EURESCOM

Study Report

Network Virtualisation – Opportunities and Challenges

D1: Network Virtualisation – opportunities and challenges for operators

Editor: Jorge Carapinha, Portugal Telecom Inovação

Suggested readers

Business units and their subsidiaries may use results and related recommendations to launch exploitation projects leading to services offers or to establish sourcing strategies, leveraging telecom networks and data centres.

Abstract

In the last few years, the concept of network virtualisation has gained a lot of attention both from research and industry communities. Network virtualisation was initially promoted as an enabler of Internet technological diversity, as well as a solution to overcome the obstacles to novelty and innovation that currently afflict the Internet evolution. Furthermore, for telecom operators, it has become clear that the potential of network virtualisation, both from economical and operational viewpoints, can be also quite relevant in multiple scenarios. The widespread adoption of virtualisation technologies in IT environments and trends like cloud computing have contributed to stimulate the interest in network virtualisation. However, it is also clear that the challenges posed by the strict requirements of carrier-grade commercial environments require thorough investigation.

This report evaluates the potential of network virtualisation from an operator's perspective, with the short-term goal of optimising service delivery and rollout, and on a longer term as an enabler of technology integration and migration. Based on possible scenarios for implementing and using network virtualisation, new business roles and models are examined. Open issues and topics for further evaluation are identified. In summary, the objective is to identify challenges but also new opportunities for operators raised by network virtualisation.

EDIN **0589-1956**
Study **P1956**
For full publication
December 2010

Eurescom participants in study P1956 are:

- Portugal Telecom Inovação
- Deutsche Telekom AG
- Síminn hf. (Iceland Telecom Ltd.)
- Türk Telekom A.Ş.

Network Virtualisation – Opportunities and Challenges
Deliverable 1: Network Virtualisation – Opportunities and Challenges for operators
Editor: Jorge Carapinha, Portugal Telecom Inovação
Study leader: Jorge Carapinha, Portugal Telecom Inovação
Study supervisor: Ádám Kapovits, Eurescom
Eurescom published study result; EDIN 0589-1956

Preface

In the last few years, significant research activities have been launched in the area of network virtualisation. In Europe, in the framework of FP7, projects such as 4WARD, Reservoir and Federica have focused in some way on network virtualisation issues. In the US, GENI, CABO and VINI initiatives have been active in this area, as well. However, these initiatives have focused mainly on architectural frameworks and on the exploration of virtualisation as a key tool to enable future Internet architectures. An operator-centric evaluation of virtualisation and a roadmap for network virtualisation deployment by operators are still largely missing.

Standardization in this field is at a very incipient stage. The "Focus Group on Future Networks" was set up to collect and identify visions of future networks. One of the deliverables to be produced by the Focus Group will be a framework of network virtualisation. In addition, an IRTF Network Virtualization Research Group is currently being setup, with a draft charter under preparation.

In summary, the full impact of the changes enabled by network virtualisation has not been fully understood yet, but it is clear that both opportunities and hurdles lie ahead for network operators. The deployment of network virtualisation imposes new requirements and raises new challenges in relation to how networks are provisioned, managed and controlled today.

Thus the study set out to deliver on the following goals:

- Assess the real potential of network virtualisation from a network operator perspective in the short/medium term, namely as a tool for optimal service delivery and a service rollout facilitator;
- Evaluate network virtualisation in the medium/long term, mainly as an enabler of technological diversity and a migration tool to new Internet architectures and network technologies;
- Describe possible scenarios for network virtualisation deployment, both in short and long term time scales;
- Analyse new business roles and new business models emerging from network virtualisation;
- Evaluate interoperability issues and identify areas requiring standardisation;
- Outline a roadmap leading to the adoption of network virtualisation by network operators.

The study started at the end of November December 2009 with the participation of Deutsche Telekom, Portugal Telecom Inovação and Síminn hf. under the leadership of Jorge Carapinha from Portugal Telecom Inovação. Türk Telekom A.Ş joined in March 2010.

In addition to this deliverable the study will issue another deliverable (D2) with identical title but in the form of a set of presentation slides.

Executive Summary

Network virtualisation (NV) is not a new concept in the telecommunications world but has been seen in a narrow range of applications, notably Virtual Private Networks. Developments now underway are likely to bring the concept to new heights. It is conceivable that a substantial part of the future's telecommunications networks may be built by using virtualisation technologies. An analogy may be drawn up between network virtualisation and cloud computing, which is heavily based on virtualisation technologies. Cloud computing and related concepts are about to radically transform the IT environment from a very distributed resources paradigm to one where resources are centralised and can be shared among a number of users accessing them through networking. A network is a distributed phenomenon in its nature and therefore networking resources are not centralised. However they are becoming so powerful that they can be shared amongst a number of virtual networks on top of a certain infrastructure.

NV in this broad context is still at a research stage and has been investigated within a number of research projects that are described in this report. They indicate that NV will bring about much needed advantages for the telecommunications industry, such as reduced OPEX and CAPEX, more dynamic and flexible service provisioning and constitute a basis for a range of new methodologies and services. Some of the possible scenarios enabled by NV are described in this report including cloud computing access, network as a service, experimentation and technology migration. An example of NV's important role is through the migration to IPv6 which can be introduced in virtual networks and expanded gradually. Gaps and open issues of NV are also treated in the report. Issues like carrier grade compliance, isolation between virtual networks, and security, to name a few, must be solved before NV becomes an adopted technical solution. Standardisation is needed for further advancement of NV. Preparatory efforts in this field have commenced under the auspices of ITU-T and IRTF/IETF. The report gives a detailed analysis of gaps and open issues of NV, from a technical, operational, business and regulatory point of view. An analysis of NV with regard to opportunities and challenges for operators is further given in the report. An exciting opportunity could emerge by developments such as Stanford's OpenFlow protocol which has the potential to convert the networking infrastructure environment from today's integrated solutions towards an open source business model with lower CAPEX, increased openness to innovation and smaller likelihood of vendor lock-in.

By actively participating in NV R&D projects and standardisation efforts, telcos can have a significant influence on the development and uptake of this exciting technology. This report gives a good starting point to learn about the new concepts and methodologies associated with NV and to ponder over the new business models and service scenarios enabled by NV.

For the long term, this study has reached three main conclusions. The telecom industry will further converge with the IT industry. This development will be powered by NV and is expected to result in a better economy for both industries. NV will help to address the main problems facing cloud computing at present, i.e. security issues, by providing flexible networking solutions that can offer good isolation. NV will be indispensable to meet demands set for the Future Internet in an economical manner.

List of Authors

Jorge Carapinha, Portugal Telecom Inovação
Peter Feil, Deutsche Telekom AG
Paul Weissmann, Deutsche Telekom AG
Saemundur E. Thorsteinsson, Síminn hf.
Márcio Melo, Portugal Telecom Inovação
Çağrı Etemoğlu, Türk Telekom A.Ş.
Ólafur Ingþórsson, Síminn hf.
Selami Çiftçi, Türk Telekom A.Ş.

Table of Contents

List of Figures

List of Tables

Abbreviations and Acronyms

AAA	Authentication, Authorisation, Accounting
AGAVE	A liGhtweight Approach for Viable End-to-end IP-based QoS Services
Amazon EC2	Amazon Elastic Computer Cloud
Amazon S3	Amazon Simple Storage Service
API	Application Programme Interface
ATM	Asynchronous Transfer Mode
B-DA	Backbone Destination Address
BGP	Boarder Gateway Protocol
BIOS	Basic Input Output System
BoF	Best of Friends
BRAS	Broadband Remote Access Server
B-SA	Backbone Source Address
BT	British Telecom
B-VID	Backbone VLAN ID
CABO	Concurrent Architectures are Better than One
CAPEX	Capital Expenditure
CDM	Code Division Multiplexing
CDN	Content Delivery Network
CPU	Central Processing Unit
CRM	Customer Relationship Management
CRS	Cisco Carrier Routing System
DC	Data Centre
DiffServ	Differentiated Services
DLCI	Data Link Connection Identifier
DPI	Deep Packet Inspection
DSLAM	Digital Subscriber Line Access Multiplexer
DWDM	Dense Wavelength Division Multiplexing
ERP	Enterprise Resource Planning
ESX	Enterprise-level virtualisation product offered by VMware, Inc
FDM	Frequency Division Multiplexing
FEDERICA	Federated E-infrastructure Dedicated to European Researchers Innovating in Computing network Architectures
FG	Focus Group
FP7	Framework Programme 7
FTP	File Transfer Protocol
FTTC	Fibre to the Curb/Cabinet
FTTH	Fibre to the Home
GENI	Global Environment for Network Innovations

G-LAB	German-Lab
GMC	GENI Management Core
GQL	SQL-like language for retrieving entities or keys from the Google App Engine scalable datastore
HP	Hewlett Packard
HTTP	HyperText Transfer Protocol
I/O	Input/output
IaaS	Infrastructure as a Service
IBM	International Business Machines
IEEE	Institute of Electrical and Electronic Engineers
IETF	Internet Engineering Task Force
InD	Infrastructure Division
InP	Infrastructure Network Provider
INTAREA	Internet Area
Intserv	Integrated Services
IP	Internet Protocol
IPTV	IP Television
IPvN	Internet Protocol version N
IRTF	Internet Research Task Force
ISP	Internet Service Provider
IT	Information Technology
ITU-T	International Telecommunication Union-Telecommunication
L2VPN	Layer 2 Virtual Private Network
L3VPN	Layer 3 Virtual Private Network
LER	Label Edge Router
LISP	Location and Identity Separation Protocol
MAC	Medium Access Control
MPLS	Multi Protocol Label Switching
MPLS-TP	MPLS Transport Profile
MSAN	Multi-Service Access Node
MVNO	Mobile Virtual Network Operator
NaaS	Network as a Service
NEC	Nippon Electric Corporation
NIC	Network Interface Card
NICT	National Institute of Information and Communications Technology
NSF	National Science Foundation
NV	Network Virtualisation
NVLAB	the Network Virtualisation Research Lab
OAM	Operation Administration Maintenance
OCS	Optical Circuit Switching

ODU	Optical channel Data Unit
OGF	Open Grid Forum
OpenVZ	An operating system-level virtualisation technology based on the Linux kernel and operating system
OPEX	Operational Expenditure
OS	Operating System
OTN	Optical Transport Network
PBB-TE	Provider Backbone Bridge Traffic Engineering
PLC	PlanetLab Central
PNP	Physical Network Provider
POP	Point of Presence
PPVPN	Provider Provisioned Virtual Private Network
QoS	Quality of Service
Quagga	A routing software suite for UNIX platforms
REST	Representational State Transfer
RFC	Request for Comments
Rx	Receiver
SDH	Synchronous Data Hierarchy
SDK	Software Development Kit
SDM	Space Division Multiplexing
SDR	Secure Domain Routers
SOA	Service Oriented Architecture
SOAP	Simple Object Access Protocol
SONET	Synchronous Optical Network
SP	Service Provider
SQL	Structured Query Language
TaaS	Telco as a Service
TCP	Transmission Control Protocol
TDM	Time Division Multiplexing
T-MPLS	Transport MPLS
TRILL	Transparent Interconnection of Lots of Links
Tx	Transmitter
UCLP	User Controlled Light path
UML	Unified Modelling Language
URL	Uniform Resource Locator
VCI	Virtual Circuit Identifier
VIf-ID	Virtual Interface Identity
VLAN	Virtual Local Area Network
VM	Virtual Machine
VN	Virtual Network

VNC	Virtual Network Customer
VNet-ID	Virtual Network Identity
VNIC	Virtual Network Interface Card
VNO	Virtual Network Operator
VNode-ID	Virtual Node Identity
VNP	Virtual Network Provider
VNRG	Virtual Networks Research Group
VoD	Video on Demand
VoIP	Voice over Internet Protocol
VPI	Virtual Path Identifier
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding
VRRP	Virtual Router Redundancy Protocol
WAN	Wide Area Network
Xen	Open source industry standard for virtualisation
XORP	eXtensible Open Router Platform
xVM	Product group from Sun Microsystems that addresses virtualisation technology on x86 platforms

Definitions

End user[**]	User of the service that is offered by the SP (or directly by the VNO in the cases where a distinct SP does not exist as such). End user nodes are not part of the virtual network topology but are attached like leaves. End users typically have to authenticate themselves towards the VNO by means of correct credentials.
Infrastructure Provider (InP)[**]	Entity that owns, controls and administers physical resources, which may be used, or offered for leasing to third parties, to build custom-tailored VNs.
Network virtualisation (NV) [*]	Networking environment that allows one or multiple service providers to compose (in a dynamic or static way) multiple heterogeneous virtual networks that co-exist together in isolation from each other and to deploy customised end-to-end services on-the-fly, as well as manage them on those virtual networks for the end-users by effectively sharing and utilising underlying network resources leased from one or multiple infrastructure providers.
Service Provider (SP)	Entity responsible for providing services to end users. Network virtualisation is not supposed to be visible from the SP perspective. In some cases, the role of service provider may overlap with the VNO, but from a functional viewpoint they should be defined as different entities.
Substrate[**]	Physical resources (typically, network nodes and links) that are owned, controlled and administered by infrastructure providers and may be virtualised to build virtual networks. Not all virtual nodes have to be virtualized, in which case network virtualisation may still be supported by means of virtual links.
Virtual Network (VN)[**]	Running instance of a slice. This implies configured and active virtual nodes as well as virtual links that are potentially in use.
Virtual Network Customer (VNC)	Customer of the VN, i.e. the entity that holds a commercial relationship with the VNO. Depending on the specific use case, it may correspond to the end user or a third party service provider.
Virtual Network Operator (VNO)[**]	Entity in charge of establishing, managing and operating VNs, as well as handling end user attachment.
Virtual Network Provider (VNP)[**]	Entity in charge of assembling a virtual network, according to a given description. The VNP composes a VN slice based on resources from one or more infrastructure providers.
Virtual Network Slice[**]	The set of reserved resources (e.g. virtual nodes and links) which belong to a virtual network. VN slices are typically reserved and assembled (but not used) by VNPs.
Virtual Private Network (VPN)[***]	Generic term that covers the use of public or private networks to create groups of users that are separated from other network users and that may communicate among them as if they were on a private network. There are two basic types of VPN: CE-based VPN, in which the shared service provider network does not have any knowledge of the VPN and all the VPN-specific procedures are performed in the Customer Edge devices (CE); PE-Based VPN, in which the service provider network is used to interconnect customer sites using shared resources and the Provider Edge device (PE) maintains VPN state, isolating users of different VPNs.

[*]Adapted from [2]. [**]Adapted from [7]. [***]Adapted from [32].

1 Introduction

Virtualisation is a potential enabler of profound changes, both in the IT and communications domains, and is expected to bridge the gap between these two worlds, traditionally living quite apart. In particular, the combination of cloud computing and network virtualisation is likely to open up an immense field of opportunities for network operators. Virtualisation of computational and storage resources is already commonplace in operational environments, but bringing virtualisation to networks has proven to be a lot more challenging.

In spite of the significant potential of the concept, it is clear that deployment of network virtualisation requires overcoming major obstacles. Network virtualisation has followed the usual development cycle, starting with research and testbed experimentation. Evaluation of the concept with a view to deployment in carrier-grade commercial operator environments is still largely unaccomplished.

The network virtualisation concept is not entirely new – to some extent, network-based Virtual Private Networks (VPNs), which are essentially separate networks sharing a common infrastructure, can also be seen as a materialisation of this idea. However, VPNs cannot be decoupled from the underlying infrastructure and should be seen more as a service, rather than a real network. VPN customers have access to something they perceive as a network cloud interconnecting their private network domains, but without any control or even visibility to the protocols running inside the VPN cloud.

Virtual networks (VN) offer a full separation and independence from the underlying infrastructure. By definition, VNs are technology agnostic, in the sense that VNs sharing a common physical infrastructure may internally run different protocols, decoupled from the infrastructure layers below. The physical infrastructure may be stretched across multiple administrative domains, i.e. may be owned by multiple infrastructure providers. The foundation for the VN is the physical infrastructure, consisting of links (e.g. fibres, copper lines, wireless connections) and nodes (e.g. routers, switches, servers and respective interfaces). The physical resources are virtualised through partitioning or slicing, thereby making the resources available to a number of isolated VNs.

The objective of this study was to evaluate the potential of network virtualisation from an operator's perspective, with the short-term goal of optimising service delivery and rollout, and on a longer term as an enabler of technology integration and migration. Based on possible scenarios for implementing and using network virtualisation, new business roles and models were examined. Challenges, open issues and topics for further evaluation were identified.

The document is structured as follows: After this brief introduction, chapter 2 contains an extensive overview on the state-of-the-art of network virtualisation. Following the description of existing concepts, important research projects and initiatives in this area are presented. Also, an overview of current commercial product offerings and standardisation activities is given. Chapter 3 details various relevant scenarios, which can benefit from network virtualisation. Chapters 4 and 5 finally identify open issues and gaps, pointing towards the way forward for operators with the challenges and opportunities network virtualisation is expected to bring.

2 State of the Art

The concept of virtual networks has been known to the telecommunications industry for a number of years, ever since VPNs were introduced. A VPN is typically built by tunnelling through the Internet, building a structure that is private in the eyes of the user but utilises a public ubiquitous network such as the Internet.

The term "virtualisation" is used in a multiple context and its meaning varies between disciplines. In computing, virtualisation usually means that a server is programmed to emulate a number of computers. A user can be allocated a single Virtual Machine (VM) which appears to him as a standalone computer. The software used to bring about the virtualisation is called "hypervisor" [1].

Network virtualisation can be defined in the following way, based on the definition given in [2]:

> Network virtualisation is a networking environment that allows one or multiple service providers to compose (in a dynamic or static way) multiple heterogeneous virtual networks that co-exist together in isolation from each other and to deploy customised end-to-end services on-the-fly as well as manage them on those virtual networks for the end-users by effectively sharing and utilising underlying network resources leased from one or multiple infrastructure providers.

Generally, network virtualisation is based on the following attributes [3]:

- Abstraction: Details of the network hardware are hidden
- Indirection: Indirect access to network elements, network nodes may be combined to form different virtual network topologies.
- Resource sharing: Network elements can be partitioned and utilised by multiple virtual networks
- Isolation: Loose or strict isolation between virtual networks must be provided

Even within the telecommunications community, many networking technologies have been termed virtualised, e.g. ATM, MPLS and DWDM. In fact, those are examples of link virtualisation, a necessary component of a Virtual Network. Another primary component is node virtualisation, where the network nodes, notably routers are virtualised based on isolation and partitioning of hardware resources [4]. Virtual switches may also appear, e.g. on virtual nodes within computer clouds where a virtual switch connects a number of virtual machines [1].

2.1 Concepts and Terminology

2.1.1 General Architecture

In Figure 1a virtualised network environment is depicted [3]. The foundation for the virtual network is of course the physical infrastructure, consisting of links and nodes. The links may be fibres, copper lines or wireless connections and networking equipment like routers, switches, servers and interfaces is located at the nodes. The physical resources are virtualised through partitioning or slicing, thereby making the resources available to a number of virtual networks. The nodes become virtual nodes and the links become virtual links through the slicing procedure. The virtualisation concept is technology agnostic, which means that the virtual networks built on the physical infrastructure may run different protocols, e.g. ATM, IPv4 and IPv6, given of course that the underlying hardware permits. The physical infrastructure may be stretched across many administrative domains.

On top of the physical infrastructure is the so-called virtualised substrate where the sliced resources are available for building up virtual networks. These building blocks may be arranged in a number of ways, a virtual node terminates one or many virtual links. The nodes can be classified according to their role in the network into edge nodes, core nodes and border nodes. The edge nodes are usually the points-of-presence of the infrastructure provider and connect the end users. The core nodes are the main forwarding engines of the network and the border nodes are the nodes on the borders between different network domains [4].

At the top in Figure 1 are the virtual networks. Those can be setup and torn down dynamically according to customer needs, the customers don't need to know how the underlying infrastructure is

built or what network domains it crosses. Additionally, network virtualisation helps to enhance the utilisation of the underlying physical infrastructure and the virtual networks need not run the same protocols. This can e.g. considerably ease the introduction of new technologies – an example could be the deployment of IPv6 into today's Internet.

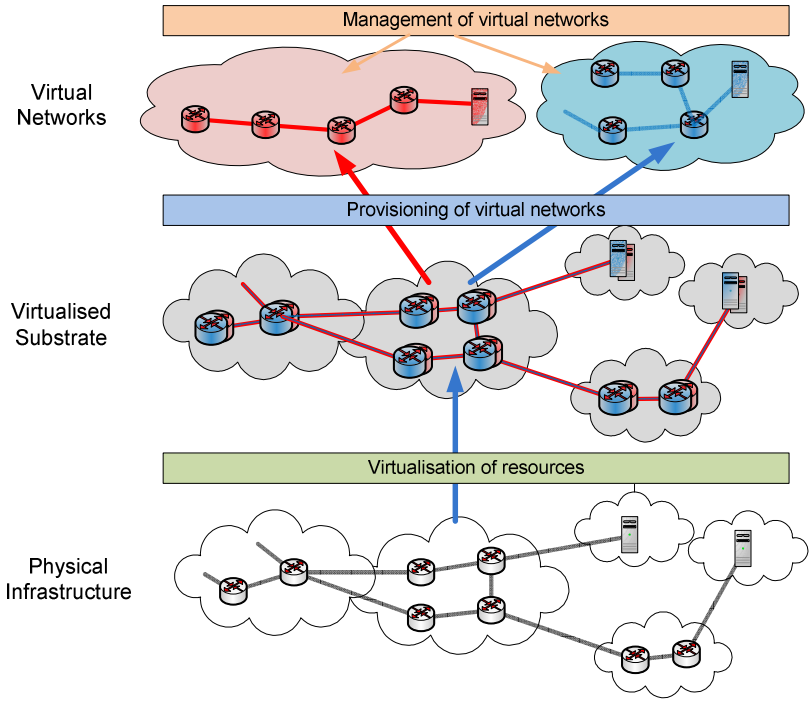

Figure 1 – Virtual Network environment and basic architecture

2.1.2 Roles and players

Figure 1 clearly shows that network virtualisation architecture is divided into three separated layers. This functional separation is a major evolution in relation to the classical monolithic approach followed in operator networks and paves the way to the definition of new business roles and players. However, one should clearly distinguish between functional roles that derive from the layered nature of the network virtualisation architecture and the business models that can be exploited based on that architecture.

There are different opinions on how to model the virtual networking business environment. In the traditional vertical telco business model, the same entity owns the physical network infrastructure, like fibres, DWDM- and SDH-systems, as well as the IP-networking infrastructure, routers and switches. In an environment based on network virtualisation, this is not necessarily the case any longer.

In [6] two roles are proposed for a network virtualisation environment, the InP and the Service Provider (SP). However, the structure proposed by the European Project 4WARD [7] has three functional roles, those of the Infrastructure Provider (InP), the Virtual Network Provider (VNP) and the Virtual Network Operator (VNO), which basically correspond to the three layers represented in Figure 1.

These three functional roles are analysed below, regardless of who is in charge of playing them, under which business case. No business models are implied at this stage and this topic will be covered in section 3. Figure 2 provides an overview of the different roles involved in NV.

Infrastructure Providers

InPs are responsible for the physical network which they deploy and manage. Their customers are VNPs who access the resources via programmable interfaces. According to the 4WARD model, they do not have any relationship with the end users, which are handled by the VNOs. The InPs can be

classified into two categories, those offering the access part of the physical network, called access providers, and those responsible for the core networking part and interconnection to other InPs, called facilities providers.

Programmability of the infrastructure networking resources is a key feature of network virtualisation. This is indispensable for allowing the VNPs to build customised virtual networks for customised services. An example of this could be a virtual network running tailor-made protocols.

According to the 4WARD model the InPs must fulfil the following requirements [7]:

- Virtualise their physical resources and provide deterministic degrees of isolation between them in order to equip virtual networks with corresponding guarantees
- Provide an interface that allows third parties to negotiate and lease virtual resources
- Monitoring and management of their physical resources and on-demand creation of virtualised resources on top of them
- Offer control interfaces for virtual resources rent to third parties in order to allow them to instantiate and manage the virtual networks built from the leased virtual resources
- Optionally, monitoring information of the Infrastructure Provider about virtual resources may be exported to the third party leasing those virtual resources
- Optionally, assist in attaching end users to virtual networks

Virtual Network Providers

Optionally, VNPs may play a mediation role between the InPs and the VNOs. They are responsible for building the virtualised substrate depicted in Figure 1 and for offering the substrate resources to the VNOs. The VNP composes a so called virtual network slice, which is a set of virtual resources consisting of physical resources from one or more InPs. The contents of the virtual slice are as requested by the VNO. After a virtual network has been set up and during its lifetime, the VNP's role becomes small. When for example an incidence occurs in the underlying infrastructure (e.g. link failure) the VNO would turn directly to the InP rather than having the VNP as middleman complicating and delaying a solution to the problem. A useful analogy for the function of the VNP is that of a travel agency. The travel agency is expert in travelling methods and routes, knows when and where trains and flights commute and has detailed knowledge on how to reach a certain destination. However, after the customer assumes the trip, he mostly contacts the on-site service providers and rarely contacts the travel agency. The VNP's tasks are summarised in [7] as follows:

- Provide an interface for VNOs allowing them to request creation, modification, or tear-down of custom-tailored virtual networks
- Request the required virtual resources from one or multiple InPs
- Assemble the requested virtual network slice from virtual resources available from one or multiple InPs, which may include assisting in the setup of virtual links between different InPs
- Handle over control of the virtual network slice to the VNO
- Allow for migration of virtual nodes between different InPs transparently to VNOs

Virtual Network Operators

VNOs have access to virtual network slices where they operate the required network architecture to offer a virtual network for a certain purpose. The VNOs set up virtual networks, manage and operate them. Their access to the virtual resources is done in such a way that allows transparent migration of virtual resources, occurring e.g. due to traffic engineering by the InPs. The VNOs handle end user attachment, check the users' authentication and delegate them to suitable virtual access points. An interesting feature of virtual networking is the possibility of a VNO to create its own "child virtual network" by partitioning of resources [6]. The VNO can lease the child virtual network to other VNOs, essentially taking a VNP role. According to [7] the VNO has to fulfil the following tasks:

- Assess the amount of virtual resources required
- Request creation, modification, or tear-down of a virtual network slice by a VNP

- Setup of the virtual network slice, i.e., installation and instantiation of network architecture in the virtual network slice and proper configuration of it, thereby transforming the virtual network slice into a functioning virtual network.
- Monitoring of the virtual network
- Management and operation of the virtual network
- Granting access to the virtual network to third parties such as service providers and end users

End users

End users are the users of the service offered by the SP, or directly by the VNO in the cases where a distinct SP does not exist as such. End user nodes are not part of the virtual network topology but are attached like leaves. End users connect to the virtual networks via virtual last mile links bridging their edge equipment and the virtual network access node. End users typically have to authenticate themselves towards the VNO by means of correct credentials. From the end user point of view, network virtualisation is supposed to be transparent.

Service Providers

Service providers (SP) are responsible for providing services to end users. Network virtualisation is not supposed to be visible to the SP – the functionality provided by a virtual network should be indistinguishable from that provided by a network based on physical resources. For this reason, the role of service provider is often excluded from the set of network virtualisation functional roles. In some cases, the role of service provider may overlap with the VNO, but from a functional viewpoint it is useful to consider them as different entities.

Virtual Network Customers

The Virtual Network Customers (VNC) is the entity holding a commercial relationship with the VNO. Depending on the specific use case, the VNC may correspond to the end user or can be of different nature, i.e. service providers with networking needs (e.g. IPTV providers), content providers or service subscribers [7].They can also be providers of cloud computing services needing to interconnect data centres and provide connectivity for their end users.

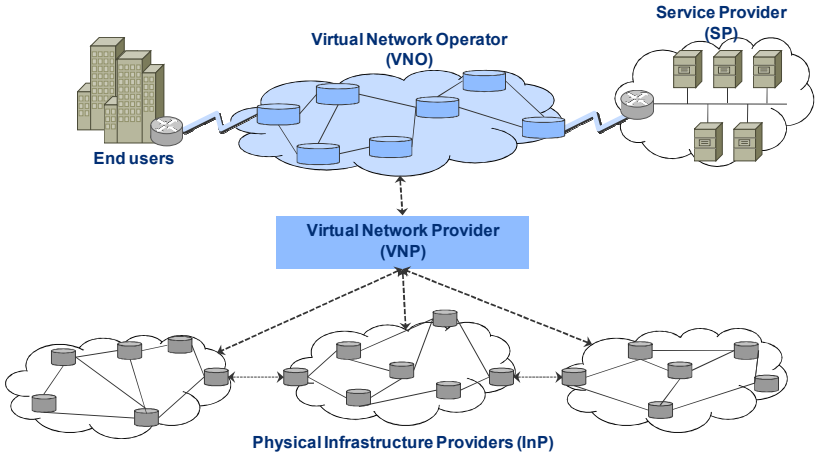

Figure 2 – Network virtualisation roles and players

2.1.3 Generic services enabled by Network Virtualisation

Network as a Service / Telco as a Service

Network as a Service (NaaS), or Telco as a Service (TaaS) as is sometimes referred to, is a concept that has been the forefront of the US GENI project. The idea is to provide a common network and IT substrate that can then be virtualised and combined as 'slices'. This type of service requires all type of virtualisation elements to be available and supported by the telco (computing, bandwidth, I/O and storage) beyond networking layer.

Furthermore, NaaS is also proposed as the Web 2.0 model to operators whereas they provide services in a Software-as-a-Service fashion, utilising internal systems and processes and opening up a range of APIs to 3rd parties – for example billing, SMS/MMS, location information and more. Examples of this include the Vodafone Betavine[1] and BT's Web21C[2], which is an SDK consisting of libraries that make it simple for developers to access BT's web APIs.

Cloud networking

Normally, cloud networks are considered to require non-monolithic network software, to better distribute information across an entire switch fabric, which in turn enables administrators to provision and administer clouds, add new services and update software. This design approach is considered better adapted to scalability, low latency, guaranteed performance, self-healing resilience and extensible management. This approach tends to promote the use of a hierarchy of non-blocking network elements, to interconnect compute and storage systems within data centres, and between data centres belonging to the same cloud infrastructure. Analogous to virtualisation in the cloud computing environment, NV is expected to bring similar advantages to the networking environment within and between data centres, allowing for optimised provisioning and administration of resources.

Figure 3 illustrates a networking architecture, displaying the networking layout between enterprise users and cloud providers. This networking is currently typically through the public Internet with associated security breaches, latency and general lack of QoS. Utilising NV in this context can transform this paradigm into one with controlled security and QoS.

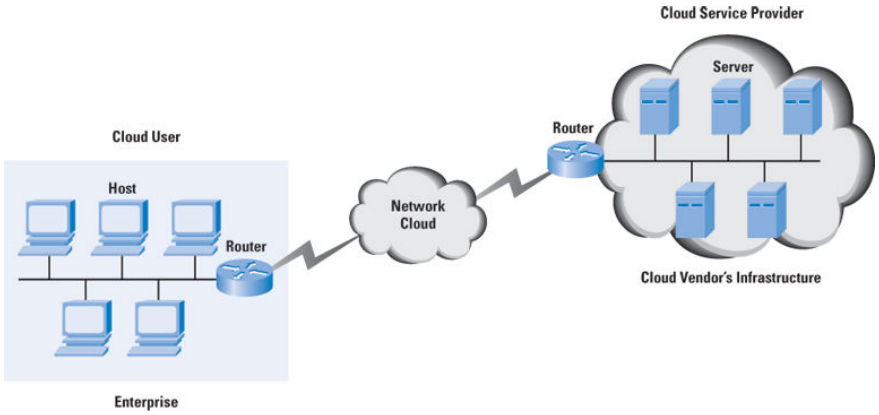

Figure 3 – Cloud computing context (Source: Cisco.com)

Instantiable services (Open Flow)

Service instantiation is a model that is considered an on-demand computing service using virtualisation techniques as close to the subscriber as possible. For example, if an enterprise needs to upload a 10 GB file today, it probably will take from a few hours to a few days depending on where the FTP server is located. Implementing an architecture that supports easy and dynamic service instantiation, allows a temporary FTP server to be instantiated on-demand at the first DSLAM/BRAS/MSAN/POP where the subscriber is connected.

FlowTable virtualisation using OpenFlow is a novel technology that is currently being developed. A number of research organisations are researching how OpenFlow virtualisation can be applied and monetised. OpenFlow is added as a feature to commercial Ethernet switches, routers and wireless access points – and provides a standardised hook to allow researchers to run experiments, without requiring vendors to expose the internal workings of their network devices. OpenFlow is used for applications such as virtual machine mobility, high-security networks and next generation IP based mobile networks.

[1] http://betavine.net

[2] http://web21c.bt.com

2.1.4 Elements of Virtual Networks

A virtualised environment has three basic elements, the virtualised server, virtual node and virtual link.

Virtualised Servers

A virtualised server is a computer environment implementing Virtual Machines (VMs). A Virtual Switch is a software based switch used to switch between VMs inside the same physical server and aggregate the traffic for connection to the external switch. Often, a Virtual Switch is implemented as a plug-in to a Hypervisor. The relationship is illustrated further in Figure 4.

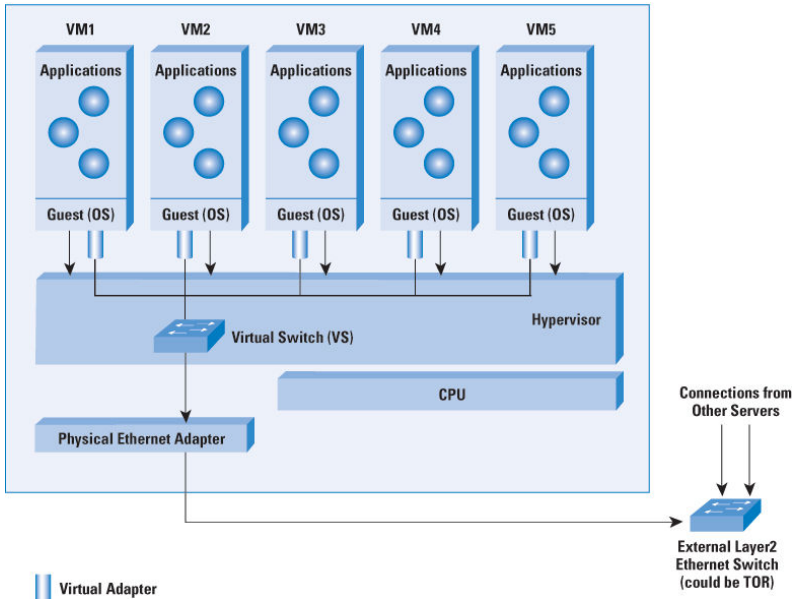

Figure 4 – Virtual Ethernet Switch in a virtualised server environment (Source: Cisco.com)

VMs have virtual Ethernet adapters that connect to a Virtual Switch, which in turn connects to a physical Ethernet adapter on the server and to the external Ethernet switch.

Virtual Nodes

Virtualisation of the nodes that constitute a physical network is a fundamental issue to network virtualisation, router virtualisation being the most notable aspect. Modern routers are built on very powerful hardware and software that allows resources to be "sliced" amongst many virtual networks passing through the node. A modern router is a complicated network element with a range of functions. It operates conceptually in two operational planes, the forwarding and the control plane. The forwarding plane functionality is to actually forward traffic from ingress to an egress interface. The control plane decides on the route a packet is forwarded to, QoS issues, and other aspects. Traditionally, routers operate on Layer 3 packets but modern devices extend their operation across layers below and above. So called switching routers are now commonplace using MPLS technology to create fast switched paths through the network, instead of the traditional hop-by-hop routing approach.

In a virtualised environment, the forwarding and control planes need to be virtualised. In Figure 5 a substrate node is shown, having two physical interfaces and two virtualised nodes (a and b) are implemented. In order to distinguish between the two virtual networks, the node needs to have a unique network identifier, called VNet-ID [5]. The main purpose of this is:

- End user attachment. The VNet-ID can be used by users anywhere to attach to the desired virtual network. For this, the VNet-ID must be globally unique.
- Accounting and billing. The virtualised environment is based on multiple infrastructure providers and the VNet-ID provides means to identify and assign resource usage. Again, the VNet-ID must be globally unique in order to allow for correct accounting and billing.

In the example in Figure 5 the physical node implements two virtual nodes for VNet#1 and one virtual node for VNet#2. An identifier is thus needed to distinguish between the virtual nodes; this identifier is called VNode-ID. The virtual nodes are connected by virtual links, either inside or outside the physical node. They are connected to virtual interfaces that also have identifiers, VIf-ID. The VNode-IDs and the VIf-IDs only need to be unique inside each physical node, since they are only used in conjunction with a globally unique VNet-ID.

Other functions that appear in Figure 5 are of importance. The Substrate Node Control is only accessible by the InP and is used for slicing of node resources and setup of virtual links. The (De-) Multiplexing and QoS Mechanism demultiplexes incoming virtual links and multiplexes outgoing virtual links via the connected physical links. It also implements the required QoS measures. The Hypervisor/Resource control creates the virtual nodes and manages their resources. The Out-of-VNet-Management Access allows VNOs to access each of their virtual nodes, permits reboot and other management functionalities. This interface is highly critical from a security point of view.

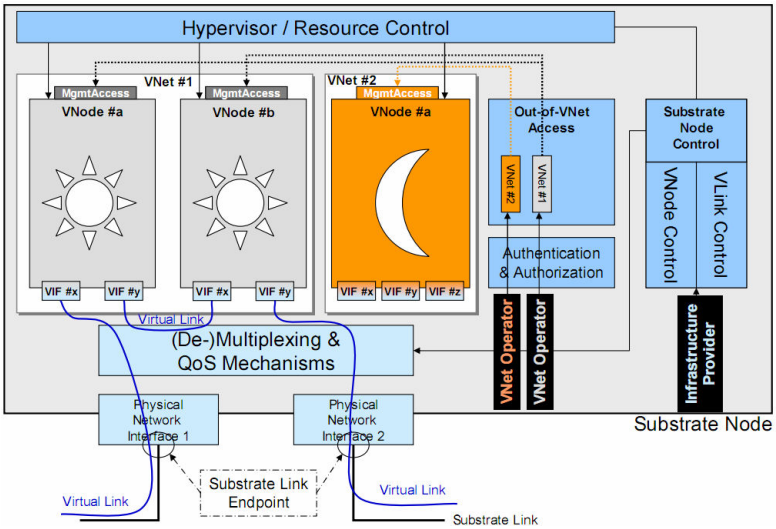

Figure 5 – Virtual nodes implemented on a substrate node [5]

Virtual Links

In several forms (for example, by means ATM, Frame Relay, MPLS or, more recently, Ethernet-based technologies), link virtualisation has been deployed in large scale operator environments for a long time. The basic purpose of link virtualisation is to divide, share and isolate the resources of physical links. A virtual link is an abstract entity that represents the functionality of a traditional physical link (i.e. bit transport between connected endpoints), but is not based on physical resources.

In the context of network virtualisation, virtual links provide the ability to flexibly connect virtual nodes with certain guarantees such as isolation, dedicated resources or QoS. Link virtualisation may have different issues, depending on the different types of physical links.

There is a wide range of available options for wired link virtualisation, from Ethernet VLANs to optical circuit switching. For each technology, a specific virtual link identifier is used to provide a separation between different virtual links. A general overview of the characteristics of relevant wired link virtualisation technologies is provided in Table 1.

Table 1 – Existing approaches for wired link virtualisation[3]

Technology	Virtual link identifier	Virtual link aggregation	Strengths	Shortcomings
Frame Relay	DLCI	Not supported	Mature technology	Circuit-based, obsolete
IP-over-IP	Source/ destination IP addresses	Not supported	Easily available (only IP connectivity is required).	Virtual links not visible to service providers; limited isolation of virtual links.
ATM	VPI/VCI	VPI	Powerful QoS mechanisms; adequate for traffic with strict QoS guarantees.	Circuit-based, not possible to integrate with IP control plane, obsolete.
802.1q	VLAN tag	Not supported	Pervasively used, easily available	Not scalable, limited applicability in WANs
802.1ad (Q-in-Q)	VLAN tag	Provider VLAN tag	Higher scalability than 802.1q. Widely used in metro networks.	For Ethernet traffic only. Scalability issues in large networks.
802.1ah PBB	[B-SA, B-DA, B-VID]	B-VID (backbone VLAN identifier)	Unlimited scalability, optimised for Ethernet traffic.	Lack of traffic engineering. Ethernet oriented.
PBB-TE (802.1Qay)	[B-SA, B-DA, B-VID]	B-VID	Simpler/more economical (lower OPEX&CAPEX) than MPLS	Requires an external management plane; unproven in large scale implementations
MPLS	Label	Label stacking (outer label)	Versatile, mature and reliable; used in many service provider backbones.	Complex, too many protocol variants, complicated inter-domain interoperability
MPLS pseudowire	[Tunnel label, VC label]	Tunnel label	Supports a wide range of protocols	Complex
T-MPLS	Label	Label stacking (outer label)	Reliable, simpler than MPLS, enhanced OAM	Incompatible with MPLS, discontinued in favour of MPLS-TP.
MPLS-TP	Label	Label stacking (outer label)	Based on a subset of MPLS, which is a widespread and proven technology.	Immature, standard still under development
SDH	Virtual container	SDH hierarchy.	Mature, reliable and stable.	Fixed resources reservation, expensive, approaching the end of lifecycle.
OTN	ODU container	Optical Transport Hierarchy	Extends SDH hierarchy to lambda-rate technologies	Fixed resources reservation.
Optical Circuit Switching	Lambda label	Waveband and fibre	Reliable, simpler than MPLS, enhanced OAM	High bandwidth links, fixed resources reservation.
Optical Burst/Packet Switching	Burst/Packet label	OCS, fibre	Optical circuits multiplexing	Not mature. Very complex.

Since many of these technologies are expected to be deployed in core networks, an important feature is the capability to aggregate multiple virtual links into a single pipe, rather than handling a potentially huge number of individual virtual links. This feature is crucial to evaluate scalability of these technologies and whether or not they can be deployed in large scale scenarios.

A particular form of link virtualisation, especially relevant in the access network segment, is wireless virtualisation. Most wireless virtualisation solutions involve the use of resource division and concurrent utilisation. As the wireless medium is a common resource for the whole physical network and does not belong to a specific node, resource sharing usually raises major challenges. It is usually a problem to schedule Tx/Rx power, frequency, time, and code or space allocation. Well known wireless transmission strategies can be used, such as TDMA (Time Division Multiple Access), FDMA (Frequency Division Multiple Access), SDMA (Space Division Multiple Access), CDMA (Code

[3] Adapted from "Virtualisation Approach: Concept", 4WARD Project Deliverable 3.1.1 [7].

Division Multiple Access), as well as hybrid approaches, involving the use of several of these schemes.

2.1.5 VN Management

In this section, we provide an overview of the basics of a Virtual Network (VN) Management processes and a detailed description of a simple VN creation without implying any business models. According to the 4WARD model [7], the VN Management processes consist of four basic phases during VN life cycle.

- **VN Design:** VNO describes the required resources to create a VN and might add some constraints such as QoS and geographical restrictions.
- **VN Provisioning:** The aforementioned VN description or parts of it is forwarded to one or more InPs, which reply whether they can fulfil the request. The InPs can then setup their substrate resources by picking a set of nodes and links that match the requirements.
- **VN Instantiation**: If the functional roles of VNP and VNO are undertaken by different operators, VNO gets access to the virtual network slice via VNP's control interface after the VN slice is created successfully. This interface offers functionalities that operate on a low level such as allowing the VNO to reboot virtual machines or to install an operating system. On the other hand, if there is no VNP between VNO and InP, VNs are ready to operate right after the successful creation of VN slice and there is no need for this kind of handover.
- **VN Operation:** Future modification of the virtual network, e.g. extension, shrinking, modification of QoS requirements or tear down of the virtual network, further functionalities such as attachment of end users to the virtual networks are the common runtime operations of virtual networks. An end user contacts the VNO for these operations. However, if the VNP is present, the VNO may require contacting again. If, for example, a new virtual node and corresponding virtual links are added to the virtual network, the VNO will have to activate the newly added resources and is responsible that they are properly added to the virtual network, e.g. that the virtual node gets an address inside the virtual network or in the case of a virtual router that its routing tables are initialised.

An exemplary creation of a very simple network depicted in Figure 6, which is also described in the 4WARD model [7] in a detailed manner, will enlighten the three phases of management processes, namely VN Design, Provisioning and Instantiation phases. According to the 4WARD model [7], all three functional roles InP, VNO and VNP exist in the VN architecture. The creation of VN consists of setting up virtual links and nodes for the new VN and giving management access of each virtual node to the operator serving the end user.

- VNO describes the desired VN topology.
- This description is then passed to a management node of a chosen VNP.
- VNP requests the list of available virtual nodes that matches the required VN topology. VNP may optimise and avoid asking for unavailable or geographically not existing nodes from InPs if VNP has any prior knowledge.
- VNP chooses the nodes and virtual links from the list which best serves his needs and does a pre-reservation of those nodes and links. If more than one InPs must be involved in the described VN topology, InPs negotiate whether and how virtual nodes can be interconnected across InPs' borders. VNP notifies the related InPs following the pre-reservation of all the required virtual resources.
- Following the notification from the VNP, the InP contacts the substrate nodes in order to set up required virtual machines, related virtual node management systems (Out-of-VN Management) and virtual links between virtual nodes are installed with the desired properties.

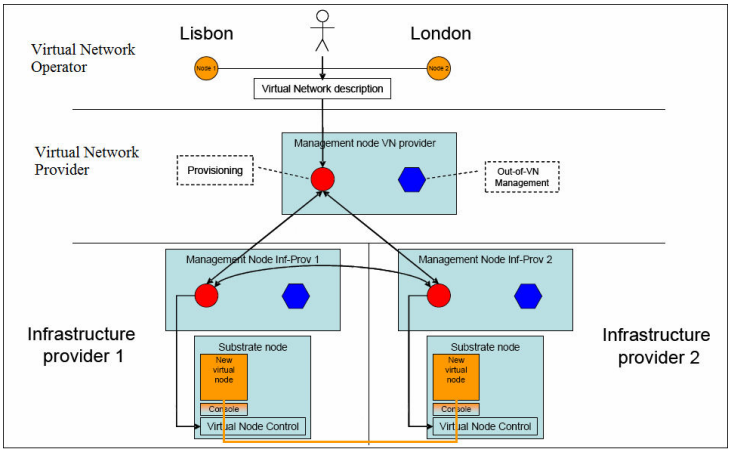

Figure 6 - Creation of a Virtual Network (I/II)[4]

At this point the VN is set up and ready to be used just like a physical network with no operation running on it. In order to set up a system on the VN, management access to virtual nodes must be given to the VNO. This access must be limited to ensure security but should fulfil every function needed to operate a VN. This type of management access is called Out-of-VN access and the node belongs to the VNP that includes Out-of-VN management interfaces for both VNP and InP, Figure 7. Assuming the virtual nodes and links have been set up successfully, the following steps are taken to give Out-of-VN management access for each node:

- The virtual node control acknowledges the successful creation of the virtual nodes and links to the InP's management node.
- The Out-of-VN Management is then informed of successful creation of the virtual node and needs to contain the current mapping from virtual node to substrate node. By this level of indirection, InPs may transparently migrate virtual nodes.
- The Out-of-VN management creates a unique identifier, which maps the virtual node local to the InP to the current location of the virtual node, to provide the VNO with a unified view of his network. This adds another layer of indirection, allowing for moving virtual nodes between InPs transparent to the VNO. The resulting unified view of the virtual network topology is then handed over to the VNO.

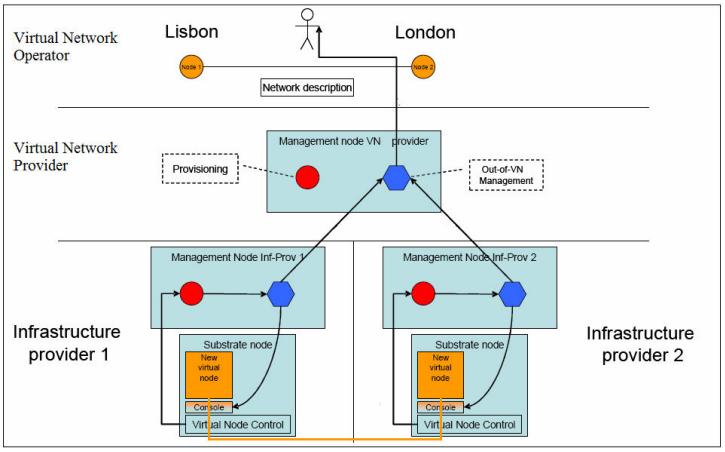

Figure 7 - Creation of a Virtual Network (II/II)[5]

[4] Source: Baucke, et al.[7]

[5] Source: Baucke, et al.[7]

The VNO is now able to access each single node via Out-of-VN management and to install the desired operating system and software in order to finally activate the virtual network topology and to provide a useful service to end-users.

VN Provisioning Mechanisms

One of the most attractive features of NV is the high degree of flexibility offered by VN provisioning mechanisms, which requires cooperation between role players. This flexibility not only improves the quality of service perceived by end users but also makes the networks resources usage more efficient. From this perspective, VN provisioning has a crucial role in VN management.

One of the VN provisioning mechanisms is initial provisioning, which composes the requested VN at the early stages of VN creation. Initial provisioning includes three consecutive steps, namely candidate discovery and matching, candidate selection and candidate binding. The first step is finding a set of available resources either by querying the service discovery framework or using a database which is created from the previous provisioning attempts, resulting in a set of discovered VN candidates. The candidate selection process consists of determining and selecting the best (or optimal) VN candidate(s) based on optimisation approaches and algorithms. Candidate binding, the last step of initial provisioning, is nothing but assigning the virtual nodes and links constituting the virtual network to the selected virtual resources in the substrate.

Following the successful deployment of the requested VN, the adaptive and dynamic provisioning and maintenance of the VN comes into play. Dynamic provisioning occurs when node splitting, migration, fix of mobility failures and maintenance of VN are needed. Dynamic provisioning and resource management take place respecting contracts and service level agreements and should maintain the topology of the VN at all times. Either binding changes and updates or new selection processes occurs during dynamic provisioning. When optimisation is needed for the VN to handle more requests, the system uses the already identified and still available candidates from the previous selection process in binding changes process and replaces the node bindings without any selection process. When substrate node or path is no longer able to support the virtual node working on top due to failures and dynamic changes in the environment, selection of new nodes may be needed. This selection process consists of selecting new substrate resources to maintain virtual resources and running the binding step. After either of these processes, modified mapping with extended matching, selection and binding is needed. Modified mapping may also be needed when InPs introduce new virtual resources which better fulfil the VN's needs.

2.1.6 Access to VN (by providers and end users)

Authentication, Authorisation, and Accounting (AAA) is a term for a framework to intelligently control access to network/IT resources, enforcing policies, auditing usage, and providing the information necessary to bill for services. These combined processes are essential for effective network management and security and are crucial in virtual networks as well. While the authentication process provides a way of identifying a user, the authorisation process determines whether a particular user is authorised to perform a given activity, typically logging on to an application or service. Accounting is the process of measuring resource consumption, allowing monitoring and reporting of events and usage for various purposes including billing, analysis, and ongoing policy management. AAA processes for Virtual Networks are considered to be established between InP and End-User as well as VNO and End-User for the attachment of End-User. In addition, mutual authentication of InP and VNP must be performed upon VNPs' each provisioning request.

Although VNs are separated and isolated from each other, they may need to cooperate with each other to better serve the End-User. Authentication and authorisation must also be performed at the intersection points between the VNs to provide secure inter-VN communication.

An end-user's VN access is to extend the virtual networks towards an end-host via a tunnel. The attachment of an end-user to a VN includes authentication and authorisation processes, the steps of the attachment process are:

1. End-users need to be authenticated at their physical attachment point in their local InP before being connected to VN.

2. End-user contacts his InP and indicates the VN he intends to join and accompanying initial credentials, if any. If there are multiple VNs, end-user considers the network parameters the VNs offer such as latency, bandwidth, etc. in VN selection.
3. The InP's authentication server contacts the VNO's authentication and authorisation server which is responsible for the VN that the end-user intends to join and also relays the accompanying credentials, if any.
4. The proper authentication and authorisation of the end-user begins.
5. If the end-user is a legitimate user, the end-user's contract with the VNP determines how it will be connected to the VN. The VN will then be extended by a virtual link towards the user.

Accounting is also one of the concerns for VNOs and InPs. VNO would like to know whether InPs provide reliable connections for their customers, which also meet their QoS expectations. This issue is crucial in terms of the agreements between VNOs and InPs and between VNOs and end-user as well [31]. Billing of the end-user depends on the contract between the end-user and the VNO and becomes more complex when dynamic resource parameters are involved. The requests of the end-user for an update in resources should be taken into account for billing.

2.2 Network virtualisation architectures put forward by research projects and initiatives;

2.2.1 European Projects

2.2.1.1 4WARD

4WARD ("Architecture and Design for the Future Internet") was a large FP7 Integrated Project, which ran from January 2008 to June 2010. Network virtualisation was one of the main focus areas of the project, with the following main objectives [7]:

- To develop a framework for the systematic and scalable provisioning of virtual networks utilising a wide range of virtualised network resources;
- To develop methods for the efficient virtualisation of diverse network resources for use in virtual networks;
- To define methods, interfaces, and protocols for the operation, and management of virtual networks;
- To provide a high degree of flexibility and enhanced resource provisioning capabilities for virtual networks with built-in security and trust properties.

Unlike most network virtualisation approaches, 4WARD considered network virtualisation not just for research purposes, but for use in commercial environments. The potential of virtualisation to enable new provider roles opens the way for new business roles and models, and may potentially lower the barriers of entry for new service providers. Moreover, it allows creating, controlling and managing coherent networks spanning multiple infrastructure providers offering more powerful control for virtual network operators. A number of constraints on technology are raised by economic and business realities.

4WARD devised a scalable framework for the systematic provisioning and management of virtual networks. This included mechanisms for the on-demand instantiation of virtual networks at scale (from heterogeneous substrates), virtualisation signalling and control methods, as well as dynamic management of virtual networks.

Another focus of 4WARD in terms of network virtualisation was the efficient virtualisation of diverse networking resources in a common framework, including mechanisms for virtualisation of routers, links, wireless resources (e.g. spectrum, modulation scheme processing etc.) and other types of network resources.

4WARD defined a network virtualisation functional model, based on three basic roles, namely the InP, the VNO and VNP, as explained before. However, the roles can be used in a number of different ways and a one-to-one relationship between roles and business entities should not be assumed.

2.2.1.2 FEDERICA

Federica was an FP7 project running from 1/1/2008 until 30/6/2010. The core objective of Federica was to support research on Future Internet by creating a Europe-wide, technology-agnostic e-infrastructure of network resources and nodes that can be 'sliced' to provide virtual Internet environments for research, as well as the mechanisms to allow researchers to control resources within these slices and conduct disruptive experiments without adverse effect on existing production networks [13]. Virtual slices of FEDERICA's infrastructure could be created, allocated and used simultaneously by researchers for testing, even with disruptive experiments, within a large production substrate. The researchers had full control of the allocated virtual nodes and network in their slice and could access specific network monitoring information.

In parallel, Federica researched virtualisation of e-infrastructures, in particular multi-domain control, management and monitoring, virtualisation services and user-oriented control in a federated environment.

The project made use of an infrastructure based on gigabit Ethernet circuits from the GÉANT2 backbone, coupled with virtualisation technologies [14]. The Federica physical topology is represented in Figure 8 [15].

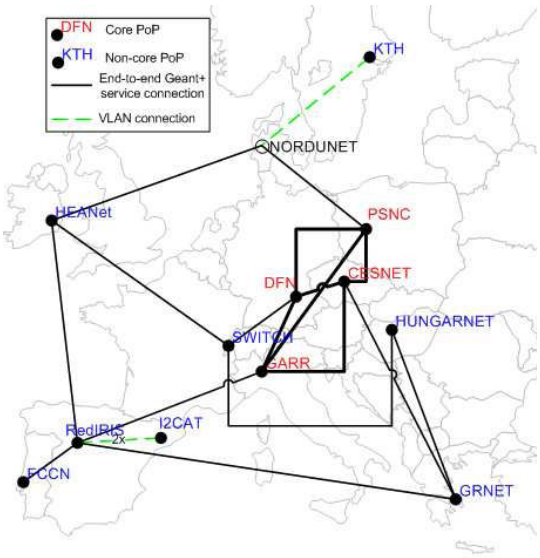

Figure 8 – Federica physical topology

The project had a phased approach for service offer - in the first phase, Federica was able to open Layer 2 and above functions for the researchers. In a later phase it gave access to the lower layer functions (Layer 1) in a particular piece of the core network.

2.2.1.3 G-LAB

G-Lab[6] (German-Lab) is a German federal research project for network test beds for Future Internet architecture, conducted in two phases. Phase 1 includes five universities and ends in 2011; phase 2 is an expansion of experimental facilities with industrial partners to conclude at the end of 2012.

G-Lab has several work packages that deal with technical issues of network architecture. One subproject, VirtuRAMA (Virtual Routers: Architecture, Management, Applications) deals with Internet network virtualisation. The aim is to have several different virtual networks on virtual routers running on consolidated routers with the aim to both reduce resource usage and provide a way to introduce services and protocols. VirtuRAMA runs from January 2009 until mid of 2011.

[6] http://www.german-lab.de

2.2.1.4 AGAVE

AGAVE (A liGhtweight Approach for Viable End-to-end IP-based QoS Services) was a 6th framework IST programme project which ran from December 2005 until May 2008. A description on the project's web page[7] is as follows:

> "AGAVE developed solutions for open end-to-end service provisioning based on the notion of Network Planes that may be interconnected across multiple providers to create Parallel Internets tailored to service requirements. The project investigated a range of Traffic Engineering techniques to realise Network Planes. A lightweight QoS approach was developed, based on the principles of differentiated routing with inherent load balancing and resilience, without requiring universal deployment of differentiated forwarding."

AGAVE focussed on network layer solutions, namely to provide end-to-end QoS-aware service provisioning over IP networks, using forwarding mechanisms such as Diffserv and Intserv. AGAVE proposed a new inter-domain architecture based on the concept of Network Planes, which allow network providers to build and provide Parallel Internets to achieve service differentiation. The Network Planes are defined within each autonomous IP network provider's domain. They can be described as slices of network resources (incl. bandwidth and routing/forwarding tables) allocated for a specific set of services with similar requirements, including QoS and availability [29]. The concept of Parallel Internets enables end-to-end service differentiation across multiple administrative domains, based on IP Network Providers' agreements. Parallel Internets are composed of interconnected Network Planes that transport traffic flows from services with common connectivity requirements. It is interesting to note that a Parallel Internet does not require all the Network Planes participating in it to be homogeneous, resulting in a high degree of flexibility.

It is interesting to compare the architecture proposed by AGAVE to that proposed by CABO, described in section 2.2.2.1. A key difference between the concept of Network Planes proposed by AGAVE and the concept of Network Substrates proposed by CABO is that a Network Plane is completely managed by the underlying InP instead of being leased to external Service Providers who have the actual control over the spliced resources. The AGAVE approach is more scalable in the sense that the network resources allocated to each Network Plane serve a set of Service Provider's services in an aggregate fashion, rather than being dedicated to any single Service Provider who has the actual control over its own substrate. Therefore, the number of Network Planes does not increase linearly with the number of requesting Service Providers. The concept of Network Planes is illustrated in Figure 9.

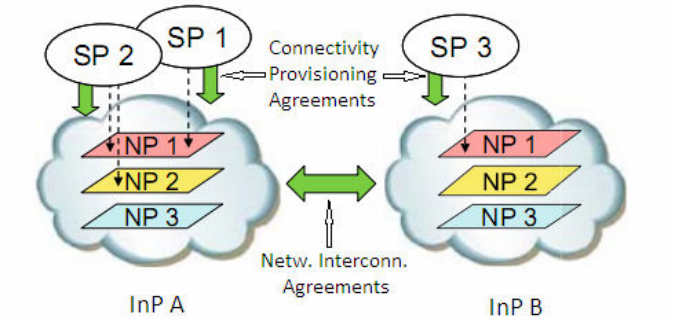

Figure 9 – The Network Planes concept proposed by AGAVE

2.2.2 North-American Projects

2.2.2.1 CABO

CABO (Concurrent Architectures are Better than One) is an architecture proposed in [28]. Cabo proposes to decouple infrastructure providers from service providers, as opposed to today's Internet where generally Internet service providers (ISP) manage the network infrastructure and provide

[7] http://www.ist-agave.org

services to their customers. Today's Internet is therefore fragmented and it is difficult to deploy new solutions on an end-to-end basis, although they make sense incrementally. Cabo is based on virtualisation, which is used by service providers to run multiple end-to-end services on infrastructure owned by others, i.e. InPs. The Cabo architecture is shown in Figure 10. Multiple InPs offer their resources to multiple Service Providers (SP), who in turn can build virtual networks across multiple infrastructure domains using slices of their network elements.

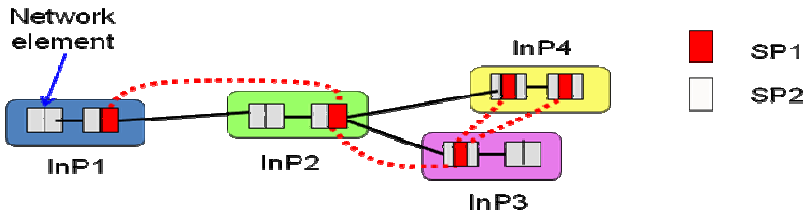

Figure 10 – Architecture proposed by Cabo

The Cabo approach is very much in line with the description on virtual networking given in section 2.1.1, except that Cabo does not define the role of the virtual network provider. The service provider in Cabo plays both the VNP and the VNO roles defined in 4WARD.

2.2.2.2 GENI

Global Environment for Network Innovations (GENI[8]) is a US initiative for network innovation, with one target being network virtualisation, with the current practical aim for network experiments across distributed networks and systems. GENI is funded by the US National Sciences Foundation (NSF) and managed by BBN Technologies. Participants in the projects include various US universities and non-profit institutions.

The GENI architecture is based on a layered design with a physical substrate, a management core (GMC) and user services for researchers on top. The bottom layer is made up of individual physical networking devices and hosts. By using OpenFlow and FlowVisor, the networking devices are virtualised and controlled by decentralised OpenFlow controllers. This control plane is in turn managed by a central GENI "clearinghouse," which experimenters contact to configure virtual slices for their services and experiments.

The host virtualisation in turn depends on OrbitLab, Emulab or PlanetLab and its technologies. For PlanetLab, any request coming to the GENI clearinghouse is routed to the PlanetLab aggregate manager (PLC). The experimenters, then, request for the virtualised network substrate on the physical layer, which connects the virtualised host systems configured by PlanetLab/PLC. The individual network and host groups form "substrates". The virtualised network and host constitutes an end-to-end slice.

2.2.2.3 OpenFlow

The OpenFlow protocol[9] is a project started in 2008 by the Clean Slate Lab at Stanford University, conducted in cooperation with partners from the networking industry. The current version of the protocol specification, released in 2009, is OpenFlow 1.0. Hardware implementations by several networking vendors for a range of their products are available and shipping in beta versions. First networks are being built with OpenFlow technology in laboratory and university settings. Some OpenFlow trial deployments are part of NSF/GENI funded programmes in the United States.

The OpenFlow project develops and specifies the OpenFlow protocol, an API to program switches. It contains a programmable control plane, flow-based forwarding and virtualisation. OpenFlow itself is an open protocol that defines an API for switches and routers, by which a central controller component can program and control these devices. Flow setup is delegated from the switches to this central controller. New services, architectures and protocols can make use of the central controller or provide

[8] http://www.geni.net/

[9] http://www.openflowswitch.org/

their own. Network virtualisation is achieved by slicing up the available network space into "slices" based on any number of properties of the underlying networking protocols, i.e. IP addresses, ports, MAC addresses, etc.

The implementation of the protocol in switching and routing products, and the development of the central controller are all independent of the protocol development itself. There is substantial support from vendors and research groups alike. Presently, vendors are developing the ver1.0 firmware with their lab teams. Larger scale tests and deployments of OpenFlow networks are targeted for 2010.

A description of OpenFlow as network technology is provided in section 2.3.2.3.

2.2.2.4 UCLP (User Controlled Lightpaths)

UCLP was a Canadian research project with the main goal of providing a network virtualisation framework upon which communities of users could build their own middleware or applications. The system was designed as a Service Oriented Architecture (SOA) where Web Services are the basic building blocks [25]. UCLP allows computer programmes and human operators to manipulate network resources as though they were software objects, thus making it possible to create and configure lightpaths within a given network and manage them as separate domains.

UCLP technology changes the management and control of optical networks by giving users the ability to create and manage their own Virtual Private Network (VPN), increasing their level of bandwidth and quality of service by drawing on resources from more than one supplier, without a network operator co-ordinating this kind of activity. UCLP allows end-users to cross connect, add or drop lightpaths anywhere in the network, as well as partition these lightpaths and offer them to other users. The resulting network can transfer large amounts of data, support real-time multimedia exchanges, and enable globally distributed broadband computing.

Argia™ is the commercial evolution of UCLP and is currently commercialised by Inocybe[10].

2.2.3 Asian Projects

2.2.3.1 Akari

AKARI[11] is a joint project of NICT (National Institute of Information and Communications Technology) of Japan and the University of Tokyo, started in 2006. AKARI's aim is to provide a "New Generation Network" with a focus on different fields, including one on network virtualisation.

The approach of AKARI is targeted towards defining high-level concepts to be included in future Internet architecture; however the project did not result in a protocol or a firm architecture specification (yet). The basic aim for network virtualisation in AKARI is similar to GENI in that it envisages a virtualised core network, on top of which virtual networks are overlaid.

2.2.3.2 NVLAB

NVLAB[12], the Network Virtualisation Research Lab, is a joint project of NICT of Japan and the University of Tokyo. NVLAB does not define a strict architecture but is a common test bed for different network virtualisation architectures.

2.2.4 Other Projects

2.2.4.1 PlanetLab

PlanetLab[13] is an international academic project for a distributed global network research test bed. The project makes use of host-based virtualisation of nodes running virtual systems that are grouped together to provide virtualised services. The basic architecture consists of *Physical nodes* in an interconnected global network which contains multiple virtual servers, *Slivers (Virtual servers)* on physical nodes. In turn, *Slices* contain one or more virtual servers distributed throughout the network, on which *Services* run, on their own (distributed) slices.

[10] http://www.inocybe.ca

[11] http://akari-project.nict.go.jp/eng/index2.htm

[12] http://www.nvlab.org/wordpress/

[13] http://www.planet-lab.org/

The network virtualisation is basically a Linux based system solution, using customised Linux kernels (linux-vserver) and custom management and control software. The network virtualisation parts interfaces between the host-system's netfilter IP filter and the Vserver virtual servers.

CoreLab[14], a collaboration between the NICT (National Institute of Information and Communications Technology) of Japan and the University of Tokyo, is a development of the PlanetLab networking test bed. CoreLab takes over the basic architecture of PlanetLab, but extends it on several parts, including the handling of the Slivers, the guest OSs etc. Deployment of CoreLab is focussed on Japan and Asia.

2.3 Network virtualisation technologies

2.3.1 Software based virtualisation solutions

Although platform virtualisation is an old technology, it was only recently that hardware and operating systems have become mature enough to make the promise of virtualisation a reality. Usually, the key component of virtualisation is the hypervisor, which is a software layer between the virtualised guest operating system and the real hardware. The virtualised guest operating system is supposed to view the underlying hardware as exclusively belonging to it and it is up to the hypervisor to provide this illusion. Virtualisation solutions fall into three basic categories – full virtualisation, paravirtualisation and operating system level virtualisation.

2.3.1.1 Full Virtualisation

In full virtualisation, the interface provided by the virtualisation system fully replicates the actual physical hardware, which allows operating systems to run as guests in a virtual machine without any adaptation, as if the virtual machine was a physical system. A complete set of hardware elements is provided by the hypervisor for running unmodified guest OS. Full virtualisation enables complete decoupling of the software from the hardware, therefore facilitates migration of applications between different physical systems. Another good property of full virtualisation is the complete isolation of different virtualised applications, which makes this approach highly secure by design. Microsoft Virtual Server and VMware ESX Server are examples of full virtualisation.

Loss of performance is usually the price to pay for full virtualisation. The hypervisor must provide the virtual machine with an image of an entire system, including virtual BIOS, virtual memory space, and virtual devices. The hypervisor must also create and maintain data structures for the virtual components, like memory page table. These data structures must be updated for every corresponding access by the virtual machines.

2.3.1.2 Paravirtualisation

In order to make virtualisation more efficient, two solutions were introduced: hardware-assisted virtualisation and paravirtualisation. Hardware-assisted virtualisation, as its name suggests, used physical hardware to take away the strain of virtualisation from software and the operating system.

In paravirtualisation, most services are provided directly from the underlying hardware, rather than an abstraction of it, which offers two main advantages – firstly, an entire hardware emulation layer between the guest operating system and the physical hardware is not needed, as the virtualisation software is just a thin layer that basically multiplexes access by guest operating systems to the underlying physical resources. Secondly, paravirtualisation avoids the dependence of device drivers in the virtualisation software, as the device drivers contained in one of the guest operating systems are used (privileged guest). Thus, it is possible to take advantage of all the capabilities of the hardware in the server, rather than being limited to hardware for which drivers are available in the virtualisation software as in hardware emulation virtualisation.

One of the drawbacks of paravirtualisation is that it requires the adaptation of guest OS – this is specially an issue with closed source operating systems, for which hardware support for virtualisation is needed to ensure that the native binary of the guest OS can still share resources with other guest OSs.

[14] http://www.nvlab.org/wordpress/

2.3.1.3 OS-level Virtualisation

A third approach, which is not a virtualisation solution in a strict sense, is OS level virtualisation, in which all guests share the same operating system as the base machine. Thus, by definition, OS virtualisation systems do not support the ability to run many different operating systems on the same physical machine. A key advantage of OS virtualisation is that a single OS instance is used, which provides enhanced efficiency in terms of resources; indeed OS virtualisation has very little overhead because it does not need to emulate hardware. Another advantage is increased flexibility offered by dynamic reconfiguration of resources allocated to each guest. On the flip side, lack of software heterogeneity is a key limitation of OS virtualisation.

Two examples of operating system level virtualisation are Solaris Containers and OpenVZ. In Solaris Containers, a virtual machine is called a zone, and a zone with resource limitations is called a container. Solaris Containers establish boundaries for consumption of resources such as memory, CPU time, and network bandwidth. Software applications and services can be isolated using flexible, software-defined boundaries, so many private execution environments can be created within a single OS instance. As processing requirements change in line with business needs, one or more of the boundaries of a Container can be expanded to accommodate a spike in resource demand. OpenVZ is another container-based virtualisation solution, for Linux. OpenVZ creates multiple secure, isolated containers on a single physical server. Each container performs and executes exactly like a stand-alone server; a container can be rebooted independently and have root access, users, IP addresses, memory, processes, files, applications, system libraries and configuration files.

2.3.1.4 Existing implementations of Network virtualisation

The combination of virtualisation with routing software (e.g. Linux-based XORP, Quagga routing suites), provides a relatively easy solution to set up virtual networks using commodity hardware. This has usually been the approach taken to build network virtualisation testbeds for research purposes. For example, in PlanetLab Linux Vserver is used to build virtual networks [20], 4WARD has used mainly Xen [21] in experimental activities, while Federica has used both Xen and VMware [22].

These approaches are appropriate for research environments as they provide open software platforms on which new networking paradigms can be experimented and validated in a virtualised environment, freed from limitations imposed by legacy technologies. However it should be noted that these solutions still lack fundamental performance, reliability and dependability properties to be used in a commercial setting.

In particular, I/O virtualisation and the ability to support fair sharing of the physical network interface cards (NIC) resources is a key requirement that is still not properly addressed by most virtualisation approaches. Crossbow is the code name for the new OpenSolaris networking stack that supports virtualisation of a physical NIC into multiple virtual NICs (VNICs). A VNIC operates like and appears to the system as a physical NIC. Each VNIC is assigned a MAC address, which can be configured to a value other than the default MAC address assigned to the physical NIC. Crossbow provides resource control features (bandwidth management and flow control) on a per VNIC basis, which enables allocation of resources to the individual VNICs. Traffic through each VNIC can be classified and separated into individual flows, based on port number, destination IP address and other parameters. These features can be used to improve system efficiency and enable differentiated services for separate VNICs [23].

Table 2 summarises the strengths and shortcomings of the three basic virtualisation approaches analysed above [18] [19].

Table 2 – Software based virtualisation approaches

	Full Virtualisation	Paravirtualisation	OS Virtualisation
Strengths	• No modification required in the guest OS • Guarantees complete isolation of virtual machines • Excellent compatibility - most operating systems supported without any modification	• Easier to implement than full virtualisation • Tends to perform better than full virtualisation • Better performance than full virtualisation for network and disk I/O.	• Best possible (i.e. close to native) performance • Dynamic resource management • Single OS installation
Shortcomings	• Requires the right hardware / software combination • Complicated to implement in the x86 architecture because of some of the privileged calls that cannot be trapped • Performance can be impacted by binary translation techniques for x86 privileged instructions	• OS running in virtual machines require adaptation – portability may be an issue • Modification of guest OS required; cannot run on native hardware or other hypervisors • Poor compatibility; not available on Windows OSes	• Supports just one OS • Isolation and security of virtual machines is not as effective
Relevant products	• Microsoft Virtual Server, Vmware ESX Server, Sun VirtualBox, Parallels Workstation, QEmu, Bochs	• Xen, UML, Solaris xVM	• Solaris Containers, OpenVZ, Linux-VServer, Parallels Virtuozzo Containers,

2.3.2 Network equipment vendors

Virtualisation is not new for router equipment vendors. For a long time, major equipment vendors such as Cisco and Juniper have supported a limited form of virtualisation - Virtual Routing and Forwarding (VRF) supports multiple instances of routing tables and can be seen as predecessor of "full-blown" network virtualisation technology, by allowing multiple instances of a routing table to co-exist within the same physical router. In the data plane, MPLS label switched paths interconnect edge routers or "Label Edge Routers" (LERs), which are at the edge of an MPLS network. These capabilities are in widespread use today, but they have some limitations.

An important limitation of "first-generation" router virtualisation is that there is no hard partitioning of resources between virtual instances. Control plane processing requirements are normally increased in proportion to the number of virtual router instances in VRF, and all such processes compete for the processor/memory resources of the same router control plane blade. On the other hand, a VPN is a service provided by an operator, not a real network – for example, customers are able to decide what kind of routing protocols run in their premises, and with the edge nodes, but not what routing protocols run inside the core network.

The potential of full-blown virtualisation has been recognised by the industry [26] [27]. It is clear that bringing router virtualisation to the core of the network offers multiple economical and operational advantages to operators and is able to enable new business models. However, carrier class requirements, such as reliability, stability, security, scalability, availability, consistency and predictability are not yet easily fulfilled by the general purpose virtualisation solutions described before. Two major vendors, Cisco and Juniper, have their own plans and roadmap for network virtualisation and commercial products have been launched, as briefly described below.

2.3.2.1 Cisco

Cisco uses the term network virtualisation in a broad sense as the architectural approach to provide separate logical networking environments. The concept can be applied in various scenarios and network segments.

The Nexus 7000 Series is targeted at data centres. The overall objective of the Nexus 7000 Series Virtualisation Architecture is to "allow IT departments to maximise existing resources by sharing a physical device among several logical functions, rather than devoting the entire device to a single function and underutilising the capacity of that physical device" [8]. Characteristics of Cisco Nexus 7000 Series include fault containment, independent management contexts per virtual device and allocation of hardware resources (e.g. ports) to specific virtual devices.

For campus networks, the Catalyst switching series (3560, 3750, 4500, and 6500) offers network virtualisation capabilities that enable partitioning of a single physical network into many logical networks across multiple locations.

For the WAN segment, Cisco offers a technology known as Secure Domain Routers (SDRs), based on Hardware-Isolated Virtual Routers, for routers CRS1/16. SDRs are defined on per-slot boundaries and provide full isolation between virtualised routing instances. Hardware-based separation of virtual networks provides a strict separation of resources between virtual networks, either at control plane or at data plane level, which guarantees fault isolation, as well as highly resilient and predictable performance and service [9].

2.3.2.2 Juniper

Juniper has played a leading role in network virtualisation, especially since the launch of the TX Matrix Plus core routing system, in February 2009. Juniper network virtualisation is targeted to core network equipment, namely the T-series router family. It is based on two fundamental building blocks:

- JCS1200 independent control plane: JCS1200 provides 12 Routing Engine slots and connects to one or more Juniper Networks routers via redundant Gigabit Ethernet connections. Up to 12 hardware-based virtualised routers can be built, each of which supports up to 16 software-based logical routers, for a total of 192 logical routers (96 if redundant) in one chassis. Because the Routing Engines are separate physical entities, complete security and isolation between each router is assured. The JCS1200 enables the scaling of control plane capacity without impacting forwarding plane performance, which is a critical requirement to virtualise core networks and services.
- High performance transport plane of T1600 and TX Matrix Plus: the TX Matrix Plus is a central switching and routing element, which can interconnect up to 16 T1600 chassis into a single routing entity that can be partitioned using the JCS1200.

With the TX Matrix Plus, Juniper aims to bring the advantages of virtualisation to core networks in a flexible multi-chassis routing system, thus lowering total cost of ownership and enabling operators to build scalable, reliable networks that can deliver innovative services with investment protection [10].

The target of Juniper network virtualisation is currently the network core, but it can be extended gradually to the edge and ultimately offer a solution to provision what could be seen as a more advanced form of today's customer VPNs [26].

2.3.2.3 OpenFlow

OpenFlow as a networking technology has two different technical scopes. In the narrow one OpenFlow is a standardised, open API to program the internal forwarding tables of networking switches and routers. In the larger scope the OpenFlow protocol and its ecosystem of controllers and software allow for a fully programmable and virtualisable networking substrate. OpenFlow is currently available as add-on to commercial LAN switches.

The OpenFlow model of programming networking devices is to centralise the control plane of these devices into a single external controller system, which handles all forwarding and processing decisions, while the switches themselves only act as execution units and forward packets. Incoming packets from new connections are forwarded by the switch to the central controller that instructs the switch(es) in the data path on how to handle this packet and subsequent data in that "flow". OpenFlow itself defines the protocol through which the switches and controller communicate.

Use cases:

- Centralised control and security
- Access management

- Traffic engineering
- Data centres
- Commoditising networking devices

If the OpenFlow protocol gains enough traction amongst switch chipset vendors, future switch implementations could be made cheaper and more powerful by using simplified hardware targeted for external control and programmability. This would remove the need for special functionalities in the switch hardware and provide much larger hardware tables and packet modification abilities while reducing the hardware cost of the switches.

Network virtualisation can be made possible by allowing the OpenFlow protocol to isolate control for different sets of flows, i.e., the set of flows can be partitioned based on packet fields to create different "flow spaces", each of which can be associated with a different controller. This allows creating "slices" from the underlying physical network substrate.

The current, first production-ready protocol version with a focus on enterprise and campus networks is OpenFlow 1.0, with follow-on versions under development targeting outstanding aspects of data centre environments and WAN requirements.

Network devices for enterprise LAN environments currently supporting OpenFlow are available by HP, NEC and Quanta-derivates, which have modified firmware with OpenFlow extensions. This allows the switches to be completely controlled by external OpenFlow controllers. Currently supported by these development efforts are switches in the NEC IP8800 series, HP ProCurve 6600, 5400 and 3500 models, and clones of the Quanta LB4G switch, sold under various brands. All of these switches are 24-48 port Gigabit and 10-Gigabit switches, with the exception of the modular HP ProCurve 5400 with up to 288 ports.

There are some performance limitations and unknowns related to the feasibility and exact configurations of complex OpenFlow networks. Currently commercially available switches were originally not designed for outside modifications of their forwarding tables and have very weak CPUs. Since all newly incoming packets are forwarded to an outside controller, this puts a bottleneck on the setup of new connections which thus depend on the performance and bus attachments of the switch CPUs. Future switch architectures might take greater switch programmability into account to make access to the forwarding table and control I/O less costly. Upcoming large-scale experimentations aim to show which practical problems this bottleneck might present and at which scale this could become a problem. Meanwhile optimisations are being worked on.

Further on, the general idea of centralised control of packet forwarding in a network needs to be investigated. Current switches and routers gained enough intelligence to operate on their own after successful configuration. In the OpenFlow model the network devices depend on an additional component to provide basic networking service – an active connection to an outside system and the active controller itself. This may present both scalability and business continuity issues, since the whole network now depends on additional, highly-centralised components. Much will depend on the possibilities of scaling controllers in more complicated networking setups and proof of concept of redundant controller configuration, which may prove to be more complicated than current redundant network setups.

In the same vein the security impact of a central networking component potentially handling and having information on all traffic in the network is not investigated yet. The controllers in larger networks would be a source of information and control of the complete network traffic traversing the network and all of its devices. A compromise of this single device, which can be caused by attack happening on the network used to communicate with the controller, could be catastrophic for the whole network.

2.4 Relevant industry activities by standardisation groups in Network Virtualisation

2.4.1 IRTF

Recently, the IRTF created the Virtual Networks Research Group (VNRG) to identify architectural challenges resulting from Virtual Networks, addressing network management of Virtual Networks, and exploring emerging technological and implementation issues [5].

The VNRG provides a forum for interchange of ideas among a group of network researchers with an interest in network virtualisation in the context of the Internet and also beyond the current Internet and will encourage the organisation of the work in smaller design teams focused on specific areas of research.

The initial set of work items includes topics such as concepts/background/terminology, common parts of VN architectures and common problems/challenges in virtual networks. VNRG takes as an input efforts of a number of IETF WGs, including encapsulated subnets (LISP at layer 3, TRILL at layer 2), subnet virtualisation (PPVPN, L3VPN, L2VPN) and aspects of managing virtual components (VRRP), as well as some work in more general areas, notably on tunnels (INTAREA).

VNRG will produce Informational and Experimental RFCs in order to document the activity of the group and to formalise the outcome of the research topics carried by the group. In addition, such documentation could become input to IETF working groups. VNRG will also encourage prototyping of virtual network technologies to validate this exploration.

The first meeting of the VNRG group took place during IETF 77, in March 2010.

2.4.2 ITU-T

Focus Group on Future Networks

The ITU-T Focus Group on Future Networks was set up to collect and identify visions of future networks based on new technologies, assess the interactions between future networks and new services, familiarise ITU-T and standardisation communities with emerging attributes of future networks, and encourage collaboration between ITU-T and Future Networks communities [12].

To accomplish these objectives, the Focus Group plans to:

- gather new ideas relevant to Future Networks and identify potential study areas on Future Networks,
- describe visions of the Future Networks,
- identify a timeframe of Future Networks,
- identify potential impacts on standards development, and
- suggest future ITU-T study items and related actions.

Network virtualisation is identified as the prominent technology that can realise the isolation of networks and can be used to build the large scale testing infrastructure for Future Network technologies. This is a key condition to develop new technologies that are not bound to the current state of the art and to overcome the limitations of the current networks.

With regard to network virtualisation, the following problem spaces are identified: Isolation, Performance, Scalability, Flexibility, Evolvability, Management and Security.

Focus Group on Cloud Computing

The ITU-T Focus Group on Cloud Computing (FG Cloud), in operation since May 2010, aims at contributing with the telecommunication aspects for flexible cloud infrastructure, security and management aspects of telecommunications, service requirements, etc., in order to support services/applications of "cloud computing" making use of telecommunication networks and service platforms [30].

3 Scenarios for network virtualisation uptake

3.1 Cloud computing scenario

3.1.1 Problem/scenario

Cloud services normally represent remote delivery of computing resources, whether hardware or software resources, most often via the Internet, either public or managed. This is especially relevant in public cloud computing environments where customers obtain cloud services from a third-party cloud provider, sometimes without any knowledge of the origin of applications or data storage/residence. Usually, this means that data crosses multiple networks before it is delivered to the end-user. Unless a somehow managed connection is used, the content delivery is mostly only a "best-effort" endeavour.

Users increasingly expect to proficiently access cloud services seamlessly, equally from fixed and mobile networks and assume a trouble-free handover of sessions between networks. Software-as-a-service applications in the cloud are becoming independent of the type of device being used to access the services. With the explosion of capable smartphones and small network computers like the iPad, users prefer to use whatever device is most appropriate or applicable for them. Additionally, users increasingly expect a problem-free and fast access to cloud services, irrespective of access networks or devices.

Currently, in a complex multi-network, multi-provider and best-effort cloud delivery model environment, the realisation of this vision still remains elusive, except for perhaps larger enterprises that can afford to provide content delivery quality assurance and increased security measures– usually through costly managed network connections (e.g. VPN/MPLS) and/or contracts with third-party CDN providers (e.g. Akamai).

With the advent of the virtualised networking technology, introducing new roles and layers of providers and operators, the cloud access and delivery model could potentially be improved, providing options for VNOs to implement end-to-end VNs and offer customised networking solutions for individual customers or customer segments. A possible scenario is shown in Figure 11. Here, a virtual network is created for each customer accessing the cloud services. The underlying virtual network resources can be based on fixed IP-networks and fixed or mobile wireless networks. This would greatly simplify the network usability for the end user such as accessing network services by different devices.

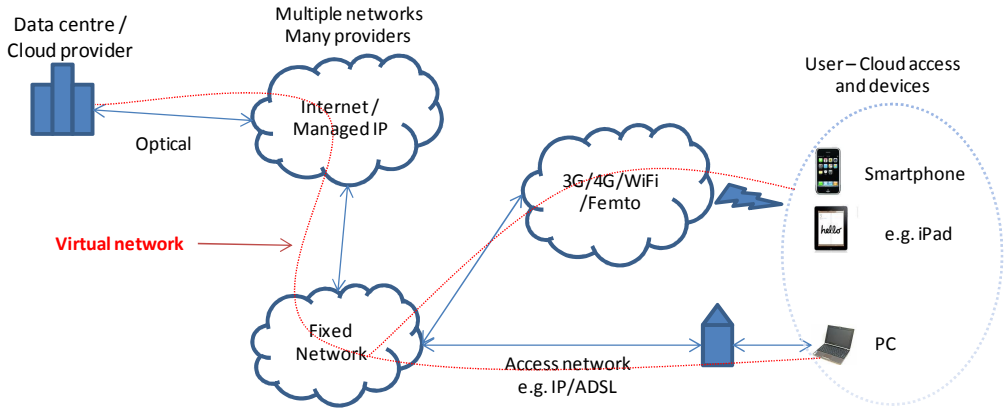

Figure 11 – A possible scenario for utilising a virtual network for cloud computing access

3.1.2 Stakeholders involved; basic business model

This scenario includes numerous stakeholders including;

- Cloud service providers that have high-capacity access connections to one or more WAN providers, through a point-of-presence (PoP) connection node or similar.

- InPs that operate physical networks and connection links for the purpose of exchanging traffic between customers of each network. Usually, these include telecom operators and large network operators.
- VNPs that manage and control virtual networks on top of multiple physical networks, through agreements with subsequent InPs.
- VNOs that provide services directly to end-users.
- End-users (customers) who obtain managed access to cloud services through the VNO, simplifying service provisioning and presentation.

3.1.3 Basic requirements

The described scenario illustrated in Figure 11 assumes a direct virtual end-to-end connection from the service/content provider to the end-user. Thus, the scenario is not "a pure" cloud computing configuration, as the service or application origin is bound to a particular location or data centre, unless the virtual network connection would be able to automatically detect multiple dynamic routes. Still, it does represent an interesting concept as it enables secure and prioritised network tunnelling and potential automatic and dynamic handover between networks, mobile and access. For this to be realised there are numerous tasks and technologies that need to be in place, including;

- Networking infrastructure components/nodes offering virtualisation and partitioning capabilities
- Networking protocols that provide for secure data transmission with QoS across public networks such as the Internet
- Intelligent systems capable of performing dynamic seamless handover from one network to another – automatically establishing and tearing down virtual network connections

3.1.4 Gaps/open issues

Although the scenario represents an interesting paradigm, there are several issues that remain unresolved or elusive, including:

- Limitations of supporting distributed cloud service provisioning, i.e. the virtual networking assumes a static end-to-end connection.
- Seamless networking handover technologies are still immature and inefficient. Despite years of efforts, the seamless handover problem has not been resolved adequately.
- Potentially complex service delivery process or business models. Requires cooperation and synchronisation between service/content providers, network providers and local operators.

3.2 Content Delivery Networks

3.2.1 Problem/scenario

A content delivery network (CDN) is a system of servers containing copies of data, placed at various points in a network to maximise the client access bandwidth. This is opposed to all clients accessing the same central server, with potential bottlenecks and latency problems. Benefits of CDNs are numerous, including;

- CDNs are used to distribute content from origin servers to users
- Avoids large amounts of same data repeatedly traversing potentially congested links
- Reduces Web server load
- Reduces users' perceived latency
- Routes data around congested networks

Although CDNs accomplish a great deal in providing enhanced content delivery, both in terms of implementing distributed caching servers for content replication and intelligent data route selection through proactive Internet traffic monitoring, CDNs are still basically "best-effort" networks. This scenario explores the pros and cons of CDNs vs. VNs and if CDNs' limitations can be somehow compensated through VNs' capabilities.

Strategically placed servers at edge locations decrease the loads on interconnects, public peers and backbones, freeing up capacity and lowering delivery costs. Instead of loading all traffic on a backbone or peer link, a CDN can offload these by redirecting traffic to edge servers. In order to reduce latency and packet loss, CDN servers are normally placed as close to the users as possible, thereby minimising network distances.

Commercial CDNs like Akamai and Limelight Networks comprise of tens of thousands of servers distributed worldwide, connected through hundreds of backbone networks. These distributed "overlay" networks possess proprietary technologies that promise to minimise the inherent performance problems or bottlenecks associated with the Internet, especially related to issues like overburdened public peering points, routing vulnerabilities in the Internet's "best-effort" inter-network routing algorithm (Border Gateway protocol – BGP) and, reducing drag caused by TCP multiple round-trips (between the communication parties) to set up and tear down connections.

In the context of cloud computing service delivery, CDNs can play an important role. Highly distributed CDNs shorten distances and enable delivery of content from network edges, avoiding as many middle mile bottlenecks as possible. Also, various unique routing, communications and application optimisation technologies are used to accelerate cloud service delivery. By real-time monitoring of the Internet's conditions, alternative paths can be identified that provide better performance than congested and default BGP-defined routes. Routing around troubles spots and finding alternative paths can provide improved connectivity and accelerated content delivery.

Some CDNs also possess proprietary transport protocols that aim to overcome TCP's and HTTP's inefficiencies by leveraging persistent connection techniques, eliminating TCP slow-start, enabling intelligent retransmission after packet loss by leveraging network latency information, allowing multiple requests to be pipelined over a single connection, using multiple routes simultaneously, and more.

3.2.2 Stakeholders involved; basic business model

The scenario represents a business model where content providers would offer their end-users (customers) premium services utilising VNs as an addition to the native CDNs. VNs would provide premium distribution of content from origin servers and/or replication servers to the relevant ISPs that connect end-users to the network.

- Content providers typically want to distribute fee-based services and are concerned with offering the highest possible end-user experience. For this purpose, providing additional service quality through deploying virtual networks in the CDN could be in their interest. Typically, the content provider would pay a VNO for implementing and managing the VNs, either directly or via the CDN provider.
- CDN providers include global providers like Akamai and Limelight networks. The CDNs sometimes have a global coverage but, depending on marketing strategy, can also provide a more condensed service level in chosen geographical areas.
- VNOs would provide additional services to end-users and content providers by deploying the needed virtual links to individual surrogates in the networks and/or original content servers.
- End-users access the content through the ISP and via virtual networking connections to the most appropriate surrogate – as determined by the CDN routing server.

3.2.3 Basic requirements

Traditional CDNs are unable to account for the level of personalisation and customisation required by many content providers today. For example, content providers that offer premium, fee-based services are concerned with offering the highest possible en-user experience while many other content providers are mostly interested in delivering content at the lowest possible cost. Normally, in CDNs that replicate content throughout the network in multiple data centres, the user requesting data is being directed to the best available server through the mechanism of a load balancer or routing optimisation server.

The scenario would be to provide virtualised load balancers that can be customised or segmented to the needs of different content providers, e.g. content delivery can be prioritised within a certain geographic domain or given premium precedence, and VNs would be established in the network

connecting the origin content servers and the most important replication servers to the end-user vicinity. Figure 12 illustrates this idea further.

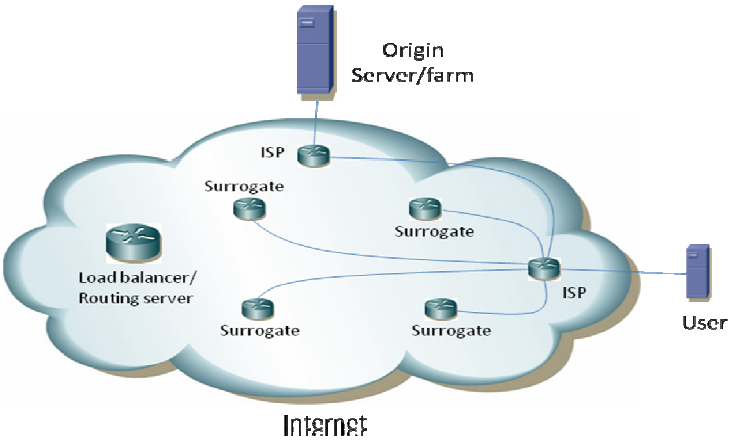

Figure 12– A possible scenario depicting VNs in a CDN

3.2.4 Gaps/open issues

There are several factors in this scenario that still remain unclear;

- Is it possible or beneficial/economical to extend a VN to multiple replication servers dispersed strategically in the Internet for optimum content delivery to the end-user?
- The required virtualisation capabilities of load balancers in CDNs are still being developed.
- The business case is unclear – e.g. if content providers would be willing to charge higher subscriber fees from their customers for this alternative service option.

3.3 Network as a Service

3.3.1 Problem/scenario

This scenario corresponds to the materialisation of the "Network as a Service" (NaaS) concept on a commercial environment. Basically, a virtual network service is provided by the owners of the network infrastructure to a third party.

3.3.2 Business model and role of stakeholders

This scenario is characterised by a clear separation of the roles of the InP, responsible for operating the underlying infrastructure and the VNO, responsible for operating the virtual network. Optionally, a VNP may be involved, particularly in the cases where a virtual network spans over multiple infrastructure domains, to find and collect the adequate network resources and act essentially as a broker between the InPs and the potential VNOs.

This model may be driven by business incentives (new revenues for the InPs, lower CAPEX/OPEX for VNOs compared to building a network based on physical resources) or regulatory measures (open access and sharing of the network infrastructure).

Three basic scenarios may be envisaged about the relationship between the VNP and the InP(s), as depicted inFigure 13:

- Scenario A: The VNO is supposed to select and establish a direct business relationship with the InP(s) that are able to fulfil the requirements, wherever they are located. This is typically the case if the virtual network is based on a single or few infrastructure domains. On the other hand, it might become complicated to handle if the virtual network extends across a considerable number of infrastructure domains.
- Scenario B: To overcome the limitations of scenario A, the VNP provides an intermediation role between the VNO and the InP. The VNP is responsible for finding the network resources at the best possible price and offering them to the VNO to configure the virtual network. The main advantage of this scenario is that it enables VNOs to roll out large virtual networks

without the need to establish business relationships with the potentially many involved InPs. The VNOs would have to deal with a single VNP who handles relations to the InPs. However, the fact that all the interactions between the VNO and the InPs have to be realised through the VNP represents an additional layer of complexity in the process.

- Scenario C: This scenario is similar to the previous one, except that in this case the participation of the VNP is restricted to the virtual network setup phase. Once the network is established, the VNP drops out of the picture and the VNO has a direct relationship with the InP. Thus, the VNP has essentially a brokerage function, facilitating the collection of the virtual resources, which makes the process more streamlined in the sense that it avoids the redundancy represented by the VNP in the operational phase in the previous scenario, but on the other hand it makes the responsibility for quality control more diffuse.

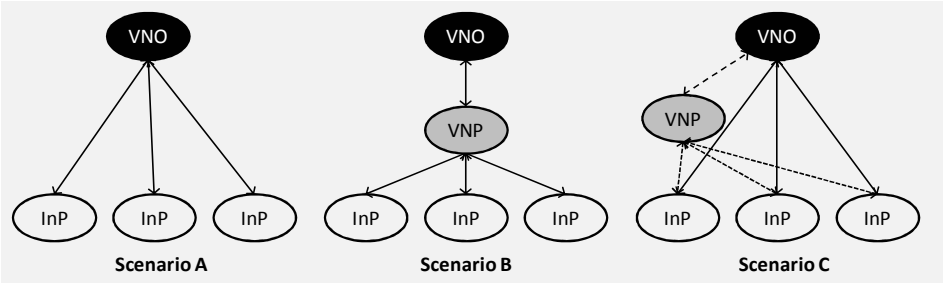

Figure 13 – Basic scenarios for virtual network provision

3.3.3 Basic requirements

A virtual network is supposed to replace the traditional network service fully based on physical resources. Therefore, all the usual requirements which are applicable to a commercial network must also be fulfilled in the case of a virtual network. This includes:

- Reliability: Commercial networks typically feature reliability levels in the order of 99.99% or 99.999%. Therefore this should also be the standard reliability level for virtual networks.
- Interoperability: A single virtual network is expected to span multiple network domains; therefore interoperability across multiple physical domains is a basic requirement.
- Security and privacy: Access to all control and management functions must be secured properly. Any possibility of eavesdropping between virtual networks must be prevented.
- Isolation and deterministic performance: To perfectly replicate the behaviour of a physical network, a virtual network should provide a performance level which does not depend on the traffic flowing on other virtual networks sharing the same infrastructure.

3.3.4 Gaps/open issues

Based on the basic requirements identified above it is straightforward to identify corresponding gaps and open issues:

- Reliability of a virtual network is ultimately determined by the reliability of the underlying infrastructure. Virtualisation introduces an additional level of complexity and represents a potential extra source of failure, which must be taken into account at the virtual network design phase and may represent an obstacle against the widespread adoption of network virtualisation in commercial environments.
- Interoperability between different heterogeneous domains is a challenge that requires standardisation. A particularly challenging scenario corresponds to the interconnection of non-contiguous network domains.
- Even if appropriate computational resources at the node level can be guaranteed, it is not trivial to guarantee strict isolation of link resources. Usually (for example in MPLS VPNs) the followed approach consists of over-provisioning to avoid a complicated control of resources; however this may be insufficient for virtual networks.

- As the definition of the three scenarios above clearly shows, it is difficult and probably impossible to find a unique model to describe inter-relationships between network virtualisation players. This in turn is likely to complicate the definition of standardised interfaces.

3.4 Virtual network as an enterprise service

3.4.1 Problem/Scenario

In the last decade, network-based VPNs have been a very successful service targeted at the enterprise market. A network-based VPN[15] is a secure, quality-assured network service which has been able to fulfil the enterprise networking requirements in a wide range of scenarios. For operators, VPNs represent an important source of revenues and, perhaps even more important, allowed them to play the role of true service providers (which enabled the inclusion of added value features like security management and hosting of data centres) rather than just pure connectivity providers. On the flip side, VPNs are a relatively static type of service, in the sense that once they are established, reconfigurations are usually complex and time consuming. This may be unacceptable in future networking environments characterised by highly dynamic conditions.

Virtual networks can be seen as an advanced form of enterprise VPN service by enabling high dynamicity and elasticity. On the other hand, virtual networks can be managed and controlled separately from the infrastructure, which in many cases represent a significant advantage over VPNs.

Emergence of cloud computing services targeted at the enterprise market could be a major driver to foster the emergence of this type of service offering. Network virtualisation combined with cloud services enables the easy distribution of applications and traffic/loads to ensure optimisation and extension of resources. Building on a highly isolated, converged, secure virtual network, a "virtual enterprise" can be established. The "virtual enterprise" concept makes doing business simpler by reducing operational costs, improving productivity and supporting new customer-centric services through marrying a consistent user experience with "anywhere, anytime" information access.

The management and control of the virtual network resources can be outsourced to the InP (in a similar fashion to today's VPNs) or executed independently by the service customer.

3.4.2 Business model and role of stakeholders

Basically, the players and roles defined for the Network as a Service scenario are fully applicable here, except that in this case the VNO corresponds to an enterprise, rather than to a virtual operator. It may also be the case that the infrastructure provider plays the VNO role as well, in the scenario where the virtual network management and control is outsourced to the operator.

Similarly to the Network as a Service use case, the virtual network may span multiple infrastructure domains, in which case there may be room for the VNP in charge of locating and selecting the virtual resources.

3.4.3 Basic requirements

The basic requirements should be the ones indicated in section 3.3.3. However, the enterprise scenario should have a few specific requirements, particularly in terms of dynamicity and elasticity. Security is another obvious concern. In the case where the customer is in charge of managing and controlling the virtual network, strong authorisation and authentication mechanisms must be put in place to ensure security. Controlled connectivity to/from external networks represents a major requirement.

3.4.4 Gaps/open issues

Scalability represents a concern in this case. Whereas in the Network as a Service use case a limited number of virtual networks (in the order of tens) is expected, in this case the number is expected to grow to the order of hundreds or thousands, based on the typical number of enterprise VPNs established in an operator domain.

Also, integration of cloud computing with network virtualisation is currently a hot research topic but many aspects of the integration of these two components are still not clear and still constitute a challenge.

[15] Not to be confused with a CPE-based VPN.

3.5 Network partitioning and dynamic resource allocation

3.5.1 Problem/Scenario

This scenario describes possible commercial telco services that can be implemented through an on-demand or predefined virtual network. The idea here is to provide the end users with a better user experience with commercial telco services on a virtual network rather than on a regular dedicated infrastructure. In order to do that, the mentioned services will be provided via dedicated virtual network slices composed on single or multiple physical infrastructure domains. Varying by each different commercial service, the VN may have a highly isolated, high performance, high security and privacy or alternating degrees of these features. VNs will be initially provisioned for each service and later can be adaptively maintained through dynamic resource allocation mechanisms. Providing a highly isolated, high performance VN for a commercial service as IPTV and VoIP can create an uninterrupted, very high quality and next to perfect user experience.

Possible commercial services that may run on a virtual network topology to achieve higher performance and quality with decreased costs are:

- IPTV network running on top of a virtual network may result in high quality, uninterrupted audio and video streaming achieved by dedicated bandwidth and isolated network resources. If possible, when handling a VoD request from a user or multiple users located in different terrains, combining with cloud services to instantiate a VoD server for that certain locations may provide increased download speeds therefore reducing download times.
- VoIP: By creating a virtual network with dedicated resources for VoIP service will ensure the highest possible voice quality.
- Enterprise services enabling the “virtual enterprise” discussed in detail in section 3.4. Virtualisation combined with cloud services enables the easy distribution of applications and traffic/loads to ensure optimisation and extension of resources.
- Online Gaming: Gaming companies started to build their business cases on online gaming since it became almost impossible to prevent the gaming software’s being cracked and therefore distributed freely over the Internet. The focus changed direction to providing the best possible multiplayer gaming experience with less bandwidth requirements. Virtual networks seem to be the most convenient solution to the problem since it may provide dedicated virtualised resources and may be adaptively, dynamically maintained to achieve highest user experience.

3.5.2 Business model and role of stakeholders

In order to explain the business model and roles of the stakeholders, it is assumed that the scenario takes place inside a typical operator with services and infrastructure divisions.

Within the telecom operator, Services Divisions (SeD) will demand the required VN topology to be prepared from the Infrastructure Division (InD), which decides whether or not it can fulfill the request. If the request can be satisfied, the InD composes the requested VN and hands over full control to the SeD, which constructs the desired service and starts serving its subscribers over the VNs.

In order to explain the SeD of an operator in detail, the services division is divided into four major sub-divisions.

- TV/Content Services Division’s main target is serving IPTV solutions and distributing VoD content over an IPTV network.
- Telephony Services Division deals with VoIP services.
- Enterprise Services Division composes private and highly secure virtual networks for companies.
- Gaming Services Division deals with online multiplayer gaming service and its objective is creating the best possible gaming environment for its customers.

3.5.3 Basic requirements

Virtual services networks running as hosts to commercial Telco services can be classified as:

- TV/Content Services require highly reliable virtual network as well as, if possible, dedicated bandwidth or virtual ports since the service will have high volume traffic for each end user.
- Telephony Services: Because delay in a VoIP call will result in major degradation in user experience, VN for telephony services must fulfil features which guarantee low delay for voice over IP calls.
- Enterprise Services: Security as well as privacy are some of the most important concerns for a company considering a virtual network. Therefore, the VN provided for a company must require a highly secured authorisation procedure and be invisible to third party users.
- Gaming Services: Online gaming probably will be the most intolerable service that can be served over a virtual network, in terms of reliability and low delay. Online games will demand an uninterrupted, low delay performance therefore highly isolated network resources for gaming platforms; otherwise the gaming experience will probably not fulfil the expectation of the customer and thereby causing churn.

3.5.4 Gaps/open issues

The gaps and issues based on the requirements identified above can be identified as;

- Dealing with the increasing number of services as well as increasing number of subscribes for each service would be challenging. Scalability and better yet dynamic resource allocation must be considered thoroughly when designing and operating the virtual network and corresponding services.
- In the near future, it is possible and highly necessary to build virtual networks supporting cloud based applications. Integrating the aforementioned telco services with the related applications hosted inside operators' clouds must be seriously considered.
- Isolation is required between each virtual network running separate telco services. However, determining the level of isolation and applying it for different types of resources will not be a trivial task. Depending on the business model and service type, a virtual network might require complete isolation while another one might share its resources when idle.
- In this section the operator is assumed to have the entire required infrastructure to provide its services. However, in some cases additional telcos acting as infrastructure providers for a virtual network might be required. In such cases dealing laterally with multiple telcos must be considered an important issue.

3.6 Experimentation

3.6.1 Problem/Scenario

In this scenario network virtualisation technology is used for running experiments on production infrastructure next to production traffic. Currently, testing requires separate networks and environments, which differ in scale and quality from production infrastructure and lack real-world characteristics desired by R&D and experimenters.

3.6.2 Business model and role of stakeholders

Running experiments on the same infrastructure alleviates the need for a large number of separate and isolated experimentation networks. At the same time, this would allow network experiments to run on a larger scale than in confined lab and research networks and to close the gap between them and the production environments.

The business model here is internal to service providers or telcos, which want to make use of their existing infrastructure for research and development purposes. The infrastructure division would provision virtual network circuits or slices on their existing production networking infrastructure to either internal customers (R&D, product development and technology) or their own uses. The legacy production network would either run as-is on the current infrastructure or be provisioned in the same way as virtual networks.

The goal is to allow internal customers/business units (BUs) to access the production infrastructure for their experimentation and testing needs. Thus the requirements for separate test networks in each of

the BUs could be reduced, leading to lower CAPEX for building and upgrading them. OPEX could be reduced through a reduced need for non-operations BUs to run their own testing networks.

This scenario thus tries to lower the bar for large-scale experiments and moves testing of network services, protocols and topologies closer to the production environment. Furthermore, it lowers network operations-related CAPEX and OPEX throughout the business.

3.6.3 Basic requirements

In order for experimentation and research to be safely allowed on virtual networks on the production infrastructure, several requirements need to be met;

- Isolation (security): The experiments must run completely isolated from the production network and traffic. Isolation is required on the data links, circuits, networking devices and management systems. Virtual networks for experimentation must not have any access to production traffic.
- Isolation (resources): Experiments on the same networking substrate will need to be provided with limits and quotas for their bandwidth, CPU, latency and QoS.
- Reliability (production/legacy traffic): The existing production traffic must not be affected by the experiments in its normal operation. This requirement encompasses both isolation requirements.
- Configuration: The customers and BUs that access virtual networks for experimentation need configuration access to their (virtual) network devices in order to run their experiments or demonstration traffic. This also applies to network managements and related systems. A whole-sale management through the InD might be feasible but could make the use of virtual networks less interesting for research, due to limited configuration and management access.

3.6.4 Gaps/open issues

Allowing experimenters access to the production backbone might pose significant technical and political questions:

- Isolation for both security and resource allocation needs to be proven in real-world scenarios and tested in edge cases. Testing this will prove to be hard to plan and execute, as experiments will try to push the envelope of available networking resources. To be considered for successful roll-out the isolation both for security (information) and resources needs to be completely secure.
- In addition to the technical isolation requirements, a political consensus is needed between all divisions in the company to allow use of VNs on production infrastructure. Given the technical requirements are met, the infrastructure division still has to allow researchers and experimenters access to the production infrastructure.
- It needs to be investigated if the virtual networks for experiments need host virtualisation on the network device level as well. Limiting the virtualisation to the links or circuits may not be sufficient for all experiments, which may need access to (logical) network devices in the path as well.
- Isolation needs and “real-world data” requirements: Some of the use cases for experiments on real-world production system depend on information about the legacy/production traffic to monitor effects and draw correlations. This conflicts with the isolation requirements of the production traffic and legacy network to not allow experiments access any other resources outside their virtual network.
- Integrated/automated configuration: The configuration/provisioning process for virtual networks needs to be automated and integrated to a large extent to lower the entry barrier for customers of this service. Since many resources will be affected by the virtual network provisioning for experiments, modifications and configurations will be required for a variety of systems. If this process ends up being too complicated or technology-heavy, the targeted BUs may simply turn to their own resources and build separate test networks.

3.7 Technology migration

3.7.1 Problem/Scenario

Network virtualisation could ease the introduction of new networking technologies into operators' backbone and access networks by making parallel or staged deployments on the existing infrastructure easier. Introducing new technologies or protocols (e.g. IPv6, Multicast) on parallel, virtualised parts of the network could lessen the impact of this transition by not affecting all stakeholders and networks at the same time.

In the example of IPv6, the new protocol could be introduced in well-defined parts (slices) of the network by limiting the transition to single virtual slices of the whole network. Current technology introduction efforts in backbone networks are sometimes hampered by scalability issues and negative impact. Slicing this network by confining the introduction to specific slices (e.g. single customers) could soften and distribute the impact.

A similar scenario would be the introduction of new routing protocols or methods into the backbone. By using virtual networks, the new routing protocol could be limited to small slices of the network first and gradually expand to include more space and traffic. Virtual networks could here reduce the impact and scope of fundamental changes to backbone technologies.

3.7.2 Business model and role of stakeholders

Using virtual networks for technology migration has several business cases by limiting the scale of the efforts to smaller scopes at first. Thus, the operator would gain a relative market advantage by allowing the introduction of potentially disruptive technologies at a higher pace with smaller and more targeted efforts. This would allow quicker reactions to innovations and market demands.

On an operational level, virtual networks could lessen or distribute over time the OPEX needed for introducing disruptive technologies; CAPEX could be reduced by alleviating the need for dedicated infrastructure in case of new technologies and conflicting customer demands – legacy networking and new protocols could run on the same infrastructure for different customers, making investments for parallel devices for new customers/technologies unnecessary.

The business case here is again internal to the company – the infrastructure division would provide other BUs with virtual network slices. These BUs could be R&D, product development or technology integration that would be involved with a roll-out of new networking technologies into the backbone. The infrastructure division could also take the role of VNO and run the virtual networks completely, or the role of VNO would be distributed to the business units involved in the technology introduction. More complex relationships are possible, where the role of VNO/VNP could be unified in a special business unit that provides virtual networking services for technology introduction and experimentation. Other BUs would cooperate with this unit or act as customers.

The end-users of the business would not be cognisant about any of the underlying relationships.

3.7.3 Basic requirements

The following requirements need to be met for migrating networking technologies using virtual networks:

- Isolation (stability and resources): Technology migration on/with virtual networks would need complete reliability for resource allocation and isolation. This includes issues such as QoS, bandwidth allocations, resource allocation on network devices in the path and other connection parameters.
- For the infrastructure division to safely open their infrastructure through virtualisation to other BUs a (resilient) process needs to be in place that defines the roles and responsibilities of all stakeholders. Infrastructure divisions could be wary of opening up their domain of responsibility to outside BUs; opening up the process of technology migration/introduction could place more pressure on them due to higher frequencies of change and innovation and the prospect of relinquishing control.
- Virtualisability: To allow introductions even in a limited scope, all systems of the network in question have to support network virtualisation.

- Reliability: The existing production network and customers must not be affected by the parallel introduction of new technologies on virtual networks on the same infrastructure. The isolation and performance metrics should be at the same level as the legacy networking infrastructure.

3.7.4 Gaps/open issues

Various open issues could arise when migrating networking technologies on virtual networks:

- Given a particular networking technology is successfully introduced in a VN – how will this technology be transferred either to other sibling VNs or the underlying backbone infrastructure? I.e. at which point will the different networks converge and how will an operator overcome technical incompatibilities between technologies in different VNs? If network virtualisation will be widespread in the provider's backbone and multiple VNs coexist for different purposes, this could lead to high OPEX maintaining these different networks and technology introduction might have to be repeated for all single separate VNs.
- Similar to the point above is the question how OPEX would be affected by running several concurrent VNs with different technologies. Technological change and migration might not happen as fast as planned, extending the life-span of VNs considerably. Each of these VNs will need to be maintained and operated, and if no streamlined solutions for automated or assisted network operation can be developed, the increase in OPEX for operating multiple VNs on a single infrastructure might be considerable.

4 Analysis of gaps and open issues

Analysis of gaps and open issues has been carried out regarding technical, operational, business and regulatory issues. Gaps and open issues mentioned in Section 3 with other relevant issues are summarised in Table 3, Table 4 and Table 5.

4.1 Technical issues

Table 3 – Technical gaps and open issues

Requirement	Explanation	Gaps, open issues	Impact
Reliability and resiliency	A virtual network should be as reliable as a physical counterpart.	Reliability depends primarily on the physical network. Network virtualisation adds an additional complexity layer, which represents an extra potential source of failure.	Virtual networks may be seen inadequate by potential virtual network providers, as carrier-class reliability is impossible to guarantee.
Scalability	Scalability must enable applicability in very large scale scenarios.	Isolation of virtual networks requires fine grained "per-virtual network" resource control. With the growth of number of virtual networks, this may become problematic.	Resource isolation may be impossible to guarantee in very large scale scenarios.
Interoperability of network virtualisation actors	When the roles of VNO, VNP and InP are played by different actors, standardised interfaces are required to enable interoperability.	Standardisation of interfaces is still missing.	Lack of standardised interfaces will make the establishment of a global network virtualisation marketplace difficult to achieve.
Interoperability of virtual networks	Virtual networks will have to interconnect with other virtual networks.	Although not strictly a network virtualisation issue, interoperability of protocols running *inside* virtual networks is required to allow interoperability of virtual networks.	Interoperability limitations will discourage early adopters, as VNs would be essentially isolated islands.
Interoperability with legacy networks	Virtual networks will have to interconnect with non-virtualised networks in multiple scenarios (e.g. virtual networks composed of non-contiguous physical segments).	Interfaces between virtualised and non-virtualised network domains to be defined.	Interoperability limitations will discourage early adopters, as VNs would be essentially isolated islands.
Heterogeneity and interoperability across multiple infrastructure domains	This requires that a single virtual network may be based on multiple heterogeneous infrastructure domains.	Building a seamless virtual network domain based on multiple heterogeneous infrastructure sub-domains is a challenging task which requires the definition of a standardised interface to abstract the physical characteristics of the network.	Lack of standardisation of inter-domain interfaces will hinder large scale deployment of network virtualisation, or limit interoperability.
QoS support	Any kind of QoS policy (e.g. definition of classes of service, resource reservation mechanisms) that might be applicable to a physical network should be equally applicable to a virtual network.	Current data plane QoS mechanisms (e.g. scheduling, policing) are based on the direct control of physical resources. Addition of the virtualisation extra layer will make resource management more complicated.	Lack of QoS guarantees will discourage adoption of NV-based solutions.
Isolation of Resources	VNs should offer the same level of isolation of any network supported by	Scalability issues depending on the level of isolation are still an open issue,	The level of isolation of resources affects the QoS policy and security mechanisms of

	physical resources.		VNs.
Programmability	Programmability is required in research environments in the cases where different and/or new network architectures and approaches are deployed. Further, it could become an effective means of VN management, enabling a high degree of flexibility and dynamicity.	Commercial network equipment vendors are not likely to provide open software platforms, or means to virtualise hardware or software.	Without flexibly programmable equipment, protocols and network architectures implemented in commercial virtual networks will be dependent on the respective vendors.
Security and privacy	Control and management functions must be properly secured.	As VN management interfaces may be exposed to external parties, VNs are vulnerable to abuses that could e.g. be used to compromise the infrastructure, specific virtual networks, or to break mutual isolation.	Security and privacy features of network virtualisation must be clearly demonstrated; otherwise its applicability in business environments will be very limited.
Functional separation of infrastructure and virtual network	A VNO should be able to run a virtual network independently of external agents, namely the providers of the infrastructure.	Standardised VNO/VNP and VNO/InP interfaces will be required.	If an effective separation of functional roles is not possible, entry of new players to play the role of VNO will be disincentivised.

4.2 Operational issues

Table 4 – Operational gaps and open issues

Requirement	Explanation	Gaps, open issues	Impact
Configuration access to virtual nodes	Access to virtual nodes by multiple VNOs with differing rights	Depending on the flexibility of virtualisation, what kind of rights to be offered to VNOs is not clear yet.	Limitation of the number of VNOs to access a virtual node and/or the rights to be offered to each VNO might be considered by InP due to its operational concerns.
Automated configuration / provisioning of virtual networks	Necessary for the dynamic creation and modification of a high number of virtual networks	Scalability issues depending on the number of virtual networks	Seamless changes in the virtual network configuration are crucial for the end-user's satisfaction.
Configuration and operation complexity	Introduction of new layers and roles for the VN management could result in an unmanageable complexity.	Clarification of the actor roles by defining the standards is necessary to ease the operation complexity.	If configuration and operation of virtual networks ends up being too complicated or technology-heavy, potential users may turn to other alternatives.
Interoperability with legacy networks	Large scale deployment is complicated because all systems must support virtualisation.	In certain cases a gateway may be required when interconnecting to legacy networks.	Interoperability with legacy networks is essential at the early stages of deploying network virtualisation but its necessity diminishes with the transformation of the non-virtual networks to virtual ones in time.
Support programmability for experimentation purposes	Programmability is required to conduct experimental projects.	Vendors are reluctant to provide open source interfaces controlling their products. Programming flexibility should not reduce the reliability of the existing infrastructure.	Experimentation reduces CAPEX and OPEX of companies by identifying the possible real-life problems before deploying a new technology commercially.

Dynamic resource allocation	Dynamic allocation of virtualised resources to VNs from the physical resources of InPs	Scalability issues depending on the number of VNs and how dynamic the network is.	Inefficient dynamic resource allocation would decrease the utilisation of the virtual networks and result in customers' dissatisfaction.
Vendor interoperability	Governing standards and open interfaces between vendors are required.	Standardisation activities have started but not reached maturity yet.	Vendor lock-in
Integration with cloud-based technologies and services	With the emergence of cloud computing, integration of virtual networks with cloud based applications will be increasingly necessary.	Standardisation of APIs that will be used by various cloud services providers.	Potential to lose an important application area and the corresponding customer (cloud service provider)

4.3 Business and regulatory issues

Table 5 – Business gaps and open issues

Requirement	Explanation	Gaps, open issues	Impact
Sound business model	An attractive business model for all players involved is a crucial requirement for any network service to be successful.	It is difficult to find a proven NV business model. For example, in the NaaS use case, it is not clear that the InP role will be attractive for operators.	Without a clear business model it is unlikely that NV will take off as a commercial service (e.g. NaaS).
Attractiveness to network operators	The role of infrastructure provider only essentially reduces the operator to the provision of a commodity service.	Lack of business incentives to deploy network virtualisation by potential InPs (today's network operators)	High possible barrier against roll-out of VNs or the main force of VNs may move else-where (free/wireless networks) and leave out operators.
Accountability	It should be possible to clearly identify responsibility for SLA violations.	Definition of roles and respective interfaces is required but so far this has not been achieved. On the other hand, given the wide range of possible scenarios, this may be difficult to accomplish.	Lack of accountability will make effective implementation of SLAs impossible, which will hinder the utilisation of virtual networks in commercial environments.
Political consensus to access virtual resources	Political will is needed within operators to open up the infrastructure via VNs to other parties.	The infrastructure owners may be reluctant to share their network resources with others. New attractive collaboration models are needed.	Opening up the infrastructure could lessen the role of infrastructure divisions to mere providers of wires and move authority to other entities.
OPEX reduction	Network virtualisation should enable OPEX reduction in order to become an attractive solution to operators.	Potential increase of OPEX to run concurrent virtual networks with different technologies (in a technology migration scenario)	Longer-term impact on OPEX not yet known. Attempts on VNs roll-out need to be carefully planned to prevent backlash due to higher OPEX.
Regulation	Regulatory environment should not discourage adoption of technology.	The introduction of Network Neutrality is unclear.	Network Neutrality could discourage Network Virtualisation depending on regulator interpretation.

4.3.1 Network Neutrality

In simple terms, network neutrality means that ISPs may not discriminate between different kinds of content and applications online. It proclaims that broadband providers should not be allowed to use e.g. their last-mile infrastructure to block or prioritise Internet applications and content and potentially obstruct desirable competition. Network neutrality primarily applies to the public Internet proclaiming a secure, fully open, unbundled and equal access to everyone – without setting any restrictions to any

kind of content being accessed/consumed. Several civil bodies have been formed to advocate for network neutrality.

Network neutrality proponents claim that network service operators seek to impose a tiered service model in order to control network service levels and thereby hinder competition and create an artificial bandwidth scarcity. Examples of prioritisation such as intentional slowing of Peer-to-Peer and FTP services, have been used to argue against data discrimination of any kind.

It is difficult to determine whether network neutrality will affect virtual network services directly. It is however obvious that the concept of network neutrality has the potential of detrimentally affecting the networking business as it prevents network operators to use any means of prioritising and impacting the usage patterns of their networking pipes in relation to Internet services/traffic. This means that operators cannot for example, impose prioritisation between their own local managed services, e.g. IPTV services and other regular free Internet services e.g. YouTube traffic. Furthermore, this means that operators are being pressured to update and reinforce their networks without necessarily any monetary gains according to the past and forecasted increase in data traffic. This is potentially a deadlock situation for operators as they are expected to provide sufficient bandwidth for increasing Internet data traffic without little or any reimbursement. Regulators have the tendency to extend their regulations as far as possible and therefore it should not come as a surprise that a potential future regulation in favour of net neutrality should limit operators' options severely to efficiently manage their data traffic prioritisation levels.

Case: Iceland

With its small population and remote location, Iceland only has a few ISPs, relying upon three submarine cables for delivering global Internet traffic. Primary peering points are in London and Toronto. With the sparse capacity it is important to prioritise Internet traffic at the peering points. For this purpose, network operators operate network elements at the network edge (Peering points) for managing individual IP traffic flows through Deep-Packet-Inspection (DPI) and differentiated classes of service (DiffServ). Prioritising and classifying data flows according to their nature, e.g. Peer-to-Peer and web traffic, enables operators to provide continuous and sufficient service levels and prevent pervasive traffic bursts and bandwidth saturation. This is depicted in Figure 14 where the MPLS network plays a central role both for Internet connectivity and for enterprise interconnections.

Although DPI technology has been used for Internet management for many years, some advocates of network neutrality fear that the technology can be used anti-competitively or to reduce the openness of the Internet.

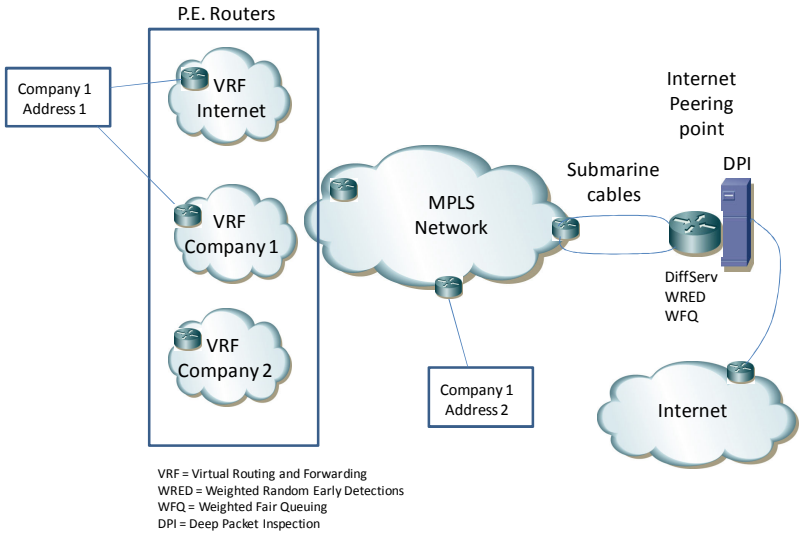

Figure 14 - Icelandic IP networking structure

4.3.2 Functional separation

How and if functional separation of service providers will impact VN provisioning is still unclear. Functional separation entails the separation of service provider business into two parts, the "Service" function and the "Network" function. Eurescom study P1754 provides detailed explanation of functional separation and its implications, especially in relation to regulatory issues and policy.

With regard to VN, it may seem that functional separation supports the vision of separated roles of InPs and VNPs where the InPs relate to the network function part and VNPs to the service function part. Clearly, both parts can belong to the same company or operator. The network side could indeed offer its infrastructure services to multiple players in the market and not just to the services side of the parent company operating a VNP business. Functional separation might appear to align well with the distinctive roles in the virtual networking environment as described in section 2.1.1 at first sight but a closer look reveals another result.

Implementing functional separation is a step towards so called service competition, where market players compete mostly in delivering services to customers. The service-companies have an equal or similar access to the company delivering the network infrastructure. Opposed to this is facility based competition where market players compete in offering not only services but also the facilities on which the services are offered. It is well known that facility based competition is more likely to foster new innovations and investments into new technologies. The virtual networking environment needs a high degree of innovation and new investments to build an ecosystem where competition at all levels thrives. Facility based competition is likely to lead to stronger and more innovative InPs than service competition and thus functional separation would yield.

Therefore, if operators will increasingly pursue functional separation particularly through regulatory pressures, it is likely to have a negative impact on network virtualisation evolution.

5 Opportunities and challenges for operators

Our aim in this section is to briefly discuss how virtualisation could take place in future networks and to give insight how feasible it is to fully implement VNs in reality, by elaborating the current challenges and the possible roadblocks.

5.1 New business models and opportunities enabled by network virtualisation.

The core of the new business models offered by the widespread use of network virtualisation is in operators' backbone networks and centres on the possible independence from infrastructure constraints and boundaries for offering new network services. New players could emerge as independent providers of virtual networking resources not only to fill some gaps but also to re-shape the existing landscape. Virtualising resources and links on the existing infrastructure could open up space for new networks and technologies, while not interfering with the legacy traffic and services.

The aim of most of the scenarios discussed in Section 3 is removing the constraints placed on service delivery and instantiation by the existing network and operating environment that makes the introduction of new services and technologies cumbersome and prohibitive. Virtualised networks on top of the existing infrastructure could free up Telcos to experiment and integrate new services without affecting their existing business. The majority of scenarios use network virtualisation as delivery model for new services ranging from creating tailored networks for individual customers for cloud services and computing access to offering complete network packages as a commercial service. Operators would give up some of the authority of their network into the hands of internal and external customers that buy these resources or offer service on them.

One of the main attractions of network virtualisation is the probable benefits in flexibility currently lacking in the legacy networks. After successful rollout and configuration, provision of networks could be as easy as a simple service request for a new network that spans multiple providers – InPs, VNOs and VNPs – that are invisible to the end-user or service providers. Reconfiguration would be made much easier if the process of configuring is confined mostly to the virtual network parts, while the task of running the physical hardware and circuits is done by the InPs or infrastructure divisions.

This would lead to a model analogous to today's electric power brokerage and distribution, where energy bought by the customers from different sources is input from the providers into the same "virtual" power grid. The customer has the freedom to choose his provider(s), whose role may include running the actual power grid or solely buying and brokering electric energy and taking care of proper input and distribution.

Aside from improving existing business and delivery models, a whole new class of operators could emerge that only acquire and resell VN resources from existing operators and offer these virtualised bundles as VNs directly to customers. The benefit for customers would be getting a large-scale network from a single source, the virtual network providers.

The special requirements of data centres could lead Telcos to implement virtualised network there as first production trial runs and commercial offerings. The proximity of computing and network systems in a single location could make a rollout of network virtualisation very interesting to DC operators and lead to an early adoption of technologies and solutions there. However, these same special demands could make a migration of network virtualisation out of the data centres into more diverse types of more demanding networks difficult, since the requirements and operational environments differ so greatly.

In the ongoing process of the large-scale rollout of fibre links close to the customers (FTTH/FTTC), the financial burden of these undertakings could lead to a closer collaboration between operators and other parties. Building and running the new fibre and wired networks wholesale cannot always be done alone by a single player anymore, and thus new ways of collaboration and sharing of resources have to be found. Network virtualisation would offer an interesting way forward for operating and managing shared infrastructure resources by multiple parties on multiple levels. Telcos could offer to run all parts of the network, from the physical infrastructure to virtual network provision, while

smaller or outside players could opt for focusing on the virtual networks on top of virtual network providers.

Having most parts of the network infrastructure based on VNs would allow operators to more easily and rapidly adapt to changing market environment and demands. If a network is already provisioned on a virtual basis, new market dynamics can be adopted with less effort. Following trends, new virtual networks and resources could be added and removed quickly without the burden of building and maintaining specialised infrastructure and network for each new service or business model.

Network virtualisation could lead to a break of the existing market monopoly of the current network equipment vendors. Today's offerings of horizontally and vertically integrated solutions preclude market entry of new vendors and technologies that may threaten the business model of existing vendors. By opening up the networking hardware with approaches such as Stanford's OpenFlow protocol, the market could be changed rapidly by new players if a wide-scale adoption and requirement of OpenFlow occurs. Having a single baseline of network functionality by providing OpenFlow could lead to much more competitive market that would free up operators to use different devices and technologies in their backbone, whose common interface would be the OpenFlow protocol. The goal is to have a similar "business model" as today's open source world of software packages and operating systems, that challenged the existing software and system vendors and cleared the path to significantly lower the CAPEX on hardware. Virtualised networks on top of commodity hardware could lead to a comparable reduction in CAPEX and increase in vendor independence in operators' networks and data centres.

5.2 Network virtualisation challenges

Deployment of network virtualisation may be hindered by a number of issues, most of which result directly from gaps and issues identified before. To properly evaluate how challenges can represent real obstacles and affect the widespread deployment of network virtualisation, it will be necessary to analyse the different use cases separately, as they are differently affected by the various challenges.

Carrier-grade reliability

Virtual networks must be reliable, at least as reliable as a physical network counterpart. Major vendors, such as Juniper, have launched products with network virtualisation capabilities; however, most of these products are mainly targeted at the high-end segment of the market. On the other hand, very promising, flexible and adaptable technologies such as OpenFlow are perceived as research tools and have not yet reached a point of maturity to enable large scale deployment. This gap is expected to be filled by a number of vendors in the coming years but it is still not clear how long this process will take and how far it will go.

Scalability

The importance of scalability as a network virtualisation requirement is particularly relevant in the cases where the number of VNs is expected to grow. This number, for example, can easily grow to the order of hundreds or maybe thousands for "VN as an Enterprise Service" scenario based on the VPN experience. Several issues have to be analysed in this kind of scenarios. For example, it is not clear to what extent traffic can be aggregated while still preserving fine grained resource control, as well as migration flexibility.

Isolation

A virtual network is supposed to offer the same level of isolation of any network supported by physical resources. Since virtual networks are by definition based on shared resources, this represents a challenge, especially in the cases where a high number of virtual networks share the same infrastructure. On the other hand, the strictness of isolation varies according to the specific use case. In a "Network as a Service" scenario, isolation will obviously be a fundamental requirement; on the other hand, isolation is not as important in the network partitioning scenario.

Security

Security issues will certainly represent a potential factor to discourage adopters of network virtualisation, at least in the initial phase. Information privacy is crucial and must be guaranteed by whatever mechanisms are available (e.g. data encryption). Also, a well-defined degree of isolation has

to be guaranteed to prevent virtual networks from mutually affecting each other adversely — either purposefully (e.g. DoS attacks) or inadvertently.

Interoperability

Many VNs will span multiple network infrastructure domains. Interoperability is a crucial requirement to enable widespread deployment of network virtualisation. Standardisation will be required to enable interoperability between VNs, as well as interoperability between virtualised and non-virtualised networks. In addition, the interface between VNOs and InPs also needs to be standardised to ensure smooth interoperability.

Operational complexity

In cases where multiple stakeholders are supposed to be involved in provisioning and controlling of VNs such as in the NaaS case, operational complexity represents a significant challenge. For example, once the VN is established and running, the role played by the VNP seems to be redundant. On the other hand, removing the VNP from the process would make the establishment of multi-domain VNs very complex.

Quality management and handling of failures

VNOs are supposed to provide services to end users, and service quality should be monitored and measured against specific SLAs. The extra level of complexity introduced by virtualisation complicates this process significantly. In such an environment, the handling of failures is another complex challenge, as the identification of the source of the problem might not be trivial.

Programmability

Programmability is supposed to represent one of the main features offered by network virtualisation in research and experimentation environments. This may be feasible using software-based routers, which are based on open source, however, it is so far not feasible in commercial equipment.

Accountability

Accountability is required to track and report utilisation of physical resources and is essential from the business and regulation perspective. The severity of this challenge would be proportional to the number of end users in the virtualised networks and amplified further with the diversity of quality of service demands.

Monetisation

Last but not least, to become a successful technology, network virtualisation must bring significant economic advantages to all the shareholders involved. Operators will only be incentivised to build virtual networks if a sound business model can be built. For example, it is not clear that the role of InP will be an attractive proposition to operators.

5.3 Possible roadblocks

Overcoming all of the challenges is time consuming and it would be a good idea to resolve the most challenging issues in use case scenarios. As explained before, each network virtualisation challenge should be evaluated in the scope of a specific use case. Table 6 summarises how important the challenges mentioned in section 5.2 are for the scenarios identified in the section 3.

Table 6 – Network virtualisation challenges vs. use cases

	Scenarios						
	Cloud computing	CDN	Network as a Service	VN as an enterprise service	Network partitioning	Experimentation	Technology migration
Carrier-grade reliability	***	***	***	***	***	*	***
Scalability	**	**	**	***	*	*	*
Isolation	***	**	***	***	**	*	**
Interoperability	**	**	***	**	**	*	***
Security	***	***	***	***	***	*	***
Operational complexity	***	**	***	***	**	*	*
Quality management	**	***	***	***	**	*	***
Programmability	-	-	-	-	*	***	*
Accountability	**	**	***	***	*	-	-
Monetisation	**	**	***	***	**	-	-

*** Crucial challenge; will surely represent an obstacle against deployment if an appropriate solution cannot be found.
** Major challenge; may represent an obstacle to widespread deployment.
* Minor challenge; should be addressed, but does not represent a major obstacle in this specific use case.
- Not relevant in the scope of this specific use case.

From use scenarios perspective, carrier-grade reliability and security are two of the most crucial challenges and seem to be the possible roadblocks for network virtualisations. In fact, realisation of VNs is not possible unless they are as secure and reliable as non-virtual ones. VNs' solutions to these challenges would be a requisite for acceptance of network virtualisation both within operators and from outside customers.

In addition to this, standardisation is also the key requirement not only to clarify the unclear points in VN architecture and regulations but also to resolve the interoperability issues between different stakeholders and between virtual networks and legacy networks. Hence, standardisation expedites the process of deployment of VNs and not having mature standards can be seen as a possible roadblock. Without standards, independent hardware and software vendors, who play key roles for VN uptake, cannot engage in this process in a unified manner. In this way, the invaluable resources of the stakeholders such as time, money and manpower are not wasted and the challenges and open issues are resolved through the collaboration of the stakeholders. Standardisation is still in a relatively early stage, but it should be noted that relevant activities in IRTF [11] and ITU-T [12] have recently started.

Last but not the least, political will is required not only for offering new services and taking on new business models but also for giving up authority and control of the existing infrastructure and traffic to new internal and external operators that run virtual networks. The possible economic advantages such as OPEX reduction and increase in profits via offering new services would attract the existing and new stakeholders and expedite the deployment of VNs.

5.4 Possible threats to telcos due to NV

The focus of this chapter has been on the opportunities and challenges regarding the introduction of network virtualisation, what issues might be in the way and what might be a positive outcome for telcos. There is also the aspect that NV could jeopardise telco operations and even their existence. Today, telcos are generally vertically integrated entities owning or ruling all major parts of the telecommunication services value chain. The mobile services constitute a good example, the telcos build, own and operate the infrastructure, provision the services and maintain customer relationships from all sides. However, the introduction of the Mobile Virtual Network Operator (MVNO) changes the telcos' role dramatically. MVNOs use telcos' infrastructure to offer their services to customers. The customer relationship is completely taken over by the MVNO and the customer does not know whose infrastructure the services are delivered on. The MVNO-telco relationship is that of a wholesaler-retailer relationship. The InP (telco) competes with other telcos for offering the infrastructure and is forced to run on low margins due to regulatory interventions or competitive reasons. The MVNO has a possible threatening position to take his business to another InP if he is not satisfied with the current InP's pricing model or other aspects of the service. The telco takes the investment risk for building up the infrastructure, whereas the MVNO has very limited or no liability.

On the other hand it can also be argued that if an MVNO wants to offer a new or improved service an advancement of the infrastructure would be needed. The InP may have many MVNOs operating on his infrastructure and only one of them has an idea for this new service. The InP would thus have little incentive to invest in the improvement.

This scenario is likely to apply for all telco business areas where NV is applied. In fact the last point demonstrates that NV where infrastructure and services are held in separate entities is likely to hinder innovation. This is exactly one of the main arguments against functional or structural separation which means that the infrastructure is split from the service company.

Another example can be taken which demonstrates that NV could be a limiting factor for innovation. If a vertically integrated telco wants to set up a new innovative service, (e.g. 3D-IPTV) and the ADSL network needs to be upgraded to accommodate more speed, the telco simply decides to upgrade his ADSL network, implement the new service and gain a competitive edge. If the new service proves to be popular, all the competitors will follow and upgrade their own networks or build new. The telco will nevertheless have a head start on his competitors and thereby gain a larger market share because he was first to market.

If this same thing would happen after NV has been implemented, a VNO would get the good idea to be the first to offer 3D-IPTV. He would ask the InP to upgrade the network but it is very unlikely that the InP would be willing to invest in the upgrade. The upgrade would probably not increase the InP's revenues, being a bit carrier is not a profitable business and the InP will get the same revenue from delivering 25 Mb/s in a year as he gets now for delivering 12 Mb/s. Therefore, there is no return on investment. Additionally, the InP has many other VNOs as customers who are not demanding a network upgrade. They would oppose loudly against any price increase. This demonstrates that NV can potentially prevent innovation. Stakeholders in telecommunications must devise ways to introduce NV without this undesirable side effect.

6 Concluding Remarks

6.1 General conclusions

Network virtualisation seems to be a promising technological solution for operators to improve their existing service models and build upon new business cases.

There is currently a strong push in the field of network virtualisation research to standardise technologies and widen the scope of solutions. Existing solutions are either commercial single-vendor solutions integrated into their vertical business portfolio or wider concepts still in the academic and industrial R&D phases. Promising projects are looking at opening up commodity networking hardware to allow greater programmability, which in turn would make larger-scale virtualisation solutions across device and vendor barriers possible.

Many of the current research projects concentrate on top-down, business and organisational-heavy concepts for architecture and operation of virtual networks and their providers. Comparably fewer academic projects are devoted to investigation of technological aspects of virtualisation, with a growing momentum building around the software-defined network theme.

Looking at the bigger picture, most of the current virtualisation solutions are either looking at the link layer or at individual hosts or network devices. Looking at the use cases this study came up with, many barriers still need to be overcome, both in functionality and in reliability. For fully integrated, wide-scale virtual networks and services, these solutions need to be integrated and proven to be carrier-ready. All use cases require reliable services for the legacy networks and additional, virtualised offerings. Adding new virtual networks to a carrier's network imposes additional requirements for security and isolation. While the technologies and potential solutions look promising, industrial research projects need to prove how they match operators' requirements for security and reliability and their ability to work in a large-scale environment, outside of closely-supervised academic tests.

Granted the new technologies and concepts gain traction and prove themselves to fulfil operators' requirements and achieve wide-spread industrial adoption, these technologies could play an important role in operators' effort to reduce CAPEX and OPEX and cope with new business models and services. It remains to be seen if these technologies will allow for completely new business fields and markets or will be another tool in running backbones and provisioning networks and services.

6.2 Long term visions

The long term vision of network virtualisation will be fulfilled mainly through achievements in three key areas.

Convergence of IT and Telecoms

Convergence presents a vital part of an overall cost-saving solution. With convergence, business processes are enhanced and time for taking decisions shortened, saving time and money. Network virtualisation is expected to bridge the gap between the IT and Telecom sectors.

With a converged solution, there is a single-point of visibility of IT and telecoms operations, providing the opportunity to manage both more effectively. There is also a practical standpoint to convergence as dealing with one single supplier for a solution rather than two separate providers. This helps to reduce costs and save time through simplified interaction.

At the same time, the company can also enjoy improved business agility and better collaboration between employees to enable innovation and allow new services to be brought to market much more quickly. In fact, convergence has contributed directly to the emergence of new solutions and better ways to serve customers as a dynamic enterprise.

Virtualisation and Cloud Computing

Security is still a major concern slowing down cloud computing adoption in the IT sector. It is critical to ensure data privacy and integrity in the cloud at a level that is at least comparable to that in current enterprise networks. The current cloud computing services do not provide isolation of computing resources and networks between customers. It is essential for cloud service providers to take advantage of resource sharing and multiplexing among customers. Virtual machines of different

customers may reside on the same physical machine, and their data packets may share the same LAN. Such lack of isolation brings security risks to users. On one extreme, cloud service providers may offer very secure systems by creating entirely separate hardware, software and administrators for special customers. However, this kind of solution is very expensive. Network virtualisation techniques can logically separate different networks on the same hardware and partition resources accordingly. This feature is useful for providing good isolation as well as network resource sharing among different users. Network virtualisation techniques may offer less expensive solutions to the security issues with cloud computing and boost the adoption of that technology.

Network Virtualisation and Future Internet

Future Internet is a general term for research activities addressing the shortcomings of the current Internet architecture. The number of end-points, quantity of data and number of services is growing such that operational and management costs have increased very significantly. Mobility, controllability and quality of service are poorly addressed and these issues have become more visible with the spread of mobile applications. Trustworthiness, security and reliability are still major concerns of the existing architecture. It is also desirable for the Future Internet to play a vital role in search for solutions to global issues such as energy, health and education.

Network virtualisation is a very promising technology for the Future Internet thanks to its high degree of flexibility and cost-effective nature. In fact, the existence of the technology totally depends upon addressing the aforementioned issues of the current Internet. Even if the solutions without network virtualisation meet the requirements of the Future Internet, they do not seem to be economically viable.

6.3 Standardisation

Standardisation is usually a crucial requirement to enable the widespread deployment of new network technologies. In the case of network virtualisation, standardisation is particularly important because interoperation of multiple actors is typically involved. In a network virtualisation environment, interoperability is required at multiple levels: between different equipment vendors; between different players/stakeholders (VNO, VNP, InP, end users); between heterogeneous virtual networks; between legacy and virtual networks.

Perhaps first of all, a general network virtualisation framework is required, including the establishment of a common terminology and the definition of reference points and interfaces. Interoperability in a network virtualisation scenario typically involves two different dimensions, as shown in Figure 15 [33] – horizontally, multiple infrastructure providers have to cooperate to support VNs spanning more than one physical network domain; vertically, a set of standardized interfaces is required to enable the establishment and control of VNs, particularly in the cases where the roles of VNO, VNP and InP are played by different entities.

Figure 15 indicates the relevant interfaces – numbered 1 to 6 – following the architectural model developed in the FP7 project 4WARD, some of which are obvious candidates for standardisation.

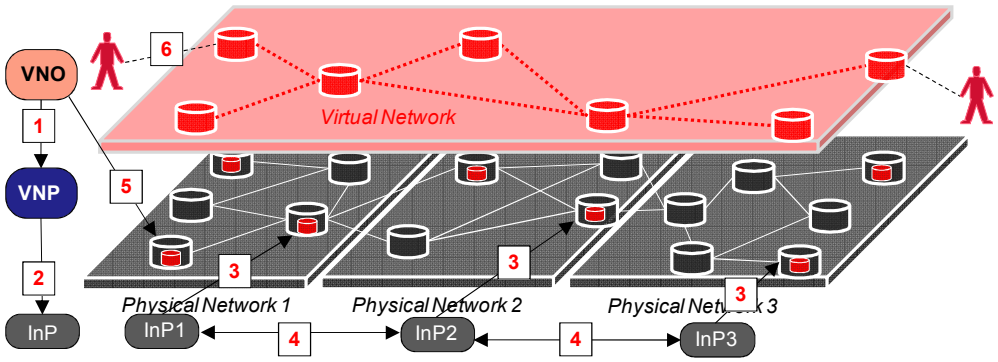

Figure 15 - Network virtualisation 4WARD model architecture interfaces

A few examples where standardisation is required include:

- VNP/InP (interface 2 in Figure 15): this interface is used to provision, setup, manage and reconfigure virtual networks, as well as to request and negotiate virtual resources and virtual networks. A standard resource description language is required to describe networks and network resources (VNPs to specify resources to be requested from InPs; InPs to describe resources provided to VNPs).
- InP/Network equipment (interface 3 in Figure 15): this interface is used by the InP to manage physical resources and setup virtual nodes and links. Lack of standards will make the process of building virtual networks cumbersome and dependent on the network equipment vendor.
- InP/InP (interface 4 in Figure 15): this interface is used to setup virtual links and virtual networks spanning multiple network infrastructure domains, including the case where two InPs are indirectly connected, i.e. through non-virtualised network domains, which is likely to be the most common case in an initial phase of network virtualisation deployment.
- Signalling for virtual link / virtual node setup: this would be required to enable automated establishment of virtual networks. Lack of standards in this area would hinder vendor interoperability and easy establishment of virtual networks, thus affecting network virtualisation scalability.
- Security and resource isolation: strict isolation of competing virtual networks sharing resources is required to prevent or mitigate the impact of DoS attacks or misconfiguration in neighboring virtual networks.

Recently, the first steps in network virtualisation standardisation have been taken, with two relevant initiatives. The ITU-T established the Focus Group on Future Networks (FG FN), in which network virtualisation is one of the fundamental topics. The following deliverables are expected from this Focus Group: future network benefits, future network vision, high-level description of future networks attributes, vocabulary, incremental terms required to address future networks [12]. The activity started in July 2009 and by August 2010, five meetings had taken place already. Although not exclusively focused on network virtualisation, relevant output is expected from this group, especially in terms of defining a general set of concepts and ideas.

In addition, the Internet Research Task Force (IRTF) created in early 2010 the Virtual Networks Research Group (VNRG), specifically focused on network virtualisation. According to the group's web page [11], "the VNRG provides a forum for interchange of ideas among a group of network researchers with an interest in network virtualization in the context of the Internet and also beyond the current Internet". The VNRG will produce Informational and Experimental RFCs in order to document the activity of the group and to formalize the outcome of the research topics carried by the group. In addition, such documentation could become input to IETF working groups. The draft "Network Virtualization Problem Statement" [34] has been the first outcome of this activity.

6.4 Areas for further study

Throughout this report, several topics which require further study have been briefly discussed. For example, each of the challenges identified in section 5.2 represent a potential topic to be further explored. This section describes some of the areas that may require attention in future research activities related to network virtualisation and is by no means exhaustive.

Migration to network virtualisation

Network virtualisation has often been heralded as a potential enabler of migration to new network architectures and solutions. However, another issue related to migration is often overlooked – how to enable migration from today's non-virtualized networks to a full-blown virtualization scenario. As a matter of fact, it is not fully clear how this migration path could be materialized in production environments, thus this should represent a topic for further study.

Fault management

Networks virtualisation poses difficult challenge for fault management, as faults and disturbances experienced in a virtual network must be correctly attributed to a root cause in the physical layer, while at the same time independence between virtual and physical networks must be guaranteed. Finding an adequate solution for this problem certainly represents a topic for further investigation.

Network virtualisation economics

So far, research in network virtualisation has concentrated mostly on technical aspects, whereas the business aspects have been relatively downplayed. Nevertheless, it is clear that finding a sound business model that suits all participating stakeholders is a fundamental requirement to enable the expansion of network virtualisation in commercial scenarios. Actually, one of the most frequent criticisms against network virtualisation is the fact that there is little incentive for network operators to invest in a technology which would mainly benefit potential competitors (prospective virtual operators).

Cloud networking

A significant number of organisations have embraced, or are in the process of evaluating, the concept of cloud computing. For IT this represents a fundamental transformation toward a utility model of computing, just like other commodities electricity, gas, or water. Extending the same concept to network services seems to be a logical next step. However, there are a number of new requirements in terms of elasticity and dynamicity which the present state-of-the-art is not able to fulfil. For operators this represents a new set of challenges.

Applicability to large scale environments

So far, network virtualisation has been deployed in small-scale research testbeds, or in operational environments but to a limited extent, only. The widespread use of network virtualisation in operators' network backbones raises a number of issues in terms of scalability that have not been fully evaluated up to now. For example, scalability tends to favour aggregation of resources and handle large chunks of networks (this is the case, for example, in today's MPLS core networks, where multiple VPN tunnels are usually aggregated in a single transport LSP). This goes against strict isolation of virtual networks, which require a fine grained control of resources.

6.5 Recommendations

The potential advantages of network virtualisation for operators are clear in several business scenarios and use cases. Because of challenges like isolation and scalability, it seems plausible that in a first stage network virtualisation will be exploited internally by operators to partition network infrastructure into multiple logical domains, to segregate different types of services (e.g. VoIP, IPTV, Internet) or to enable coexistence of different network technologies (e.g. IPv4, IPv6). Virtual networks are likely to be offered as a commercial service only at a later stage, when technology maturity issues are finally sorted out.

In summary, based on the analysis performed in the previous sections, the following recommendations can be made:

- The decoupling of networks and infrastructure (and consequently the roles of network providers and infrastructure providers) has a potentially disruptive effect on operators' business. For incumbent operators, this brings new opportunities, already analysed before, but also new threats. Easier establishment of networks will enable increased competition from new entrants; in addition, the economic attractiveness of providing network infrastructure, which basically becomes a commodity, may be questionable in many scenarios. Thus, commercial pros and cons of network virtualisation should be carefully evaluated.
- The potential of network virtualisation should be analysed together with other relevant emergent trends, namely cloud computing and the convergence of IT and networking. The combination of these ongoing transformations is likely to raise an immense field of opportunities for operators. Network virtualisation will surely have a role to play in this scenario.

- Important lessons about running isolated virtual networks over a common physical infrastructure should be learnt from the experience with MPLS VPNs, which most operators have been offering for more than 10 years. Although VPNs should be seen as a limited form of network virtualisation, similar challenges are to be faced in several problem spaces, such as scalability, reliability and security.
- Standardisation, still in a relatively incipient stage, will surely play a key role to enable interoperability between different vendors and, no less importantly, between network operators. Network virtualisation-ready commercial products from major vendors have been launched. Standardisation is therefore essential to guarantee interoperability and avoid vendor lock-in.
- Most network virtualisation initiatives have focused on "full-blown" virtualisation of network resources, both nodes and links. One of the problems with this approach is that vendors have to expose the internals of network devices. The OpenFlow initiative proposed an alternative, a more lightweight approach to implement network virtualisation. OpenFlow is already supported by commercial Ethernet equipment and its evolution must be closely followed by operators.
- Network virtualisation has so far been successfully demonstrated in small-scale research test beds, and has proven to be a great tool for experimentation, but it is clear that there is still a way to go before the technology can be considered mature enough for large-scale commercial deployment. The evolution of this technology should be closely followed, particularly with regard to reliability, security and isolation.

References

[1] T. Sridhar," Cloud Computing – A Primer, Part 1: Models and Technologies", Internet Protocol Journal, vol 12 no.3, September 2009

[2] Mosharaf Chowdhury, "Network Virtualization: Present and Future", Presentation of the , eNVy Project, University of Waterloo – eNVy, May 21 2008

[3] Roland Bless, Stephan Baucke, "Network Virtualization within FP7 EU Project 4WARD Network of the Future"

[4] Jorge Carapinha, Javier Jiménez, "Network Virtualization – A view from the Bottom", VISA'09, August 17 2009, Barcelona, Spain.

[5] Roland Bless, Christoph Werle, "Control Plane Issues in the 4WARD Network Virtualization Architecture", Workshop der Wissenschaftlichen Konferenz Kommuniation in Verteilten Systemen 2009, (WowKiVS 2009)

[6] Mosharaf Kabir Chowdhury, Raouf Boutaba, "Network Virtualization: State of the Art and Research Challenges", IEEE Communications Magazine, vol. 47 No. 7, July 2009, pp 20-26.

[7] S. Baucke et al. "Virtualisation Approach: Concept", 4WARD Project Deliverable 3.1.1.

[8] "Cisco Nexus 7000 Series Virtualization Architecture – Multi-Degree Virtualization Enabling Resource Consolidation", http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9402/ps9512/brochure_cisco_nexus_7000_series_virtualization_arch.pdf

[9] Router Virtualization in Service Providers, http://www.cisco.com/en/US/solutions/collateral/ns341/ns524/ns562/ns573/white_paper_c11-512753_ns573_Networking_Solutions_White_Paper.html

[10] "Juniper Networks Brings Virtualization to the Core with Industry's Most Flexible Multi-chassis Routing System", Juniper Press Release, http://www.juniper.net/us/en/company/press-center/press-releases/2009/pr_2009_02_02-x17_19.html

[11] Internet Research Task Force -Virtual Networks Research Group (VNRG), http://www.irtf.org/charter?gtype=rg&group=vnrg

[12] Terms of Reference of ITU-T Focus Group on Future Networks, http://www.itu.int/dms_pub/itu-t/oth/3A/02/T3A020000010001PDFE.pdf

[13] "European Virtualisation Project is Looking for Future Internet Researchers", Federica Project, http://www.fp7-federica.eu/documents/20081124-Federica_lyon_poster_v6.1.pdf

[14] "The FEDERICA project - A federated infrastructure for Future Internet research", Mauro Campanella, Eurescom mess@ge 2/2008

[15] FEDERICA Project, Deliverable DSA1.1: "FEDERICA Infrastructure Version 7.0", http://www.fp7-federica.eu/documents/FEDERICA-DSA1.1.pdf

[16] Jeanna Neefe Matthews, Wenjin Hu, Madhujith Hapuarachchi, Todd Deshane, Demetrios Dimatos, Gary Hamilton, Michael McCabe, James Owens, "Quantifying the Performance Isolation Properties of Virtualization Systems"

[17] "Solaris™ Containers: Server Virtualization and Manageability", Sun Microsystems, A Technical White Paper, September 2004

[18] David E. Williams, Juan Garcia, "Virtualization with Xen", Syngress, 2007

[19] "Understanding Full Virtualization, Paravirtualization, and Hardware Assist", VMware White Paper, www.vmware.com/files/pdf/VMware_paravirtualization.pdf

[20] PlanetLab User's Guide, http://www.planet-lab.org/doc/guides/user

[21] Liang Zhao et al. "D-3.2.0 Virtualisation Approach: Evaluation and Integration", 4WARD Project, 2009

[22] Cristina Cervelló-Pastor et al. "Deliverable JRA1.1: Evaluation of current network control and management planes for multi-domain network infrastructure", Federica Project, 2008

[23] Sunay Tripathi, Nicolas Droux, Thirumalai Srinivasan, “Crossbow: From Hardware Virtualized NICs to Virtualized Networks”, VISA’09, 2009, Barcelona, Spain

[24] User-Controlled LightPaths (UCLP) software, Communications Research Centre Canada http://www.crc.gc.ca/en/html/crc/home/info_crc/publications/technology_showcase/uclp

[25] Eduard Grasa, Sergi Figuerola, Albert López, Gabriel Junyent, Michel Savoie, Bill St Arnaud, Mathieu Lemay, “Articulated private networks in UCLP”, 2007

[26] Michael Morisy, “Network virtualization, led by Juniper, promises efficiency boost”, http://searchtelecom.techtarget.com/news/article/0,289142,sid103_gci1353321,00.html

[27] “Virtualization in the Core of the Network”, Juniper White Paper, 2009

[28] Nick Feamster, Lixin Gao, and Jennifer Rexford, "How to lease the Internet in your spare time," in the Editorial Zone of ACM SIGCOMM Computer Communications Review, p. 61-64, January 2007 (slides). A longer version appears as Georgia Tech Technical Report GT-CSS-06-10, August 2006.

[29] M. Boucadair, P. Georgatsos, N. Wang, D. Griffin, G. Pavlou, M. Howarth, A. Elizondo, “The AGAVE approach for network virtualization: differentiated services delivery” Ann. Telecommun. (2009) 64:277-288, Institut TELECOM and Springer –Verlag France 2009

[30] Focus Group on Cloud Computing (FG Cloud), http://www.itu.int/ITU-T/focusgroups/cloud/

[31] E. Keller, R. Lee, and Jennifer Rexford, “Accountability in hosted virtual networks” Proceedings of ACM SIGCOMM Workshop on Virtualized Infrastructure Systems and Architectures (VISA), New York, NY, USA: ACM, 2009

[32] “Provider Provisioned Virtual Private Network (VPN) Terminology”, Internet Engineering Task Force, IETF RFC 4026, March 2005

[33] 4WARD Project Work Package 3, “Network Virtualisation: the 4WARD perspective” http://docbox.etsi.org/Workshop/2010/201003_FNTWORKSHOP/3_NETWORK_VIRTUALIZATION/CARAPINHA_4WARD.pdf, ETSI Future Network Technologies Standardization Workshop 10-11 March 2010

[34] S. Jeong, M-K. Shin, T. Egawa, H. Otsuki, “Network Virtualization Problem Statement”, http://www.ietf.org/id/draft-shin-virtualization-meta-arch-02.txt, July 2010

Annex A Cloud computing

A.1 The implications of Cloud Computing

Cloud computing is a broad term that encompasses many things in the IT sector, but it is most commonly understood as computing resources that a provider hosts and provisions to customers over an Internet connection. Today, most cloud service providers, such as Amazon AWS, Google, SalesForce, etc., provide cloud computing services, including Infrastructure as a Service (IaaS), Platform as a Service (PaaS) or Software as a Service (SaaS), that are distributed and based on best effort network delivery. As service delivery is primarily through the Internet, only best effort can be promised in terms of service quality. Due to this actuality, and other reasons like security issues, services and applications in the cloud tend to be those that tolerate minor delays or are not mission critical. Generally, this includes services like email, CRM systems, collaboration suites, social networking, etc. Although an elusive definition, few company core systems, like ERPs, accounting and salary systems, have been migrated to the cloud. However, by providing managed network connections and added security options, it is foreseen that continued migration to the cloud will occur.

Cloud computing will undoubtedly bring new challenges to operators – increased resilience (since a higher number of crucial applications will depend on the network), security, privacy, mobility, dynamicity, elasticity. Strictly speaking, these issues are not new to operators, but cloud computing requires a new approach to handle some of these issues. In some sense, a change of paradigm is required – from static networking, in which networks are planned, provisioned and configured for relatively long periods of time, we are moving to elastic networking, in which network resources are allocated, reconfigured and finally deallocated in a much more dynamic way, in much shorter periods of time. That is where the network virtualisation should prove to be an important asset for network operators – in particular the capability to build, provision, reconfigure and eventually remove networks "on-demand", in a simple and flexible way.

A.2 Cloud Computing platforms

A.2.1 Microsoft Azure

Azure Services Platform is an application platform in the cloud that allows applications to be hosted and run at Microsoft data centres. It provides a cloud operating system called Windows Azure which serves as a runtime for the applications and provides a set of services that allows development, management and hosting of applications off-premises. Windows Azure takes advantage of the benefits of virtualisation. Typically, each instance of an application runs in a separate virtual machine on Windows Server 2008. The VMs run on a hypervisor that Microsoft designed specifically for the cloud computing environment. The most important aspect of Windows Azure is that it provides on-demand computing power and storage that can scale as required by the load at any given moment in time.

Windows Azure has three core components: Compute, Storage and Fabric. As the names suggest, Compute provides computation environment while Storage focuses on providing scalable storage (Blobs, Tables, Queue, Drives) for large scale needs and the Fabric provides secure connections between servers, switches and other components of the Microsoft data centre.

The hosting environment of Windows Azure is called the *Fabric Controller* - which pools individual systems into a network that automatically manages resources, load balancing, geo-replication and application lifecycle without requiring the hosted apps to explicitly deal with those requirements. Azure Services Platform provides an API built on REST, HTTP and XML that allows a developer to interact with the services provided by Windows Azure. A client-side managed class library is also provided that encapsulates the functions of interacting with the services. It also integrates with Microsoft Visual Studio so that it can be used as the IDE to develop and publish Azure-hosted applications.

A.2.2 Amazon Web Services (AWS)

The Amazon Web Services (AWS) are a collection of remote computing services (also called web services) that together make up a cloud computing platform, offered over the Internet by

Amazon.com. Amazon Web Services' offerings are accessed over HTTP, using REST and SOAP protocols. All are billed on usage, with the exact form of usage varying from service to service. The most central and well-known of these services are Amazon EC2 and Amazon S3.

Amazon Elastic Compute Cloud (EC2) is a central part of Amazon.com's cloud computing platform, Amazon Web Services (AWS). EC2 allows users to rent virtual computers on which to run their own computer applications. EC2 allows scalable deployment of applications by providing a web service through which a user can boot an Amazon Machine Image to create a virtual machine, which Amazon calls an "instance", containing any software desired. A user can create, launch, and terminate server instances as needed, paying by the hour for active servers, hence the term "elastic". EC2 provides users with control over the geographical location of instances which allows for latency optimisation and high levels of redundancy. For example, to minimise downtime, a user can set up server instances in multiple zones which are insulated from each other for most causes of failure such that one backs up the other.

Amazon S3 (Simple Storage Service) is an online storage web service offered by Amazon Web Services. Details of S3's design are not made public by Amazon. S3 stores arbitrary objects up to 5 gigabytes in size, each accompanied by up to 2 kilobytes of metadata. Objects are organised into buckets (each owned by a Services or AWS account), and identified within each bucket by a unique, user-assigned key. Amazon Machine Images (AMIs) which are modified in the Elastic Compute Cloud (EC2) can be exported to S3 as bundles. Buckets and objects can be created, listed, and retrieved using either a REST-style HTTP interface or a SOAP interface. Additionally, objects can be downloaded using the HTTP GET interface and the Bit Torrent protocol. Requests are authorised using an access control list associated with each bucket and object. Bucket names and keys are chosen so that objects are addressable using HTTP URLs.

A.2.3 Google App Engine

Google App Engine is a platform for developing and hosting web applications in Google-managed data centres. It virtualises applications across multiple servers and data centres. Google App Engine is free up to a certain level of used resources. Fees are charged for additional storage, bandwidth, or CPU cycles required by the application.

Compared to other scalable hosting services such as Amazon EC2, App Engine provides more infrastructures to make it easy to write scalable applications, but can *only* run a limited range of applications designed for that infrastructure. App Engine's infrastructure removes many of the system administration and development challenges of building applications to scale to hundreds of requests per second and beyond. Google handles deploying code to a cluster, monitoring, failover, and launching application instances as necessary.

While other services let users install and configure nearly any *NIX compatible software, App Engine requires developers to use only its supported languages, APIs, and frameworks. Current APIs allow storing and retrieving data from a Big Table non-relational database; making HTTP requests; sending e-mail; manipulating images; and caching. Most existing Web applications can't run on App Engine without modification, because they require a relational database. Google App Engine's data store has a SQL-like syntax called "GQL".

Per-day and per-minute quotas restrict bandwidth and CPU use, number of requests served, number of concurrent requests, and calls to the various APIs, and individual requests are terminated if they take more than 30 seconds or return more than 10MB of data.

A.2.4 IBM cloud initiatives

As IBM is a large provider of cloud computing infrastructure, it has established a different approach towards providing cloud services than many others cloud providers, including Microsoft and Google. IBM's strategy is to avoid creating a undesirable competitive position with many of its customers, or potential customers, that use or want to use IBM infrastructure technology for implementing their own cloud services. This has led IBM to create a private cloud service mainly directed at organisations for performing development and testing projects and to realise faster application deployment with reduced capital and operational costs.

IBM recently started to offer their CloudBurst business platform that enables organisations to simplify cloud computing acquisition and deployment as a pre-packaged and self-contained service delivery platform intended for fast implementation in data centres. The CloudBurst platform is built on IBM's System x BladeCenter technology.

Finally, IBM has started to provision its collaboration solutions in a cloud fashion, perhaps similar to Google Apps and Microsoft Live. This includes IBM LotusLive for social networking and online collaboration, IBM LotusLive iNotes for enterprise email, scheduling and contact management and LotusLive Integrated Apps which is a partner program that provides several cloud based applications, e.g. SalesForce, through the IBM platform.

