end devices.

Data Centers based on virtual networking embedded in the servers such as Network Virtualization and Security Software (NSX) or Open Virtual Network (OVN) terminate the VXLAN tunnels at the hypervisor, while other solutions such as Cisco ACI provides support for both VTEPs: At the hypervisor level and the physical NVE.

VXLAN forwards multidestination traffic using Protocol Independent Multicast (PIM). Each network receives a unique Virtual Network Identifier (VNI) that is tied to a multicast group[12] whose members are the VTEPs associated with the specific Network.

[11] MAC-in-IP/UDP

[12] Multicast is used to forward traffic to specific multiple destinations[57] preventing hosts in a network to use computational resources receiving packets they are not interested in. Multicast groups are identified with a Class D IP address. If a device is interested in listening to the messages in a multicast group, it

The endpoints can be identified using the F& L method which floods messages to all the hosts in the network creating a MAC Address Table with IP to MAC information. However, this mechanism might negatively impact the scalability as the data center become larger. This limitation is addressed with the introduction of BGP EVPN to handle the control plane while VXLAN is dedicated to the data plane. When a new endpoint joins the network, BGP EVPN distributes its MAC, and IP addresses to all the VTEPs, which preserve the information in a state table until they receive an update notifying that the host has moved to another VTEP or has left the network[31].

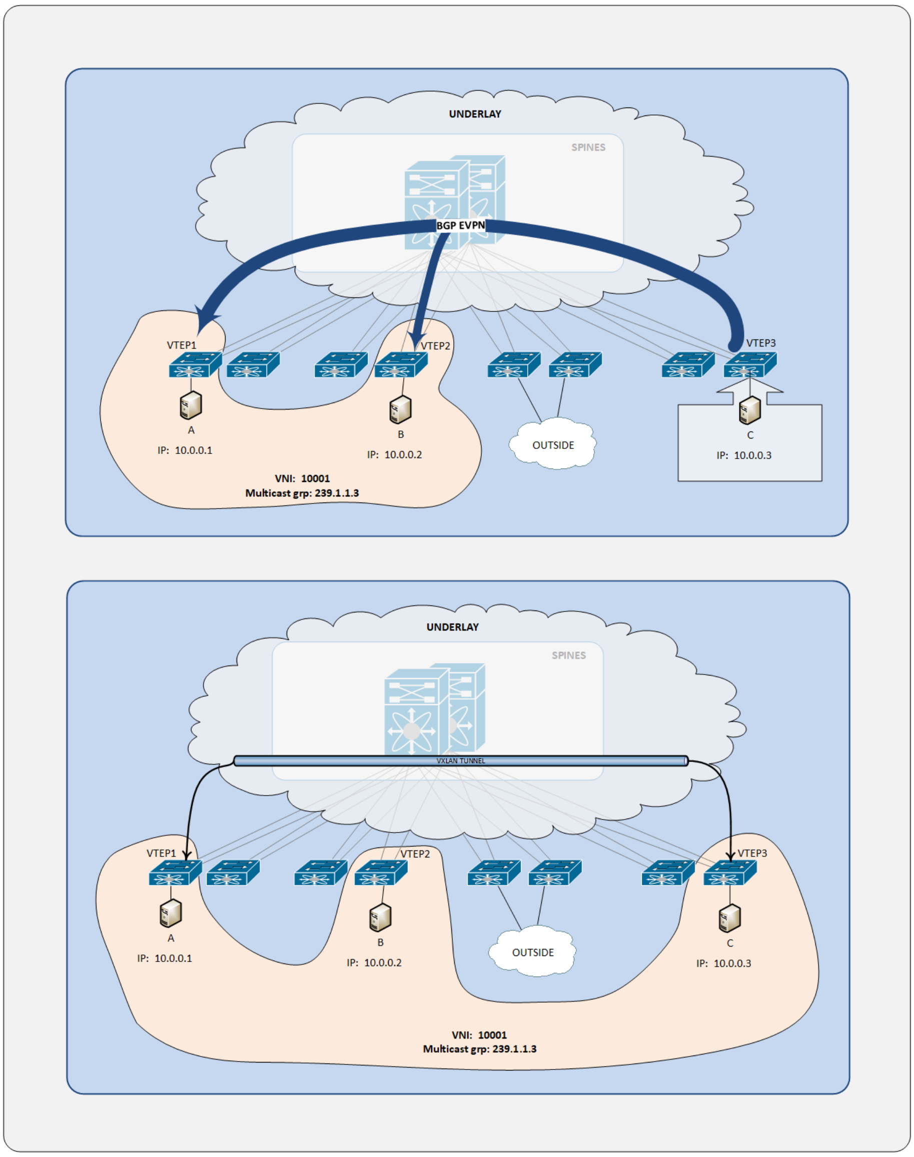

Figure 10: VXLAN communication in an ACI fabric

The communication between two devices in the network is achieved through VXLAN

has to join the group and listen to the traffic sent to the multicast address *ibid*.

tunnels terminated at the VTEPs contained by the leaves. A simplification of the process illustrated in the figure 10 is described below.

1. Host C joins the network, and the VTEP 3 sends the information to all other VTEPs using BGP EVPN. All VTEPs update their state tables with the new information. VTEP 3 gets assigned the VNI 10001 and will join the multicast group 239.1.1.3

2. Host A wants to talk to host C. VTEP 1 receives the message and establish a VXLAN tunnel to VTEP 3 that decapsulates the frame and forward it to host C. The VXLAN frame depicted in figure 11, is created by VTEP 1 by doing changes to the original frame: Removes the 802.1Q header from the original frame; adds the VNI and the original 802.1Q header to the VXLAN field, adds the UDP header containing layer 4 (protocol) destination port which is 4789 and source port which is calculated to provide information about the path; adds an outer IP header containing the VTEP source IP and the VTEP destination IP; adds an outer MAC header containing the VTEP source MAC and the VTEP destination MAC.

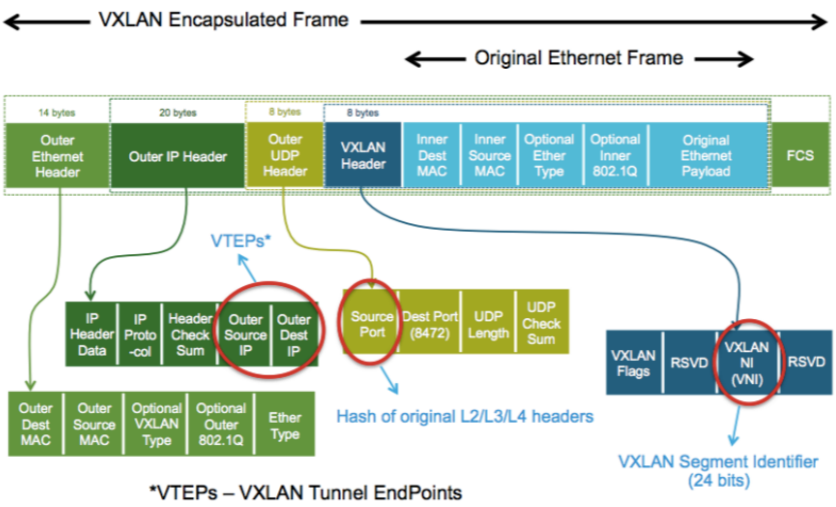

Figure 11: VXLAN Frame Format, taken from [31]

4.4 Software-Defined Networking (SDN)

SDN is a vendor-independent network architecture developed by the Open Networking Foundation (ONF) as an initiative to address the requirements of next-generation networks. SDN is defined by the ONF in [7] as:

> *"The physical separation of the network control plane from the forwarding plane, and where a control plane controls several devices."*

The management plane, control plane, and data plane were separated functions in legacy network equipment, but they were restricted to a specific device, i.e., network administrators had to log in to each device to manage, configure and troubleshoot network connections. The purpose of SDN is to extract the management and control plane from network devices and centralize these functions in a set of controllers that administrate all the devices in the network[28].

The SDN architecture is built on the three layers depicted in figure 12: the infrastructure layer, the control layer, and the application layer [7].

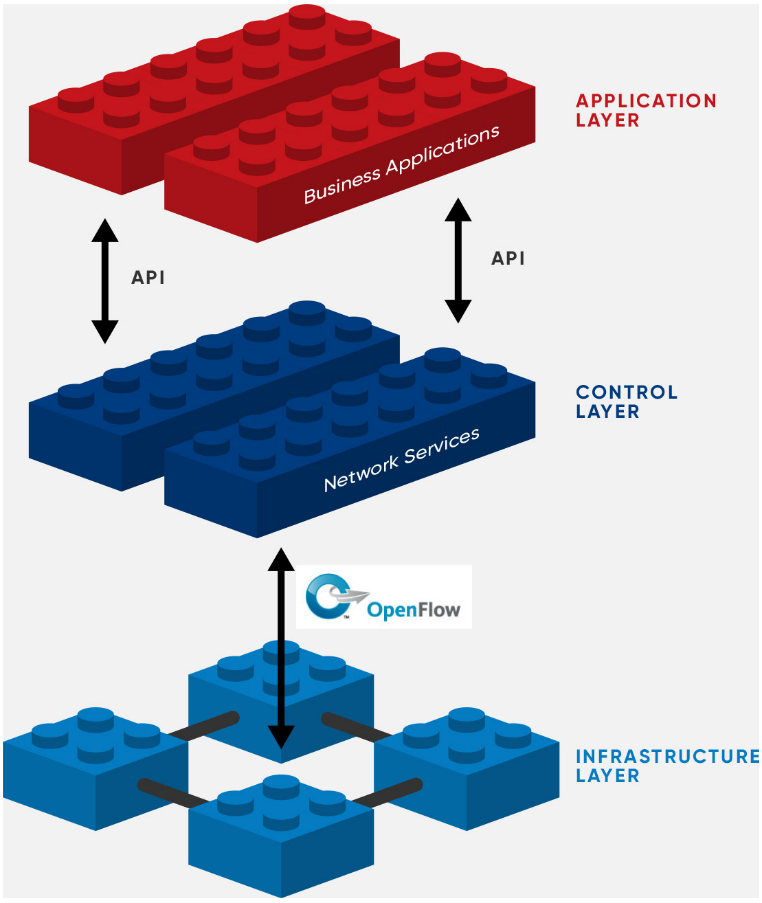

Figure 12: SDN Architecture, taken from [7]

The infrastructure layer consists of network devices such as switches and routers that have as a function to encapsulate and forward traffic based on a set of policies defined in the controller.

The control layer is the interface to the network [13] and performs the management and control plane functions, viz: *Configuration and forwarding state*. The controllers have an

[13]This interface is also called SDN Control Data Plane Interface (CDPI) or SDN Southbound interface

overview of the whole network and communicate with network devices using the OpenFlow protocol to provide information that allows switches and routers to update their flow table according to the network policies [41].

The use of centralized management simplifies the way the network is administrated and reduces the likelihood of human failure as the configuration is done and verified at the controller and then pushed to all the network devices. It also provides integration with virtual environments such as VMware or Red Hat. The advantages of this integration include:

- *Visibility increase:* the controller can identify virtual networking inside the hosts and provides an overview of all endpoints, including virtual machines and containers[26].
- *Functionality enhancement:* The network configuration and policies can be defined in the controller and propagated to the physical and virtual infrastructure, providing consistency across the environments and reducing the risk of misconfiguration[26].
- *Security improvement:* The devices can be segregated at the endpoint level, providing the possibility to group, isolate or limit the access to endpoints, even if they share a physical host or a broadcast domain. It also allows the use of Layer 4 to Layer 7 (L4-L7) devices, such as load balancers or firewalls[26].

The application layer includes programs and business applications that interact with SDN controllers to specify the desired network features and behavior [42]. These applications include GUIs embedded in the controller, scripts, third-party programs, and others.

The interface between the application and control layer [14] is based on API and provides programmability and automation capabilities, allowing scripting and rapid replication of configuration. Moreover, the programmability features of SDN simplify the migration from legacy networks, the implementation of new deployments, and allow the collaboration between development and operations to the point that developers can be granted the privileges to self-provision network resources.

[14] This interface is also called SDN North-Bound Interface (NBI)

5 ConocoPhillips Solution

ConocoPhillips is the world's largest independent exploration and production company, based on proved reserves and production of liquids and natural gas. Headquartered in Houston, Texas, ConocoPhillips has operations in 17 countries and approximately 11,400 employees [6].

ConocoPhillips is one of the largest foreign operators on the Norwegian continental shelf. Headquartered in Tananger outside Stavanger, the company has around 1850 employees. The company's core activities in Norway are petroleum exploration and production (ibid).

ConocoPhillips holds a strong position in big fields on the Norwegian continental shelf. The company is the operator of the fields in the Greater Ekofisk Area, which is the mainstay of the company's activities in Norway. The company has a 35.112 percent interest in the Ekofisk, Eldfisk and Embla fields and 30.658 percent in the Tor field (ibid).

The company also has assets in fields operated by co-venturers, including Heidrun, Visund, Oseberg, Troll, Grane and Alvheim (ibid).

ConocoPhillips -Phillips Petroleum at that time- found oil on Ekofisk in 1969 and started production 18 months after. In 1992 the company invested in a reliable and robust communication's network that could help to reduce the workforce and move several functions from offshore to onshore without any safety or production impact [54]. In 2013 ConocoPhillips installed Ekofisk 2/4L; an accommodation platform and communications hub for the Ekofisk complex [6].

All ConocoPhillips North Sea installations are connected to onshore using owned Fiber Optics (FO) cables and leased lines presentation; this facilitates the collaboration between offshore and onshore organizations and the performance of real-time operations.

ConocoPhillips is a technology company that recognizes the value of innovation developing and adopting digital solutions in different areas such as remote-controlled subsea installations; production and drilling optimization using real-time data; container logistic; heat recovery, and power generation management to reduce emissions; exploration and reservoir data; production control and monitoring; and telemedicine to support health services offshore.

In 2013, ConocoPhillips put together a multidisciplinary team to study cloud computing capabilities and challenges. This team included different BUs and departments and had the purpose of developing a cloud strategy. In 2014 the company started a cloud adoption, putting in place policies, procedures, and legal negotiations with CSPs. Later on, the application migration was slowed down due to the decline of the commodity prices, and ConocoPhillips decided to focus on the deployment of Office 365.

In 2017, ConocoPhillips established the goal to migrate 10% of Houston and Bartlesville virtual environments to public clouds. Further, the company formed a team to operate IaaS solutions with the purpose of moving workloads to the cloud and build the required expertise for future operations.

In 2018, ConocoPhillips announced the "all-in cloud philosophy" encouraging Data Architects to consider cloud computing as the preferred solution rather than on-premise deployments.

This section will analyze the current ConocoPhillips Norway data center network deployment, translate the company business requirements into technical specifications, and explore the alternatives to replace the existing solution.

5.1 Previous Design

ConocoPhillips data center network was based on legacy Data Center technologies built upon the Cisco NX-OS platform in a three-tier topology, illustrated in figure 13.

The core layer consisted of two Nexus 7000 Layer Three (L3) switches with ten slots; eight Input/Output (I/O) slots with a forwarding capacity of 550 Gigabits per second (Gbps) and two supervisors[15] [13]. The switches were configured in HA with a module arrangement illustrated in figure 14.

Table 1 shows a brief description of the chassis type and the modules installed in each Nexus 7000.

Table 1: Nexus 7000 Module Arrangement

QTY	Part Nr.	Function	# Ports	Capacity per port
1	N7K-C7010	10 Slot Chassis	NA	NA
2	N7K-SUP1	Supervisor Module	NA	NA
3	N7K-M132XP-12	Fiber Optic Module	32	10 Gbps
1	N7K-M132XP-15	Fiber Optic Module	32	1/10 Gbps
4	N7k-M148GT-11	Cobber Ethernet module	48	10/100/1000 Mbps

[15] For more information about Nexus C7000 ref. to the data sheet in [13]

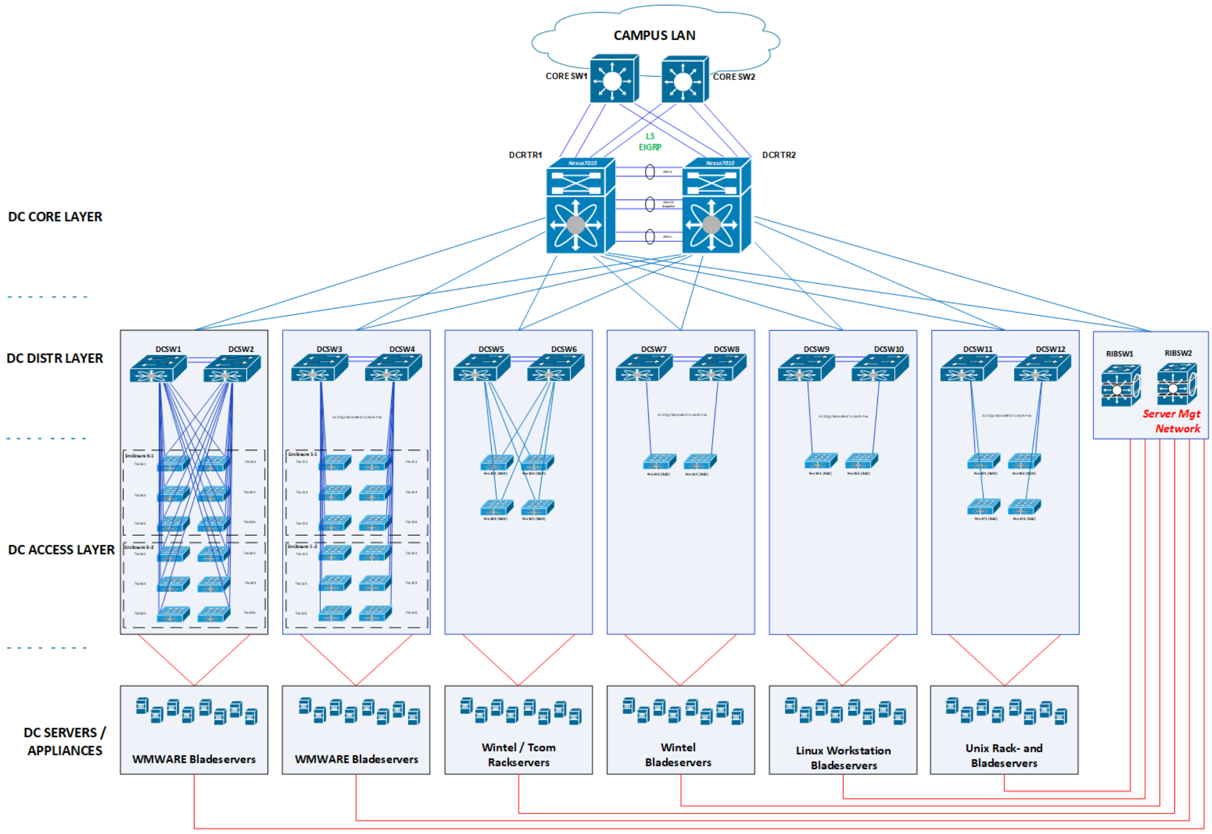

Figure 13: ConocoPhillips Original Data Center Network Topology 2018-2019

The core switches were installed in 2011 and were due to refresh at the end of 2017 when several modules were out of support; see table 2. i.e., the manufacturer no longer provided support and parts.

Table 2: *End of Life (EoL) dates for modules installed in core and distribution switches. Source: [15, 16]*

Part Number	End of Support
N7K-SUP1	August 31, 2019
N7K-M132XP-12	December 31, 2017
N7K-M132XP-15	August 31, 2019
N7k-M148GT-11	December 31, 2017
N5K-C5596UP	May 31, 2024

In the distribution layer, the COP's data center had 12 Nexus 5000 switches configured in HA using vPC, i.e., even if there are two physical switches, end hosts perceived them as one logical device. The distribution switches provided Layer 2 connectivity to the access layer with physical segregation for the different types of end devices hosted by the data center:

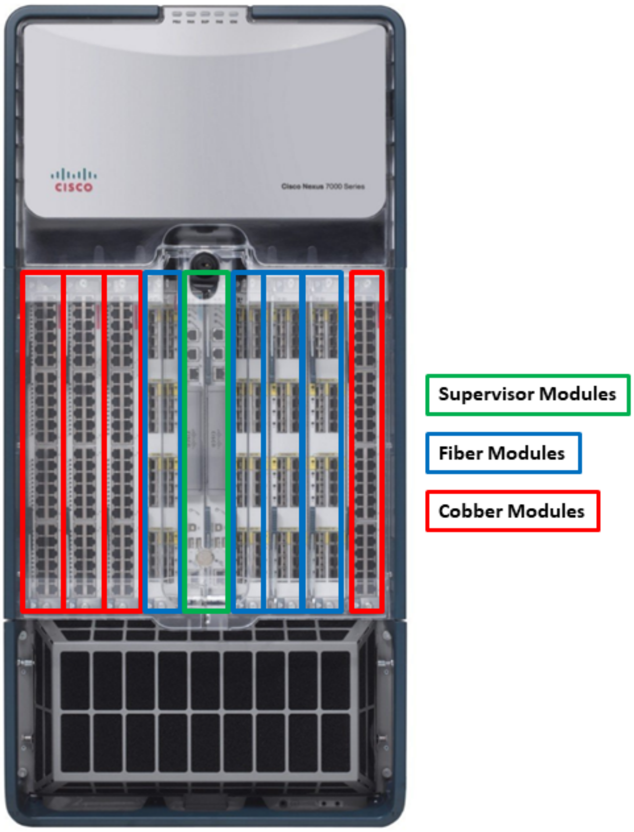

Figure 14: Slot arrangement of COP's Nexus 7000

- Two pairs of distribution switches dedicated to VMware blade servers with integrated Nexus 2000 (N2K) B22HP FEXes that acted as an access layer.
- One pair of distribution switches with two pairs N2K FEXes to connect Wintel Rack servers using copper Ethernet interfaces.
- One pair of distribution switches dedicated to Wintel blade servers.
- One pair of distribution switches to connect Linux Workstation Blade servers with integrated B22 Blade Fabric Extender in the access layer.
- One pair of distribution switches with two pairs N2K FEXes in the access layer to connect Unix Rack and Blade Servers.

The distribution switches were built on Nexus 5596 Chassis with 48 x 10Gbps ports. These switches were installed in 2014 and the last day of support is May 31, 2024.

The access layer relied mainly on N2K FEXes, which were preferable single homed. However, some of the environments hosting single connected devices required dual-homing in the access layer to provide resiliency to the end-points. Figure 15 depicts the FEX connections used in the Data Center.

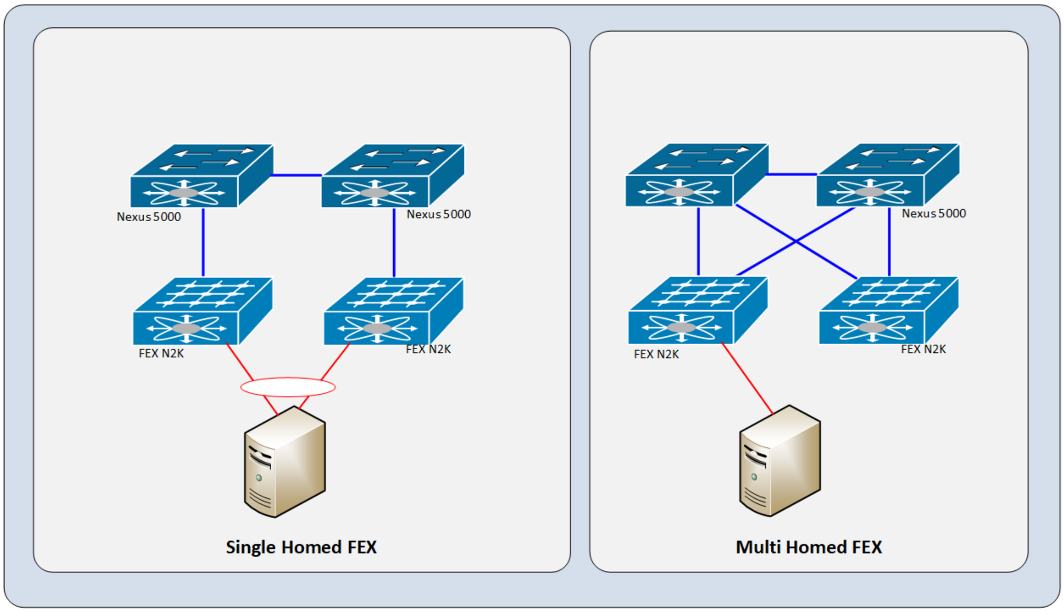

Figure 15: Single vs Multihomed FEX configuration

In brief, the legacy data center of ConocoPhillips Norway was a robust Cisco NX-OS deployment whose features included:

- High availability and redundancy based on vPC.
- Three tier topology with End of Row (EoR) distribution switches and both EoR and Top of Rack (ToR) access FEXes.
- The boundary between the Data Center and the campus network was a set of routed interfaces based on the dynamic routing protocol Enhanced Interior Gateway Routing Protocol (EIGRP).
- The management network was run in dedicated layer 3 switches routed to the campus via the Data Center Core.
- The deployment included both single homed and multihomed FEXes.
- The data center supported copper and FO interfaces with up to 10Gbps forwarding capacity.
- The design provided a dedicated distribution and access layer to specified server environments.

Some of the challenges the network administrators experienced during the lifetime of the deployment were:

- Most of the administrators were not aware of the network design philosophy; this made the network difficult to maintain restricting the administration to one or two senior engineers.
- The dedicated environments restricted the type of devices that could be connected to the switches; as a result, the Data Center was over-dimensioned with a large number of unused ports.
- Servers that did not belong to any of the defined environments had to be connected directly to the core or the distribution switches.
- With the boost of virtualization, the horizontal traffic increased and the server team required 40Gbps interfaces that could not be offered with the previous design.
- The Data Center did not support the extension of layer 2 networks across physical locations; this complicated server operations such as VMotion and storage mirroring.
- The network did not provide programmability and automation capabilities, making an integration with cloud environments difficult.

5.2 New Data Center Network

As described in the previous section, the core layer of COP's Data Center Network is built upon the Nexus C7010 Chassis, consisting of several modules that were EoL or were about to reach the EoL date, see table 2. It was time for COP to replace some of the network equipment; this represented an opportunity for the company to evaluate how the requirements had evolved since the last implementation and whether a redesign was necessary to address those requirements.

The project was framed using the Prepare, Plan, Design, Implement, Operate, and Optimize (PPDIOO) Lifecycle framework introduced by Cisco in [21] and illustrated in figure 16, with the purpose to design an agile and robust Data Center Network that aligned with business requirements and reduced the overall costs of the solution.

Figure 16: *Prepare, Plan, Design, Implement, Operate, and Optimize (PPDIOO) Lifecycle Framework. Source:[21]*

5.2.1 PREPARE PHASE: Business and Technical Requirements

The methodology used to assess the business requirements and technical specifications was based on action research using the methods and activities described below to collect the data:

- Identification of corporate goals and regional strategy.
- Interviews with server and applications team to gather information about their needs, future projects and desired features.
- Interviews with the network team in Norway to understand the benefits and challenges of the current design.
- Study of the current design using existing documentation, configuration and getting information from the installed devices.
- Virtual meetings with network teams in other BUs in ConocoPhillips to streamline the Data Center Network philosophy, get information about trends and analyze the support models.
- External consultancy from Cisco and Atea to understand the technical features and constraints.

The business requirements and technical specifications of the new Data Centers are summarized in table 3.

1. **Supported End-Points and Applications**

   The new data center network shall support all the existing server, network and storage infrastructure, including:

   - VMware
   - Wintel rack servers
   - Wintel blade servers
   - Linux rack and blade servers
   - Unix rack and blade servers
   - Netapp storage
   - N2K B22 FEXes

Table 3: Technical Specifications based on business requirements

End-Points and Applications	Required	Preferred
VMware Hypervisor	X	
Bare Metal Servers	X	
Netapp	X	
Firewalls	X	
Netscaler	X	
Capacity and Interfaces		
600Gbps fwd capacity	X	
2.8Tbps fwd capacity		X
100Gbps interfaces		X
40Gbps interfaces	X	
25Gbps interfaces	X	
10Gbps interfaces	X	
10/100/1000 Mbps copper ethernet	X	
Non-blocking		X
No oversubscription		X
Scalability in the core		X
Scalability in the access layer	X	
Multisite and Cloud Integration		
Programmability and Automation		X
Multisite support		X
L2 over L3		X
Security and Reslience		
High Availability	X	
Monitoring	X	
Traffic Filtering and QoS		X
End Point visibility		X
Access Control	X	
Support		
Easy provisioning		X
COP standard compliant	X	
NetFlow		X
SPAN ports		X

- Check Point, Cisco, and Palo Alto firewalls
- Netscaler load balancers

Moreover, the data center should be ready for integration with cloud applications.

2. **Traffic pattern and access interfaces**

Applications have changed, and many of them are not hosted by a single server; instead, they are distributed across several machines leading to a significant increase in the horizontal traffic in the data center. Moreover, new servers provide higher I/O throughput with interfaces up to 40 and 100 Gbps. The requirements regarding the traffic and access interfaces are:

- Increase bandwidth in the access and core layer
- Copper ethernet interfaces supporting 10/100/1000 Mbps
- FO interfaces supporting 10, 25, 40 and 100 Gbps
- Non-blocking architecture and no oversubscription
- Possibility to expand the capacity in the core and access layer, i.e., scalability.

3. **Multisite and Cloud Integration**

   The company has a hybrid cloud strategy, and the Data Center should support a future cloud integration with orchestration, automation and programmability capabilities. Additionally, the Data Center Network should support multisite configuration as the BU is planning a reengineering of the Disaster Recovery Site and several Data Centers offshore. The requirements are:

   - Automation and programmability capabilities to facilitate the orchestration with cloud computing.
   - Multisite architecture.
   - Possibility to extend broadcast domains across sites to allow an effortless failover between locations.

4. **Security improvements, resilience and Network services**

   The data center shall support the current security infrastructure and configuration. Moreover, it is desired to study the Data Center SotA that provides security enhancements. The requirements regarding security, resilience and Network Services are:

   - The availability of the Data Center should be equal or better than the current design.
   - Monitoring capabilities

- Support for traffic filtering and QoS
- End-Point visibility and improved security
- Access control based on existing solutions, *vis:* Tacacs and Radius[16]

5. **Management and Support**

   The Data Center hosts mission-critical servers and applications that run continuously and require high uptime. ConocoPhillips takes advantage of globalization and uses the FTS model to provide support in all time zones. To adhere to the FTS support model, it is necessary to design a network based on a global philosophy ensuring that qualified personnel is always available to assist in case of failures or outages. The requirements regarding the support model are described below:

   - Simplify the provisioning, management and troubleshooting processes, reducing the operating expenses.
   - Collaborate with other BUs to design a solution according to global standards
   - Support for external NetFlow logging
   - Support for Switched Port Analyzer (SPAN) ports.

[16] Tacacs+ and Radius are the most used protocols to control access to networking equipment. They provide Authentication (who has access), authorization (what the user is allowed to do) and accounting (logging) services.

5.2.2 PLAN PHASE

5.2.2.1 Design Alternatives

ALTERNATIVE 1: New Cisco Nexus Core Switches in NX-OS mode, same topology

This alternative requires the lowest up-front investment and shortest implementation time. However, it does not fulfill all the requirements and does not generate significant additional value for ConocoPhillips.

The purpose of this alternative is to replace the core switches with devices in the same family and gradually replace the distribution and access layer depending on the devices' EoL dates: Cisco Nexus 7000 or Nexus 9000. The setup of the new devices is done by replicating the configuration from the existing hardware, and the installation is a simple process that does not cause significant downtime. There is no changes in the existing topology, the support model or the network design.

The approach supports all the required end-points, applications and interfaces. Moreover, it provides an upgrade in the switching and forwarding capacity. However, it relies on a legacy topology that does not provide flexibility and scalability in the core layer, resulting eventually in oversubscription.

ALTERNATIVE 2: Cisco NX-OS as underlay with NSX SDN solution

This option uses the solution described in alternative 1, adding NSX in the VMware environments to provide an SDN implementation that facilitates integration with other sites, Layer 2 (L2) over L3 support for applications as vMotion and integration with public clouds.

The NSX alternative performs VXLAN overlay over the NX-OS network as underlay, meaning that the Data Center could benefit from SDN relying on the existing infrastructure. However it requires a high investment as it is based on two implementations: The NX-OS network delivered by cisco and the NSX SDN delivered by VMware.

Moreover NSX is dedicated to the virtual environment excluding the bare metal servers and other physical appliances from the SDN solution.

ALTERNATIVE 3: Application Centric Infrastructure

Application Centric Infrastructure is Cisco's SDN solution for Data Center Networks, uses a spine-leaf two-tier topology with VXLAN as an overlay, the configuration is centrally managed from a controller and escalates well at the spines and leaves level.

ACI provides SDN features to all the server environments, storage, and network devices hosted today by the data center. It integrates with most of the hypervisors, load-balancers, and firewalls, increasing visibility, programmability and automation capabilities. Moreover, ACI supports several multisite configurations and integration with public clouds, providing features such as extension of broadcast domains across sites, Disaster Recovery (DR) site implementations, and redundancy.

The ACI implementation is only supported by Cisco Nexus 9000 series switches; i.e., it requires a complete redesign of the data center network and the replacement of the equipment in the core and distribution layer.

The migration from the existing NX-OS data center network to an ACI deployment demands to have both implementations in production at the same time. This is a time consuming process that requires an exhaustive planning, and good management of change to minimize the impact on production systems.

5.2.2.2 Life Cycle Cost (LCC) Analyzis

This section will analyze the cost of alternatives 1 and 3 throughout the complete lifecycle, including initial costs, operational costs, and maintenance costs. Alternative 2 is not analyzed, as it was discarded early in the technology selection due to the limitations regarding the integration with existing physical appliances and bare metal servers.

The initial costs are primarily engineering and procurement. Engineering cost was calculated based on the man-hours expended on the selection of technologies, topology design, and migration planning. Procurement refers to hardware investment and is calculated based on the quotes received from the vendor.

The installation costs refer to the man-hours expended on the installation of the equipment and the arrangement of the facility, i.e., Rackspace, new FO cables, and patch panels.

Energy costs were calculated based on the equipment's data sheets. The value used is typical consumption.

Operational costs are divided into monitoring and configuration. Monitoring refers to the periodic checks of the equipment health. These checks usually are performed using Cisco Prime for alternative 1 (NX-OS) or the Application Policy Infrastructure Controller

Cost Element	Subcost Element	Year	0	1	2	3	4	5	6	7	8	9	10
		disc	1	0.95238	0.90703	0.86384	0.8227	0.78353	0.74622	0.71068	0.67684	0.64461	0.61391
Initial Cost	Engineering		57.6	0	0	0	57.6	0	0	0	0	0	0
	Procurement		170	0	0	0	180	0	0	0	0	0	0
Installation			27.2	0	0	0	0	0	0	0	0	0	0
Energy Cost	Electricity		0	12.3582	12.3582	12.3582	12.3582	7.00362	7.00362	7.00362	7.00362	7.00362	7.00362
Operations	Monitoring		0	31.025	31.025	31.025	31.025	31.025	31.025	31.025	31.025	31.025	31.025
	Configuration		0	17.68	17.68	17.68	17.68	17.68	17.68	17.68	17.68	17.68	17.68
Maintenance	Preventive		0	13.6	13.6	13.6	13.6	13.6	13.6	13.6	13.6	13.6	13.6
	Corrective			13.6	13.6	13.6	13.6	13.6	13.6	13.6	13.6	13.6	13.6
Total			**254.8**	**88.2632**	**88.2632**	**88.2632**	**325.863**	**82.9086**	**82.9086**	**82.9086**	**82.9086**	**82.9086**	**82.9086**
NPV - Left Axix 1000USD			254.8	84.0602	80.0573	76.245	268.088	64.9611	61.8677	58.9216	56.1158	53.4436	50.8987
cumulated NPV - Right Axix 1000U			254.8	338.86	418.917	495.163	763.251	828.212	890.08	949.001	1005.12	1058.56	1109.46
TOTAL NPV 1000USD			1109.46										

	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
Initial costs	227.6	0	0	0	195.474	0	0	0	0	0	0	423.074
Installation and commissioning	27.2	0	0	0	0	0	0	0	0	0	0	27.2
Operations and Energy costs	0	58.1554	55.3861	52.7487	50.2368	43.6492	41.5706	39.5911	37.7058	35.9103	34.2003	449.154
Maintenance Costs	0	25.9048	24.6712	23.4964	22.3775	21.3119	20.2971	19.3305	18.41	17.5334	16.6984	210.031
												1109.46

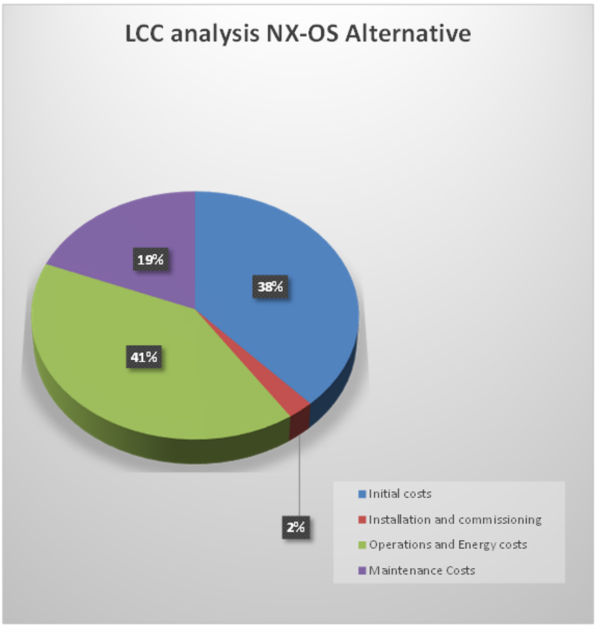

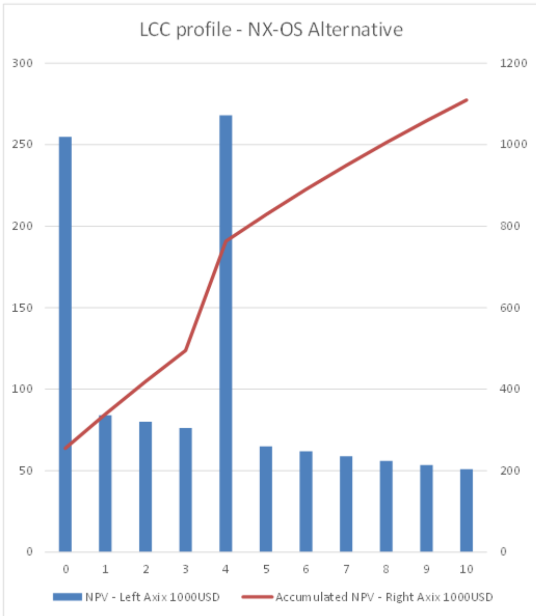

Figure 17: Life Cycle Cost (LCC) for Alternative 1

(APIC) GUI for alternative 3 (ACI). Configuration costs are mainly the provisioning of new equipment that requires connectivity to the data center.

Preventive maintenance costs are the staff hours expended on software upgrades. These upgrades are released by the manufacturer to address vulnerabilities, fix bugs and add features to the deployment. Based on early experience, it is assumed that these upgrades are typically done once or twice a year.

Corrective maintenance is the hours that network staff expend on the troubleshooting and

Cost Element	Subcost Element	Year	0	1	2	3	4	5	6	7	8	9	10
		disc	1	0.952380952	0.90702948	0.863837599	0.822702475	0.783526166	0.746215397	0.71068133	0.676839362	0.644608916	0.613913254
Initial Cost	Engineering		172.8	0	0	0	0	0	0	0	0	0	0
	Procurement		396	0	0	0	0	0	0	0	0	0	0
Installation Cost			163.2	0	0	0	0	0	0	0	0	0	0
Energy Cost	Electricity		0	6.479115	6.479115	6.479115	6.479115	6.479115	6.479115	6.479115	6.479115	6.479115	6.479115
Operations	Monitoring		0	15.5125	15.5125	15.5125	15.5125	15.5125	15.5125	15.5125	15.5125	15.5125	15.5125
	Configuration		0	4.42	4.42	4.42	4.42	4.42	4.42	4.42	4.42	4.42	4.42
Maintenance	Preventive		0	4.08	4.08	4.08	4.08	4.08	4.08	4.08	4.08	4.08	4.08
	Corrective			6.8	6.8	6.8	6.8	6.8	6.8	6.8	6.8	6.8	6.8
Total			**732**	**37.291615**	**37.291615**	**37.291615**	**37.291615**	**37.291615**	**37.291615**	**37.291615**	**37.291615**	**37.291615**	**37.291615**
NPV - Left Axix 1000USD			732	35.51582381	33.8245941	32.21389915	30.67990395	29.21895614	27.82757728	26.50245455	25.24043291	24.03850753	22.89381669
cumulated NPV - Right Axix 1000U			732	767.5158238	801.3404179	833.5543171	864.234221	893.4531772	921.2807544	947.783209	973.0236419	997.0621494	1019.955966
TOTAL NPV 1000USD			1019.955966										

	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
Initial costs	568.8	0	0	0	0	0	0	0	0	0	0	568.8
Installation and commissioning	163.2	0	0	0	0	0	0	0	0	0	0	163.2
Operations and Energy costs	0	25.15391905	23.95611338	22.81534607	21.72890102	20.69419145	19.70875376	18.77024168	17.87642065	17.02516252	16.2144405	203.94349
Maintenance Costs	0	10.36190476	9.868480726	9.398553072	8.951002926	8.524764691	8.118823515	7.732212872	7.364012259	7.013345008	6.679376199	84.012476
												1019.95597

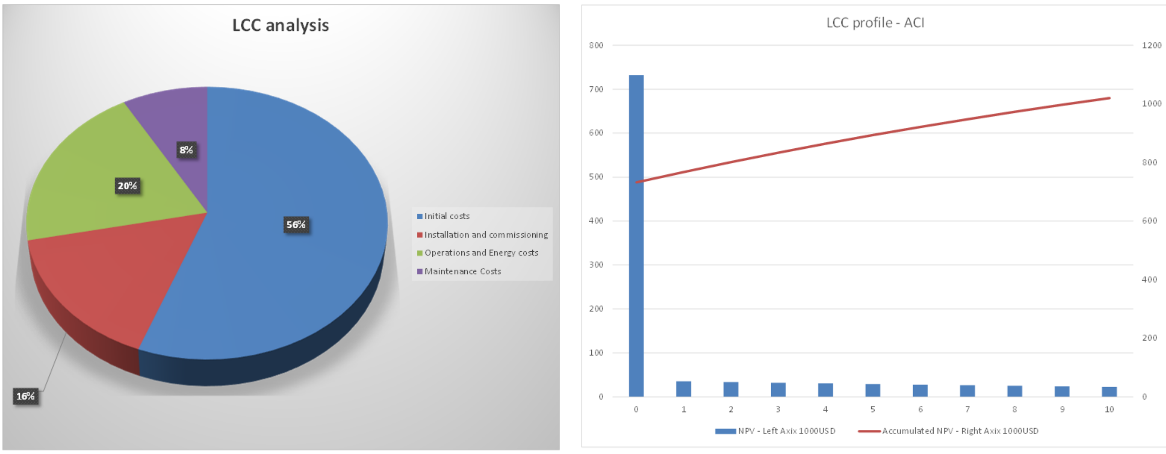

Figure 18: Life Cycle Cost (LCC) for Alternative 3

failure remedy.

Figure 17 illustrates the cost's distribution and the Life Cycle Cost (LCC) profile for alternative 1. The second equipment investment in year 4 refers to the replacement of the distribution and access layers as most of the devices are EoL in 2022.

The LCC Analysis and Profile for alternative 2 is illustrated in figure 18.

The complete LCC analysis is included in Appendix F.

5.2.2.3 GAP analysis and Comparison between alternatives

Table 4 shows a comparison between the three alternatives based on the criteria explained below:

- *CapEx:* Is the up-front investment. This includes the purchase of the equipment, preparation of the Data Center facility -electricity, Rackspace-, installation, consul-

Table 4: Comparison of the Alternatives to Replace the Data Center Network

	ALT 1	ALT 2	ALT 3
CAPEX	low	high	high
OPEX	high	medium	low
IMPLEMENTATION TIME	6 months	1 year	1 year
IMPLEMENTATION APPROACH	replace box by box	replace box by box and implement SDN	Parallel deployment
REQUIRED TRAINING	low	high	high
RESOURCES	in-house	network resources in-house, server and SDN resources had to be trained or outsourced	no resources in-house. Need to train network engineers or outsource the job.
INNOVATION	low	medium	high
SDN CAPABILITIES	no	partial	full
REQUIREMENTS GAP	no gap	no gap	no gap
PREFERENCES GAP	medium	medium	very low

tancy, migration, test, and documentation.

- *OpEx:* Refers to the operational expenditures. This includes the resources and man-hours required for provisioning, and operation activities, e.g., provisioning, software upgrades, and troubleshooting.
- *Implementation time:* The duration of the deployment from the day the equipment is powered up.
- *Implementation approach:* Describes briefly how the network team will proceed to the installation including the physical requirements, e.g., Rackspace, floor space, electricity, and cooling.
- *Required training* of network staff to install, operate and maintain the new implementation.
- *Resources:* Competencies and skills of ConocoPhillips' network staff and the need for external consultancy.
- *Innovation:* Refers to the solution alignment with ConocoPhillips goals and innovation strategy.
- *SDN Capabilities:* Automation and programmability capabilities and integration

with existing endpoints.

- *Requirements and preferences GAP:* Level of compliance with the requirements and preferences described in section 5.2.1 and summarized in table 3.

Alternative 1 -*Cisco NX-OS*- allows for a progressive replacement of the hardware based on EOL dates reducing the upfront investment. On the other side, this alternative results in higher operational expenditures, as day to day operations such as provisioning and software upgrades cannot be automated, i.e., network engineers have to log in and do the configuration in a box per box basis; this process is time-consuming and increases the risk for human failure.

The implementation approach of this alternative is relatively simple and several engineers in the network staff are qualified to perform the installation without additional training or consultancy. The configuration of the new equipment is almost identical to the existing one, i.e., the configuration file can be replicated to the new boxes. Moreover, the equipment redundancy allows for the replacement of one box at the time eliminating the need for extra physical space and utilities at the data center.

Because of the reduced scope, the implementation simplicity, and the in-house competencies, the implementation time of alternative one is estimated to be approximately six months. Furthermore, this alternative fulfills all main requirements explained in section 5.2.1 and summarized in table 3. However, it does not align with company goals and strategy as it does not provide significant value translated into innovation, programmability features, cloud integration, and multisite support.

Alternative 2 -*Cisco NX-OS with NSX SDN*- Requires the same hardware investment than alternative 1 in addition to the VMware SDN software. The up-front investment for this solution could be high because it relies on two separate implementations from different vendors, and requires dedicated resources from the network and server team. The operational expenditures of this solution might be slightly reduced in the virtual environment; however, the NSX software does not integrate with legacy and physical appliances, resulting in two operational modes: one for the VMware environment, and the other for the traditional NX-OS network which is equivalent to alternative 1, *viz:* Network configuration is done in a box per box basis.

The physical implementation of this deployment is identical to alternative 1, while the NSX SDN has to be designed, deployed and tested in collaboration with VMware experts, probably outside the company. Because of the NSX SDN scope, the need for external resources and training requirements for server and network staff, this implementation is estimated to take one year.

This alternative fulfills all main requirements, but the poor innovation in the physical network and the lack of integration between NSX and physical appliances, limits the scalability, automation capabilities, cloud integration, and multisite support.

Alternative 3 *-Cisco ACI-* Relies upon the Cisco Nexus 9000, and requires the replacement of the whole Data Center Network. This increases the scope of the project and the up-front investment considerable. However, this solution extends the SDN capabilities to all endpoints in the data center -virtual environment, physical appliances, and network services- providing endpoint visibility, allowing for a high level of automation, simplifying the maintenance activities, and reducing the operational expenditures.

For the implementation part, this alternative requires a parallel deployment. Meaning both data center networks need to be operational at the same time. This introduces additional requirements with regards to the facilities such as floor space, rack space, electricity, and cooling. Additionally, network staff requires thorough training in the new technology in order to implement, operate and maintain the new data center network.

However, this alternative fulfills all the requirements and preferences described in section 5.2.1 and summarized in table 3. Moreover, it aligns with ConocoPhillips goals and innovation strategy, providing high scalability, flexibility, automation, programmability, visibility, analytics, cloud integration, and full support for multisite deployments.

Table 5 provides a GAP analysis that summarizes the requirements compliance for the three alternatives.

5.2.2.4 Decision

The ACI Cisco solution was chosen based on the LCC, GAP analysis, the project's budget and the experience the company had with this technology in the USA. The next section will go through the detailed analysis of the Data Center facility and the design of the new network.

Table 5: GAP Analysis of the Alternatives to Replace the Data Center Network

End-Points and Applications	Alt 1	Alt 2	Alt 3
VMware Hypervisor	Compliant	Compliant	Compliant
Bare Metal Servers	Compliant	Compliant	Compliant
Netapp	Compliant	Compliant	Compliant
Firewalls	Compliant	Compliant	Compliant
Netscaler	Compliant	Compliant	Compliant
Capacity and Interfaces			
600Gbps switching capacity	Compliant	Compliant	Compliant
1.2Tbps switching capacity	Compliant	Compliant	Compliant
100Gbps interfaces	Compliant	Compliant	Compliant
40Gbps interfaces	Compliant	Compliant	Compliant
25Gbps interfaces	Compliant	Compliant	Compliant
10Gbps interfaces	Compliant	Compliant	Compliant
10/100/1000 Mbps copper ethernet	Compliant	Compliant	Compliant
Non-blocking	Partially Compliant	Partially Compliant	Compliant
No oversubscription	Not-Compliant	Not-Compliant	Compliant
Scalability in the core	Not-Compliant	Not-Compliant	Compliant
Scalability in the access layer	Compliant	Compliant	Compliant
Multisite and Cloud Integration			
Programmability and Automation	Partially Compliant	Partially Compliant	Compliant
Multisite support	Partially Compliant	Partially Compliant	Compliant
L2 over L3	Not-Compliant	Not-Compliant	Compliant
Security and Reslience			
High Availability	Compliant	Compliant	Compliant
Monitoring	Compliant	Compliant	Partially Compliant
Traffic Filtering and QoS	Compliant	Compliant	Compliant
End Point visibility	Not-Compliant	Not-Compliant	Compliant
Access Control	Compliant	Compliant	Compliant
Support			
Easy provisioning	Not-Compliant	Not-Compliant	Compliant
COP standard compliant	Compliant	Compliant	Compliant
NetFlow	Compliant	Compliant	Compliant
SPAN ports	Compliant	Compliant	Compliant

X	Compliant
X	Partially Compliant
X	Not-Compliant

5.2.3 DESIGN PHASE

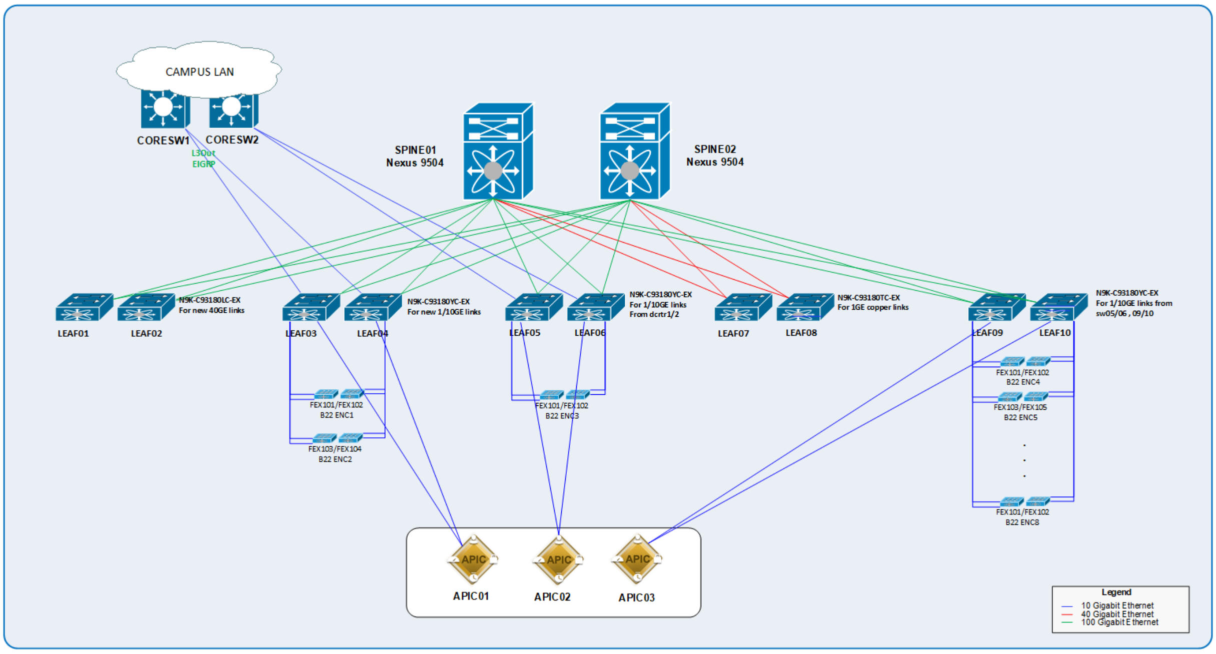

Figure 19: COP ACI Topology

The design depicted in figure 19 is based on a spine-leaf two-tier Next Generation Data Center (NGDC) architecture with 40 and 100Gbps non-blocking fabric ports that provide high-bandwidth, and low latency as the end-points are maximum one hop away from each other.

The spines provide connectivity to the leaves, maintain an endpoint table and serve as a BGP route reflector. In future implementations of multisite, the spines will be the connection point to the WAN and will do the name translation between domains.

The spines are built on the Cisco modular chassis Nexus 9504 with 4x100 fabric modules and 4 slots for payload line cards and 1 slot for up to two supervisors. The solution has one line card that provides 32 x 40/100Gbps QSFP28 ports. Appendix G shows the comparison of the different chassis and modules for the Nexus 9500 Switch Series.

The topology contains 10 leaves from Cisco Nexus 9300 series grouped in pairs to provide redundancy to the endpoints. The first pair (LEAF01/02) supports up to 32 x 40/50Gbps QSFP+ ports or 18 x 100Gbps QSFP28 ports; it also supports fiber and copper breakout cables, but for this design, it will only be used as 40/100Gbps interfaces. The second, third, and last pair of leaves (LEAF03/04, 05/06, 09/10) provide 48 x 1/10/25Gbps fiber ports and 6 x 40/100Gbps QSFP28 ports. The fourth pair (LEAF07/08) provides 48 x

100M/1/10Gbps copper ports and 6 x 40/100Gbps QSFP28 ports; this pair will replace the copper FEXes from the previous design. Appendix H shows a summary of the Cisco 9300 series.

The APIC cluster APIC-CLUSTER-M2 consists of three controllers *"with medium-size CPU, hard drive, and memory configurations (up to 1000 edge ports)[14]"*.

All leaves are uplinked to both spines using 40/100 Gbps ports. The APICss are connected to three different pair of leaves to increase the resilience of the fabric. Connections outside the fabric are routed via an L3 EIGRP router consisting of 4 x 10 Gbps connections between border leaves and campus core switches.

Two new racks with power circuits are required to place the spines beside the old Data Center core. The leaves are EoR and are collocated or close to the legacy distribution and access switches. The FO infrastructure is increased considerably to facilitate the end-point connectivity to the new Data Center Network.

Management routing is moved out of the Data Center Network and Out of band IP addresses are assigned to all nodes for support and monitoring purposes.

The solution is summarized in table 6.

Table 6: Summary of COP's Solution

SUMMARY	
Spines	
Ports per spine	34 x 40/100Gbps
Capacity per spine	3200 Gbps
40 Gbps Leaves	
Fabric/Access Ports per leaf	3200 Gbps
Capacity per leaf	3200 Gbps
10 Gbps Leaves	
Fabric Ports	6 x 40/100Gbps
Access Ports	48 x 1/10/25Gbps
Capacity	3600 Gbps
Copper Leaves	
Fabric Ports	6 x 40/100 Gbps
Access Ports	48 x 100 Mbps/1/10Gbps
APIC	
Edge Ports	1000
Total 40/100 Gbps ports	60 Access Ports
Total 10Gbps Ports	288 Access Ports
Total Copper Ports	96 Access Ports

5.2.4 IMPLEMENTATION PHASE

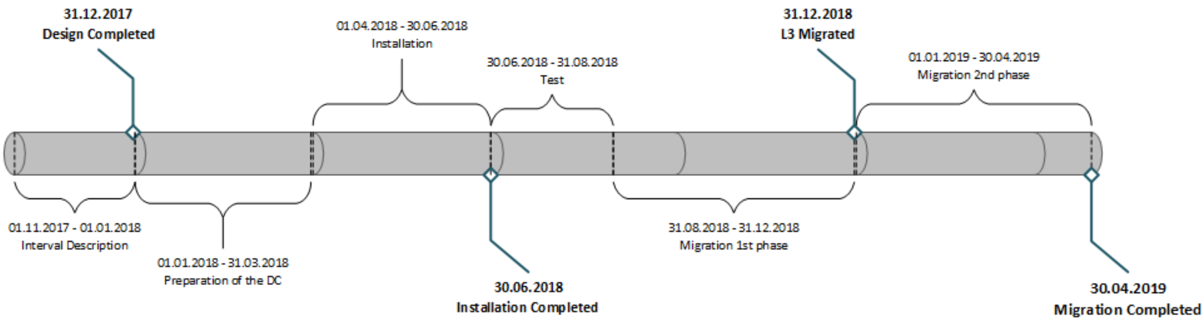

Figure 20: Project's Timeline

After the design completion and approval, the next step is the implementation phase. The project's activities and timeline are illustrated in figure 20.

5.2.4.1 Preparation of the Data Center

The preparation of the Data Center includes placement of purchase orders, installation of racks, FO cables, patch panels, power circuits, mounting of network equipment, and cabling.

5.2.4.2 Installation

The installation activities include equipment power-up, initial setup, software upgrade, access control, monitoring, fabric policies, network configuration, and tenant configuration.

During the initial setup Out-of-Band (OOB) [17] IP addresses are assigned the controllers; then a setup wizard is started through the KVM connection [18] of the first APIC. The wizard asks for some entries before it starts discovering the fabric nodes using the LLDP protocol. When the controller finishes the discovery, OOB IP Addresses are assigned to spines and leaves through the GUI.

APIC nodes are upgraded to the ConocoPhillips' standard firmware version, and access control is configured to use TACACS+ servers.

Fabric policies include NetFlow, backup, SNMP, and Syslog. The monitoring servers are

[17] Addresses that are not routed within the Data Center Network. This guarantees the access to the devices in case of an outage

[18] KVM is a connection to the server attaching directly a screen and a keyboard, it gives access to the initial setup utility

configured to ping the OOB IP addresses of all fabric nodes, *viz:* APICs, Spines, and Leaves.

The network configuration includes the definition of the VLAN range, virtual and physical domains, interface profiles, and leaf profiles.

For this deployment, two tenants are defined: *A production and a test tenant.* The tenant configuration includes the Virtual Routing and Forwarding (VRF), Bridge Domains, Endpoint Groups (EPGs), contracts, and association with interface profiles. Finally, the routing to the outside -campus core- is configured using 4 x 10G routed ports.

5.2.4.3 Test

When the installation is completed, the ACI fabric is connected to the old Data Center Network using an L2 interface or L2out. This interim topology illustrated in figure 21 allows the migration of endpoints from the legacy data center network to the ACI fabric without significant outages.

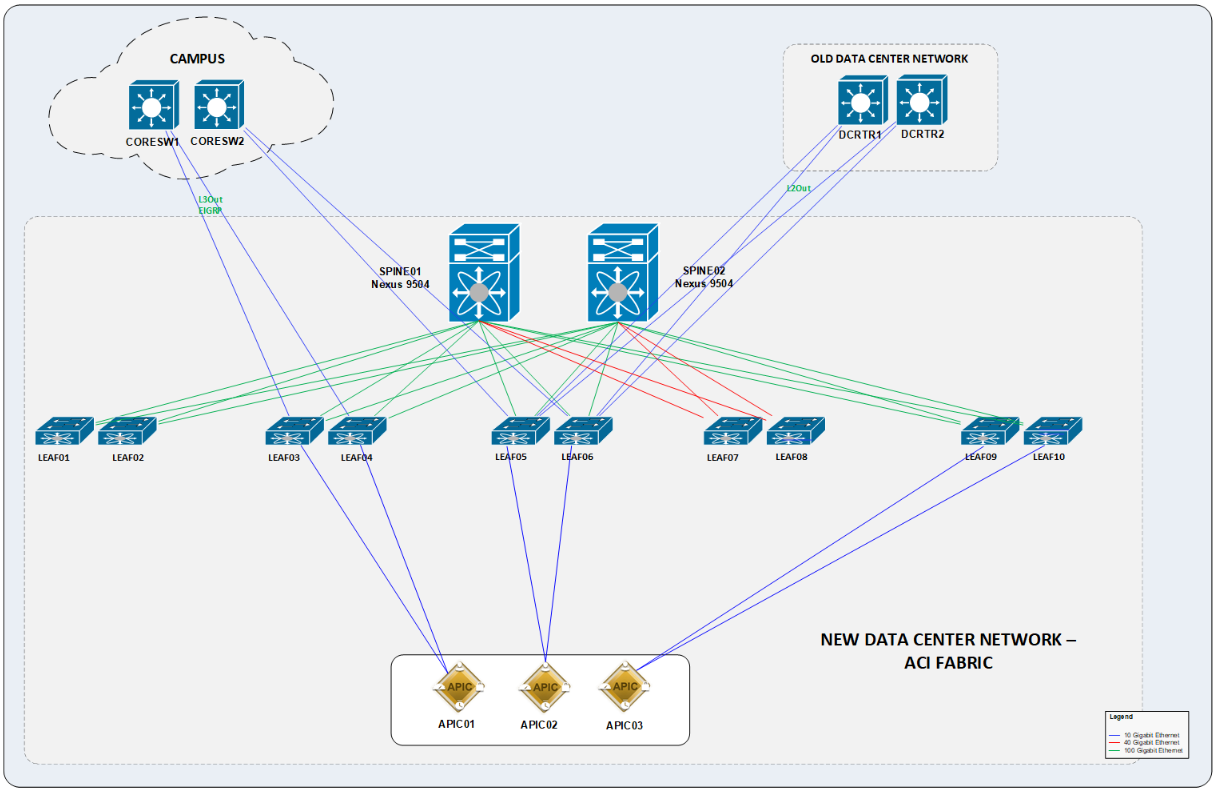

Figure 21: Interim Topology during the migration phase

During this phase, each EPG is tested to assure that there are no errors in the configuration or software bugs that can impact the Data Center during the migration.

5.2.4.4 Migration

During the migration phase, the endpoints are moved from the old data center to ACI. A detailed migration plan is elaborated in collaboration with system owners, and change requests are submitted through ConocoPhillips' Management of Change (MOC) platform to remind the owners about the agreed date and time and to inform the final users.

The migration plan includes a risk analysis of the job, activities required before the change, test plan, the back-out plan and activities required after the change. During the risk analysis, the applications running at the endpoint are identified; their functions are correlated with other systems to find interdependencies, and the impact is estimated based on previous experience, knowledge of system owners and existing documentation. Finally, the criticality of the application is assessed together with system owners and users.

Depending on the results of the risk analysis, mitigating actions are defined, e.g., move applications to another host in the case of VMs, perform changes during less critical time windows and inform end users. Activities required before the change are typically documented in the form of a checklist and include the execution or confirmation of mitigating actions, configuration, health check on devices, and required instruments or test equipment.

The test plan is elaborated in collaboration with system owners and users. The first part of the test plan consists of traffic and health inspection of network equipment, e.g., interface link, amount of traffic, errors or packet drops, and others. The second part is the functionality test; this is performed by system owners or application users. If one or more tests fail, the migration is rolled back and analyzed to find the issue before a new change is scheduled.

The back-out plan contains the activities that need to be performed if the change is not successful, i.e., if one or more tests fail. The activities are documented in order, and configuration scripts are predefined to simplify the rollback process and reduce the risk of human failure.

Activities required after the change include the cleanup in old network equipment, documentation updates, information to network administrators, system owners, and end users.

When all the endpoints in a system are physically moved to the new data center network, the L3 interface is shut down in the legacy network and configured in the ACI fabric. During this activity, the endpoints experience a short outage; typically some seconds.

After all endpoints and L3 interfaces are moved to the ACI fabric, the L2out configuration can be removed and the old data center network is ready to be decommissioned.

Some of the configuration tasks were scripted in postman and pushed to the APIC using Representational State Transfer (REST) Application Programming Interface (API) calls. This streamlined the implementation, simplified the provisioning process and minimized the duration of the outages. Among the activities that are useful to automate in an ACI implementation, it is worth to mention the migration of B22 blades and configuration of new EPGs with their respective Bridge Domains (BDs).

6 Discussion

The overall aim of this work was to Design a Data Center Network solution for ConocoPhillips Norway, that aligned with its global cloud strategy positioning the company for the adoption and integration with emerging technologies, such as Cloud Computing, Internet of Things, Artificial Intelligence, and Analytics. The specific research objectives were:

1. Understand next generation Data Center Network's concepts, services, value and its role in the inter-cloud journey.
2. Assess COPNO requirements and specifications for the data center network
3. Evaluate the technologies critically and architect a solution for the replacement of the data center network in COPNO.
4. Describe the Data Center Network implementation process, achievements, challenges, and topics for future research.

This section presents a reflection of the research objectives listed above and outlines the principal findings throughout the research study. Further, it provides recommendations for the operation and optimization of the Data Center Network and possible future works. Finally, in the self-reflection part, main challenges and lessons learned are summarized.

6.1 Research Objectives: Summary of Findings

6.1.1 Objective 1: Understand next generation Data Center Network's concepts, services, value and its role in the inter-cloud journey

The literature review describes how the digital transformation or fourth industrial revolution is impacting traditional business models and value chains. Even though this impact might have adverse effects; it also offers opportunities for consumer and provider companies. ConocoPhillips, as a consumer has access to a broader portfolio of technologies and services on-demand, that can quickly be provisioned, without the burden of long-term contracts and high up-front investments.

However, to take advantage of the features and benefits of emerging technologies, ConocoPhillips should achieve certain technological maturity. Integration between on-premises and public cloud computing deployments is a crucial milestone in this matter.

For a company like ConocoPhillips, with significant existing data center infrastructure, the Data Center Network has an essential role in the integration and migration of assets to public clouds. The network technology should be chosen to provide agility, scalability, programmability, automation, and orchestration with public clouds.

Legacy data center networks have become an obstacle for the development of highly virtualized Data Centers and their integration with clouds. New Generation Data Center Networks or SDN, proposes a new architecture and adopts different technologies and protocols such as VXLAN, BGP and OPFLEX, to centralize the management, increase the resilience, and offer automation capabilities.

The main conclusion from this study is that the Data Center Network is a critical element in the integration with public clouds. In the transition between on-premises and hybrid cloud deployments, it is beneficial to centralize operations through a controller or cloud orchestrator. This feature, provided only by SDN networks, increases visibility, control, agility, and reduces the risk of human error.

6.1.2 Objective 2: Assess COPNO requirements and specifications for the data center network

ConocoPhillips is an exploration and production company that has operated in Norway since 1969. Its main office is located in Tananger where it operates an on-premises Data Center hosting a large number of services including mission-critical applications.

The company has developed a cloud strategy and encourages all the BUs to prepare their facilities for "all-in" cloud migration. However, this is a time-consuming process that will probably take several years.

In 2017, ConocoPhillips Norway decided to upgrade the data center network to comply with corporate policies regarding vendor support and to address operational requirements such as capacity, layer two extension, and programmability.

The conclusion from the research regarding this objective is that the company in order to align with its cloud strategy and business needs, should increase the data center capacity, enhance endpoint visibility, add automation capabilities, centralize network operations, provide multisite and cloud integration.

6.1.3 Objective 3: Evaluate the technologies critically and architect a solution for the replacement of the data center network in COPNO

This work addresses this objective by evaluating three alternatives: Cisco NX-OS, NSX, and Cisco ACI.

Cisco NX-OS provides addresses the capacity requirements, but it does not provide centralized management, layer two extension, multisite deployment, and automation capabilities. This option is not recommended as it does not align with the ConocoPhillips IT strategy.

NSX is a software solution that implements VXLAN overlay over an IP network underlay. In ConocoPhillips' case, it requires the replacement of the EoL equipment in the data center. This solution is not recommended as it is limited to the VMware environment, excluding the remaining physical appliances and Bare Metal servers.

ACI is the Cisco SDN solution. It is based on a two-tier topology that provides high capacity and scalability. The network is centrally administrated from a controller; this increases visibility, simplifies management, provides programmability and automation capabilities. ACI uses VXLAN as an overlay and has the ability to terminate the tunnels in the hypervisor level or at the network equipment extending the SDN network to the bare-metal hosts and physical appliances. This solution also supports multisite support and cloud deployments that can be managed from an orchestrator.

Based on ConocoPhillips business strategy and technical requirements, the recommendation is to adopt the ACI solution. This solution represents a higher up-front investment and longer migration time. However, this is the alternative that addresses most of the company needs, and it also has the lowest cost through its whole life-cycle.

6.1.4 Objective 4: Describe the Data Center Network implementation process, achievements, challenges, and topics for future research

To minimize the impact on production services, the new SDN and the legacy Data Center Network had to operate parallel. This increased the scope of the project with additional requirements of power, cabling, floor, and rackspace.

The initial configuration of the ACI fabric was performed from the APIC controller; this was a comprehensive process done in-house that required a detailed survey of configuration guidelines, standards, and best practices.

Repetitive and extensive tasks such as the configuration of B22 blade servers was automated using postman scripts. This was an opportunity to explore the benefits of a centrally managed SDN, mitigating the impact on production systems, reducing the implementation time, and minimizing the configuration complexity.

Finally, it is worth mentioning the importance of proper project management, MOC, excellent communication with applications' owners, users, a thorough handover to system administrators, and regular documentation updates.

6.2 Self-reflection

I have dedicated most of my work life to highly technical assignments, and I probably did not appreciate the value of a management background before I started to study this master degree.

As a highly qualified staff, engineers must often participate in decision making. However, without management foundation, we could easily ignore or overweight decision factors.

In the specific case of this master thesis, I decided to evaluate the available data center network alternatives against ConocoPhillips' goals and needs, including also future implementations. The literature review and the State of the Art provided a theoretical fountain, that combined with a thorough analysis of technical specifications served as the principal source of information for the final decision.

The implementation phase comprised coordination and interaction with different people and systems. This was carried out with minimum failures, thanks to organized and systematic project management.

I understood during the execution of this project that the human factor could sometimes be the main enemy of change. Pushing experts out of their comfort zone and moving them away from what has made them unique during many years, was one of the biggest challenges.

However, I am pleased to appreciate the results after over a year of work. We have upgraded our data center network to one of the last available technologies; we are operating the implementation without significant issues; and finally, we appreciate the value of our data center architecture as we automate processes and easily integrate with other sites.

6.3 Limitations

I started working on this project at the beginning of 2018, and some months after I decided to adopt it as my master degree dissertation. I expended most of the first semester researching the available literature and exploring training opportunities.

My biggest mistake throughout this master study was waiting too long to define the framework, the overall aim, specific objectives, and structure. As a result, I expended much time on an unsystematic literature review that did not represent a valuable contribution to the final product.

I also underestimated the value of the research methodology and I decided to postpone this section. I was surprised to find out how the study of the different qualitative methods could help to accomplish a structured data analysis and literature review. It was then, I changed my data collection tool from OneNote to NVivo and started to expend less time in data classification focusing more on its analysis.

I believe that beginning with a comprehensive review of research methods and establishing a strategy from the beginning, followed by a clear definition of the overall aim and specific objectives, could have been a better foundation for this work.

6.4 Recommendations

The Information Technologies' field has always been under continuous research and development and catching up with state of the art, and emerging technologies is not always viable.

Many companies, like ConocoPhillips, have as a part of their technology strategy, the cyclic replacement of obsolete components and solutions; this might be an opportunity to evaluate the benefits and drawbacks of the different alternatives out in the market.

There is not a standard solution that fits the needs of all companies. Hence, the importance of a comprehensive study of existing literature, state of the art, vendor documentation, and current trends.

With all the options on the table, the next step is to evaluate them against the company's strategy and technical specifications. This evaluation should contain the whole life cycle of the assets.

In the case of Data Center Networks, an SDN solution represents a high up-front investment and might not be the right alternative for many organizations. However, if the company in question has a hybrid cloud or multi-site strategy, an SDN solution is an option worth to evaluate in detail. Moreover, new topologies and distributed intelligence across network devices, allow for smaller deployments that can easily scale from two leaves to around 30 for the smallest fabric solution.

The implementation phase is a critical part of the chosen alternative and could be a tedious and time-consuming stage. Appropriate project management is recommended to handle activities such as viability studies, preparation of facilities, schedules, communication with customers and vendors, management of change, documentation, and handovers.

It is recommended to consider the participation of system administrators in the implementation phase of the project. In this study, this was a critical decision that added significant value to the company as the involved individuals built knowledge and understanding of the deployment building blocks.

The final recommendation is to maintain continuous documentation throughout all phases of the project. Document everything: meetings, e-mails, brain-storming sessions, expert recommendations, quotes, topology, configuration, changes, and lessons learned.

6.5 Future Implementations

The focus of this project was to replace and upgrade COPNO's data center network and to prepare the company for future cloud integration and multisite implementations. Some suggestions for future studies include:

1. *Application approach and integration with VMware:* ACI offers the possibility to segregate endpoints based on their functions, e.g., databases, web servers, application servers, storage, and others. This might improve security as the environments are isolated from each other and only required protocols are allowed between hosts. Integration with VMware increases endpoint visibility, facilitating the identification and classification of applications.

   However, multiple questions need to be answered to determine the real value of this approach. E.g., How the implementation improves visibility, security, agility, and automation? What are the implications of the implementation on the support model? How to define roles and boundaries between network and server teams? What is required in terms of equipment and resources to perform the implementation?

2. *Implementation of multisite and extension of L2 across sites:* ACI is based on

VXLAN. This overlay protocol simplifies the implementation of multisite deployments significantly.

I believe it could be worth to study the advantages and disadvantages of a multisite implementation that comprises the main data center, disaster recovery sites, offshore locations, and other BU's. This could be analyzed from the operational and functional point of view.

3. *Integration with public clouds:* One of the principal characteristics of SDN Data Centers is the way they simplify the integration process with cloud computing. A topic that needs to be explored include the alternatives for doing that integration and the advantages that ACI provides compared with other Data Center Network models.

   An practical implementation or a PoC and the analysis of the results could be a beneficial contribution.

7 Conclusion

In addition to being a crude exploration and production company, ConocoPhillips also seeks to develop and utilize technology to optimize operations, reduce costs, improve safety, and reduce their carbon footprint. In 2018, the company started a project to replace its Data Center Network in Norway. Through this thesis, the available Data Center Network technologies were studied and correlated with the company's IT strategy and business requirements.

Considering ConocoPhillips' "all in cloud" philosophy, their ambition to orchestrate all sites in the region, the life cycle cost of the different alternatives and the innovation goals, it was decided to implement the Cisco Software-Defined Networking (SDN) Data Center, better known as Application Centric Infrastructure (ACI).

The initial installation and the migration from the legacy to the new Data Center Network were performed in over one year. Comprehensive project management, management of change, risk analysis, and decision making were essential in the execution of the project.

The programming features of the ACI implementation were explored, simplifying the configuration process and minimizing the outages.

Currently, there is an ongoing project to extend the data center implementation to a disaster recovery site. This will be an opportunity to explore the multisite features offered by the new data center network. Cloud implementation is an extension of the data center to one or several public clouds, and the administration is done from a central orchestrator.

This project can be used as a guide to plan and design SDN data centers and hybrid clouds. However, the individual requirements and specifications of different companies may vary and need to be analyzed independently.

References

[1] : *Amazon Web Services (AWS) - Cloud Computing Services.* – URL https://aws.amazon.com/

[2] : *The Customer Success Platform To Grow Your Business.* – URL https://www.salesforce.com/

[3] : *Dropbox.* – URL https://www.dropbox.com/

[4] : *How to Secure Cisco Routers and Switches.* – URL https://www.globalknowledge.com/us-en/resources/resource-library/articles/how-to-secure-cisco-routers-and-switches/

[5] : *Microsoft Azure Cloud Computing Platform and Services.* – URL https://azure.microsoft.com/

[6] : *Our History.* – URL http://www.conocophillips.no/who-we-are/our-history/

[7] : *Software-Defined Networking (SDN) Definition.* – URL https://www.opennetworking.org/sdn-definition/

[8] : *Web Hosting, Lightning Fast Hosting and One Click Setup, GoDaddy IN Pro.* – URL https://in.godaddy.com/hosting/web-hosting

[9] : *Welcome to Office.* – URL https://www.office.com/

[10] : *What Is Amazon EC2?.* – URL https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html

[11] : *Cloud Computing Service Models and their Benefits.* Sep 2015. – URL https://www.msigeek.com/7357/cloud-computing-service-models-benefits

[12] : *Understanding Rapid Spanning Tree Protocol (802.1w).* Sep 2017. – URL https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/24062-146.html

[13] : *Cisco Nexus 7000 Series Switches Data Sheet.* Mar 2018. – URL https://www.cisco.com/c/en/us/products/collateral/switches/nexus-7000-series-switches/Data_Sheet_C78-437762.html

[14] : *Cisco Application Policy Infrastructure Controller (APIC) - Cisco Application Policy Infrastructure Controller Data Sheet.* Feb 2019. – URL https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/application-policy-infrastructure-controller-apic/datasheet-c78-739715.html

[15] : *Cisco Nexus 5000 Series Switches - End-of-Life and End-of-Sale Notices.* Jan 2019. – URL https://www.cisco.com/c/en/us/products/switches/nexus-5000-series-switches/eos-eol-notice-listing.html

[16] : *Cisco Nexus 7000 Series Switches - End-of-Life and End-of-Sale Notices.* Mar 2019. – URL https://www.cisco.com/c/en/us/products/switches/nexus-7000-series-switches/eos-eol-notice-listing.html

[17] Ahdoot, Albert: *The Connection between Data Centers and the Web.* 2011. – URL https://www.colocationamerica.com/blog/history-of-internet-and-colocation.htm

[18] Ali Khoshkholghi, Mohammad ; Abdullah, Azizol ; Latip, Rohaya ; Subramaniam, Shamala ; Othman, Mohamed: Disaster Recovery in Cloud Computing: A Survey. In: *Computer and Information Science* 7 (2014), 09

[19] Asay, Matt: *Capital One's 'all-in' cloud strategy is much more than a tech decision.* Dec 2018. – URL https://www.techrepublic.com/article/capital-ones-all-in-cloud-strategy-is-much-more-than-a-tech-decision/

[20] Biggam, John: *Succeeding with your master's dissertation: a step-by-step handbook.* McGraw-Hill Education (UK), 2015

[21] Bruno, Anthony ; Jordan, Steve: *CCDA 200-310 Official Cert Guide: Exam 61 Offic Cert ePub _5.* Cisco Press, 2016

[22] Bryson, Teren: Cisco ACI in the Real World. In: *Enterprise Networking Planet* (2015)

[23] Cherkaoui, O. ; Menon, R.: *Virtualization, Cloud, SDN, and SDDC in Data Centers.* Wiley Blackwell, 2014. – 389–400 S. – ISBN 9781118937587

[24] Columbus, Louis: Roundup Of Cloud Computing Forecasts, 2017. In: *Forbes* (2017)

[25] Cunningham, J B.: Strategic considerations in using action research for improving personnel practices. In: *Public Personnel Management* 24 (1995), Nr. 4, S. 515–529

[26] Dagenhardt, Frank ; Moreno, Jose ; Dufresne, Bill: Deploying ACI: The complete guide to planning, configuring, and managing Application Centric Infrastructure. (2018)

[27] Gonsalves, Antone: Cisco ACI software hasn't found a place in enterprise data centers. In: *Search Networking* (2017)

[28] Halabi, Sam: *Hyperconverged Infrastructure Data Centers: Demystifying HCI*. Cisco Press, 2019

[29] Jackson, Chris ; Preston, Hank ; Steve Wasko, publisher=Cisco Press: *CCNA Cloud CLDFND 210-455 Official Cert Guide*

[30] Kogias, Dimitrios G. ; Xevgenis, Michael G. ; Patrikakis, Charalampos Z.: Cloud federation and the evolution of cloud computing. In: *Computer* (2016), Nr. 11, S. 96–99

[31] Krattiger, Lukas ; Kapadia, Shyam ; Jansen, David: *Building Data Centers with VXLAN BGP EVPN: A Cisco NX-OS Perspective*. Cisco Press, 2017

[32] Lavalle, Brian: *Undertaking the challenge to reduce the data center carbon footprint*

[33] LogicMonitor: Cloud Vision 2020: The future of the cloud / Logic Monitor. 2017. – Forschungsbericht

[34] Marr, Bernard: *What is Industry 4.0? Here's A Super Easy Explanation For Anyone*. Sep 2018. – URL `https://www.forbes.com/sites/bernardmarr/2018/09/02/what-is-industry-4-0-heres-a-super-easy-explanation-for-anyone/#770dd6ec9788`

[35] Mell, Peter ; Grance, Tim u. a.: The NIST definition of cloud computing. In: *National institute of standards and technology* 53 (2009), Nr. 6, S. 50

[36] Merriam, Sharan B. ; Tisdell, Elizabeth J.: *Qualitative research: A guide to design and implementation*. John Wiley & Sons, 2015

[37] Michael Lenox, Sonja Rueger Ching Fong O.: *Digital Transformation*. 2019. – URL `https://www.coursera.org/learn/bcg-uva-darden-digital-transformation/home/info`

[38] MOORE, Geoffrey A.: Dealing with darwin: How great companies innovate at every phase of their evolution, portfolio. In: *New York, USA* (2005)

[39] MOOZAKIS, Chuck ; LINTHICUM, David ; WATSON, Ray ; KNAPP, Kristin: A New Network Architecture: Send in the Clouds. In: *Search Networking* (2015)

[40] NUTANIX: Enterprise Cloud Index. 2018. – Forschungsbericht

[41] ONF: *Software-Defined Networking: The New Norm for Networks.* April 2012. – URL https://www.opennetworking.org/images/stories/downloads/sdn-resources/white-papers/wp-sdn-newnorm.pdf

[42] ONF: *SDN Architecture Overview.* December 2013. – URL https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/SDN-architecture-overview-1.0.pdf

[43] PEACOCK, Melanie J.: *The human resource professional's guide to change management : practical tools and techniques to enact meaningful and lasting organizational change.* 2017

[44] PEÑA-LÓPEZ, Ismael u. a.: Cloud Computing: The Concept, Impacts and the Role of Government Policy. (2014)

[45] PORTER, Michael E.: The five competitive forces that shape strategy. In: *Harvard business review* 86 (2008), Nr. 1, S. 25–40

[46] SALAM, Abdul: *Deploying and managing a cloud infrastructure : real world skills for the CompTIA cloud+ certification and beyond.* 2015

[47] SANTANA, Gustavo A.: *CCNA Cloud CLDFND 210-451 Official Cert Guide: Exam 55 Offic Cert Guide.* Cisco Press, 2016

[48] SCARPATI, Jessica: Cloud network strategy: One size doesn't fit all. In: *Search Networking* (2015)

[49] SCHWAB, Klaus: *The fourth industrial revolution.* Currency, 2017

[50] SHAMSEE, Navaid ; KLEBANOV, David ; FAYED, Hesham ; AFROSE, Ahmed ; KARAKOK, Ozden: *CCNA Data Center DCICT 200-155 Official Cert Guide: Exam 67 Offic Cert ePub _1.* Cisco Press, 2017

[51] SILVA, Peter: *What is Load Balancing?.* – URL https://devcentral.f5.com/articles/what-is-load-balancing-24740

[52] Simmon, Eric: Evaluation of Cloud Computing Services Based on NIST SP 800-145. In: *National institute of standards and technology* (2018)

[53] smartsheet: *How Your Business Can Benefit from Value Chain Modeling.* – URL https://www.smartsheet.com/value-chain-model

[54] Soerensen, NetPower Web S.: *Kulturminne Ekofisk.* – URL http://www.kulturminne-ekofisk.no

[55] Trailhead: *The Four Industrial Revolutions.* – URL https://trailhead.salesforce.com/en/content/learn/modules/learn-about-the-fourth-industrial-revolution/meet-the-three-industrial-revolutions?trail_id=learn-about-the-fourth-industrial-revolution

[56] Varia, Jinesh ; Mathew, Sajee u. a.: Overview of amazon web services. In: *Amazon Web Services* (2014), S. 1–22

[57] Wallace, Kevin: *CCNP Routing and Switching ROUTE 300-101 Official Cert Guide: Exam 37 Cert Guide.* Cisco Press, 2014

[58] Weinman, Joe: *Cloudonomics : the business value of cloud computing.* 2012

[59] Yoo, Christopher S.: Cloud Computing: Architectural and Policy Implications. In: *Review of Industrial Organization* 38 (2011), Jun, Nr. 4, S. 405–421. – URL https://doi.org/10.1007/s11151-011-9295-7

A Collection of questions for server staff interviews

1. Based on the DC inventory could you identify the servers you own/operate?
2. Which of those servers/storage are going to be decommissioned or replaced during the next year?
3. Do you know of new servers being installed in the data center during the next year?
4. What are the requirements of the new servers from the network?
   (a) Types of interfaces and capacity
   (b) L4-L7 services such as firewalling or load balancing
   (c) QoS
   (d) Interface monitoring
   (e) Others
5. Can any of your assets be moved to a public cloud?
6. Are you planning to move any of your Data Center assets to a public cloud? What is the timeline?
7. What are the main challenges you have with regards to the network in the data center?
   (a) Capacity issues
   (b) Support
   (c) Downtime
   (d) Provisioning
8. How do you think the Data Center will evolve during the next five years?
   (a) Growth on premises
   (b) Move more assets to the cloud
   (c) Growth on-premises and cloud
   (d) More virtualization

9. Are there any features you would like to have that the current data center network does not provide? What are those feature and how they could improve the performance?

B Collection of questions for network staff interviews

1. Are you familiar with the current Data Center Network deployment?
   (a) Hardware
   (b) Software
   (c) Features
   (d) Configuration
2. What are the main challenges you experience with regards to the Data Center Network?
   (a) Difficult to understand
   (b) Provisioning
   (c) Troubleshooting
3. What features do you consider a new data center network should provide?
   (a) Programmability
   (b) Provisioning model
   (c) Capacity
   (d) Multisite
4. Would you feel comfortable with a new Data Center Network approach?
   (a) If the approach deviates from the traditional support and configuration model via CLI
   (b) If some of the functions are automated
   (c) If other teams have access and can configure certain features in the Data Center Network
5. Do you have any background as a programmer of development? Which languages?
6. Would you like to learn to programme?
7. What do you think about configuring the network using API?
8. Would you like to help with the implementation of the new data center network?

9. Are you willing to take a course or certification in the new data center network technology?

C Collection of questions USA network interview

1. What is your current solution for the Data Center Network and when was implemented?
2. Did you run any POCs before the implementation? What were the results and conclusions?
3. Have you had any challenges with your current solution?
4. Do you recommend the same solution for ConocoPhillips Norway?
5. How many persons are qualified to support the Data Center Network Implementation?
6. Could you provide access to the documentation?
7. How was your experience with regards to the integration with Vmware and the application-centric approach?
8. Did you deploy any L4-L7 services in the Data Center network? How did you experience the solution? How does it work?
9. Have you had any challenges with the implementation that affected the availability and reliability of the deployment?
   (a) Software bugs
   (b) Outages
   (c) Security Breaches
   (d) Others
10. How did you do the configuration?
    (a) Everything from the GUI
    (b) Did you use scripts? What type? What did you script?
11. Did you outsource any part of the implementation or required vendor support at any stage of the project? What did you outsource and why?

D Collection of questions for BA's interviews

1. You own (list of systems) in the data center, could you please assess the criticality and the downtime tolerance of those systems?
2. Can you provide a description, function and affected users for each of your systems?
3. Could you take any measures to reduce the downtime of the systems?
   (a) Redundancy
   (b) Moving VMs to other servers
   (c) Migrating at certain time
4. Do you know if any of your systems are going to be decommissioned or moved to the cloud?
5. Have you received any requests, that you were not able to fulfill due to data center limitations? Can you provide an overview of those requests?

E Collection of questions for external expert's interviews

1. What are the differences, advantages, and disadvantages of the following data center implementations:
   (a) Traditional NX-OS
   (b) ACI
   (c) NSX
2. What is the best alternative for the migration phase with regards to:
   (a) Minimize downtime
   (b) Minimize costs
3. What are the advantages of integration with Vmware?
4. Should we consider the application-centric approach? What is the value of that implementation for ConocoPhillips?
5. What are the trends in the market? What are companies like ConocoPhillips doing in the data center front?
6. Validation of the topology and equipment choice
7. What are the requirements for a multisite implementation?

F Life Cycle Cost Analysis

LCC Analyzis for NX-OS Alternative

Cost Element	Subcost Element	Year	0	1	2	3	4	5	6	7	8	9	10
		disc	1	0.95238	0.90703	0.86384	0.8227	0.78353	0.74622	0.71068	0.67684	0.64461	0.61391
Initial Cost	Engineering		57.6	0	0	0	57.6	0	0	0	0	0	0
	Procurement		170	0	0	0	180	0	0	0	0	0	0
Installation			27.2	0	0	0	0	0	0	0	0	0	0
Energy Cost	Electricity		0	12.3582	12.3582	12.3582	12.3582	7.00362	7.00362	7.00362	7.00362	7.00362	7.00362
Operations	Monitoring		0	31.025	31.025	31.025	31.025	31.025	31.025	31.025	31.025	31.025	31.025
	Configuration		0	17.68	17.68	17.68	17.68	17.68	17.68	17.68	17.68	17.68	17.68
Maintenance	Preventive		0	13.6	13.6	13.6	13.6	13.6	13.6	13.6	13.6	13.6	13.6
	Corrective			13.6	13.6	13.6	13.6	13.6	13.6	13.6	13.6	13.6	13.6
Total			**254.8**	**88.2632**	**88.2632**	**88.2632**	**325.863**	**82.9086**	**82.9086**	**82.9086**	**82.9086**	**82.9086**	**82.9086**
NPV - Left Axix 1000USD			254.8	84.0602	80.0573	76.245	268.088	64.9611	61.8677	58.9216	56.1158	53.4436	50.8987
cumulated NPV - Right Axix 1000U			254.8	338.86	418.917	495.163	763.251	828.212	890.08	949.001	1005.12	1058.56	1109.46
TOTAL NPV 1000USD			1109.46										

	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
Initial costs	227.6	0	0	0	195.474	0	0	0	0	0	0	423.074
Installation and commissioning	27.2	0	0	0	0	0	0	0	0	0	0	27.2
Operations and Energy costs	0	58.1554	55.3861	52.7487	50.2368	43.6492	41.5706	39.5911	37.7058	35.9103	34.2003	449.154
Maintenance Costs	0	25.9048	24.6712	23.4964	22.3775	21.3119	20.2971	19.3305	18.41	17.5334	16.6984	210.031
												1109.46

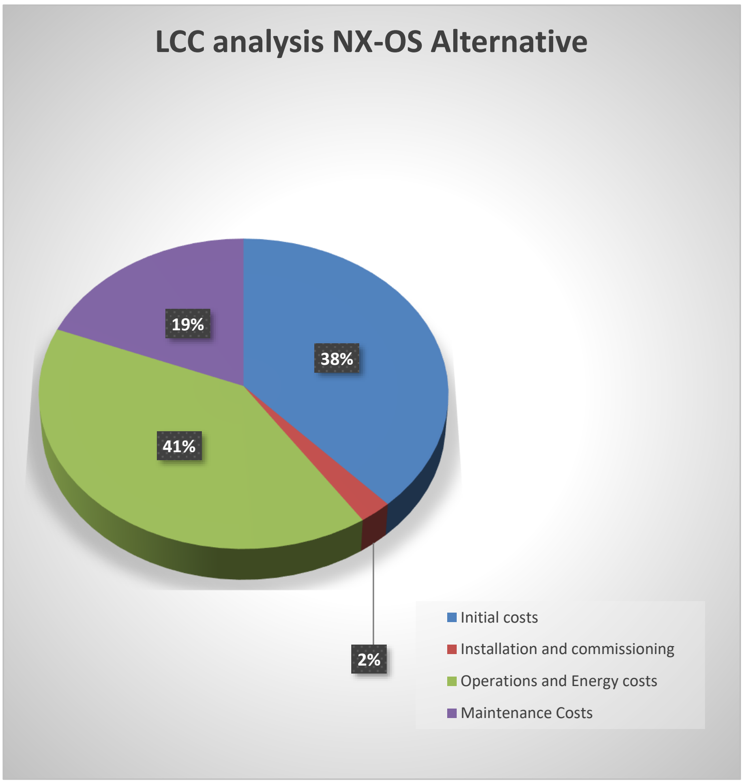

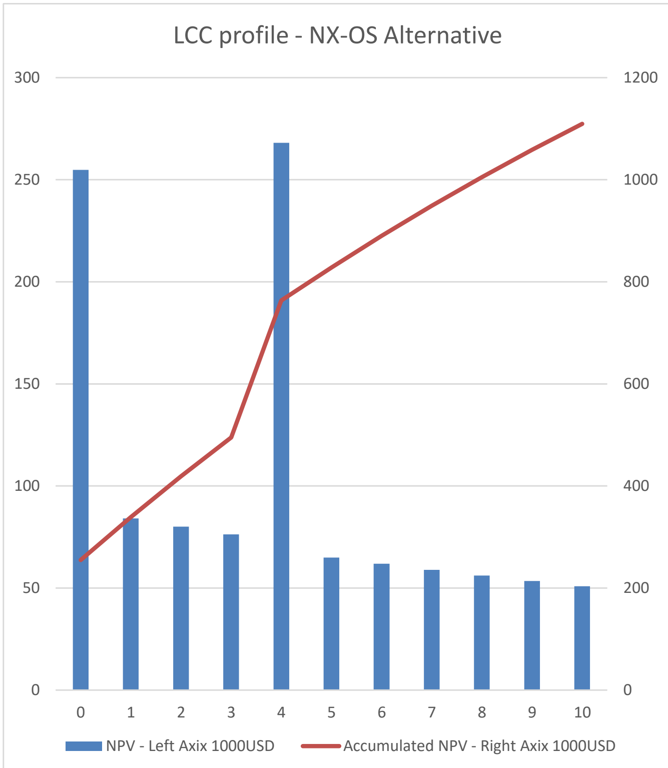

Data and assumptions for LCC calculations		
Lifetime	10	Years
Discount rate	5%	%
Engineering Year 0	320	MHR
Installation & Commissioning Year 0	160	MHR
engineering manhour	180	USD/MHR
Onshore manhour	170	USD/MHR
Procurement Year 0	170	1000 USD
Electricity	0.125	USD/kwh
power consuption	11.286	KW/hour
Power after year 5	6.396	KW/hour
uptime/year	8760	Hours
availability	0.997	Percentage
monitoring	0.5	MHR/day onshore
Configuration	2	MHR/week onshore

Maintenance			
Preventive			Costs/year 1000 USD
6 month checks	40	MHR/6 month in addition to motnhly che	13.6
Corrective	80	MHR/year	13.6

after 4 years (access)		
Engineering	320	MHR
Installation & Commissioning	480	MHR
Procurement	180	1000 USD
Power		

POWER CONSUMPTION YEAR 1-4

Qty	Power Consumption Core			
2	Supervisor	89	178	Watts/hr
12	Fabric Modules	234	2808	Watts/hr
3	Fans	95	285	Watts/hr
1	Line cards	845	845	Watts/hr
		Total	**4116**	Watts/hr

	Power Consumption Access			
10	N5K	660	6600	Watts/hr
		Total	**6600**	Watts/hr

	Power Consumption Fexes			
6	N2K	95	570	Watts/hr

Total Consumption		**11286**	Watts/hr
Total Consumption/year		**17606.16**	KW

POWER CONSUMPTION YEAR 5-10

	Power Consumption Core			
2	Supervisor	89	178	Watts/hr
12	Fabric Modules	234	2808	Watts/hr
3	Fans	95	285	Watts/hr
1	Line cards	845	845	Watts/hr
		Total	**4116**	Watts/hr

	Power Consumption Access			
2	40/100G Leaves	220	440	Watts/hr
6	10G leaves	210	1260	Watts/hr
2	Copper Leaves	290	580	Watts/hr
		Total	**2280**	Watts/hr

Total Consumption		**6396**	Watts/hr
Total Consumption/year		**9977.76**	KW

Equipment Procurement			
Qty	**Equipment**	**unit price**	**total price**
2	spine	85	170
2	leaf 40	20	40
6	leaf 10	14	84
4	leaf copper	14	56
			180

LCC Analyzis for ACI Alternative

Cost Element	Subcost Element	Year	0	1	2	3	4	5	6	7	8	9	10
		disc	1	0.95238	0.90703	0.86384	0.8227	0.78353	0.74622	0.71068	0.67684	0.64461	0.61391
Initial Cost	Engineering		172.8	0	0	0	0	0	0	0	0	0	0
	Procurement		396	0	0	0	0	0	0	0	0	0	0
Installation Cost			163.2	0	0	0	0	0	0	0	0	0	0
Energy Cost	Electricity		0	6.47912	6.47912	6.47912	6.47912	6.47912	6.47912	6.47912	6.47912	6.47912	6.47912
Operations	Monitoring		0	15.5125	15.5125	15.5125	15.5125	15.5125	15.5125	15.5125	15.5125	15.5125	15.5125
	Configuration		0	4.42	4.42	4.42	4.42	4.42	4.42	4.42	4.42	4.42	4.42
Maintenance	Preventive		0	4.08	4.08	4.08	4.08	4.08	4.08	4.08	4.08	4.08	4.08
	Corrective			6.8	6.8	6.8	6.8	6.8	6.8	6.8	6.8	6.8	6.8
Total			**732**	**37.2916**	**37.2916**	**37.2916**	**37.2916**	**37.2916**	**37.2916**	**37.2916**	**37.2916**	**37.2916**	**37.2916**
NPV - Left Axix 1000USD			732	35.5158	33.8246	32.2139	30.6799	29.219	27.8276	26.5025	25.2404	24.0385	22.8938
cumulated NPV - Right Axix 1000U			732	767.516	801.34	833.554	864.234	893.453	921.281	947.783	973.024	997.062	1019.96
TOTAL NPV 1000USD			1019.96										

	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
Initial costs	568.8	0	0	0	0	0	0	0	0	0	0	568.8
Installation and commissioning	163.2	0	0	0	0	0	0	0	0	0	0	163.2
Operations and Energy costs	0	25.1539	23.9561	22.8153	21.7289	20.6942	19.7088	18.7702	17.8764	17.0252	16.2144	203.943
Maintenance Costs	0	10.3619	9.86848	9.39855	8.951	8.52476	8.11882	7.73221	7.36401	7.01335	6.67938	84.0125
												1019.96

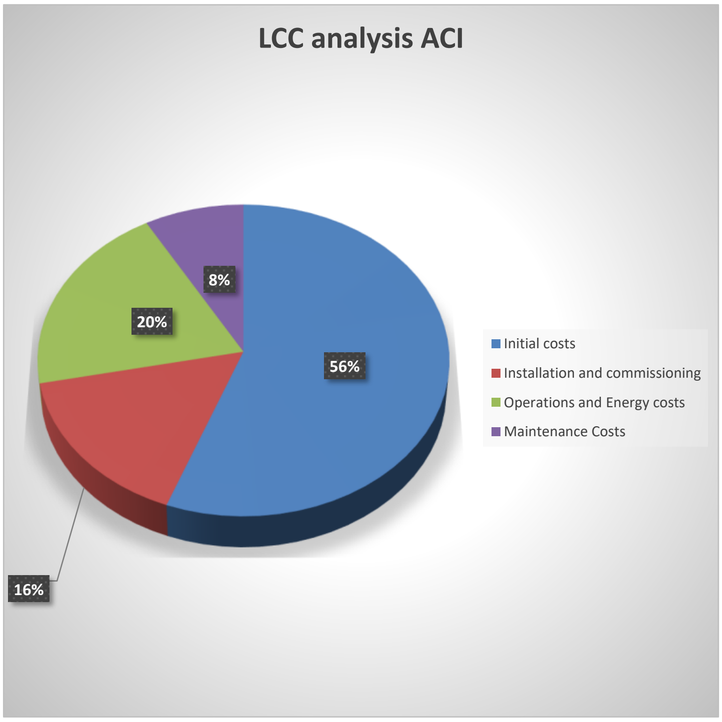

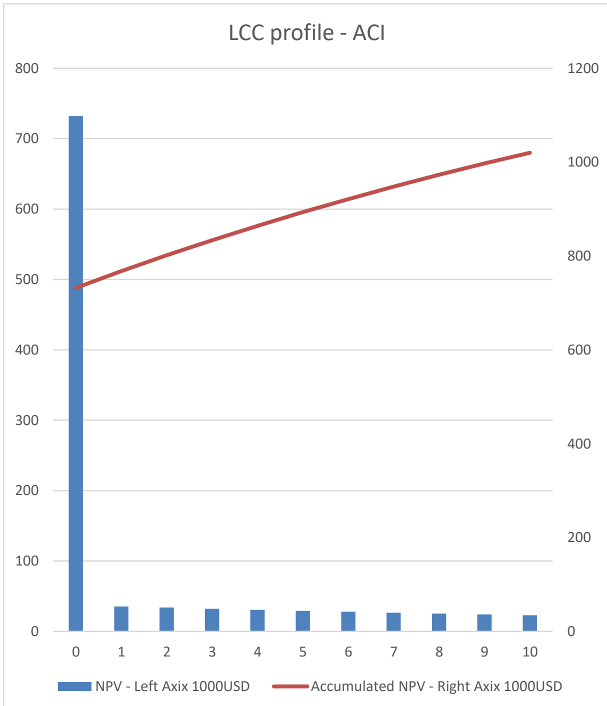

Data and assumptions for LCC calculations		
Lifetime	10	Years
Discount rate	5%	%
Engineering	960	MHR
Installation & Commissioning	960	MHR
engineering manhour	180	USD/MHR
Onshore manhour	170	USD/MHR
Procurement	396	1000 USD
Electricity	0.125	USD/kwh
power consuption	5.917	KW/hour
uptime/year	8760	Hours
monitoring	0.25	MHR/day onshore
Configuration	0.5	MHR/week onshore

Maintenance			
Preventive			Costs/year 1000 USD
6 month checks	12	MHR/6 month	4.08
Total Preventive			**4.08**
Corrective	40	MHR/year	6.8

POWER CONSUMPTION

Qty	**Power Consumption Spines**	Power/u	Total power	
2	Supervisor	75	150	Watts/hr
8	Fabric Modules	234	1872	Watts/hr
3	Fans	95	285	Watts/hr
1	Line cards	430	430	Watts/hr
		Total	**2737**	Watts/hr

	Power Consumption Access			
2	40/100G Leaves	220	440	Watts/hr
6	10G leaves	210	1260	Watts/hr
2	Copper Leaves	290	580	Watts/hr
		Total	**2280**	Watts/hr

	Power Consumption Controller			
3	UCS-M4	300	900	Watts/hr

Total Consumption		**5917**	Watts/hr
Total Consumption/year		**9230.52**	KW

Equipment Procurement

Qty	Equipment	unit price	total price
2	spines	80	160
2	leaf 40	30	60
6	leaf 10	22	132
2	leaf copper	22	44
		total	396

G Cisco Nexus 9500 Switch Series

Cisco 9500 Chassis

The tables included in this overview are taken from:
https://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/models-comparison.html

Nexus 9500 Chassis

	Cisco Nexus 9504 Switch	Cisco Nexus 9508 Switch	Cisco Nexus 9516 Switch
Form factor	7 RU	13 RU	21 RU
Line card slots	4	8	16
Supervisor slots	2	2	2
Fabric module slots	6	6	6
ACI support	Yes	Yes	Yes
Bandwidth per slot (Tbps)	3.84	3.84	3.84
Bandwidth per system (Tbps)	15	30	60
Maximum number of 1/10G BASE-T ports	192	384	768
Maximum number of 10 GE ports	576	1152	2304
Maximum number of 40 GE ports	144	288	576
Maximum number of 100GE ports	144	288	576
Airflow	Front to back	Front to back	Front to back
Power supplies (3-kW AC/DC)	Up to 4	Up to 8	Up to 10
Fan trays	3	3	3

Nexus 9500 Supervisors

	Supervisor A	Supervisor B	Supervisor A+	Supervisor B+
Processor	4 core, 1.8 GHz CPU	6 core, 2.2 GHz CPU	4-Core/8-Thread 1.8 GHz	6-Core/12-Thread 1.9 GHz
Memory (RAM)	16 GB	24 GB	16 GB	24 GB
Storage (SSD)	64 GB	256 GB	64 GB	256 GB

Nexus 9500 100GE Modules

SKU	N9K-X9736C-FX	N9K-X9732C-EX	N9K-X9732C-FX	N9K-X9736C-EX	N9K-X9636C-RX	N9K-X9636C-R	N9K-X9432C-S
Front Panel Ports	36 x 100GE	32 x 100 GE	32 x 100 GE	36 x 100 GE	36 x 100 GE	36 x 100 GE	32 x 100 GE
Multispeed Ports	1x1GE, 4x10GE, 4x25GE, 1x40GE, 2x50GE	1x1GE, 4x10GE, 4x25GE, 1x40GE, 2x50GE	1x1GE, 4x10GE, 4x25GE, 1x40GE, 2x50GE	1x1GE, 4x10GE, 4x25GE, 1x40GE, 2x50GE	4x10GE, 1x40GE , 1x100GE	4x10GE, 1x40GE , 1x100GE	1x1GE, 4x10GE, 4x25GE, 1x40GE, 2x50GE
NX-OS Mode	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ACI Spine	Yes	Yes	Future	No	No	No	No
ACI Leaf	No	No	No	No	No	No	No
Chassis Support	Nexus 9504, 9508, and 9516#	Nexus 9504, 9508, and 9516	Nexus 9504, 9508, and 9516	Nexus 9504, 9508, and 9516	Nexus 9504 and 9508	Nexus 9504 and 9508	Nexus 9504 and 9508
Switching Performance	3.6Tbps, Line rate for 170+ byte packet sizes	3.2Tbps, Linerate for 64+ byte packet sizes	3.2Tbps, Line rate for 170+ byte packet sizes	3.2Tbps, Linerate for 64+ byte packet sizes	3.6Tbps, Line rate for 120+ byte packet sizes	3.6Tbps, Line rate for 92+ byte packet sizes	Line rate greater than 250 Bytes
FM Required for Max Bandwidth	X9736C-FX: 5x N9K-C95xx-FM-E/E2	X9732C-EX: 4x N9K-C95xx-FM-E/E2	X9732C-FX: 4x N9K-C95xx-FM-E/E2	X9736C-EX: 4x N9K-C95xx-FM-E/E2	X9636C-RX: 6x N9K-C950x-FM-R	X9636C-RX: 5x N9K-C950x-FM-R	X9432C-S: 4x N9K-C950x-FM-S
Fabric Redundancy	Bandwidth is spread across all FMs	Bandwidth is spread across all FMs	Bandwidth is spread across all FMs N+1 FM redundancy with 5 FM-E/E2 FMx	Bandwidth is spread across all FMs	No	Bandwidth is spread across all FMs N+1 FM redundancy with 5 FM-R FMs	Bandwidth is spread across all FMs
Interface type	QSFP28	QSFP28	QSFP28	QSFP28	QSFP28	QSFP28	QSFP28
Buffer	160MB	160MB	160MB	160MB	16GB	24GB	32MB
Line Rate Encryption	Yes	No	Yes	No	No	No	No
Streaming analytics export off chip	Yes	No	Yes	No	No	No	No

Requires N9K-C9515-FM-E2 fabric module for ACI spine.

Nexus 9500 40GE Modules

	9636PQ	X9636Q-R	9536PQ	9432PQ	9736PQ
Usage	40 GE high-performance aggregation and FEX*	40 GE aggregation	40 GE aggregation and FEX*	40 GE aggregation and FEX*	ACI spine
Front-panel ports	36 x 40 GE	36 x 40 GE	36 x 40 GE	32 x 40 GE	36 x 40 GE
Standalone	Yes	Yes	Yes	Yes	No
ACI leaf	No	No	No	No	No
ACI spine	No	No	No	No	Yes
Chassis support	4 and 8, slot	8 slot	4, 8, and 16 slot	4, 8, and 16 slot	4, 8, and 16
40 GE ports	36	36	36	32	36
Switching performance	Nonblocking	Nonblocking	1.5:1 oversubscribed	Line rate more than 200 byte	Nonblocking
Buffer (MB)	36	12GB	104	24	42
Typical power (watts)	260	329	360	240	197
Minimum software version	6.1(2)I1(1)	7.0(3)F1(1)	6.1(2)I2(2a)	6.1(2)I2(2a)	ACI-N9KDK9-11.0
SKUs	N9K-X9636PQ	N9K-X9636Q-R	N9K-X9536PQ	N9K-X9432PQ	N9K-X9736PQ

Nexus 9500 10GE Modules

	N9K-X97160YC-EX	N9K-X9788TC-FX	N9K-X9564PX	N9K-X9464PX	N9K-X9564TX	N9K-X9464TX2
Front Panel Ports	48 x 10/25 GE + 4 x 40/100G GE	48 x 1/10 GBaseT + 4 x 40/100G GE	48 x 10 GE + 4 x 40 GE	48 x 10 GE + 4 x 40 GE	48 x 10 GE + 4 x 40 GE	48 x 10 GE + 4 x 40 GE
NX-OS Mode	Yes	Yes	Yes	Yes	Yes	Yes
ACI Spine	No	No	No	No	No	No
ACI Leaf	No	No	No	No	No	No
Chassis Support	Nexus 9504, 9508, and 9516	Nexus 9504, 9508, and 9516	Nexus 9504, 9508, and 9516	Nexus 9504, 9508, and 9516	Nexus 9504, 9508, and 9516	Nexus 9504, 9508, and 9516
100 M/1 GE/10 GE copper ports	NA	48	NA	NA	48	48
1 GE Fiber downlink fiber ports	48	NA	48	48	NA	NA
10 GE downlink fiber ports	48	NA	48	48	NA	NA
25 GE downlink fiber ports	48	NA	NA	NA	NA	NA
40 GE uplink ports	4	4	4	4	4	4
100 GE uplink ports	4	4	NA	NA	NA	NA
Switching Performance	1.6 Tbps	880 Gbps	640 Gbps	640 Gbps, Line rate for 200+ byte packets	640 Gbps	640 Gbps, Line rate for 200+ byte packets
Buffer	40MB	40MB	104MB	12MB	104MB	12MB

H Cisco Nexus 9300 Switch Series

Cisco 9300 Switch Series

The tables included in this overview are taken from:
https://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/models-comparison.html

Nexus 9300 400 GE Switches

	Nexus 9316D-GX	Nexus 93600CD-GX
Usage	ACI Spine & Leaf*	ACI Spine* & Leaf
Form factor	1 RU	1 RU
Throughput (Tbps)	6.4 Tbps	6.0 Tbps
40 GE ports	16	28
100 GE ports	16	28
400 GE ports	16	8
Latency (microseconds)	TBD	TBD
Buffer (MB)	80 MB	80 MB
Power supply (watts) 2 per switch	1100 W	1100 W
Typical power (watts)	800 W	590 W
Minimum software version	TBD	TBD
Orderable	Q12019	Q12019
SKU	N9K-C9316D-GX	N9K-C93600CD-GX
Operating System	NXOS/ACI	NXOS/ACI
*** Roadmap**		

Nexus 9300 1/10GBaseT Switches

	Nexus 93120TX	Nexus 93108TC-EX	Nexus 9348GC-FXP	Nexus 93108TC-FX	Nexus 93216TC-FX2
Usage	ACI leaf, top of rack, end of row	ACI leaf, top of rack, powered by Cloud Scale	ACI leaf, top of rack, powered by Cloud Scale	ACI leaf, top of rack, powered by Cloud Scale	ACI leaf, top of rack, powered by Cloud Scale
Form factor	2 RU	1 RU	1RU	1RU	2 RU
Throughput (Tbps)	2.4	2.16	0.696	2.16	7.2
100 M copper ports	96	48	48	48	96
1 GE copper ports	96	48	48	48	96
10 GE copper ports	96	48	-	48	96
10/25 GE ports			4		None
40 GE ports	6	6	2	6	12
100 GE ports	No	6	2	6	12
Latency (microseconds)	~4	~2.5	~2.5	~2.5	~1
Buffer (MB)	52	40	40	40	40
Power supply (watts); 2 per switch	1200	650	350	500	1200
Typical power (watts)	542	290	178	276	210
Minimum software version	NXOS-703I2.2A/ACI-N9KDK9-11.2	NXOS-703I4.3/ACI-N9KDK9-12.0	NXOS-703I7.1/ACI-N9KDK9-13.0	NXOS-703I7.1/ACI-N9KDK9-12.2A	NXOS-9.3.1/ACI-14.1.2/4.1.2
Orderable	Yes	Yes	Yes	Yes	Yes
SKU	N9K-C93120TX	N9K-C93108TC-EX	N9K-C9348GC-FXP	N9K-C93108TC-FX	N9K-C93216TC-FX2
Operating System	NX-OS, ACI	NX-OS, ACI	NX-OS, ACI	NX-OS, ACI	NX-OS, ACI

Nexus 9300 1/10/25GE Fiber Switches

	Nexus 93180YC-EX	Nexus 93180YC-FX	Nexus 93240YC-FX2	Nexus 92160YC-X	Nexus 93360YC-FX2
Usage	ACI leaf, top of rack, FEX aggregation, powered by Cloud Scale	ACI leaf, top of rack, FEX aggregation, powered by Cloud Scale	ACI leaf, top of rack, FEX aggregation, powered by Cloud Scale	Top of rack	ACI leaf, top of rack, FEX aggregation, powered by Cloud Scale
Form factor	1 RU	1 RU	1.2	1 RU	2 RU
Throughput (Tbps)	3.6	3.6	4.8	3.2	7.2
1 GE ports	48	48	48	48	96
10 GE Fiber ports	48	48	48	48	96
25 GE ports	48	48	48	48	96
100 M/1 GE/10 GE copper ports	NA	NA	NA	NA	None
40 GE ports	6	6	12	6	12
100 GE ports	6	6	12	4	12
Latency (microseconds)	~1	~1	~1	Less than 2	~1
Buffer (MB)	40	40	40	20	40
Power supply (watts); 2 per switch	650	500	1100	650	1200
Typical power (watts)	210	260	298W	10-Gb - 150, 25-Gb - 170	210
Minimum software version	NXOS-703I4.2/ACI-N9KDK9-11.3	NXOS-703I7.1/ACI-N9KDK9-12.2A	NXOS-703I7.3/ACI-N9KDK9-14.0	NXOS-703I3.1	NXOS-9.3.1/ACI-14.1.2/4.1.2
Orderable	Yes	Yes	Yes	Yes	Yes
SKU	N9K-C93180YC-EX	N9K-C93180YC-FX	N9K-C93240YC-FX2	N9K-C92160YC-X	N9K-C93360YC-FX2
Operating System	NX-OS, ACI	NX-OS, ACI	NX-OS, ACI	NX-OS	NX-OS, ACI

Nexus 9300 40/100 GE Switches

	Nexus 9364C	Nexus 9336C-FX2	Nexus 9332C
Usage	ACI spine, powered by Cloud Scale	ACI leaf, powered by Cloud Scale	ACI spine, powered by Cloud Scale
Form factor	2 RU	1 RU	1 RU
100 M/1 GE/10 GE copper ports	NA	NA	NA
1/10 GE fiber ports	2	NA	2
40 GE ports	64	36	32
100 GE ports	64	36	32
Latency (microseconds)	~1.3	~1.3	~1.3
MACsec & Cloudsec	Yes on last 16 ports	Yes, all ports	Yes, on last 8 ports
Buffer (MB)	40	40	40
Power supply (watts); 2 per switch	1200	1100	1100
Typical power (watts)	429	337	296
Minimum software version	NXOS-703I7.2/ACI-N9KDK9-13	ACI-N9KDK9-13.1.2/ NXOS-703I7.3	ACI-N9KDK9-14.0/NX-OS Release 9.2(3)
Orderable	Yes	Yes	Yes
SKU	N9K-C9364C	N9K-C933C-FX2	N9K-C9332C
Operating System	NX-OS, ACI	NX-OS, ACI	NX-OS, ACI