esa CloudSat

ESA ARTES 1
ESA/ESTEC Contract No. 40000110995/14/NL/AD
SCENARIOS FOR INTEGRATION OF SATELLITE COMPONENTS IN FUTURE NETWORKS

Final Report

Edited by G. Gardikis (Space Hellas, Athens, Greece)

Contributors S. Costicoglou, G. Darladimas (Space Hellas, Athens, Greece)

F. Arnal, C. Baudoin (Thales Alenia Space France, Toulouse, France)

L. Contreras, P. Aranda (Telefónica I+D, Madrid, Spain)

H. Koumaras, V. Koumaras, G. Xilouris, A. Kourtis (NCSR "Demokritos", Athens, Greece)

Revision 1.1

Date May 1st, 2016

Document Ref.No. SPH/RD/CloudSAT/FR/001

Name of ESA Study Manager Maria Guta (TIA-TFA)

SPACE HELLAS ThalesAlenia Space Telefónica ΔΗΜΟΚΡΙΤΟΣ

ABSTRACT

The CloudSat project studies the applicability of emerging virtualisation and softwarisation technologies to satcom platforms and determines the benefits and the challenges associated with the integration of satellite infrastructures into future software-based networks.

To this end, the CloudSat study:

- Reviews state-of-the-art virtualisation and softwarisation technologies, focusing on Software Defined Networking (SDN) and Network Functions Virtualisation (NFV);
- Determines the applicability of these technologies to satcom;
- Identifies specific use cases/integration scenarios and studies their techno-economic efficiency;
- Defines integrated virtualised satellite/terrestrial architectures and validates them in a lab environment;
- Produces a roadmap and recommendations for future virtualisation-capable satellite networks.

The satcom community is expected to derive significant benefits from the adoption of the SDN/NFV model and the interoperability/integration with terrestrial software-based networks. Especially satcom service providers are seen to receive the most benefits, associated with CAPEX/OPEX reduction thanks to virtualisation, as well as increased revenues stemming from novel added-value service offerings.

Table of Contents

1. INTRODUCTION

1.1. CloudSat objectives

The foundation of current networking infrastructures (wired/wireless and also satellite) on fixed, hardware components with vendor-specific management interfaces, although achieving satisfactory performance and reliability, significantly constrains management flexibility and resource federation, while also hampering the rapid introduction of new network services. This “ossification” is even more visible in the case of satellite networks, where the resource-demanding procedure of hardware prototyping of network technologies and protocols into on-board processors, as well as the delay and costs associated with satellite manufacturing and launch, introduce considerable delays in the adoption of new technologies.

In an effort to overcome these limitations, the telecom/network community is pursuing during the last years a paradigm shift towards the virtualisation/ ”softwarisation” of infrastructure components, enabling a novel “cloud networking” model, which allows the flexible management of network resources and functionalities in a cloud-like manner. Future networks are envisaged to consist of heterogeneous wireless and wired physical infrastructures, whose resources are abstracted via virtualisation mechanisms, unified, dynamically pooled and offered as-a-Service to multiple tenants.

In order to be able to benefit from such a progress and also seamlessly integrate with future networks, satellite communication platforms need to follow this transformation which is currently occurring in the terrestrial segment. The CloudSat study, implemented within the frame of the ESA ARTES 1 programme, focuses on this issue, studying the applicability of cloud networking technologies to satcom platforms and determining the benefits and the challenges associated with the integration of satellite infrastructures into future cloud networks.

1.2. Document scope

This document constitutes the Final Report of the CloudSat study. It is structured as follows:

Chapter 2 presents a comprehensive survey of cloud networking techniques and technologies

Chapter 3 discusses the suitability of cloud networking techniques and technologies with respect to their integration with satcom platforms.

Chapter 4 proposes some integration scenarios/use cases which show the added-value of the interplay between cloud networking and satcom.

Chapter 5 presents integrated cloud networking architectures.

Chapter 6 presents the implementation of the proposed architectures in an experimental testbed and the evaluation of the system against selected use cases.

Chapter 7 presents a techno-economic analysis of the gains and constraints associated with the integration of cloud networking in satcom.

Chapter 8 proposes recommendations for future work, technology development and standards evolution.

Finally, Chapter 9 concludes the document.

2. CLOUD NETWORKING TECHNIQUES AND TECHNOLOGIES REVIEW

2.1. The Cloud Network Model

2.1.1. Brief introduction to Cloud Computing

This section briefly introduces some concepts referring to Cloud Computing for providing context to the reader. Although this document mainly focuses on Cloud Networking, knowledge about Cloud Computing concepts is desirable to understand the cloud ecosystem.

According to the definition by the National Institute of Standards and Technology (NIST) [NIST], it is stated that "*Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction*". Five essential characteristics of cloud services are further identified, namely:

- On-demand self-service: It is possible for a user to provision resources at any time without human interaction.
- Broad Network access: The resources are available through a network to various standard platforms.
- Resource pooling: The resource pools of the provider are serving multiple customers and can be dynamically assigned as demand changes.
- Rapid elasticity: The customer can rapidly provision more resources or release some of the provisioned resources at any time.
- Measured Service: The usage is automatically metered at some level of abstraction to provide a transparent usage reporting for both the user and the provider.

Three main types of cloud services can be distinguished:

- Software as a Service (SaaS): This layer provides direct access to applications for users. The SaaS model is used for enabling access to software services on demand hosted in the cloud. The SaaS model is "built" on top of the IaaS and the PaaS models, hence processing and storage of data both happen in the cloud.
- Platform as a Service (PaaS): In this category of services the provider allows users to develop their own SaaS applications, by providing different development platforms and environments to test and execute the constructed applications. PaaS allows the end-user running custom built applications,

eliminating the expense and the complexity associated with configuration and management of the hardware and software needed for running them.

- Infrastructure as a Service (IaaS): This is the lower layer of the underlined architecture of clouds. It is responsible of providing and managing resources for the upper layers, i.e. PaaS and SaaS. In this case, the cloud services are deployed in such a way that the user can exploit the computing resources offered by the cloud via machine virtualization (VM).

Cloud computing (specially the IaaS type) highly relies on virtualization technologies, which enable to decouple the physical from the logical infrastructure.

With regard to infrastructure ownership, the deployment models suggested are:

- Private cloud - the infrastructure is not shared outside of the organization employing the private cloud and is managed internally, serving the organization's needs.
- Community cloud - the infrastructure is shared among different organizations, under a federated management entity. This approach naturally raises several concerns associated with security as well as compatibility.
- Public cloud - the infrastructure is made available to the public and is owned by an organization selling cloud services.
- Hybrid cloud - the infrastructure is a composition of two or more clouds of the above variations. The clouds are unique but there is data and application portability between them. Under this schema a private cloud is able to migrate workloads to a public cloud when the demand increased and additional external resources are needed (this approach is called "cloudbursting").

2.1.2. Benefits and Challenges of Cloud Services

Virtualisation and especially cloud services are emerging as an essential component of the enterprise IT infrastructure and, consequently, one of the fastest growing business opportunities for Internet service providers and telecom operators [Verchere11]. The Internet has evolved over time into a three layer structure: at the top layer sit the applications driving the capacity and the ultimate requirements of the lower layers. These applications can be consumer-oriented, like video, audio, gaming, file-sharing, communication, social networking, consumer cloud access, etc. [Sandvine], or business applications such as backup, inter-site connectivity or various data-center-to-data-center interactions, such as distributed search or VM migration.

Cloud services are also driving chances to the telecom operator infrastructure. On-demand virtual machine (VM) creation and new services in the cloud enables the reduction of IT resources, but requires to dynamically configure the network infrastructure in order to optimally use their resources.

The versatile consumption of IT resources and the distinct nature of the applications running on it produce very variable traffic patterns on the connections reaching the data centers. The transport network becomes the key point to efficiently connect

users to services and applications, which are now consumed independently of where either the resource or the user is located.

The flexibility provided by the cloud computing dynamically changes both the overlay service topology and the corresponding traffic demand, affecting the traditional planning and dimensioning rules of network operators. Network utilization becomes time-varying and less predictable. The efficient integration of cloud-based services among distributed DCs, including the interconnecting network, becomes then a challenge to provide performance guarantees, localization and high availability properties.

[Contreras12] introduces the idea of *cloud-ready transport network*. Such a kind of transport network must support procedures to *allow elastic on-demand connectivity* as a way to configure the network based on the changing demands, to *provide an automated connectivity control* to use dynamically the network resources, and to *enhance the network configuration* based on the *cloud information*. A cloud-ready network can achieve these goals with the support of three technological pillars: (1) a flexible transport network, able to guarantee the required capacity on-demand; (2) a multi-layer oriented network management, able to tackle the traffic demand in a cost-effective way; and (3) a set of cross-strata capabilities, able to allow a joint optimization of the resources of both the cloud-based application and the underlying network providing connectivity.

2.1.3. Cloud Networking – Concepts and Terminology

In the context of this study, the term “cloud networking” refers to the application of the main benefits of the cloud model, especially on-demand provisioning and resource elasticity, to network services. Cloud networking implies the virtualization and abstraction of network resources (i.e., links, nodes and functionalities) and their provision to the end-user as-a-Service, in a cloud-like manner, featuring dynamic resource pooling and elasticity. In this context, the physical infrastructure is partitioned into several autonomous, logically isolated *virtual networks (“slices”)* which are offered to customers/tenants. Network slices may span across several heterogeneous network domains (wired, wireless), yet the specificities of the underlying infrastructure are hidden from the tenant, who “sees” and manages a unified end-to-end virtualised service.

This concept is shown in Figure 1, as described in ITU-T Rec. Y3011 [ITUY3011], which uses for these slices the term “Logically Isolated Network Partition” (LINP). As shown, physical resources are abstracted to virtual resources, which are then aggregated to form virtual networks (LINPs). More details on network virtualisation concepts and technologies can be found in Sec. 2.2.1.2. .

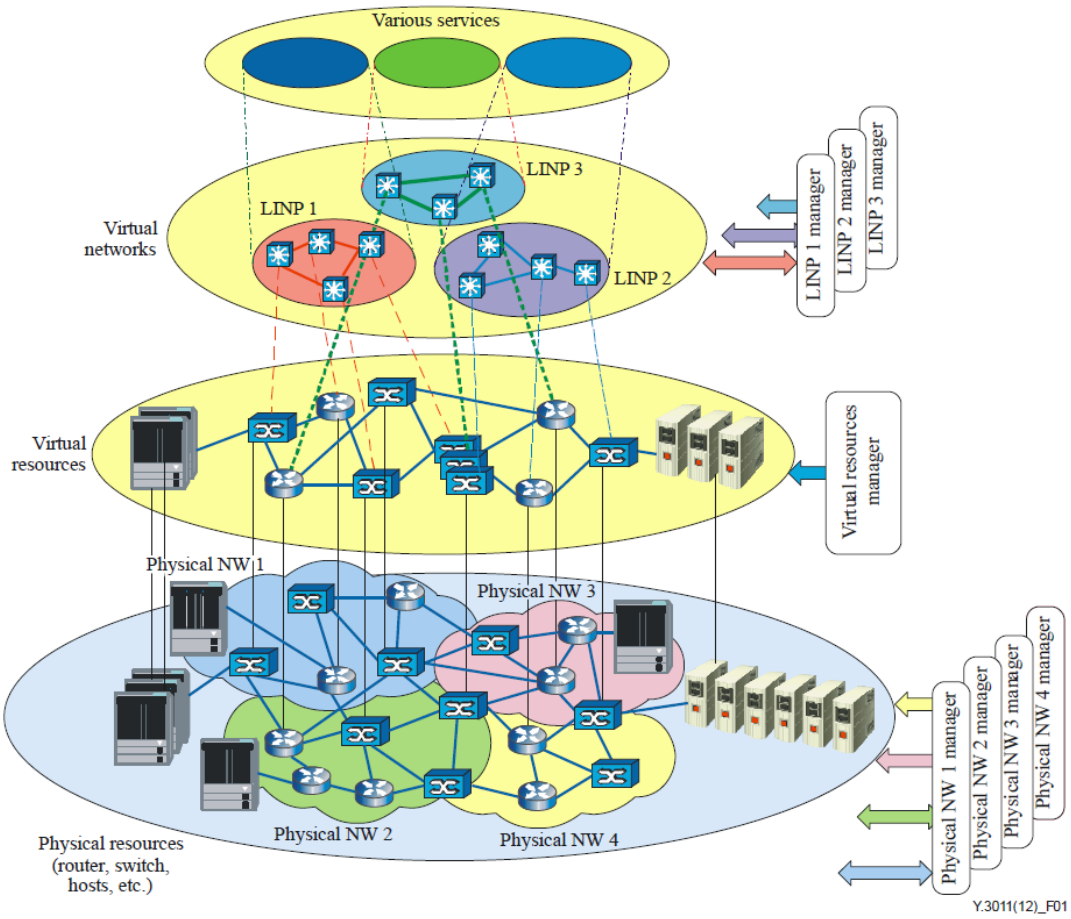

Figure 1. The concept of Network Virtualisation (source: [ITUY3011])

Moreover, in addition to the establishment of virtual slices, an added-value feature of the cloud network model, which has been specifically highlighted during the last years with the advent of *Network Functions Virtualisation (NFV),* is the ability to insert in the network slice traffic processing services in the form of software virtual network appliances (or –more commonly- *Virtual Network Functions/VNFs).* In this context, a network slice can be further enhanced with VNFs such as virtual firewalls, caches, media processors, deep packet inspectors etc. More information on NFV benefits, trends and challenges can be found in Sec. 2.2.3.

Overall, it can be summarized that the cloud network model exploits novel infrastructure management paradigms based on resource virtualization and federation across heterogeneous physical infrastructures, in order to offer next generation virtualized end-to-end *Cloud Network Services*. These services consist of a connectivity component (virtual network slice), optionally enhanced with on-demand virtual network functions. A simple visualization of the cloud network concept applied in hybrid satellite/terrestrial infrastructures is shown in Figure 2.

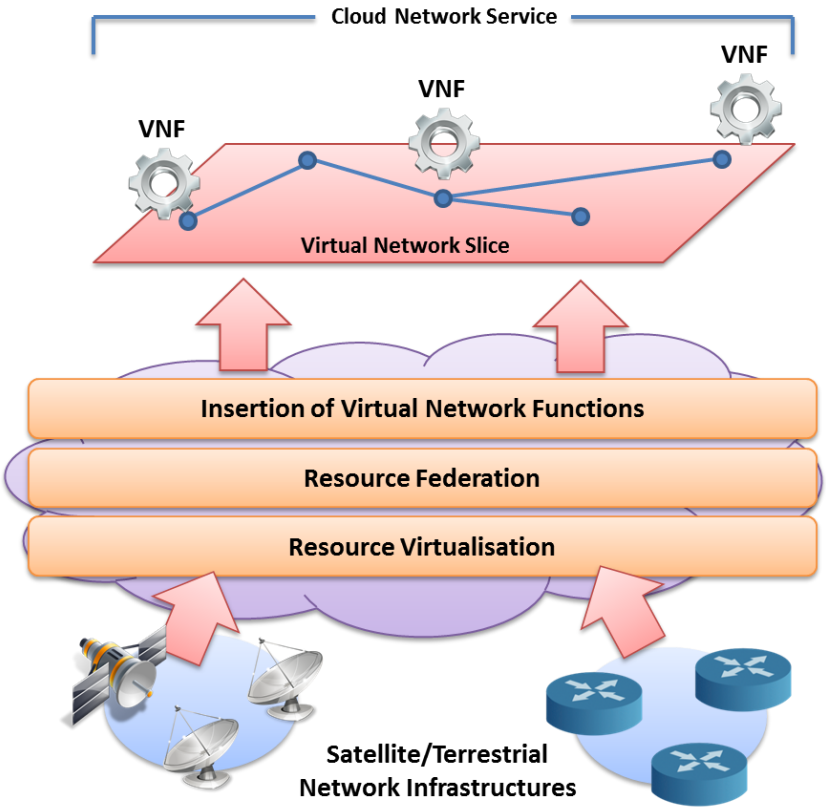

Figure 2. Simplified view of the Cloud Network model applied to a hybrid satellite/terrestrial network infrastructure

It must be noted that services offered by cloud networks, are far richer in comparison to existing terrestrial or satellite VPN bundles. Thanks to state-of-the art technologies involved (such as network programmability and network functions virtualization to be discussed later in this deliverable), Cloud Network Services feature full resource elasticity (i.e., up/down scaling) and can thus support flexible Service Level Agreements (SLAs) and billing models according to usage. Moreover, as, aforementioned, they can natively support a set of rich in-network functions (VNFs) in addition to connectivity and QoS.

This enhanced service offering capability is directly analogous to contemporary computing Infrastructure-as-a-Service platforms (see Table 1 for a brief comparison between computing and networking cloud services). In an IaaS service, users are able to demand and acquire Virtual Machines (VMs) with pre-defined computing, memory and storage capabilities. In addition, modern cloud computing technologies allow VM resources to be dynamically up- and down-scaled according to their utilization.

In the same sense, in a future satellite/terrestrial cloud network platform, users/tenants will be able to select the virtual topology which best match their need in terms of endpoint/Point-of-Presence location, capacity, QoS and in-network functionalities, which will be offered as logically isolated service, transparently spanning across the terrestrial and satellite domains. This Cloud Network Service is managed and monitored by the tenant as if it was an independent physical unified

network; the details and topology of the underlying physical infrastructure are hidden.

Table 1. Comparison of computing and networking cloud services

	Computing Cloud (IaaS model)	Networking Cloud
Physical infrastructure	Data centre (server farm)	Satellite/Terrestrial Network
Service offered	Virtual Machine	Cloud Network Service (Virtual Network Slice + VNFs)
Virtualised resources	• CPU frequency • Memory • Disk space	• Link capacity • Routing/Forwarding engines • Network Functions (VNFs)
Resource up-/down- scaling	Supported	Supported
User-side management	Per-VM, cloud-based	Per-service, cloud-based
Billing model (common)	Pay-as-you-go	Pay-as-you-go

2.1.4. Actors and Roles in a Cloud Network environment

The business value chain in a Cloud Network environment is based on the four-role model [Abarca13] [Carapinha09] which has been proposed for platforms providing virtualised network services. An additional role, this of the Virtual Network Function provider (VNFP) [TND21] is added for NFV-enabled services using third-party VNFs.

Overall, the following five business roles are identified (see Figure 3):

Infrastructure Providers (InPs) or Physical Infrastructure Providers (PIPs) are either Satellite or Terrestrial operators who own the physical network infrastructure, optionally accompanied with in-network computing resources required for VNF deployment. InPs possess the mechanisms able to virtualise these resources and advertise them to the Cloud Network Service Providers.

Cloud Network Service Providers (CNSPs) are responsible for finding and composing the adequate set of virtual resources from one or more InPs in order to offer virtualised services. CNSPs lease slices of underlying InPs and assemble them to form single- or multi-domain slices, also instructing the insertion of the required VNFs. While CNSPs and InPs are discrete roles, in most cases they are undertaken by the same business actor. That is, InPs also undertake the role of CNSP in order to ensure better and more efficient control of their own physical resources without confronting stability and/or privacy issues by exposing them to third-party CNSPs.

Virtual Network Function Providers (VNFPs) are the vendors of the virtual network appliances. The term Function Developer is often used alternatively, since the VNFPs have developed the software, which implements the function. VNFPs actually lease

instances of their VNFs to Customers, commonly not directly but via agreements with the CNSPs.

Customers/Tenants are the users and at the same time the "operators" of the cloud network service. They establish SLAs with the CNSP, holding the requirements and the constraints of the service. Customers have specific management, control and monitoring rights on the provisioned services, as if the latter were autonomous infrastructures. They also have a unified view of the provisioned network service, regardless of the multiple InP domains on which it may be built, whereas they do not have any awareness of the physical infrastructure assets which are involved in the service. Since multi-tenancy is an inherent feature of the cloud network model, the term *Tenants* can also be used.

Customers may exploit the network service for own internal use (e.g., in the case of an enterprise establishing a corporate VPN). Alternatively, Customers may also in turn act as Service Providers themselves and exploit the virtual network for offering a specific application to their own customers. That would be e.g. the case of a content provider exploiting the cloud network service to distribute IPTV streams to end users with specific QoS constraints. In this case, the model also includes *End-Users (EUs)* which access a specific application over the Customer's virtual network. The provision and configuration of the network service is totally transparent to the EUs, who interact only with the application offered by the tenant.

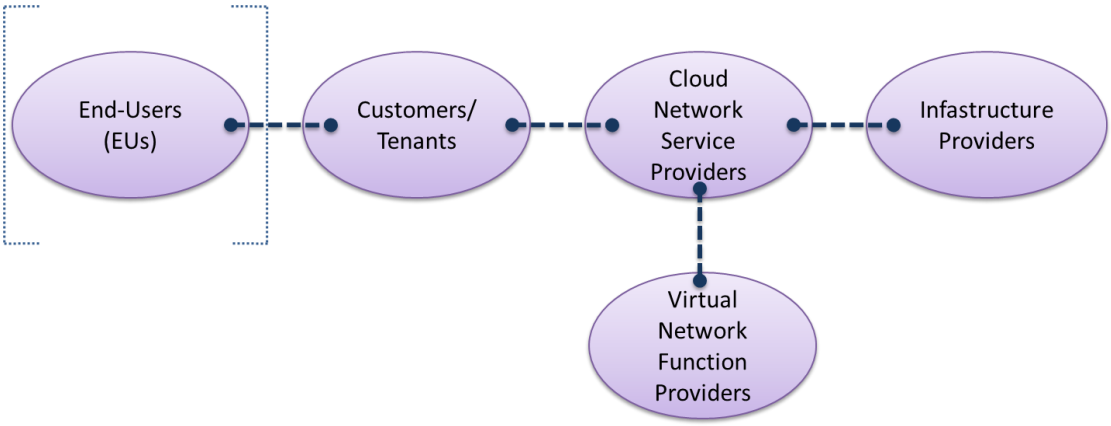

Figure 3. Roles and value chain in the cloud network model

2.2. Enabling Technologies

The cloud computing paradigm has allowed a new model for service delivery where the Information Technology (IT) resources deployed in Data Centres (DCs) form a pool able to attend multiple service demands by means of a dynamic resource assignment, like CPU or storage capacity, either physical or virtual (by using abstraction mechanisms). The computing resources can be provided on-demand depending on the user requests.

This elasticity on the resource consumption allows an agile adaptation to the business requirements and efficient resource utilization. In this multi-tenant model, the sharing of resources among users reduces costs and maximizes utilization, leveraging the economies of scale.

Furthermore, the virtualization concepts are being extended to the network side. The virtualization in the IT side implies ubiquity, service independence of the real location, and then flexibility. Existing connectivity technologies have to be revisited, towards new flexible infrastructures that can accommodate the dynamic customer demands. These new infrastructures will be composed of distinct virtualization technologies built on top of a high capacity transport based solution, able to provide enough bandwidth capacity in a dynamic manner. Software-defined networking (SDN) appears as a promising technology towards this goal. Moreover, the emerging Network Functions Virtualisation (NFV) paradigm offers the potential to augment network connectivity with virtualized network appliances for traffic processing on-demand.

Finally, in order to address these challenges and to be able to provide end-to-end services across heterogeneous domains, cloud Data Centres connecting to high-performance network infrastructures also need to leverage orchestration mechanisms allowing the joint management of the DC (storage and processing) and the network, at both infrastructure- and service-level.

In this context, cloud networking should not considered as a single technology. Instead, the realization of the cloud network concept relies on the synergistic application of specific enabling technologies, namely:

- Infrastructure (DC and Network) Virtualisation,
- Programmable and Software-Defined Networking (SDN),
- Network Functions Virtualisation (NFV),
- Federated Resource Management and Orchestration.

These enabling technologies for cloud networking are briefly overviewed in the sections to follow.

2.2.1. Infrastructure Virtualization

2.2.1.1. IT infrastructure virtualization

Virtualization of IT resources

Server virtualization is a technology that partitions the physical machine into multiple *Virtual Machines (VMs),* each capable of running applications just like a physical machine. The virtualization technology allows a flexible management of IT resources, distributing them as needed for a certain service either among distinct servers into a data center, or even spreading them across several data centers connected to the network. By separating logical resources from the underlying physical resources, server virtualization enables flexible assignment of workloads to physical machines. This not only allows workload running on multiple virtual machines to be consolidated

on a single physical machine, but also enables *VM migration*, which is the process of dynamically moving a virtual machine from one physical machine to another.

The software that controls the virtualization is called *hypervisor*. A VM simulates the behavior of a physical, dedicated server (i.e., like a "machine within the machine"). It accesses to processing units, RAM memory, hard disks and devices. However, those elements are provided by the hypervisor, which is executed in the physical machine hosting the VM.

Reference [Krishnan11] presents the run-time behavior of several concurrent Virtual Machines, achieving a good reference model useful for describing VM workloads. In particular, the authors introduced several models for characterizing CPUs, RAMs, Disk and I/O within VMs under different working conditions (i.e. percentage of load, power consumption, etc.).

As shown in Figure 4, any operating system can be installed on a VM and is unaware that it is being executed on a virtual rather than a physical machine. The user programs interact with the operating system as usual. The hypervisor ensures that the VMs are isolated among them and have the necessary resources.

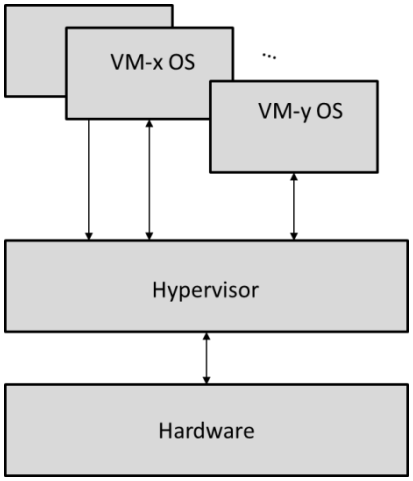

Figure 4. IT Virtualisation hierarchy

The virtualization of IT resources [Smith05], which encompass processing units, storage devices, and even working memory, implies an on-demand creation and configuration of virtual machines that provide a full computerization capability.

VM placement (i.e. the decision about which physical node will host a VM) and also migration can be a complex issue in some cases. In contrast to a small number of warehoused-sized data centers (DC) commonly used in public clouds, a telecom operator cloud may support a large number of small, distributed DCs to reduce traffic in the core network. By encapsulating workloads in VMs a cloud resource manager can migrate workloads from one DC to another, towards improving the perceived quality of experience (QoE), reducing energy consumption, or even in response to situations such as network failures or high-demand events. In addition, placing DCs

closer to end-users enables the development of services and applications that can take advantage from very low latency.

Management of virtualized IT resources

Cloud management platforms are integrated tools that provide management of cloud environments. These tools incorporate self-service interfaces, provision system images, enable metering and billing, and provide some degree of workload optimization through established policies. Through the self-service interface, the user can request virtual infrastructure. This request is issued to a Cloud Controller, which provisions this virtual infrastructure on available resources within the DC. The Cloud Controller provides the central management system for cloud deployments as shown in Figure 5.

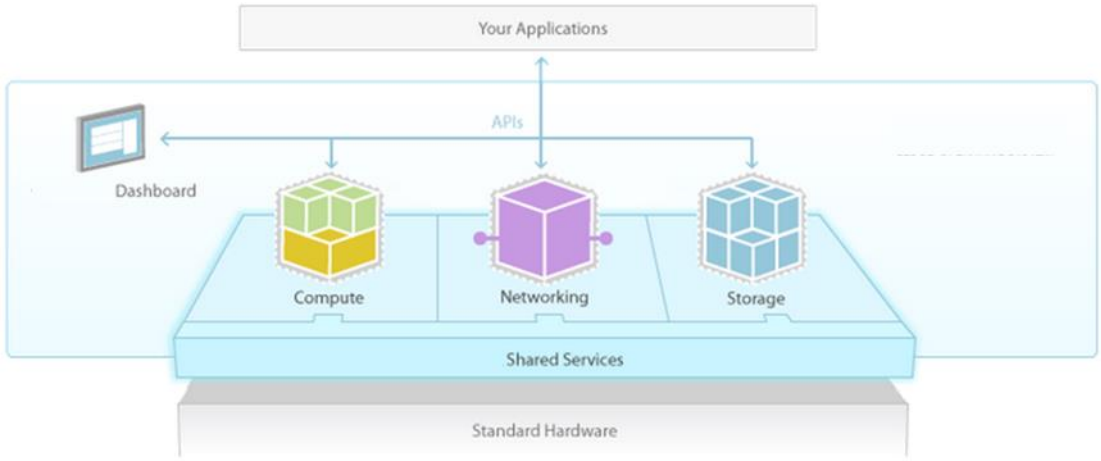

Figure 5. Cloud Management system architecture

The most popular cloud management platforms include open source solutions such as OpenStack, CloudStack and Eucalyptus, as well as commercial solutions from Microsoft and VMware. These solutions are briefly introduced here below.

OpenStack

OpenStack [Openstack] is a "Cloud Operating System" that controls large pools of compute, storage, and networking resources throughout a DC, all managed through a dashboard that gives administrators control while empowering their users to provision resources through a web interface. As an open source solution, OpenStack is developed and supported by a global collaboration of developers and cloud computing technologists. The project seeks to deliver solutions for all types of clouds by being simple to implement, scalable, and feature rich. The technology consists of a series of interrelated projects delivering various components for a cloud infrastructure solution. All OpenStack source code is available under an Apache 2.0 license.

OpenStack has a modular design that enables integration with legacy and third-party technologies[OStackArch]. It is built on a distributed, messaging-based architecture

with modular components, each of which manages a different service; these services, together, constitute an IaaS Cloud.

The components of Openstack (Figure 6) are:

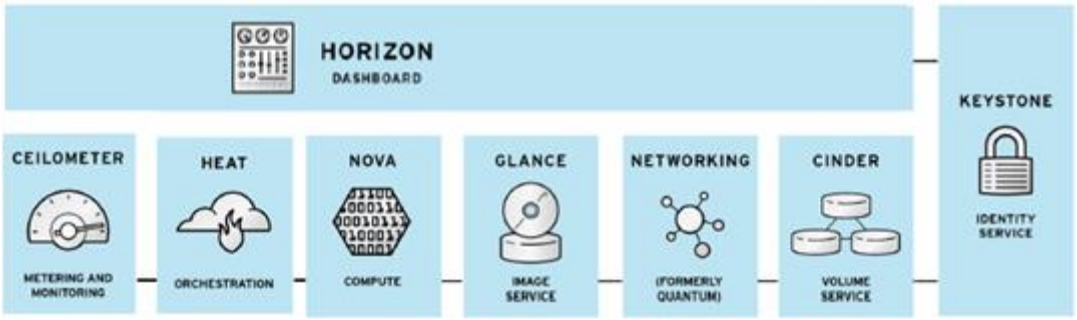

Figure 6. Openstack components

- Horizon: Service and administration Web-GUI dashboard.
- Keystone: Identity (authentication and authorization) service for users, tenants and roles and also for the platform services.
- Nova: Compute management service, which provides virtual servers upon demand. Includes “nova-api” as end-point for all API queries (EC2/OpenStack). It initiates most of the orchestration activities (such as running an instance) and also enforces some policy (mostly quota checks).
- Neutron: Provides network connectivity between interface devices managed by other OpenStack services. Neutron has different agents depending on the service to be provided (Neutron-L3-agents, Neutron-dhcp-agents, Neutron-metadata-agent, ...)
- Glance: Provides a catalogue and repository for virtual disk images.
- Heat: Provides orchestration services into the platform.
- Ceilometer: Collects infrastructure and instances measurements and provides a monitoring system. It aims to deliver a unique point of contact for billing systems to acquire all of the measurements they need to establish customer billing, across all current OpenStack core components with work underway to support future OpenStack components
- Cinder: Manages persistent block storage (data volumes) that can be attached to VM instances.

Figure 7 depicts the interactions among the different Openstack components, a deeper insight of which can be found in Annex I.

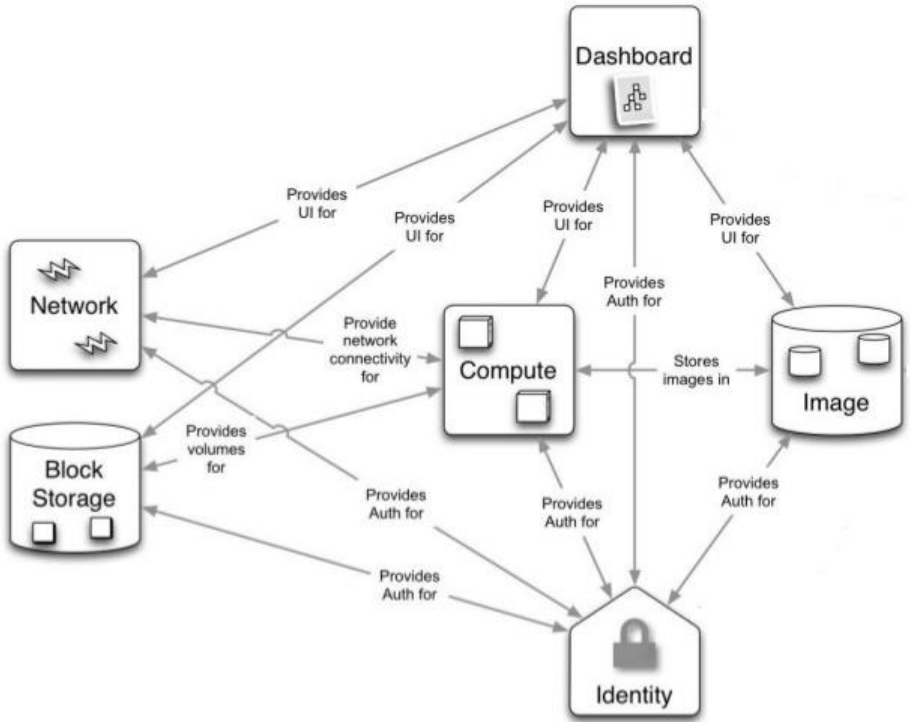

Figure 7. Interactions among Openstack components

Eucalyptus

Eucalyptus (Elastic Utility Computing Architecture Linking Your Programs To Useful Systems) [Eucalyptus] is an open-source Cloud that provides on-demand computing instances and shares the same APIs as Amazon's EC2 cloud. Eucalyptus was designed as a highly-modular framework in order to enable extensibility with minimal effort. The Cloud Controller (CLC) in Eucalyptus acts as the Cloud entry-point by exposing and managing the virtualised resources. The CLC offers a series of web services oriented towards resources, data and interfaces (EC2-compatible and Query interfaces). In addition to handling incoming requests, the CLC acts as the administrative interface for cloud management and performs high-level resource scheduling and system accounting. The CLC accepts user API requests from command-line interfaces like euca2ools or GUI-based tools like the Eucalyptus Management Console and manages the underlying compute, storage, and network resources.

Cloudstack

Apache CloudStack [CloudStack] is open source software designed to deploy and manage large networks of virtual machines, as a highly available, highly scalable Infrastructure as a Service (IaaS) cloud computing platform. CloudStack is used by a number of service providers (e.g. BT) to offer public cloud services, and by many companies to provide an on-premises (private) cloud offering, or as part of a hybrid cloud solution. CloudStack is a turnkey solution that includes the entire "stack" of features most organisations want with an IaaS cloud: compute orchestration, Network-as-a-Service, user and account management, a full and open native API, resource accounting, and a first-class User Interface (UI).

CloudStack is a framework that allows pooling of computing resources in order to IaaS cloud services that can be used to provide IT infrastructure such as compute nodes (hosts), networks, and storage as a service to the end users on demand. CloudStack Management Server is the main component of the framework, consisting of managing resources such as hosts, storage devices and IP addresses. The Management Server runs on a dedicated host in a Tomcat container and requires a MySQL database for persistence. The Management Server controls allocation of VMs to hosts and assigns storage and IP addresses to VM instances. This component also controls or collaborates with the hypervisor layers on the physical hosts over the management network and thus controls the IT infrastructure.

VMware vCloud Suite

VMware's vCloud Suite [vCloud] - is a comprehensive, integrated cloud platform for building and managing cloud environments. Tools for cloud management are delivered through VMware vCenter Server, a centralised and extensible platform for managing virtual infrastructure. The tools included in the vCenter Server framework support: configuration of ESX servers and VMs, performance monitoring throughout the entire infrastructure, using events and alerts. The objects in the virtual infrastructure can be securely managed with roles and permissions.

2.2.1.2. Network infrastructure virtualization

Survey on Network Virtualization mechanisms

The primary aim of network virtualization is to enable multi-tenancy in such a way that privacy, isolation and reliability are assured in the same degree as in a physical dedicated infrastructure.

The logical separation of services achieves protection from other tenants' services in the sense of not interfering with the operational procedures and the transported traffic (e.g., by not having IP address dependencies from the rest of the other services in the network).

There exist several methods and technologies for providing a logically separated network per tenant. Network virtualization technologies may apply either to intra-DC networks or inter-DC (wide area network) connectivity. A complete survey of the internals of DC network virtualization can be found in [Bari13], providing comparison of different approaches in terms of scalability, multipathing, bandwidth sharing and guarantee capabilities, etc.

For WAN (transport network) virtualization, the basic concept is to set up an overlay network to separately transport the information between service endpoints. The overlay network can be basically layer-3 or layer-2 based. For DC interconnection, this overlay network will typically interconnect nodes in the border of the DCs, giving access to an IP/MPLS (or even optical) wide area network.

In this context, a virtualized networking environment supports the coexistence of multiple virtual networks over the same physical infrastructure. The concept of multiple coexisting networks was supported first, in the context of Virtual Local Area Networks (VLANs) and Virtual Private Networks (VPNs).

A VLAN [IEEE802] is a group of hosts logically brought together under a single broadcast domain. VLANs have become a widely used standard with well-defined use cases [Garimella07]. In an intra-DC scenario, the idea behind is that the traffic from the VMs of the different tenants within the DCs is segmented by using VLANs. Unfortunately, VLAN-based solutions cannot meet the requirements for dynamicity, flexibility and scalability, needed both for configuration and proper operation of virtual networks [Sun10], [Greenberg08], [Kim11]. In addition, some researchers and industry vendors are attempting to adapt and extend existing network paradigms to accommodate new requirements brought by the virtualized use cases. Examples of such are academic works as VL2 [Greenberg11], Portland [Mysore09], Seattle [Kim11], vendor products, such as Cisco VN-Link [Cisco09] and network standards amendments, such as TRILL [Touch09] and 802.1Qbg [IEEE802bg].

A Virtual Private Network (VPN) [Ferguson98], [Rosen06] is a dedicated communications network of one or more enterprises that are distributed over multiple sites and connected through tunnels over public communication networks, forming an overlay network. The VPN guarantees private communication between the end-points, constituting a Closed User Group (CUG). These VPN connections can serve complex connectivity requiring not only connectivity among the participants in the VPN but also external connectivity (e.g., to the Internet). Those scenarios of complex connectivity are typically built on top of layer-3 VPNs (L3VPN) where specific virtual routing functions (VRFs) are defined per each of the border routers (or provider edges, PEs) participant in the VPN service to properly route the IP packets interchanged within the VPN. The IP addressing within the VPN has to be unique (i.e., not duplicated), but different customers can use the same addressing scheme in different VPNs. On other hand, when simple layer-2 connectivity is needed in point-to-point or point-to-multipoint, layer-2 VPNs (L2VPN) can be configured, working in base of the MAC addresses instead of the IP addresses.

However, VPNs are too rigid, and cannot support a network virtualization environment where dynamics, flexibility and scalability are highly required attributes. Managed network VPN (e.g., BGP/MPLS), which represents a widely deployed network service for enterprises, is a significant example. This type of services has been conceived to work in a relatively stable network environment (which is the case with most enterprise networks today), but is not appropriate to cope with the typical dynamics of cloud services. The traditional VPN model is not able to handle essential cloud properties such as elasticity and self-provisioning, which means that those properties should be also extended to network resources. Quite often, expanding or reducing cloud resource capacity, or provisioning new cloud resources, requires a corresponding reconfiguration of network resources, e.g., bandwidth assigned between two data centers, whether they are in the same geographical place or not, or between the data center and the end user. In order to cope with the cloud, future network services will certainly require on-demand and self-provisioning properties.

Today the network can provide static connectivity to cloud resources, to what we call conventional networking. The next evolutionary step is to make the network elastic and adaptable according to the cloud dynamics.

Lately, it has become clear that the overlay based approach is the correct answer for achieving independency from the physical networking infrastructure [Narten12], [Laserre12] and [Kreeger12]. An overlay network can be created on top of an existing network, by generating logical communication links between hosts within the service domain. Overlay networks enable the design of modular networking protocols and services in which logical functions are separated from the underlying physical infrastructure.

Multiple vendors have put significant efforts in order to achieve efficient overlay solutions built upon different tunneling protocols. Recent solutions building upon overlays to achieve benefits of scale in multitenant virtual networks are VXLAN [Maha12], NVGRE [Sridharan11] and STT [Davie12]. Reference [Bitar13] introduces different DC interconnection techniques existing today. Among them VXLAN emerges as the most used technology. The basic concept behind VXLAN is the encapsulation of an original Ethernet frame on top of an UDP packet sent between two corresponding Network Virtualization Edges (NVEs). The role of NVE is played by the virtual switches where VMs are attached for internal communication in the DC. When it is required to get connectivity between DCs, the node in the border of the DC will play the role of NVE as well, and stitching that traffic to the inter-DC overlay network. VXLAN header includes a 24-bit long field named VXLAN Network Identifier (VNI) that allows per-tenant network differentiation. Then with VXLAN it is possible to create a virtualized end-to-end layer-2 network on top of a layer-3 overlay transport.

IBM's SDN for Virtual Environments is an architecture based on VXLAN that enhances VXLAN scalability and allows for integrated overlay networks across hypervisors. VMware's NSX is another network virtualization solution, mainly based on distributed virtual switches (vSwitches) and a NSX controller. It supports VXLAN, STT and GRE tunnelling. NetLord [Mudigonda11] is a novel multitenant network architecture that uses an overlay encapsulation method.

The major improvement provided by overlay networks is their separation from the underlying infrastructure, and from each other. This separation facilitates independent address spaces, ensures isolation, and allows different virtual networks to be managed by different administrators.

In order to effectively support cloud services, multiple distributed high-performance datacenters are interconnected by high-bandwidth dynamic optical networks. In such a Cloud environment, optical network virtualization [ADVA10] plays a key role in interconnecting geographically distributed virtual IT resources (i.e. computing and storage) with high-capacity virtual optical network connectivity. The research studies on optical network virtualization are still on the initial stage [Peng13]. However, optical network themselves have undergone significant evolutions [Jinno09][Nag10] and are being empowered by more and more elastic, flexible, programmable optical devices/components. On the other hand, in order to support end-to-end connectivity services, the virtualization functionality should also be enabled over multiple domains with cross-layer and cross-technology characteristics.

Multi-site connectivity of Virtualized Network Environments

The conventional approach for connecting computing resources to the network uses network segmentation based on VLANs to separate end-user or tenant services. This way of segmenting the network provides a limited number of configurable networks per DC (determined by the theoretical 4096 available different VLAN tags), requiring a careful planning of resources either locally or remotely, when more than one DC is involved in a service. This approach has the burden of having to manually reconfigure multiple switches and routers every time a new service has to be deployed, creating inefficiencies and costs.

Multisite L2 services can be built under the concept of virtual patch-panels. In each location, ports on multiple switches (spread across the network) can be programmatically connected among them to set up extended point-to-point connections, in a dynamic and automated way. It is also important to implement mechanisms capable of providing a dynamic on-demand scaling (up and down) of the resources offered to those services, and capable of automatically propagating any network change that could affect those services.

The virtual patch panel function will consist on a stitching of the per-user generated VLAN in a DC with the corresponding pre-provisioned connection that connects with the remote DC. The same is done in the other end to build the end-to-end point-to-point layer2 connection.

2.2.2. Programmable and Software-Defined Networking (SDN)

2.2.2.1. Software Defined Networking and Openflow

The term “Network Programmability” refers to the capability provided by L2/L3 physical network elements to arbitrarily program their switching, forwarding and routing logic on-demand. Older visions of network programmability, such as the one promoted by the IETF ForCES (Forwarding and Control Element Separation) working group [FORCES] introduced the separation of forwarding and control modules within the network elements, allowing the control logic to be developed as a set of separate custom software components [Yang04].

Currently, the most popular paradigm for vendor-neutral network programmability, is Software Defined Networking (SDN) [McKeown08], a model for network control which separates the control and forwarding logic, migrating the traffic handling decisions from the network elements themselves to centralised software controllers. Conceptually, in SDN networks forwarding (physical) devices have minimal intelligence, while the control logic is implemented on top of a so-called SDN controller. The controller is a logically centralised entity which is responsible for a set of tasks, including the extraction and maintenance of a global view of the network topology and state, as well as the instantiation of forwarding logic appropriate to a given application scenario. In practice, the controller manages connections to all

substrate network elements and installs, modifies and deletes forwarding entries into the forwarding tables of the connected switches by using protocol specific control messages.

This communication between controllers and network elements in SDN is commonly based on the Openflow protocol, which originated in Stanford university and is currently maintained by the Open Networking Foundation [ONF]. To date, Openflow is the dominant driving standard for SDN. Using OpenFlow, the Controller can dictate specific rules to SDN-enabled switches (Figure 8). These rules define whether flows which match specific characteristics should be forwarded, re-routed, altered, dropped or QoS-shaped.

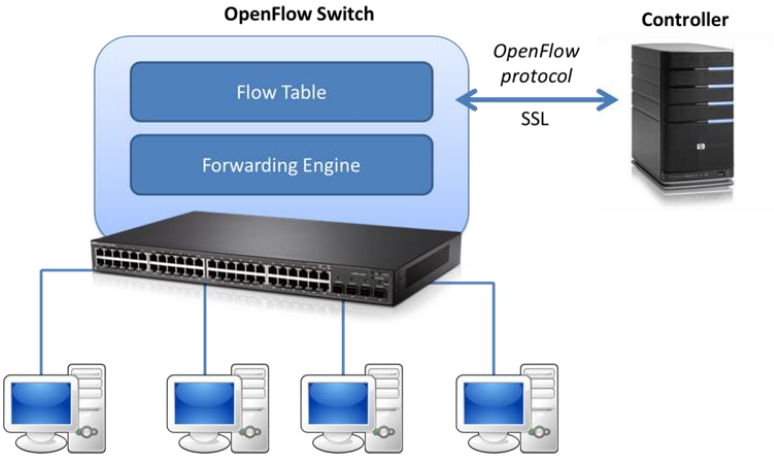

Figure 8. SDN and Openflow

SDN opens new perspectives in network management and is considered a key enabler for cloud networking, since it can facilitate centralized per-flow control across the network and orchestrate virtualization procedures.

2.2.2.2. SDN Controllers

While the OpenFlow protocol itself is quite low-level, several Controller Application Programming Interfaces (APIs) have been made available in order to facilitate high-level programming of networking applications. What these controllers do, is to abstract the OpenFlow protocol to a programming language that the network application is written in. In this context, management applications for cloud networking can be easily developed using a common set of architectural patterns as well as common means to query data flows from one or more network elements and supporting framework functions.

The NOX controller was the first widely available OpenFlow controller [NOX]. NOX was originally developed by Nicira and released as open-source software. Due to its early availability and its simplicity, NOX quickly become the de-facto reference design for

OpenFlow controllers. As a result, it has been used to test new OpenFlow features, novel controller ideas and it has been employed extensively in research and feasibility studies. NOX applications – called modules – are implemented using the C programming language. NOX is event based; each module essentially consists of a collection of callback functions, triggered by the arrival of specific OpenFlow protocol messages. A spin-off of NOX called POX [POX] enables the use of Python for programming modules. While NOX/POX is extremely versatile it is not primarily aimed for production use, as it is not optimised for performance and stability and lacks resilience features.

Other controller frameworks aimed at deployment in production environments, include Beacon [Beacon], Maestro [Maestro] and FloodLight [Floodlight], all of which are implemented in Java. FloodLight is the open source basis for Big Switch's commercial OpenFlow controller.

Apart from the aforementioned frameworks, there also exist SDN management platforms, which are more complete in terms of offered services, so that they can be considered as integrated stand-alone solutions for the management of SDN infrastructures. Most of them also leverage the SDN capabilities in order to offer network virtualisation services, supporting multi-tenancy (often called "Network-as-a-Service"). These SDN management platforms are overviewed in the next sections.

2.2.2.3. SDN-based Network-as-a-Service platforms

FlowVisor

FlowVisor is the ON.LAB network slicer, which allows multiple tenants to share the same physical infrastructure [Flowvisor]. A tenant can be either a customer requiring his own isolated network slice; a sub-organisation that needs its own slice; or an experimenter who wants to control and manage some specific traffic from a subset of endpoints. FlowVisor acts as a transparent proxy between OpenFlow switches and various guest network operating systems. It supports network slicing and allows a tenant or an experimenter to control and manage some specific traffic from a subset of end points. This approach enables multiple experimenters to use a physical OpenFlow network without interfering with each other.

FlowVisor enables network virtualisation by dividing a physical network into multiple logical networks ensuring that each controller touches only the switches and resources assigned to it. It also partitions bandwidth and flow table resources on each switch and assigns those partitions to individual controllers.

FlowVisor slices a physical network into abstracted units of bandwidth, topology, traffic and network device CPUs. It operates as a transparent proxy controller between the physical switches of an OpenFlow network and other OpenFlow controllers and enables multiple controllers to operate the same physical infrastructure, much like a server hypervisor allows multiple operating systems to use the same x86-based hardware. Other standard OpenFlow controllers then operate their own individual network slices through the FlowVisor proxy. This arrangement

allows multiple OpenFlow controllers to run virtual networks on the same physical infrastructure.

FlowVisor, originally developed at Stanford University, has been widely used in experimental research and education networks to support slicing where multiple experimenters get their own isolated slice of the infrastructure and control it using their own network OS and a set of control and management applications. FlowVisor has been deployed on a Stanford production network and sponsors, such as GENI, Internet2, NEC and Ericsson, have been contributing to it and using it in their research labs. The SDN research community considers FlowVisor an experimental technology, although Stanford University has run FlowVisor in its production network since 2009. FlowVisor lacks some of the basic network management interfaces that would make it enterprise-grade. For example it currently does not support any CLI or Web-based administration console but requires users to make changes to the technology with configuration file updates.

OpenVirteX

OpenVirteX is a network hypervisor that can create multiple virtual and programmable networks on top of a single SDN-based physical infrastructure [OpenVirteX]. Each tenant can use the full addressing space, specify their own topology, and deploy the network OS of their choice. Networks can be reconfigured at run-time, and OpenVirteX can automatically recover from physical failures.

OpenVirteX is actually a network hypervisor that enables operators to provide networks whose topologies, management schemes, and use cases are under the full control of their tenants. More specifically OpenVirteX builds on OpenFlow as protocol and FlowVisor for design. In this respect they share some common properties i.e. act as proxies between tenants and the underlying physical infrastructure. Unlike FlowVisor however, OpenVirteX provides each tenant with a fully virtualised network featuring a tenant-specified topology and a full header space.

Openstack Neutron

OpenStack Neutron [Neutron], historically known as Quantum, is an OpenStack project focused on delivering Networking as a Service (NaaS), especially tailored for cloud environments.

Neutron provides a way for organisations to make it easier to deliver networking as a service in the cloud and provides REST APIs to manage network connections for the resources managed by other OpenStack services.

It is designed to implement a “plugin” mechanism that will provide an option for network operators to enable different technologies via the Neutron API making it technology agnostic. However, currently Neutron is able to deliver all its core features only above an SDN-enabled infrastructure.

Neutron provides native multi-tenancy support (isolation, abstraction and full control over virtual networks), letting tenants create multiple private networks and control

the IP addressing on them, and exposes vendor-specific network virtualisation and SDN technologies.

As a result of API extensions, administrators and users have additional control over security and compliance policies, QoS monitoring and troubleshooting, the ability to build sophisticated networking topologies, as well as the ability to easily deploy advanced network services, such as a firewall, L2-in-L3 tunnelling, end-to-end quality of service support intrusion detection or VPN.

The core Neutron API includes support for Layer 2 networking and IP Address Management (IPAM), as well as an extension for a Layer 3 router construct that enables routing between Layer 2 networks and gateways to external networks. It is based on a simple model of virtual networks, subnet, and port abstractions to describe networking resources. Network is an isolated layer-2 segment, analogous to a VLAN in the physical networking world. More specifically, it is a broadcast domain reserved for the tenant that created it or explicitly configured as shared. Neutron includes a growing list of plugins that enable interoperability with various commercial and open source network technologies, including routers, switches, virtual switches and SDN controllers.

Starting with the Folsom release, Neutron is a core and supported part of the Openstack platform. However, it is a standalone and autonomous service that can evolve independently to Openstack.

OpenNaaS

OpenNaaS (Network-as-a-Service) [OpenNaaS] is an open-source framework, which provides tools for managing the different resources present in any network infrastructure, particularly focusing on SDN infrastructures. The software platform was created in order to offer a neutral tool to the different stakeholders comprising an Open Access Network (OAN). It allows them to contribute and benefit from a common NaaS software-oriented stack for both applications and services. It is based on a lightweight, abstracted, operational model, which is decoupled from actual vendors' specific details, and is flexible enough to accommodate different designs and orientations.

Figure 9 depicts the layered architecture of the framework, with the platform layer, the resource abstraction layer with the NaaS manageable units, and the upper layer, where the network intelligence resides, as well as the integration of the framework with third-party components. Besides, the resource abstraction, the core platform concepts are also depicted. Different OpenNaaS deployment examples can be found in the following list of European projects extending the OpenNaaS framework: OFERTIE, CONTENT and SODALES. Furthermore, authors in [Riera14] used OpenNaaS in order to build a first proof-of-concept pilot for provision of NFV services.

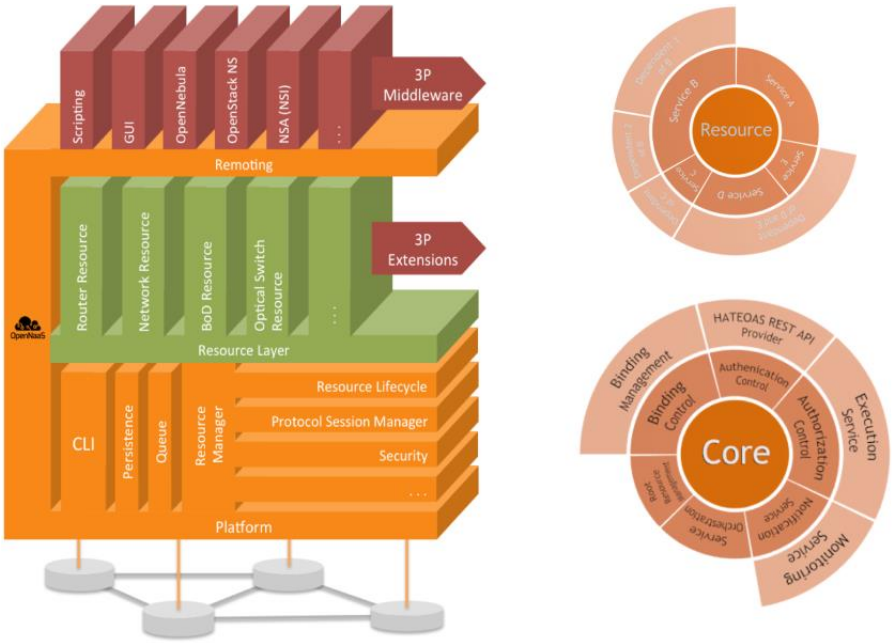

Figure 9. OpenNaaS Architecture (left), NaaS Resource Abstraction (right)

OpenDaylight

OpenDaylight [ODL] is currently the newest and also the most feature-rich SDN controller/management platform. It is backed by the Linux Foundation and developed by an industrial consortium, which includes Cisco, Juniper and IBM, among many others. OpenDayLight includes numerous functional modules which are interconnected by a common service abstraction layer. Further, OpenDayLight provides a flexible northbound interface using Representation State Transfer APIs (REST APIs), and includes support for the OpenStack cloud platform.

In specific, OpenDaylight, as seen in the architecture diagram of Figure 10 (depicting the latest “Helium” release, is built upon four “layers”, i.e.:

- technology-specific plug-ins, for managing SDN and non-SDN devices with various network configuration protocols
- a Service Abstraction Layer, unifying the capabilities of the underlying technology-specific plug-ins
- a core of basic network services, such as topology management, host tracking etc.
- a set of northbound APIs (REST-based) for communicating with network management applications

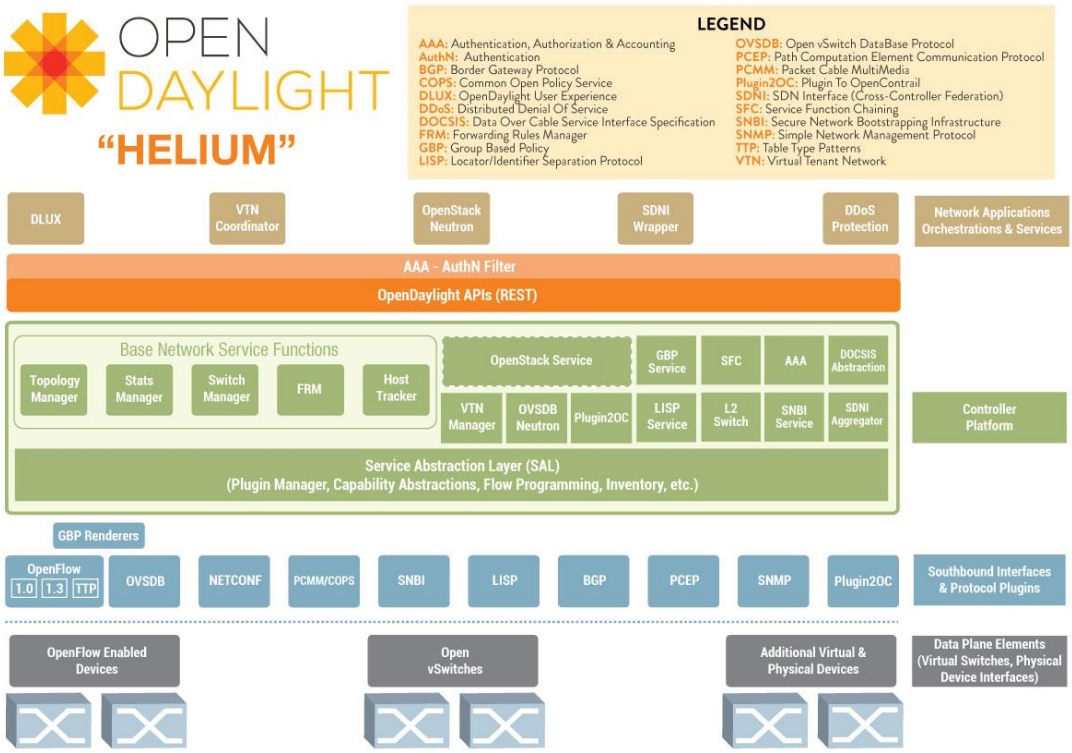

Figure 10. OpenDaylight Architecture ("Helium" release)

OpenDaylight Virtual Tenant Network (VTN)

OpenDaylight Virtual Tenant Network (VTN) [VTN] is an OpenDaylight project, initiated by NEC. VTN provides a multi-tenant virtual networking service on a physical networking infrastructure controlled by OpenDaylight.

VTN allows users to design and deploy virtual networks without needing to know the physical network topology or underlying operating characteristics. The VTN also allows the network designer to construct the virtual networks using common L2/L3 network semantics. That is, it allows the virtual network to be created using virtual links (vLinks), virtual interfaces, as well as virtual network elements (vBridges, vRouters). In this context, VTN allows the users to define the network with a look and feel of conventional L2/L3 network.

Once the network is designed on VTN, it is automatically mapped onto the underlying physical network, and then configured on the individual switches leveraging an SDN control protocol. The definition of the logical plane makes it possible not only to hide the complexity of the underlying network but also to better manage network resources. It achieves a reduction in the reconfiguration time of network services and minimising network configuration errors.

OpenDOVE

Similarly to VTN, OpenDOVE (Distributed Overlay Virtual Ethernet) [OpenDOVE] is also an OpenDaylight project, acting as a network virtualisation platform that provides

isolated multi-tenant networks on any IP network in a virtualised DC. DOVE provides each tenant with a virtual network abstraction providing layer-2 or layer-3 connectivity and the ability to control communications using access control policies. Address dissemination and policy enforcement in DOVE is provided by a clustered directory service. It also includes a gateway function to enable virtual machines on a virtual network to communicate with hosts outside the virtual network domain.

Users interact with Open DOVE to create and manage virtual networks through the Open DOVE Management Console (DMC), which provides a REST API for programmatic virtual network management and a basic graphical UI. The DMC is also used to configure the Open DOVE Gateway to provide connectivity to non-virtualised networks.

The Open DOVE Connectivity Server (DCS) supplies address and policy information to individual Open DOVE vSwitches, which implement virtual networks by encapsulating tenant traffic in overlays that span virtualised hosts in the data centre. The DCS also includes support for high-availability and scale-out deployments through a lightweight clustering protocol between replicated DCS instances. The Open DOVE vSwitches serve as policy enforcement points for traffic entering virtual networks.

The DOVE technology was originally developed by IBM Research and has also been included in commercial products.

2.2.3. Network Functions Virtualization (NFV)

After the success of the cloud computing/storage model, where computation and data are moved from end-user devices to dedicated servers, currently there is an important interest to move the telco operators' infrastructure of network functions to the Cloud too [Verchere11],[Chang12]. The motivation of this initiative is to decrease the CapEx and OpEx of such infrastructure, by decoupling the hardware and software of such network elements, substituting the former with commercial off-the-shelf (COTS) devices. In such new architecture, a Virtualized Network Function (VNF) can be deployed over hypervisors, the actual hardware below being totally transparent. This way, a VNF can be executed over any hardware platform compatible with the hypervisor, as it provides a unified interface to access virtual computing, storage, and network elements.

Both network operators and equipment manufacturers had been working on technologies related to network virtualization for some years, and in particular on the ideas around running network functions on general-purpose servers and cloud infrastructures, mostly inspired by the impact of cloud technologies in all fields of IT. By the beginning of this decade, encouraging results on the performance of these solutions for real-world network workloads were attained, facilitating industry progress in this area.

2.2.3.1. ETSI NFV initiative

The European Technical Standards Institute (ETSI) uses Industry Specification Groups (ISGs) to provide a quick path for the creation of industry fora on specific topics. One relevant example of these groups is the Network Functions Virtualization (NFV) ISG. It aims at addressing the problems faced by network operators and created by an ever growing number of network functions that are implemented in specialized appliances, namely need to find space and power to accommodate them, need for specialized device handlers, short life-cycle, etc. NFV aims at solving this problem by leveraging standard IT virtualization technology to consolidate as many network functions onto standard equipment found in todays' datacenters (Figure 11). NFV is complementary to Software Defined Networking (SDN): while network functions can be virtualized without the need of an underlying SDN infrastructure, both are mutually beneficial.

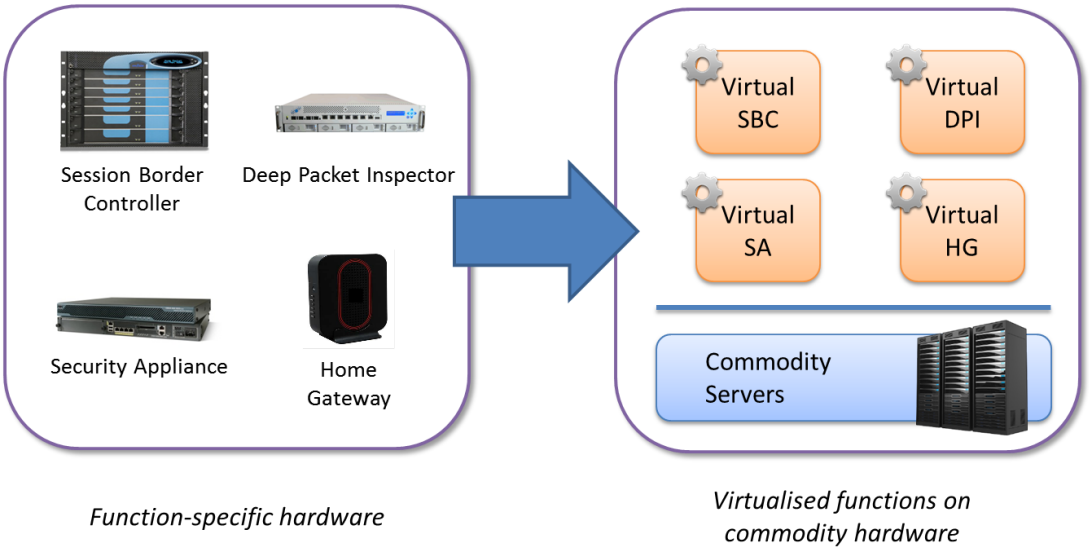

Figure 11. The Network Functions Virtualisation (NFV) concept

NFV is a technology (or set of technologies) aimed to build network infrastructure services the same way IT services are constructed relying on current cloud infrastructures. NFV advocates a homogeneous supporting infrastructure providing computing, storage and connectivity mechanisms, which are expected to be accessed through a common virtualization interface by involved software elements that implement the actual network functions. It is important to note the double role of network facilities. There is a layer of homogeneous, virtualized network mechanisms used to support the interconnection of the elements (hardware and software) required by the software modules implementing the second, an upper layer of network functions running on the infrastructure.

In the first year of its existence the ETSI NFV ISG has produced a series of documents, including: (1) a use cases document [NFV1] that describes the application field addressed by the ISG, (2) a requirements document [NFV2] that handles the

requirements for the NFV framework, (3) the NFV architectural framework [NFV3], (4) the NFV terminology document [NFV4], and (5) a Proof of Concept (PoC) framework document as a way to demonstrate the work of the ETSI NFV ISG through real-life and multi-party PoC implementations [NFV5].

The architectural framework, presented in Figure 12, provides the blueprint for vendors to implement NFV compatible products and is made of a series of building blocks vendors can choose from. The main blocks are:

- the **NFV Infrastructure (NFVI)** providing the virtual resources for executing Virtualized Network Functions (VNFs), including Commercial Off-the-Shelf (COTS) hardware and a virtualization layer to abstract and virtualize it,
- the **VNF** as a software (SW) implementation of a network function,
- the **Element Management System** (EMS) to manage a particular type of VNF,
- the **Management and Orchestration** (MANO) to orchestrate and manage the lifecycle of the VNFs. This block includes the **Orchestrator**, the **VNF Manager**(s) and the **Virtualized Infrastructure Manager**(s).

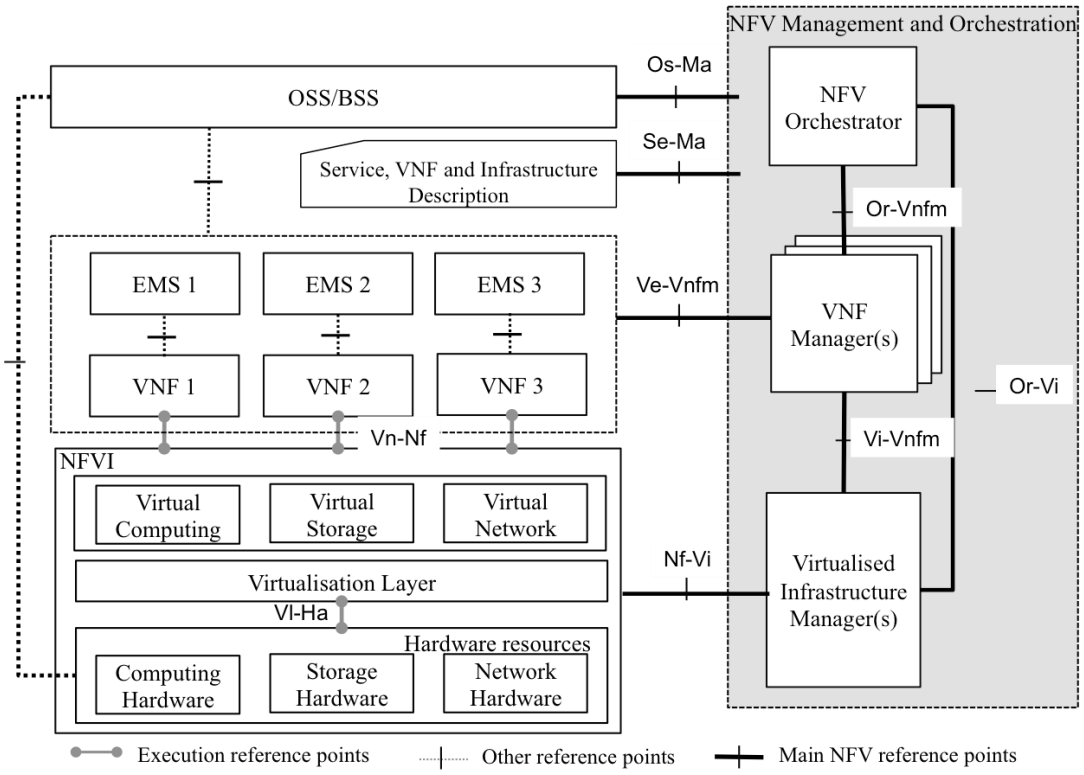

Figure 12. NFV Reference Architecture

While this was probably an initial guess in many cases, there are three essential aspects that distinguish NFV from the direct application of cloud technologies to provide network infrastructure services, and therefore require going beyond carrier clouds to implement NFV.

First of all, the kind of workloads that NFV implies is completely different from the kind of workloads considered by the current cloud practice. VNFs are extremely

dependent on direct I/O and memory operations, and much less on direct computing or storage access. And this not only has impact on VNF performance when deployed directly following “classical cloud” mechanism, but also (and even most significantly) on the portability of VNF instances across the cloud infrastructure. Available experimental evidence shows high performance deviations among workload distributions that were considered completely equal by classical cloud VM placement mechanisms in [NFVPerform]. To properly achieve performance and portability goals, it is necessary to improve cloud orchestrators, hypervisors, kernels and even hardware drivers to support finer-grained placement policies, provide better control of direct memory communication among software instances, and override the virtualization layer for direct I/O to network interfaces.

Secondly, network services need to adapt to network shape. While the classical cloud applications are endpoints in a communication (the archetypal web server in many cases), most network infrastructure services are middle-points (for example, a router or a firewall) and many of them are subject to stringent delay requirements and/or similar constraints. That implies that infrastructures and VM placement strategies must adapt to the network shape and support both highly centralized and consolidated datacentres in the cases they can be used and their economies of scale applied, and much more decentralized schemas. The important point here is not only to support both kinds of deployments but also to be able to seamlessly integrate them.

Finally, when it comes to the orchestration and management of the resources, it is worth noting again that we are dealing with networks at two layers: the supporting infrastructure already present in the current clouds, and the upper network service layer provided by VNFs and their composition into services. To guarantee performance, upper network services may need to directly manipulate the underlying network infrastructure well beyond the limits of usual northbound interfaces exposed by the SDN controllers that are being deployed within current cloud datacentres.

Cloud computing is clearly an essential enabler for NFV, and it is at the root of the NFV concept itself. NFV has to leverage on technologies that are currently applied in cloud computing. At the core of these technologies are hardware virtualization mechanisms by means of hypervisors, and the usage of virtual Ethernet switches for transferring traffic between virtual machines and physical interfaces (though other possible virtualization mechanisms could be applicable, the current focus of the NFV community is on these techniques). Furthermore, current cloud approaches provide methods to enhance resource availability and usage by means of orchestration and management mechanisms, applicable to the automatic instantiation of VNFs, resource management, re-initialization of failed VMs, creation of VM state snapshots, migration of VMs, etc.

2.2.3.2. OPNFV project

The Open Platform for NFV (OPNFV) [OPN] is a new open source project focused on accelerating the evolution of NFV. OPNFV is a Linux Foundation Collaborative Project and it has been launched on October 2014.

The focus of OPNFV is on NFVI (NFV Infrastructure) and VIM (Virtualised Infrastructure Management) components targeting the single specification northbound and southbound APIs to enable interoperability with the rest of the ETSI NFV components.

An initial release of a first implementation is expected for 2015.

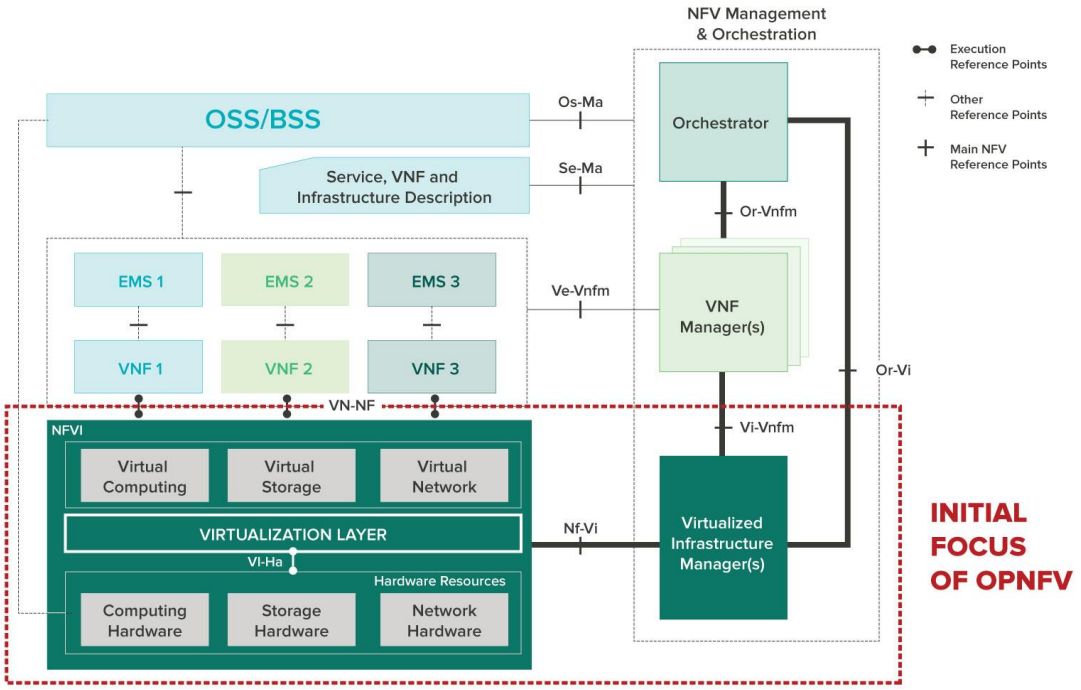

Figure 13. OPNFV scope

2.2.4. Federated Resource Management and Orchestration

There are a variety of deployment models where a number of virtualized assets need to interact in a distributed fashion. The typical cases are the private and the public cloud (see Sec. 2.1.1) Sometimes the private cloud can complement its assets by interacting with a public cloud. When both models are simultaneously in place, emerges the concept of hybrid cloud in which both private and a public cloud temporary cooperate for providing a given service. For a hybrid cloud to be realized, a federated management mechanism must be in place. This mechanism achieves the deployment and management of the service, as a result of a number of control and management actions in each of the involved domains, interfacing with the specific control mechanisms present in each of the administrative domains.

The federation of resources among various computing and network domains is a very relevant case because of the need for coordination to provide a consistent service between the heterogeneous environments existing in each of the domains. For instance, the conventional mechanisms of segmenting the network are hard to scale considering the administrative silos defined for each infrastructure domain. The time required for providing a service increases when several actions need to be

coordinated between domains. Additionally, every change to the service impacts on the separated administrative domains.

The conventional approach for connecting domain resources normally involves the burden of requiring a manual reconfiguration every time a new capability is deployed within a domain, creating inefficiencies and extra costs in the chain. The result is an evident deviation between business-level requirements (in terms of service connectivity and associated Key Performance Indicators) and the provision of the capabilities to be offered.

Although the situation is approachable when considering a domain controlled by a single infrastructure administrator, the methodology turns unfeasible in the case of configuring and operating a distributed environment, where multiple administrative entities are enclosed in a common federation, providing cross-domain services.

From a technical perspective, three types of infrastructure capacities shall be considered:

- Transport capacities: infrastructures that provide connectivity to support service provisioning and access to/from data and users (e.g. via National Research and Education Networks).
- Computing capacities: infrastructures that provide hosting capacities for provisioning software resources (e.g. Data Centers);
- Data capacities: infrastructures that provide data sources that can be connected to applications (e.g. Smart Cities or Sensor Networks);

Such heterogeneous environments can highly benefit from the existence of automated control mechanisms which can coordinate autonomously the resources present in every domain. This can allow performing unified service operation, control and management as it was operated by a single organization.

Even if vendor-agnostic management technologies such as Openflow are adopted, the challenge is only partially met - a common control point based on a single SDN controller is hard to achieve. The heterogeneity of the resources and the underlying technologies can compromise the scalability of such controller. Also, the fact of having different administrative domains prevents from delegating all the control of the infrastructure to an external entity.

However, it is yet required to keep an end-to-end view of the federated resources. The way of doing that is by means of orchestration capabilities on top of separated control domains. The orchestration will allow conjugate resources from different infrastructures composing true end-to-end services while keeping the local complexities being handled by the controller in each domain. The orchestrator will maintain the service awareness and will interact with each of the controllers in the federation responding to the customers' demands.

Actually, local controllers such as SDN controllers and cloud hypervisors should provide the orchestrator with an abstract view of the infrastructure. Technological details should be hidden and, in turn, the offered capacity should be characterized by SLA parameters (performance: guaranteed QoS, reliability, legal: location constraints…). Indeed, the effect of the infrastructure being software-driven means that the underlying hardware resources are not directly visible or addressable, rather

the resources are described and presented using abstract virtual elements. This approach allows the orchestrator to efficiently allocate the available resources independently from the underlying technology. Dynamic and automatic harmonization of heterogeneous domains enables the efficient use of the network resources.

Within a single administrative domain, the modules and components will have specific well-defined functionality, interacting with the other modules and components using task specific APIs. For the inter-administrative domain activity, there will be a set of these components that support various negotiation, control, and information exchange and functions migration operations between administrations. Actually the cost of transport resource is more critical than the cost of computing and storage, thus the optimization of the resources should be considered.

From a networking point of view, four key services should be provided by the federation orchestrator:

1. Provisioning: this capability enables the set-up, release and modification of connections in the network. Its most basic feature is to set up a point-to-point connection between two locations. However, there are other characteristics that a client interface can have like (a) excluding or including some nodes, (b) defining the protection level, (c) defining its bandwidth or (d) defining its disjointness from another connection
2. Topology discovery: this functionality requires unique identifiers for the exported network topology information. Network identifiers (such as IPv4 or datapath-IDs) help to carry out path computation and to integrate the nodes for an end-to-end scenario. Further, the local controllers can provide information about the links in the domain (physical or virtual) or even their utilization. It is clear that the more information is shared, the less abstracted the network appears. So it is important to highlight that the North Bound Interface (NBI) of the orchestrator should be more abstract than the NBI of the controller.
3. Monitoring: to collect the status of the connections that have been created is very useful in a multi-controller scenario, where after a failure in one domain, the domain's controller may request another connection.
4. Path computation: Global path computation is an important feature because individual controllers in each domain are only able to share abstracted information that is local to their domain. An orchestrator with its global end-to-end view can optimize end-to-end connections that individual controller cannot configure. Without a path computation interface, the orchestrator is limited to carrying out a crankback[1] process.

Next sub-sections present two cases for DC federation.

[1] I.e. a trial-and-error approach without detailed knowledge of the overall topology, in which alternative paths are tried every time the creation of a path is denied by one of the involved domains.

2.2.4.1. XIFI

The XIFI [XIFI] infrastructure consists of a pan-European open federation of computing resources hosted in DCs accounting a total of 19 nodes spread across 12 countries. The XIFI federation is a compendium of test infrastructures, also known as nodes or regions, belonging to independent administrative organizations, which in a collaborative manner provide multi-domain cloud-based computing services.

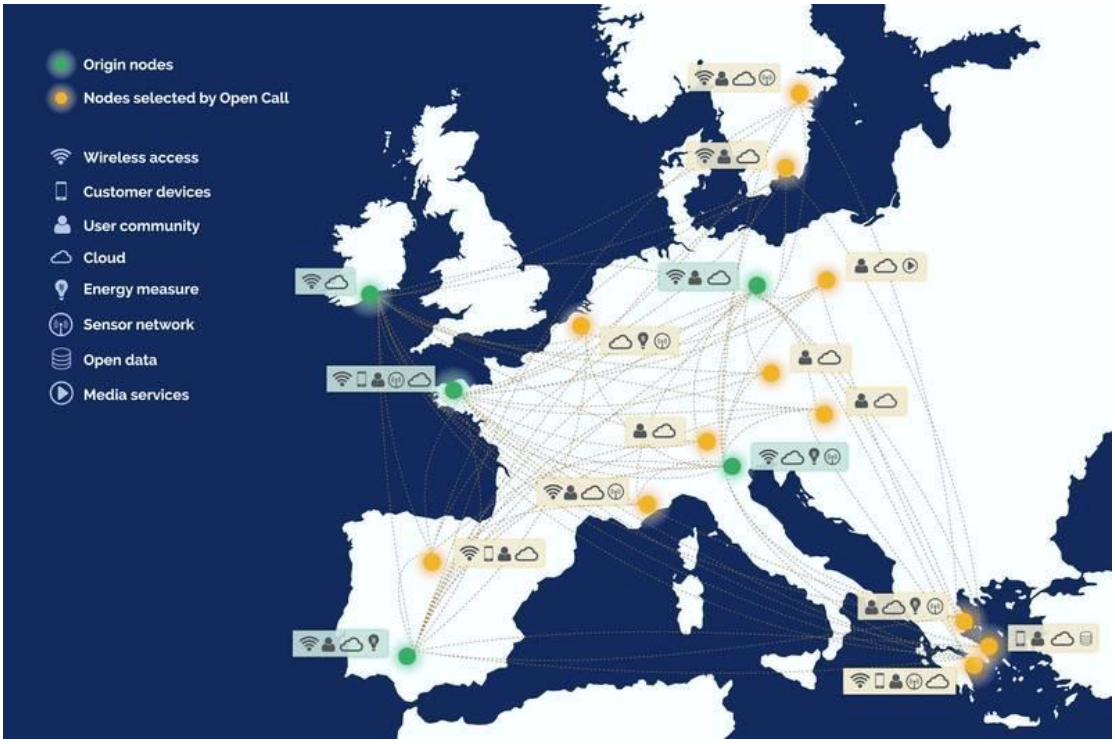

Figure 14. The XIFI federation

In principle, the connectivity between the DCs will be provided by the NRENs (National Research and Education Networks) since the project instigation departs from the European research community. However other forms of connectivity are also feasible if new participants to the XIFI network are not connected to the NRENs.

Each involved DC makes available a number of servers and storage devices to the federation. The usage of these resources by the end users are requested by means of a central portal to the federation which will expose the available capabilities for being selected by the users according to their needs.

Both the DC owners hosting the computing resources and the NRENs providing the inter-DC capacity are running additional services to that offered by XIFI in their infrastructures. That is why service isolation is fundamental in order to avoid impacts from one service in the other.

The computing resources are attached to OpenFlow-enabled switches that constitute a demarcation point for the XIFI Network Controller. Those OF switches are the ones to be instructed by the XIFI Network Controller for implementing the end to end

service. From the demarcation point to the WAN, a number of transport services are pre-provisioned in order to be dynamically selected by the XIFI Network Controller algorithms according to the end user requests. These pre-provisioned connections traversing the NRENs are dynamically stitched to the connectivity service created in the internals of the DC, comprising the VM, the Open vSwitch and the OF switches up to the demarcation point.

In particular, the selected management solution for the computing capabilities of each DC is based on OpenStack. Current versions of OpenStack allow to create the L2 networks in a single Data Center with a variety of technologies (GRE tunnels, VxLAN, VLANs, etc). However, the networking capabilities of an OpenStack instance nowadays are restricted to the directly managed devices, typically located in a single datacenter.

In order to allow the end to end service composition a XIFI Network Controller is defined to orchestrate the service by interacting with the local SDN controllers.

Each data center has its own independent OpenFlow controller. However, such OpenFlow controller will only be used for provisioning purposes (pushing the rules in the switches) and to feed topological information. The routing and orchestration is provided by the ABNO implementation, which is done inside the project. Initially, Floodlight controller has been selected to statically push the OpenFlow rules and provide the local topology inside the data center. However, other controllers can also be used (see Sec. 2.2.2.2.)

This XIFI Network Controller architecture is based on the Application Based Network Operations Architecture (ABNO) framework, defined in the IETF [ABNO]. This framework groups together a number of standard components and protocols able to control and manage the network in a coordinated way, taking into consideration requests from external applications. The ones relevant for this experiment are: (i) the *ABNO Controller* which is the module executing the workflows and coordinating the individual actions of the rest of ABNO components; (ii) the *Path Computation Element* (PCE) which is in charge of providing the best path in a domain, according to specific constraints; (iii) the *Topology Module* (TM) that maintains an updated view of the topology graph; and, (iv) the *Provisioning Manager* (PM) which directly interacts with the network elements (either directly or via an intermediate control plane element).

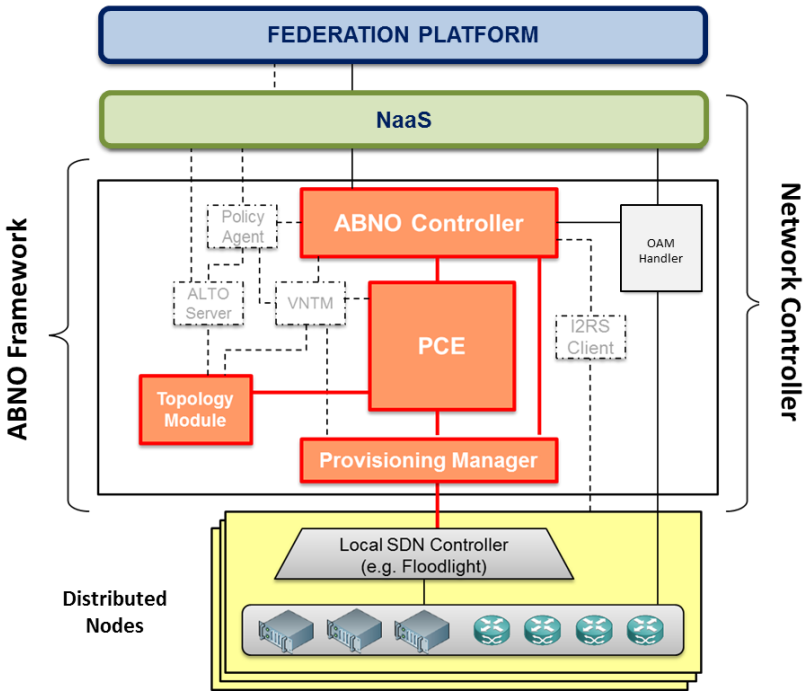

Figure 15. XIFI Network Controller

The NaaS module is required to provide connectivity service awareness across the federation. It is in charge of collecting and maintaining the information about the resources committed to the end user from the connectivity point of view, either if the resources are local to just one federation node or if they are spread in different nodes in the federation. One example of NaaS implementation is the OpenNaaS framework, as already described.

2.2.4.2. ExoGENI

GENI (Global Environment for Network Innovations) [GENI] is a US program to develop and deploy integrated network testbeds, providing a virtual laboratory for networking and distributed systems research and education. GENI allows the establishment of large-scale experiment infrastructure based on layer-2 connectivity. GENI makes use of infrastructure slices to allow multiple experimenters run multiple experiments at the same time.

ExoGENI [ExoGENI] is a testbed part of GENI that conjugates networking and cloud computing, by offering Infrastructure-as-a-Service (IaaS) cloud model capabilities together with orchestrated provisioning across sites. Some functions are delegated to the federation, such as identity management, authorization, and resource management.

In the networking part OpenFlow is used within each site, connecting (through direct Layer 2 site connectivity) to circuit backbone fabrics (through host campus networks) for inter-site connectivity, including international connectivity, by means of programmable exchange points. The flexible networking operations use traditional VLAN-based switching and OpenFlow.

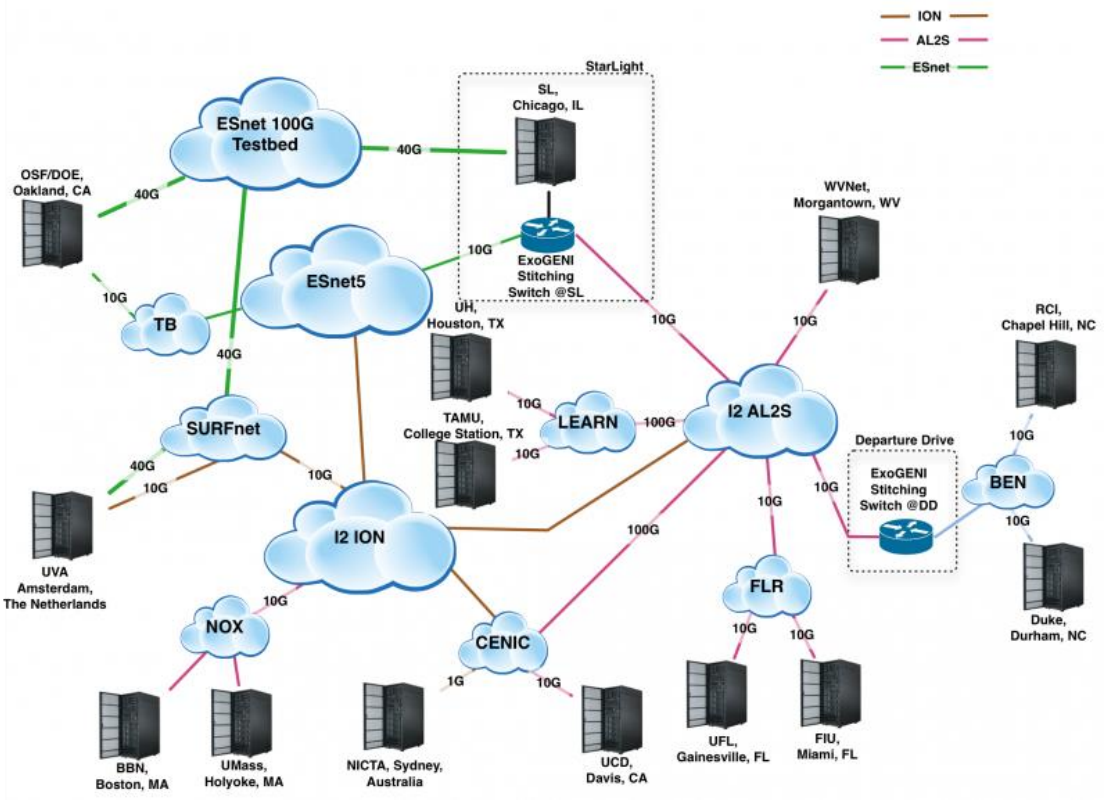

Figure 16. ExoGENI Federation

Through the ExoGENI testbed is possible to create slices with user-driven packet forwarding control via OpenFlow. OpenFlow slices are restricted to VLANs provisioned within and between racks.

ExoGENI (Figure 17) is managed by a control framework software named Orca which is used to create integrated slices of compute elements and Layer2 links, with optional OpenFlow integration. ORCA is tightly integrated with OpenStack and Eucalyptus via special extensions to both for provisioning virtual machines.

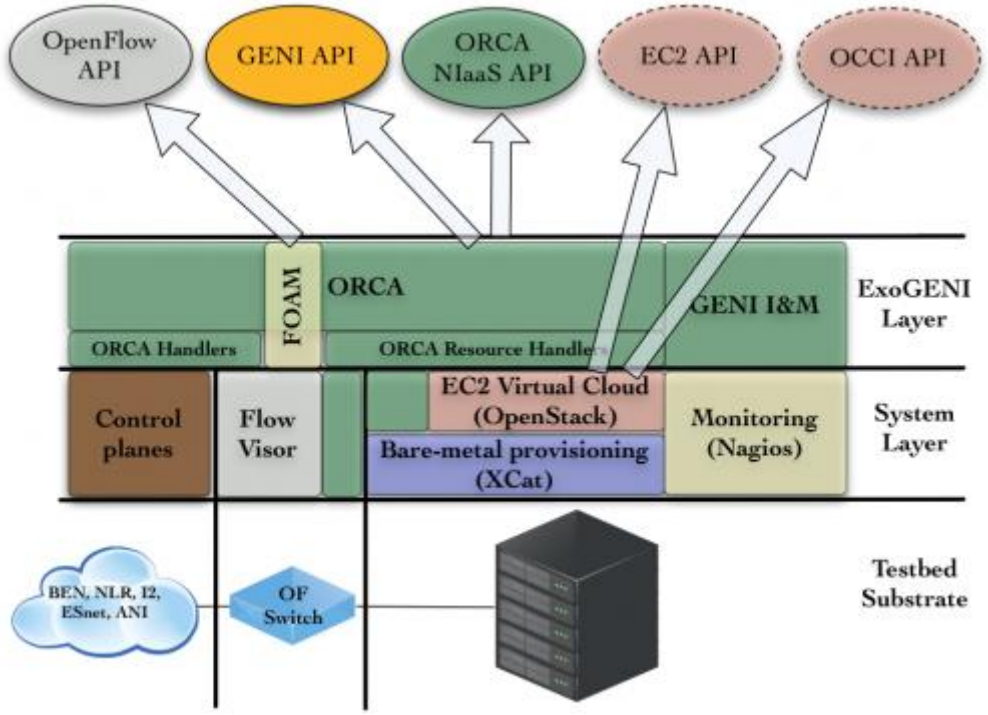

Figure 17. ExoGENI software structure

2.3. Applicability to Wireless Environments

The softwarisation and virtualization of networking infrastructures are also emerging in wireless and mobile environments. The motivation for that can be found in the ever new powerful terminals (e.g., smartphone and tablets) and the proliferation of data-hungry mobile applications (e.g., mobile health, mobile education, context-awareness applications, cloud communication etc.). In order to cope with such a demand, some network operators are now following a cloud computing paradigm, enabling the reduction of the overall costs by outsourcing communication services from specific hardware in the operators' core to server farms scattered in datacenters. These services have different characteristics if compared with conventional IT services that have to be taken into account in this cloudification process.

No less significant is the continued transformation of cellular networks, whose objective is to offer to portable devices bit-rates (bandwidth) and Quality of Service (QoS) comparable to those traditionally made available only through fixed networks.

Some of the mechanisms that are being considered and already adopted by operators include: sharing of network infrastructure to reduce costs, virtualization of core servers running in data centers as a way of supporting their load-aware elastic

dimensioning, and dynamic energy policies to reduce the monthly electricity bill. However, this has proved to be tough to put in practice, and not enough. Indeed, it is not easy to deploy new mechanisms in a running operational network due to the high dependency on proprietary (and sometime obscure) protocols and interfaces, which are complex to manage and often require configuring multiple devices in a decentralized way.

On these grounds, the Network Function Virtualization (NFV) and Software Defined Networking (SDN) emerge as the latest incarnation of technological promises for reaching the necessary cost efficiency, not only in fixed, but also in wireless and mobile networks. The following sections better illustrate this trend.

2.3.1. Software Defined Wireless Networks

Figure 18 shows an Software-Defined Wireless Network (SDWN)-based architecture [Bernardos14] of a mobile network operator, where a solid line in the figure denotes a user plane connection, and a dashed one is used for control plane. We take the 3GPP Evolved Packet System as reference architecture to link the proposed concepts with a well-established and understood system architecture.

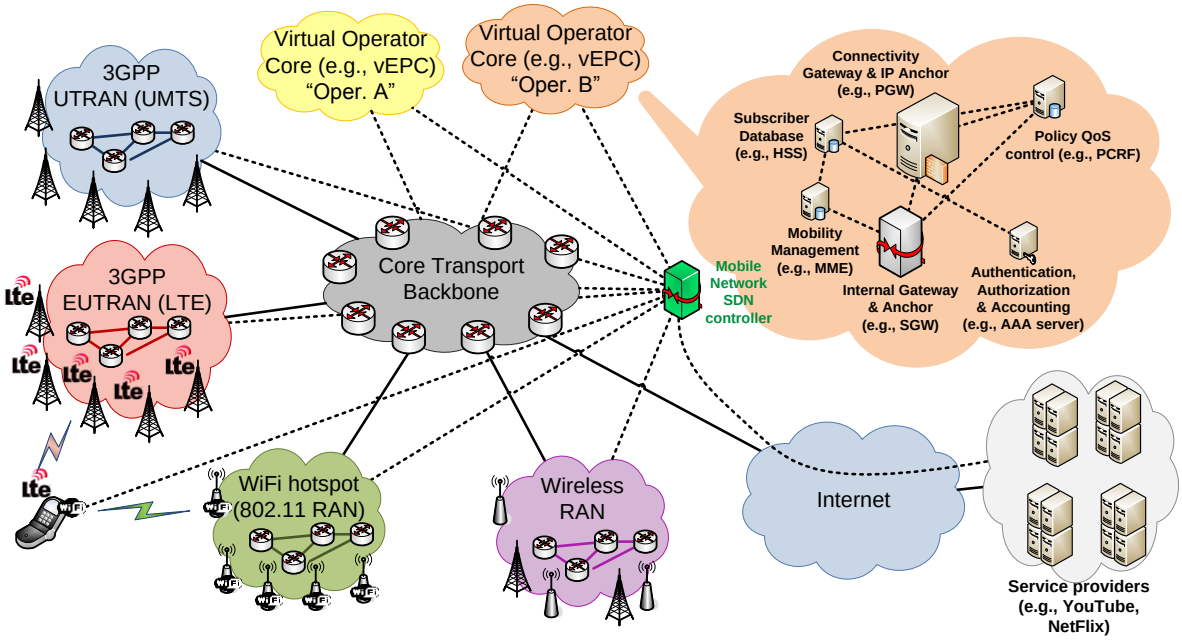

Figure 18. Software Defined Wireless Network concept

A mobile network typically exhibits multiple heterogeneous Radio Access Networks (RANs) connected to a common transport core network. Note that the connection between the last network entity providing radio access and the core transport network might involve a wired or wireless backhaul network (shown as part of the RAN in the figure) by using a combination of technologies (e.g., fiber optic, microwave) and topologies (e.g., ring structure, daisy chain) in the backhaul segment. Three well-known examples of RANs are shown in Figure 18: the UTRAN (for UMTS), E-UTRAN (LTE), and a WiFi hotspot. However, note that the proposed architecture is generic enough to support other RAN technologies as well, both already existing ones (such as WiMAX), or future ones.

In the SDWN architecture, radio access networks are enhanced with programmability (as introduced in more detail below), supporting multiple functionality levels to allow for incremental deployments. The core transport is composed of programmable L2 switches and L3 routers, allowing to set-up unicast and multicast forwarding at flow level (as supported for example by OpenFlow). Multiple (virtual) operators might share part of the radio, backhaul and transport core network, which requires the interconnection of the core control plane entities – in charge of functions such as authentication, authorization, charging, subscriber management, mobility management, QoS provisioning or connection to external services/networks – with the programmable network.

Two different models can be adopted to implement an SDWN architecture: “evolutionary" and “clean slate.” The evolutionary model allows for incremental deployment in existing networks: legacy control plane entities from the operators can connect to the transport core network without modifying the existing interfaces. In this model, the SDN controller implements standardized interfaces to support the

internetworking with existing legacy entities, even if they run on a virtualized environment (what is known as virtual EPC, vEPC).

In the clean slate model, the control plane functions are directly programmed on the SDN controller or on top of it as applications, using a software API between the virtual operators and the SDN controller. While this approach does not allow for an easy incremental deployment, it brings since day one all the advantages of programmable network architectures. For example, the deployment of new network functions and services is much easier and faster, as it can be directly implemented on the controller and does not need to impact on multiple interfaces and equipment from different vendors. We can just take the simple, but very representative, example of IPv6 support on a mobile network. With the clean slate approach, adding IPv6 support would just require additional code on the SDN controller, as compared to defining new interfaces, procedures on the different control and user plane entities, which require software/firmware updates (if not even replacing some hardware).

The brain of the architecture, the SDN controller, is connected to each programmable entity. Note that the SDN controller is a logical entity, which might also be decentralized into different physical boxes to improve scalability and performance, though this is currently the subject of extensive research.

In order to allow for third parties (such as service and application providers) to influence/control the behavior of the network, an API is enabled. This API effectively enables external players to get access to the network resources, similar to what an OS does with the access of applications to computational resources and peripherals. The API offered by the SDN controller supports different access levels to the external parties, so personalization can vary on different dimensions: per application, per user, per (virtual) operator, per access network, or a combination of them.

The following interfaces are of interest:

- *A northbound interface to the (virtual) operators* sharing the same physical set of network resources allowing them to dynamically change the share of resources, for example to adapt to network load or to the number and profile of users attached to the physical shared network at any given moment of time. This interface should be able to implement richer SLAs as compared to the ones available nowadays, as a more dynamic and almost real-time reconfiguration of the network would be possible. Each (virtual) operator should have access to an abstracted view of its assigned resources, so they can program that "virtual" network as a physical one.
- *A northbound interface to the external parties* (service and application providers) authorized to influence the network behavior. This interface should be properly secured, granting access with different granularities and permissions. The interface should be powerful enough to allow an application provider to influence how its traffic is handled, even taking into consideration the virtual operator its users are getting access from. Note that this is possible because of the centralization achieved by the use of the SDN approach, though this may introduce scalability issues (e.g., up to per flow signaling, need for frequent network monitoring, etc) that need to be taken into account.

- *A southbound interface to the physical user-plane network entities* in the core transport backbone. This interface is used by the SDN controller to implement the different behavior policies according to the requests from the external parties, the virtual operators associated to the different users attached to the network, and the network conditions. Given the logical centralization provided by SDN, close to maximum utilization of the capacity of the network links can be achieved. This interface also allows for effective sharing of a common backbone and backhaul network by different operators, which may even connect to the Internet via different gateways.
- *A southbound interface to the physical user-plane entities in the RAN*. This interface allows for effective virtualization of the access network, therefore sharing the same physical resources among different operators. Besides, this interface should allow programming the wireless access technologies to provide the expected behavior, depending on the specific needs and characteristics of the mobile terminal, the requests from the external providers and the different SLAs that the virtual operators may have in place with their users.
- *A southbound interface with the mobile node.* This interface provides the network with certain programmability capabilities on the mobile node. This can be used for example to improve the mobility experience, by better exploiting the simultaneous use of available wireless access networks, e.g., helping in access network and interface selection.

As summary of the proposal, the following table captures the major benefits identified in [Bernardos14] and the corresponding challenges to reach such goals.

Table 2: The case for SDN in Mobile Networks

Key Benefits	Key Challenges
Easier deployment of new services	Specification of the interfaces
Reduced management and operational costs of heterogeneous technologies	Need to integrate scheduled-based and contention-based systems
Efficient operation of multi-vendor infrastructures	Harmonization of the standardization efforts
Increased accountability and service differentiation	Verifiable security and privacy architecture
Continuous and transparent enhancement of network operation	Operation and management of wireless networks is more complex

2.3.2. Open Networking Foundation proposition for Mobile and Wireless

Wireless and Mobile networks have unique challenges that maybe addressed effectively by SDN solutions. These challenges include aspects like the support of the massive growth mobile data, the communication handover and simultaneous operation over multiple wireless medium, the rapidly evolving mobile service ecosystem, or the dynamic and changing physical environment and medium where the communication takes place.

In this context, last year, the Open Networking Foundation (ONF) chartered a Wireless & Mobile Working Group (WMWG)[2] to foster the adoption of OpenFlow-based SDN technology to the mobile and wireless networks. This group studies and proposes simplified reference mobile architectures for transport by leveraging OpenFlow-based SDN. The OpenFlow protocol (being standardized by the ONF) is a key component of the SDN concept to model flow abstractions.

The WMWG group aims to simplify the interaction between wireless networks and packet networks. Moreover, it is also responsible for proposing extensions to the OpenFlow protocol specification (e.g. extensions required to transfer the GTP tunnel-ids to the switches), and the parameters for interaction between the control and user planes. Current work in ONF WMWG is structured in three projects: the Mobile Packet Core, the Wireless Transport, and the Enterprise unified wired/wireless access projects.

The Mobile Packet Core project explores the application of OpenFlow to 3GPP Evolved Packet Core (EPC), considering user/data plane separation in GW, mobility management and mobile flow steering for offload.

The Wireless Transport project deals with the problematic of the wireless backhaul links, where a central SDN controller can help to optimize radio parameters in data plane using OpenFlow.

Finally, in the Unified Access project the idea to develop a unified access network that uses a common controller to manage both wireless access points (AP) and wired switches.

The ONF has publicly released a solution brief related to mobile and wireless networks. Two use cases are covered: (1) inter-cell interference management; and (2) mobile traffic management.

In the first case, the application of the SDN-based control is applied to the radio resource management. The logically centralized control layer enables radio resource allocation decisions to be made with global visibility across many base stations. Scalability is improved because as new users are added, the required compute capacity at each base station remains low.

[2] ONF Wireless and Mobile Group, https://www.opennetworking.org/working-groups/wireless-mobile

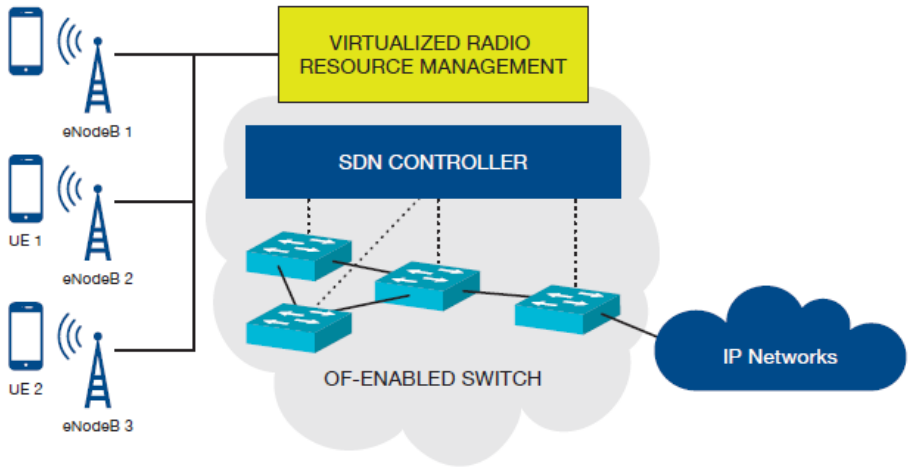

Figure 19. SDN-based wireless Resource Management

In the second case, a finer mobile traffic management can be achieved through mobile traffic offloading and roaming, content adaptation (such as adaptive streaming), and mobile traffic optimization. All of those situations can be greatly benefit from centralized SDN control platform.

Among the improvements foreseen it can be remarked the more granular network control that can be achieved by software driven mobile and wireless networks. It is expected to allow mobile operators to efficiently apply policies at a very granular level—including the session, user, device, and application levels—in a highly abstracted, automated fashion. This control enables mobile operators to support multi-tenancy while maintaining traffic isolation, security, and elastic resource management when customers share the same infrastructure.

2.3.3. Advanced services in the path to 5G – The C-RAN concept

Mobile access networks are of particular interest for network operators because the high capacity and capillarity they require to satisfy end-user expectations; this will become even more evident with the advent of 5G wireless networks. This scenario forces to explore new ways of deploying the necessary infrastructure to fulfil end-users requirements in a cost effective manner.

One of the recent trends in the mobile industry is the centralization of some functions of the Radio Access Network (RAN), named as the Centralized-RAN (C-RAN) or Cloud RAN approach. C-RAN can be considered as a special application of NFV paradigm to embrace the radio access front-end of a mobile network. The concept was first introduced by researchers at the China Mobile Research Institute [CRAN11]. Since then, the topic has attracted huge attention by the research community and the mobile network operators, which foresee significant benefits.

While traditional base stations and access points aggregate all radio front-end functionalities in a single hardware platform, the C-RAN concept applies the generic NFV softwarisation paradigm by migrating the L2 and L1 operations to software modules deployed in data centers. More specifically, C-RAN envisages the functional

split of existing eNodeBs separating the Radio Units (RU) from the Base Band Unit (BBU) and linking them by a remote connection by means of high-speed Common Public Radio Interface (CPRI). C-RAN proposes allocating common radio-access processing resources, namely Base-Band Units (BBUs) in centralised nodes with powerful computing capabilities, typically a DC, while just keeping remote only the infrastructure strictly needed to provide the wireless connectivity, i.e. the Radio Remote Units (RRUs) or Remote Radio Heads (RRHs).

In this setup, depicted in Figure 20, BBUs perform all front-end functionalities in software, over generic hardware, and produce the baseband signal which is transmitted encapsulated in a data network to the RRUs in the actual transmission sites. RRUs act only as “dumb” radio interfaces, performing only the digital-to-analog conversion, upconversion, amplification and transmission. Received signals are down-converted, digitized, and sent over the backhaul network to the BBU pool for processing.

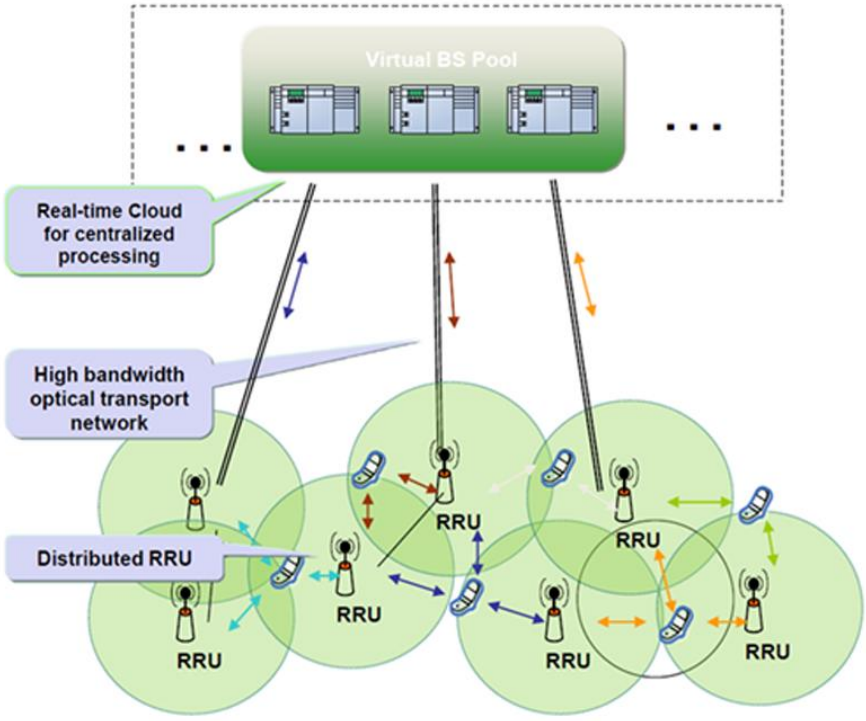

Figure 20. The C-RAN approach

The benefit of C-RAN is that L1/L2 operations can be carried out in commodity hardware, significantly reducing CAPEX and enabling easy upgrade of the network with new radio protocols, just via a software update. Moreover, resources can be dynamically pooled and assigned on-demand, eliminating the risk of resource shortage on the radio front-end. Furthermore, the C-RAN concepts allows virtualizing the RAN part of the network and splitting it across multiple tenants, each of which gains access to a separate “slice” of the BS. This is the concept of “Radio Access Network as-a-Service” (RANaaS).

Several recent research activities adopt the C-RAN concept, mostly applied to cellular networks (base stations) and small cells.The FP7 iJoin project (see Sec. 2.4.10) focuses on the cloud-based virtualization of the cellular radio front-end (eNodeB virtualization) aiming at efficient resource management and RANaaS. Similarly and with a more extended scope, Mobile Cloud Networking (see Sec. 2.4.4) extends mobile network virtualization to also reach the base station part towards providing an end-to-end virtual mobile network platform.

The C-RAN concept can be benefited of the deployment of computing facilities in DCs towards the access. Mobile-Edge Computing (MEC) is a possible solution to this problem. By providing a cloud server running at the cell site, specific tasks can be performed directly at the network edge, something that cannot be achieved with traditional network infrastructure [ETSIMEC]. Applications include active device tracking, augmented reality application delivery, RAN-aware content optimization, distributed content and DNS caching, and application performance optimization, to name a few. Also, machine-to-machine gateway and control functions can leverage such an approach. MEC expected benefits include:

- Optimized user Quality of Experience by bringing the content and compute capabilities closer to the customer
- Reduced content transport cost with increasing revenue from data mobility services
    - Backhaul traffic savings – up to 50%
    - Increased radio utilization – around 25%
- Web application delivery optimization
    - DNS & Content Caching close to cell tower
    - Adding network intelligence to the RAN
    - Eliminate redundant and repetitive content transmission
    - Support of multiple mobile technologies, with transparent handovers
- New monetization options for both internet and mobile players
- Totally transparent to the mobile network and terminals

Furthermore, C-RAN can be coupled with SDN-based management for more efficient federation and allocation of resources. In [Yang13] the authors propose an SDN controller is the control plane of heterogeneous RANs by abstracting and combining control functions of the access elements. Such controller is used not only to configure rules for the traffic flows but also to optimize spectrum, computing and storage resources.

2.4. Related Research Projects

This section surveys a number of research projects related to cloud networking technologies.

2.4.1. 4WARD

The FP7 4WARD project [4WARD] aimed to radically transform the existing Internet architecture by introducing specific innovations. Network virtualisation was one of the primary targets of the project, while others (e.g. in-network management and network of information) were much less relevant to cloud networking.

In the area of network virtualisation, the goal of 4WARD was to develop a systematic and general approach, focusing on three main areas:

1. Virtualisation of Network Resources: virtualisation of both wireless and wireline resources; performance optimisation of shared resources and secure separation of virtual networks sharing a resource; development of standardised interfaces for management and control of the virtualised resources.
2. Provisioning of Virtual Networks: instantiating complete virtual networks using the virtual resources, allowing the on-demand deployment of new virtual networks on a potentially large scale; establishment of a virtualisation framework including the discovery of available physical and virtual resources; scalable provisioning, control, and aggregation of resources to form complete networks.
3. Virtualisation Management: deployment, control, and dynamic re-allocation of resources on demand during the lifetime of the virtual network; dynamic management of volatile and mobile resources that may enter or leave the virtual network at any time.

2.4.2. SAIL

The FP7 SAIL project (Scalable and Adaptive Internet Solutions) [SAIL] aimed at designing enabling technologies and architecture for Future Networks. Among these technologies are:

- Cloud Networking[3] (CloNe) facilitates on-demand management and control of computing, storage and connectivity resources in the network, by automatically moving or scaling up or down the resources required to distribute content and applications. The multiprovider approach of CloNe allows data centres and network provider to cooperate to offer a service end-to-end. Applications can benefit of placing their data and running software at the most appropriate place in the network.
- Open Connectivity Services (OConS) provide a toolbox to extend current networks to interconnect applications and end-users, and to deliver content the best way, optionally combining multiple physical network links in a multi-path concept. OConS offers services to extend IP and L2 networks. It provides mechanisms to deliver content the best way, using a large set of connectivity technologies. OConS Multi-P, which stands for multipoint, multi-path, multi-

[3] In SAIL, the term “Cloud Networking” mostly referred to the management of in-network computing clouds

protocol, enables to deliver data using the best connection available, taking advantage of the fact that many devices use multiple connectivity technologies, with or without wires. OConS offers services that integrates the different layers and domains to get a more efficient data plane.

2.4.3. ALICANTE

The FP7/ICT ALICANTE project [ALICANTE] worked towards the deployment of a "Media Ecosystem", by proposing a comprehensive architecture with innovations and evolutions for the common actors of the Networked Media Value Chain: The Content Providers, the Service Providers, the Network Providers, and of course, the End-Users. ALICANTE aimed to provide Content-Awareness to the Network Environment, thanks to the proposed CAN (Content-Aware Network) Layer on top of the classical Internet infrastructure, and the main element constituting it: the Media-Aware Network Element (MANE), new edge router with Content-Aware capabilities.

The ALICANTE approach for network-as-a-service involved the instantiation of virtual Content-Aware Networks (VCANs) on top of the network substrate. VCANs are virtual networks, implemented via MPLS tunnels with specified QoS, which feature inherent content-awareness i.e. they differentiate between traffic flows and assign different QoS levels to different applications. This feature is supported by the MANE.

The management entity which allows VCAN instantiation is the CAN Manager, which accepts requests from Content/Service Providers and configures the MANEs in order to set-up the VCANs.

2.4.4. Mobile Cloud Networking (MCN)

The Mobile Cloud Networking (MCN) [MCN] is a 3-year European co-funded (FP7) project, which started in November 2012 and will end in October 2015. The project approaches the integration between the Cloud and Telco worlds, making Operators benefit from the principles of virtualization.

The project focuses, in particular, on mobile operators. For this reason, leveraging cloud and NFV concepts, the main target is to fully cloudify the whole components of a mobile network operation, namely (see Figure 21):

- the access (RAN - Radio Access Network);
- the core (EPC – Evolved Packet Core);
- the services (IMS – IP Multimedia Subsystem, CDN – Content Delivery Networks, DSS – Digital Signage);
- the Operational Support Systems (OSS) (Provisioning, Monitoring, SLA Management);
- the Business Support Systems (BSS) (CRM – Customer Relationship Management, Charging, Billing).

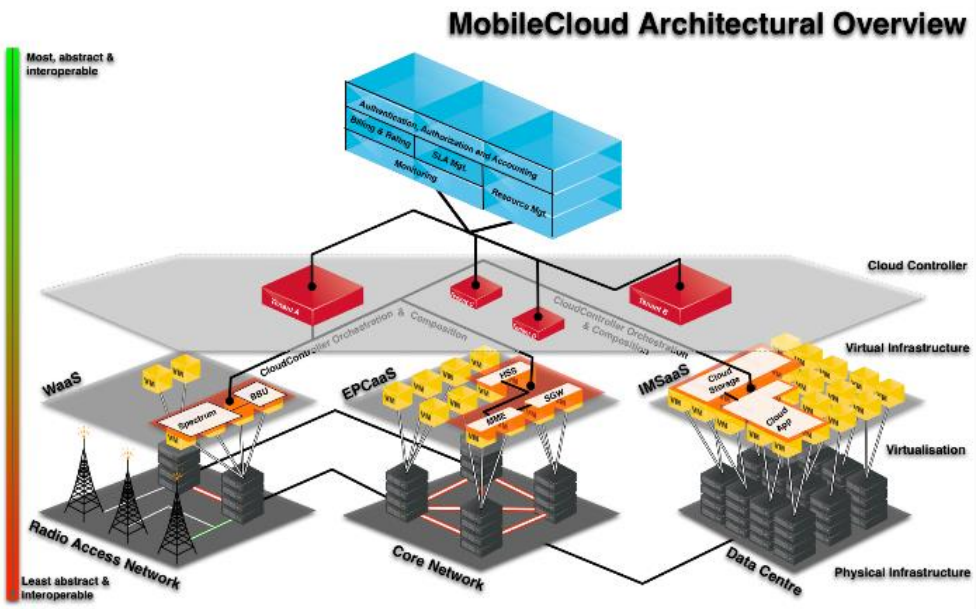

Figure 21. Mobile Cloud Networking concept

Beyond the virtualization, the project explores the "as a Service" (XaaS) concept, where the functions are provided as a full operational service. That means the customer of the service does not need to worry about implementation, deployment and dimensioning details of it.

In specific, the following main MCN services are envisaged:

The RANaaS (Radio Access Networks as-a-Service) service involves the virtualization of the RAN head-ends. RAN can be split into the RRU (Remote Radio Unit), which transmits the physical signals through the antennas, and the BBU (Base Band Unit), which processes the signal and make all necessary computations at baseband level. The RANaaS virtualizes the BBU part, by softwarizing it and allocating computing resources for BB operations.

The EPCaaS (Evolved Packet Core as-a-Service) service provides LTE connectivity to the end users, by forwarding traffic from the RAN to the Internet and vice-versa. The ECPaaS virtualizes the core as a whole by virtualizing its internal components (e.g. P-GW, S-GW, MME, etc.).

The IMSaaS (IP Multimedia Subsystem as-a-Service) service virtualizes the IMS internal components (e.g. P/I/S-CSCF, HSS, etc.) in a similar way as the EPC.

The CDNaaS (Content Delivery Network as-a-Service) service provides a caching mechanism that can be used to cache popular content close to the content consumer.

Finally, the DSSaaS (Digital Signage System) service provides a digital signage service. The virtualization allows the adaptation of the service to the usage performance.

In addition, easy creation of "end-to-end (e2e)" services by composition of basic services. As an example, it can be considered the creation of an MVNOaaS service by the composition of RANaaS+EPCaaS+IMSaaS.

2.4.5. T-NOVA

FP7 T-NOVA specifically focuses on the aspects of Network Functions Virtualisation (NFV). It aims to introduce a novel enabling framework, allowing operators not only to deploy virtualized Network Functions (VNFs) for their own needs, but also to offer them to their customers, as value-added services. Virtual network appliances (gateways, proxies, firewalls, transcoders, analyzers etc.) can be provided on-demand as-a-Service, eliminating the need to acquire, install and maintain specialized hardware at customers' premises.

For these purposes, T-NOVA designs and implements a management/orchestration platform for the automated provision, configuration, monitoring and optimization of Network Functions-as-a-Service (NFaaS) over virtualised Network/IT infrastructures. T-NOVA leverages and enhances cloud management architectures for the elastic provision and (re-) allocation of IT resources assigned to the hosting of Network Functions. It also exploits and extends Software Defined Networking platforms for efficient management of the network infrastructure. In other words, T-NOVA combines IT/cloud virtualisation and Network-as-a-Service concepts to offer a complete end-to-end Cloud Network service.

Furthermore, in order to facilitate the involvement of diverse actors in the NFV scene and attract new market entrants, T-NOVA establishes a "NFV Marketplace", in which network services and Functions by several developers can be published and brokered/traded. Via the Marketplace, customers can browse and select the services and virtual appliances which best match their needs, as well as negotiate the associated SLAs and be charged under various billing models. A novel business case for NFV is thus introduced and promoted.

As Application Scenarios, T-NOVA considers four VNFs for proof-of-concept purposes: a Deep Packet Inspector (vDPI), a Security Appliance (vSA), a Session Border Controller (vSBC) and a Home Gateway (vHG).

As shown in Figure 22, the overall The T-NOVA architecture can be hierarchically organised into four architectural layers:

- The NFV Infrastructure (NFVI) layer includes the physical and virtual nodes (commodity servers, VMs, storage systems, switches, routers etc.) on which the services are deployed.
- The NFVI Infrastructure (NFVI) Management layer comprises the infrastructure management entities (VIM, TNM). In the sections as well as the deliverables to follow, the NFVI and management layers are conceptually grouped under the name Infrastructure Virtualisation and Management (IVM).
- The Orchestration layer is based on the T-NOVA Orchestrator and also includes the NF Store
- Finally, the Marketplace layer contains all the customer-facing modules which facilitate multi-actor involvement and implement business-related functionalities.

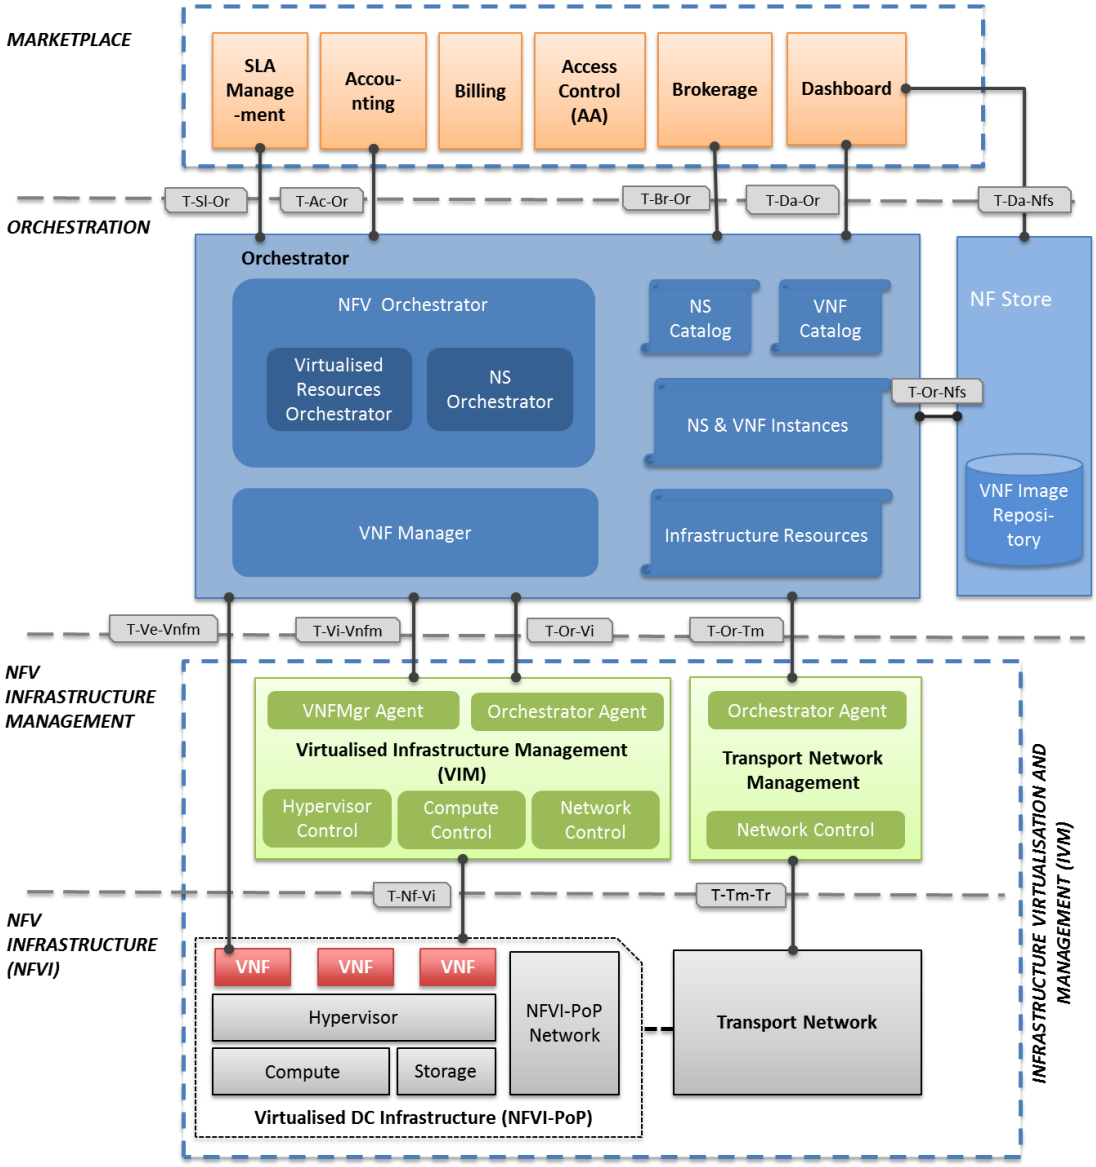

Figure 22. Overall high-level T-NOVA Architecture

More details on the T-NOVA Architecture and the functionalities of the different subsystems can be found in [TND221].

2.4.6. XIFI

XIFI [XIFI] is a FI–PPP project that is motivated by the need to establish a federation of infrastructures and an associated marketplace for large trials of Future Internet applications and services. XIFI addresses the infrastructure element of this initiative, by paving the way for a unified European market for FI facilities.

The creation of a sustainable pan-European open federation of test infrastructures is a key milestone, overcoming the current fragmentation of European infrastructure into isolated test beds that are individually unable to support large-scale trials, and

opening the way for widespread and replicable exploitation of FI services and applications.

XIFI concentrates on the coordination and federation of large test infrastructures using OpenStack and OpenStack-alike technologies. Therefore, it only has a very high-level and abstracted view of the underlying network resources. In SDN environments, it will even have an abstract view of the controllers and provide high-level configuration requests to them. An overview of the XIFI approach for federated resource management was given in Sec. 2.2.4.1. .

2.4.7. FI-WARE

FI–WARE [FIWARE] is a FI–PPP project that aims at providing the Future Internet Core Platform. It presents an ecosystem of services that will boost the economy of business in europe. Among the set of services envisioned by FI–WARE, there is a number of components to facilitate the integration between services and devices, the so called ”Interface to Networks and Devices (I2ND)” Generic Enablers. Among these components, the ”Network Information Control (NetIC)” Generic Enabler is intended to provide abstract access to heterogeneous open networking devices, including OpenFlow appliances. It exposes network status information and it enables a certain level of programmability within the network (depending on the type of network and the applicable control interface). This programmability may also enable network virtualization, i.e., the abstraction of the physical network resources as well as their control by a virtual network provider.

2.4.8. ALIEN

ALIEN [ALIEN] aims at providing a Network Operating System (NOS) running on top of an hybrid, heterogeneous network infrastructure.

Such NOS will be based on the control and management framework of the OFELIA FIRE facility. ALIEN will extend such framework to support devices that are alien to the OpenFlow technology such as optical network elements, legacy layer–2 switches, network processors and programmable hardware.

The technical enabler is a novel hardware description language as well as functional abstraction mechanisms for uniform representation of any type of network hardware that doesn’t support OpenFLow.

However, while ALIEN aims at plugging non–OpenFlow devices into an OpenFlow control framework. On the other hand ALIEN provides an OpenFlow abstraction for networking gear for which an SDK is not available.

2.4.9. OFELIA

OFELIA [OFELIA] is a collaborative project within the European Commission’s FP7 ICT Work Programme.

The project aims at creating a unique experimental facility that allows researchers to not only experiment on a test network but to control and extend the network itself precisely and dynamically. The OFELIA facility is based on OpenFlow, a currently emerging networking technology that allows virtualization and control of the network environment through secure and standardized interfaces.

The OFELIA project aims at providing researchers and practitioners with a virtualized and fully programmable platform on top of which novel protocols and algorithms can be tested.

2.4.10. iJoin

The goal of iJOIN [iJoin] is to design new network operation and management algorithms in the context of Radio-Access-Nework-as-a-Service (RANaaS), assessing their implications on the 3GPP LTE architecture.

The introduction of the RANaaS concept enables opening the RAN/backhaul market for new players, such as vendors and providers of cloud infrastructure platforms. The adoption of RANaaS also provides the technological foundation for shorter and more efficient product development cycles and for a significant reduction of costs for operators. It allows for classical functionalities which are usually processed within a small cell to be partially (or fully) deported in a cloud platform. Therefore, RANaaS benefits not only from the computing power but also from the centralisation gain.

2.4.11. CROWD

The EU CROWD FP7 project [CROWD] proposes a comprehensive architecture based on SDN for cellular and WLAN technologies. The control plane is hierarchically organized into two tiers where "districts" have a local knowledge for short time scale decision and "regions" have global information for more coarse-grained long time scales. In addition, two Northbound APIs type are proposed namely: Technology Specific where optimization is performed on details specific to a given technology (e.g. InterFrame Space for WiFi AP) and Technology Agnostic which enable generic optimization on a broad range of technologies, based on set of abstracted primitives. Finally, reconfigurable backhauling is considered with special emphasis on the opportunity given by SDN for traffic-proportional backhaul configuration through the controller.

2.4.12. EU projects' relevance to enabling technologies

Table 3 below consolidates the aforementioned research projects and maps them to the enabling cloud networking technologies mentioned in Section 2.2, showing which project makes use of which technology.

Table 3. EU projects' relevance to enabling technologies

Project	Infrastructure Virtualisation	Programmable and Software-Defined Networking (SDN)	Network Functions Virtualisation (NFV)	Federated Resource Management and Orchestration
4WARD	✓	-	-	-
SAIL	✓	-	-	✓
ALICANTE	✓	-	-	-
Mobile Cloud Networking	✓	✓	✓	✓
T-NOVA	✓	✓	✓	✓
XIFI	✓	✓	-	✓
FI-WARE	✓	-	-	✓
ALIEN	✓	✓	-	-
OFELIA	✓	✓	-	✓
iJoin	✓	✓	✓	-
CROWD	✓	✓	-	-

2.5. Consolidation of Terrestrial Cloud Networking Architectures

2.5.1. Technical analysis

Cloud networking, as a paradigm, aims to offer dynamic rich end-to-end Network Services on-demand across heterogeneous infrastructures, by virtualising, abstracting and orchestrating network and IT infrastructure resources.

While the benefits of this new paradigm are numerous and are addressed to all actors of the value chain, as highlighted in the previous sections, cloud networking is associated with a number of issues and challenges. A brief overview of these challenges is as follows:

Transport and setup delays – While current network services such as VPNs may take hours or days to be established, cloud networking promises to reduce setup time down to the order of seconds. However, in a complex network service involving numerous virtual links and VNFs, as well as physical appliances, at several physical locations, setup time is still considerable. In addition to service instantiation, the resource mapping calculation, i.e. the determination of the exact embedding of the service to available resources, is usually a demanding procedure and further contributes to setup delay. After service instantiation, transport delays are

usually increased, compared to traditional network services, due to encapsulation/ decapsulation/ tagging procedures at service endpoints as well as flow processing operations (filtering, rewriting, steering etc.) within the network. In NFV services, the traffic may be diverted off its "natural" path in order to be directed to the data centres hosting the VNFs and then injected back to the network, which significantly increases end-to-end transport delay.

Data and signalling overhead – Data overhead is often caused by encapsulation, which is a common technique for logical isolation in network virtualisation and naturally reduces the MTU and consequently the available bandwidth. Apart from that, the use of proprietary or non-standard headers prevent the network devices in the cloud hosts to employ the on-board off-load processing, thus introducing more latency in the packet forwarding. The impact of signalling overhead is far less controllable, since the centralised management of the network service requires significant signalling exchange. This issue is more intense in SDN scenarios, where management and monitoring information down to flow-level granularity needs to be exchanged between each SDN node and the controller.

Networks federation and coupling – A great challenge for associated with the provision of virtualised end-to-end services is how to couple heterogeneous infrastructure domains, as well as how to jointly manage these domains in a unified manner. End-to-end service virtualization requires a framework that handles interactions between such contrasting underlying infrastructures while providing a generic and transparent interface for service providers to easily compose and manage services. The exposure of northbound management interfaces by each infrastructure domain, based on open APIs, is essential to achieve federation. In multi-domain scenarios there are also issues associated with the privacy and independency of each infrastructure owner; high-level orchestrator mechanisms need not to be too "intrusive".

Isolation of services – While logical isolation at data plane, at least at L2/L3 is achieved via tagging and/or encapsulation, control plane isolation still seems challenging, especially in the case of multiple network applications controlling the same substrate. In the case of virtual network appliances, VNFs which are co-located on the same physical node, although isolated, might slightly affect one another ("noisy neighbour" effect).

Reconfigurability and programmability – One of the benefits of programmable networking is the ability to offer to the customer the ability to arbitrarily program the control logic of his/her network service. However, the elevation of this programmability paradigm in "production" environments i.e. ISP networks, poses significant privacy and stability issues yet unresolved. Moreover, the offered programming interfaces, mostly in the form of an SDN controller, are quite low-level for the average user; higher-level abstractions are required in order to offer the benefit of programmability to a wider customer group.

Resource mapping – Pooling virtualised resources in order to deploy a service with specific requirements can often be formulated as an optimisation problem, which becomes significantly complex in large infrastructure topologies, especially when both IT and network resources need to be combined, e.g. in the case of NFV.

Scalability – While cloud networking frameworks are being tested with success on small scale networks (e.g. in lab or campus environments), the scaling up in commercial ISP networks, possibly involving multi-domain services with hundreds or thousands of users is not considered feasible in short term, especially due to the excessive signalling overheads involved as well as SDN limitations in flow and routing tables.

Mobility – The mobility of end users from one physical location to another requires the reconfiguration of the virtual topology as well as the live migration of virtual network appliances, which can induce significant overhead, especially in highly dynamic environments.

Resource elasticity – While today's virtualised network services such as VPNs have statically assigned resources, resource dynamicity is considered a key added-value of cloud networking. However, an effective and efficient service up/down and in/out scaling mechanism needs to achieve several goals; i) the scaling response time must be satisfactory (i.e. in the order of minutes or even seconds) in order to cope with the change in the resource demand, ii) the stability of the overall platform needs to be secured even after abrupt changes, iii) other tenant services must not affected and iv) the new allocation of resources must still be close to the optimal i.e. avoiding overprovision.

Security and trust – Although virtualisation in principle achieves the isolation among tenant services thus protecting tenants, the physical infrastructure can be still vulnerable to attacks. Moreover, the network programmability paradigm poses a number of security issues associated with the authorisation to deploy and execute arbitrary network applications, since a faulty or a malicious application may severely impact the stability of the services and under certain circumstances also affect other tenants.

Accounting, billing and SLAs - In traditional network economics, bandwidth is the primary billable item. But in cloud networking, the usage of virtualised resources, especially virtual network functions, is important as well. Furthermore, resource dynamicity and up/down scaling need to be supported by the accounting and billing procedures. In addition, new SLA templates need to be proposed, taking into account not only bandwidth and QoS, but also highly dynamic usage of in-network resources as well as scaling policies and permissions.

Resilience and availability – The "softwarisation" of infrastructure involved in cloud networking poses several challenges when it comes to availability, since software-based network services and appliances are generally more vulnerable to faults and outages than traditional hardware-based networking, which (especially in the case of fixed networks) is generally considered a service with high availability. While a periodic interruption of a software application is generally considered natural and acceptable, this is not the case with a network service, where outages are much less tolerated.

Performance – Tunnelling/encapsulation and flow processing mechanisms involved in cloud networking have often an impact in network performance, decreasing available bandwidth and increasing end-to-end delay. The performance of virtual network functions is also a significant challenge, since, in order to secure

commercial viability, software appliances should achieve performance comparable to their hardware counterparts with a reasonable allocation of resources.

Standards and technologies – Although in some domains e.g. in network virtualisation there are already several well-established standards available (e.g. MPLS, VXLAN, NVGRE etc.), in other domains e.g. in NFV or in federated resource management there is a complete lack of standards, apart from some high-level recommendations (not normative) issued by bodies such as ITU, IETF and ETSI. This gap is partially bridged by several open-source initiatives, such as Openstack or OpenDaylight, which attempt to establish de facto standards and interfaces, bypassing "official" standardisation bodies.

2.5.2. Capabilities mapping

In this section, the four enabling technological domains of Cloud Networking, as surveyed in Chapter 2.2, are overviewed in terms of their relevance to the aforementioned challenges. In specific, the tables in the following subsections identify the trends, advances and techniques in each technological domain, which try to address the issues associated with Cloud Networking. Only the issues which are relevant with each domain and are adequately addressed are included in each corresponding section.

2.5.2.1. Infrastructure Virtualisation

Challenge	How it is addressed
Transport and setup delays	Tailored resource mapping algorithms for virtual networks (see below on "Resource mapping") significantly reduce setup time.
Data and signalling overhead	Emerging tunneling protocols such as STT (Stateless Transport Tunneling) use fragmentation techniques to minimize data overhead per packet.
Networks federation and coupling	Tunelling mechanisms such as VXLAN, NVGRE or STT facilitate infrastructure federation by interconnecting remote data centres. Widespread APIs for the management of virtualized infrastructures, such as the Amazon EC2 API, also adopted by Openstack, greatly promote unified management.
Isolation of services	Techniques for traffic marking e.g. VLAN provide satisfactory logical isolation for network traffic within network elements.
Reconfigurability and programmability	IT Virtualisation infrastructures such as Openstack expose a northbound API allowing (partial) programmability i.e. control by high-level applications.
Resource mapping	Although the optimal virtual network embedding onto a physical infrastructure is considered an NP-hard optimization problem, there exist in the literature several algorithms which achieve quite

	satisfactory approximations.
Scalability	Modern IT virtualization infrastructures support various segregation methods (cells, regions, availability zones, aggregates etc.) in order to improve scalability.
Mobility	Modern IT virtualization infrastructures support VM migration to a different physical host with minimal downtime (even "live" migration under certain circumstances)
Resource elasticity	Modern IT virtualization infrastructures support VM resource up/down scaling ("resizing") with minimal downtime
Security and trust	In L2 virtual networks such as VLANs, network traffic is segregated within network elements and thus well isolated. In L3 traffic, encryption can be added on top of encapsulation for increased security. In virtualized IT assets, already established security/trust mechanisms can be leveraged for access control, such as certificates etc.
Accounting, billing and SLAs	Already established billing and accounting mechanisms for IaaS cloud platforms are being extended to better support elasticity and dynamic usage models.
Resilience and availability	Already established techniques for cloud resilience can be exploited, such as VM migration and dynamic network reconfiguration in case of physical failure. As for virtual networks, several reconfiguration mechanisms have been proposed in order to rapidly react to dropped links.
Performance	Virtual network performance can be improved by adopting a lightweight virtualization scheme and allocating more processing resources to tunnel endpoints. VM performance in cloud environments can improve via exploiting virtualization-capable hardware (especially CPUs, GPUs and network interfaces)
Standards and technologies	Network virtualization techniques are adequately covered by standards. In IT virtualization, the lack of standards is partially covered by industry initiatives and open technologies which are widely adopted.

2.5.2.2. Programmable and Software-Defined Networking

Challenge	How it is addressed
Transport and setup delays	SDN achieves very low delays for rule communication and enforcement, significantly reducing setup time. SDN does not add transport delays per se.
Data and signalling overhead	SDN can achieve network virtualization/partitioning without encapsulation and thus with no data overhead, thanks to per-flow manipulation.
Networks federation	The adoption of a single control protocol by multiple vendors significantly facilitates federation of heterogeneous

and coupling	infrastructures.
Isolation of services	Although the traffic may not be logically isolated, isolation of control plane can be achieved via specialized middleware enabling multi-tenancy in SDN networks (such e.g. as OpenVirteX/VTN).
Reconfigurability and programmability	SDN by nature allows the programmability of the control plane logic by arbitrary user applications.
Scalability	Scalability in SDN can be partially addressed by distributing the control plane functionality by means of distributed controllers and/or partitioning the network into segments and organizing the controllers in a hierarchical structure.
Mobility	Via modification of per-flow rules, user traffic can be redirected to the new user location (traffic steering) with very low response times.
Security and trust	Multi-tenancy enabler middleware platforms such e.g. as OpenVirteX can logically separate (up to a certain degree) the control applications.
Accounting, billing and SLAs	The Openflow protocol already provides a rich set of monitoring metrics, including per-flow statistics, which can be exploited for fine-grained accounting and billing.
Resilience and availability	In case of link or node failures, SDN per-flow control capability can switch traffic to alternate routes, achieving very low downtime.
Performance	SDN performance is mostly affected by controller capacity and response time. In order to improve controller performance, distributed architectures are proposed, where the controller load is split in multiple instances working simultaneously.
Standards and technologies	For the controller-switch communication, Openflow is already a well established standard. With regard to higher-level programmability frameworks, the absence of standards has led to a multitude of technologies.

2.5.2.3. Network Functions Virtualisation

Challenge	How it is addressed
Transport and setup delays	Transport delays are reduced by minimizing the traffic steering/diversion path. This can be achieved by deploying multiple NFVI-PoPs close to the customer access points (the “Edge Cloud” concept). Local caching of VNF images at the NFVI-PoP significantly accelerates image deployment and thus decreases setup delay.
Networks federation and coupling	VNFs are assumed to expose a uniform management interface regardless of the underlying physical infrastructure on which they are hosted, which significantly facilitates federated management.
Isolation of services	VNF isolation can be increased by following best practices such as

	i) selection of virtualization-enabled hardware components when building the infrastructure and ii) strategic placement of VNFs, e.g. grouping VNFs of the same tenant into the same virtual machine.
Reconfigurability and programmability	NFV directly addresses the network programmability challenge by allowing arbitrary virtual appliances/traffic processors to be inserted into the network. Moreover, emerging NFV orchestrator platforms expose a northbound API which allows (partial) control by high-layer applications.
Resource mapping	Resource mapping in NFV is optimized i) by grouping the compute resources in clusters/NFVI-PoPs, enabling two stage mapping (first select the PoP and then assign the resources within the PoP) and ii) by exploiting features such as Enhanced Platform Awareness (EPA) in order to assign specialized VNFs requiring specific hardware accelerations to the physical hosts supporting this features.
Mobility	By exploiting cloud mechanisms for VM mobility, VNFs can follow the end user and be re-deployed close to the user location.
Resource elasticity	Cloud techniques such as VM resizing can be exploited to achieve VNF resource elasticity.
Security and trust	Emerging NFV security frameworks mandate VNF code to be digitally signed and verified by a trusted entity prior to deployment.
Accounting, billing and SLAs	NFV takes advantage on already established mechanisms for billing of pay-per-use cloud IaaS services. NFV-oriented SLAs are currently under development.
Resilience and availability	Emerging NFV frameworks foresee a validation / qualification procedure for VNFs prior to deployment, in order to verify that the VNF application is as stable as required. Moreover, NFV leverages techniques for cloud resilience, such as VM migration in case of physical failure.
Performance	Performance of VNFs can be improved by exploiting hardware acceleration techniques for several operations, mainly for traffic capture and processing. Additionally the use of packet handling acceleration frameworks in conjunction with specific HW capabilities (i.e Intel DPDK (Data Packet Development Kit) with SR-IOV (Single-Root I/O Virtualisation)), might provide enhanced performance at VM level.
Standards and technologies	While formal standardization processes have not been initiated yet (apart from the recommendation guidelines issued e.g. by ETSI NFV ISG), emerging open platforms such as OPNFV (Open Platform for NFV) have the potential to evolve to de facto standards.

2.5.2.4. Federated Resource Management and Orchestration

Challenge	How it is addressed
Transport and setup delays	Centralised management generally achieves better transport delays due to the capability to select optimal paths across multiple domains.
Data and signalling overhead	A hierarchically organized federated management scheme based on vertical rather than horizontal communication, generally minimizes signaling overhead and optimizes management, at the expense of sacrificing the independence of infrastructure domains
Networks federation and coupling	Emerging high-level orchestrators operating on top of infrastructure managers, provide a promising solution for feasibly federated management, while preserving the independence of the infrastructure domains.
Reconfigurability and programmability	Emerging federated management platforms are exposing an API which allows control by high-level applications, enabling the high-level control of the service to be undertaken by arbitrary software modules.
Resource mapping	Resource mapping in a federated environment is quite challenging. It is partially addressed by performing it in a two-tier procedure, first mapping the service to clusters of resources (e.g. domains or distributed data centres) and then performing a more fine-grained mapping within the selected clusters.
Scalability	Scalability in management procedures is achieved by splitting infrastructure management in two or three tiers, assigning only high-level decisions to federated orchestrators and offloading fine-grained procedures to local managers of each infrastructure segment.
Mobility	A federated management approach better addresses user mobility due to the ability to dynamically re-allocate and possibly migrate service resources across different segments of the infrastructure.
Accounting, billing and SLAs	Federated management can be further augmented via multi-actor Marketplaces, facilitating
Resilience and availability	Federated management generally achieves higher resiliency, since failover choices are more, across different infrastructure segments.

2.5.3. Technology Readiness

Technology readiness is a very important factor to be considered along with the functional capabilities of a specific technology. Especially when considering the integration of specific technologies into satellite infrastructures (and in particular satellite payloads), the maturity of the technology is crucial.

In this context, the tables below aim to briefly overview some of the most prominent of the technologies overviewed in this document in terms of maturity and readiness. The ESA TRL categorisation is used.

2.5.3.1. Infrastructure Virtualisation

Technology	Hypervisor technology
TRL	TRL 9 (Actual system proven through successful operations)
Comments	Commercial hypervisors, such as XenServer, VMWare ESX/ESXi and Hyper-V are well-established in the market, most of them powering production infrastructures for more than a decade.

Technology	OpenStack
TRL	TRL 9 (Actual system proven through successful operations)
Comments	As a cloud controller platform, Openstack is considered mature (although still rapidly evolving). Several commercial cloud services (either public, private or hybrid) running on Openstack already exist. Furthermore, companies such as HP or RedHat offer Openstack as part of a complete cloud package platform (RHEL Openstack, HP Helion respectively)

2.5.3.2. Programmable and Software-Defined Networking

Technology	Openflow
TRL	TRL 6 (System Demonstration in a Relevant Environment)
Comments	Openflow technology is considered quite mature (although still evolving). Openflow-enabled switches exist in the market for the last years by vendors such as HP, Pica, NEC, Brocade, Big Switch etc. However, the protocol has not yet been deployed in wide-area “production” networks, where traditional control mechanisms are still being used.

Technology	OpenDaylight
TRL	TRL 6 (System Demonstration in a Relevant Environment)
Comments	OpenDaylight has been released as a second stable version (Helium) and is extensively used in many lab and experimentation testbeds as well as proof-of-concept demos. Also, it has been used in some derivative commercial products, such as Brocade Vyatta controller. However, OpenDaylight it is still not known to have been validated in large-scale production environments.

Technology	OpenVirteX
TRL	TRL 6 (System Demonstration in a Relevant Environment)
Comments	OpenVirteX started as an experimental platform, yet it is now considered quite stable and has already been integrated with Openstack. However, it is lacking the industry support of OpenDaylight and there are no known commercial deployments so far.

Technology	OpenNaaS
TRL	TRL 6 (System Demonstration in a Relevant Environment)
Comments	OpenNaaS, although being developed for several years, has not yet attractive a critical mass of experimenters and adopters, thus it has not yet undergone significant validation in diverse environments.

2.5.3.3. Network Functions Virtualisation

Technology	OPNFV
TRL	TRL 2 (Technology Concept and Application Formulated)
Comments	OPNFV is currently in its first stage, and the core projects are still under development, however the basic enabling technologies have been identified. The first release is to be expected by the third quarter of 2015.

Technology	Proprietary NFV Orchestrators
TRL	TRL 7 (System prototype demonstration)
Comments	Already several commercial NFV orchestrators exist (HP NFV Director, Nokia Cloud Network Director, Overture Ensemble Service Orchestrator, Cyan Planet Orchestrate etc.) they have very diverse capabilities, and they are not known to have been validated in large-scale production environments.

Technology	ETSI-compliant NFV Orchestrators
TRL	TRL 2 (Technology Concept and Application Formulated)
Comments	For ETSI-compliant NFV Orchestrators, the TRL level should be much lower than proprietary ones, since not only the platforms themselves are under development, but also the specifications which should govern them have not been finalised yet (ETSI has not produced strictly technical normative documents so far)

Technology	Mobile Edge Computing

TRL	TRL 1 (Basic principles observed and reported)
Comments	While the MEC scope has been well defined, the ETSI group is currently working on a first set of documents, which will define the basic concepts, the terminology and the architectural framework.

2.5.3.4. Federated Resource Management and Orchestration

Technology	XIFI platform
TRL	TRL 6 (System Prototype Demonstration in a Relevant Environment)
Comments	The XIFI management platform is currently operating with considerable stability, yet it is confined in an experimental environment. The technology behind it has not been tested in a production environment.

Technology	GENI/ExoGENI platform
TRL	TRL 6 (System Demonstration in a Relevant Environment)
Comments	GENI has been active for several years, successfully supporting a large number of experimenters. Nevertheless, its technology has not been tested in a production environment in order to determine whether it meets carrier-grade requirements.

3. DIMENSIONS OF SUITABILITY FOR INTEGRATION WITH SATELLITE NETWORKS

The purpose of this chapter is to review in depth, from a technological-push approach (bottom-up), the different levels – or "dimensions" – for which the identified Cloud Networking technologies would bring added-value to satellite communications. The most promising technologies from the domains of Infrastructure Virtualisation, Programmable and Software Defined Networking, Network Function Virtualisation, Federated Management and Orchestration, are identified for satcom.

3.1. Review and Criticality evaluation of dimensions relevant to satcom

In this section we identify the dimensions of suitability of integrating Cloud Network technologies and define the criticality level for each dimension.

Dimensions are categorised as:

"Functional" dimensions refer to functionalities added or enhanced/facilitated via Cloud Networking.

"Integration" dimensions refer to the issues and aspects associated with the integration to satellite platforms, including the impact to product lifecycle and the TRL.

"Business" dimensions refer to stakeholders' interests, market aspects and high-level costs.

With regard to criticality, the following categories are used:

"Low" indicates that Cloud Networking seems applicable but complex and without any evidence of new feature or important cost saving, taking into account that non-Cloud networking technologies are also progressing and achieve satisfactory results.

"Moderate" refers to dimensions where Cloud Networking may indeed be of help to ease the support of existing service or functionality and shall probably contribute to cost saving. Typically, this can include non-negligible CAPEX that can be amortized over OPEX after a period of time.

"High" corresponds to dimensions where Cloud Networking addresses important expectations in terms of additional value, applicability for concrete and existing use cases where demands is known or can be anticipated. Even in the cases where no new real services/features are introduced, it is recognized that software/virtualisation can bring important cost reduction (via e.g. economies of scale).

An important note is that the identified "dimensions" are often inter-related with each other. Hence, the analysis may sometimes introduce partial overlaps among dimensions.

3.1.1. Cloud Networking Functional dimensions for satcoms

3.1.1.1. Networks federation and coupling

In spite of different interests expressed by various stakeholders, (standardization bodies, manufacturers, project consortiums, satcom operators, space agencies, etc…), the effectiveness of integration, federation or coupling between satcom and terrestrial systems remains rare up to now. The technical complexity involved in each domain, the rapid evolution of terrestrial systems, the certification and standardization models to define, and the difficulty to find win-win business models, constitute likely the main explanation for this situation. However, it shall be recognized that satcoms have borrowed much from the terrestrial domain.

The abstraction of telecom and network technologies shall allow two distinct classes of federation models:

- federation/coupling[4] between satellite and terrestrial systems (such as: fixed and/or mobile networks). This is in line with the "5G" vision, which encompasses federation of heterogeneous access networks in a transparent manner.
- federation/coupling between heterogeneous satellite systems:
    - Fixed Satellite Services (FSS) +Mobile Satellite Services (MSS) systems
    - Fixed Satellite Services (FSS) + Broadcast Satellite Services (BSS) systems
    - MSS+BSS systems
    - Even two FSS systems (e.g. operating in distinct bands or with non-interoperable technologies or implementations)

In addition, the implementation of such federation could apply to different components:

- Satellite Gateway and other centralized components (i.e. resource managers if externalized, routers, management systems, network appliances such as firewall, etc…
- And/or User Terminals: technology abstraction could even be implemented at user terminal only
- Even at satellite level, at long term, a satellite OBP technology could even host a SDN-compatible switch

In terms of impacted features, federation and coupling may address the following topics, for which significant differences can occur between systems (and even between equivalent systems operated by different providers):

[4] Federation and coupling can refer to different levels of integration or interworking (there might not be a common understanding on these terms). Generally speaking, federation shall imply more integrated systems (more complex interfaces) that in "coupling" model.

- Forwarding
- Addressing, Routing and switching
- Resource allocation (Layer 2)
- SLA and QoS implementation (Layer 3 and above)
- Network Management

Finally, it is worth to note that Network federation and coupling will require "joint decisions". This means that whatever the Cloud Networking technologies to be used, some entities have to implement decisions or policies that take into account information available from all involved systems. Such decision modules will represent the intelligence of the federation and coupling model that controls the federated system.

The decision functions will typically be defined and implemented i) by the provider of one of the involved systems (and that shall reflect the implementation of agreements with other providers) ii) by a third party, acting as federated manager or iii) less possibly, by the Customers themselves (case of "transparent coupling").

The expected development of network federation and coupling in satcom seems very promising in the quest of providing additional resources, features and services to satcom customers. It is important to notice here that this model would directly benefit to users (through the perceived Quality of Experience, reduction of cost to access services, etc…) and shall not be considered as a "simple" IT or network tool for network and system providers.

Therefore the criticality level of this dimension is assessed as High.

3.1.1.2. Isolation of services

Providing service and/or providers isolation in large-scale system environments such as satcom are of prime importance in situations where services are provided via Virtual providers. The requirements are:

- To guarantee a clean resource and traffic separation between providers (case of global provider reselling (part of) the system capacity to other entities (e.g. virtual network providers)),
- To cope with different and possibly heterogeneous addressing and configuration plans,
- To secure the differentiation of processing for flows belonging to very different domains (e.g. use of common platform components for very heterogeneous services with different types of constraints and purposes),
- To interconnect with multiple/inconsistent management systems while preserving isolation as needed,
- To transport data frames from different parties (e.g. aggregated services from operators) with the possibility to enforce customized policies (e.g. related to QoS & accounting associated to the user SLAs. For example, providers could define different types/number of QoS classes, billing based either on volume, rate, or combination of both, etc…).

However, Virtualisation for isolating services has already been introduced in satcom in different ways, as presented below. These approaches utilize either available in-band signalling in protocol headers, or headers from an additional tunnelling layer.

Support of VLAN and extensions in satcom

Before the advent of Cloud Networking technologies (as described in Chapter 2) several forms of Virtualisation in satellite networks have been supported, at different levels. Virtual Local Area Networks (VLANs) (IEEE 802.1Q) defined in the context of LAN services, has been widely adopted in terrestrial networks. Actually virtual 802.1x technology has been extended later (such as in Carrier-grade network, with Provider Bridge (IEEE 802.1ad) [802.1ad] and Provider Backbone Bridge (IEEE 802.1ah) [802.1ah] and/or supported in satcom standards (DVB-S2/RCS(2) networks). So far, satcom systems remained more or less transparent to this technology, except on the user plane domain (with very marginal impact at encapsulation level to identify the type of transported user frame).

Service isolation at DVB layers

More recently, the DVB-S2 within DVB-RCS2 standard [DVBRCS2] introduced its own form of Virtualisation isolation, decoupling the roles of Satellite Network Operator (SNO) and Satellite Virtual Network Operator (SVNO). Their relationships between the roles of the value chain are illustrated in Figure 23 below.

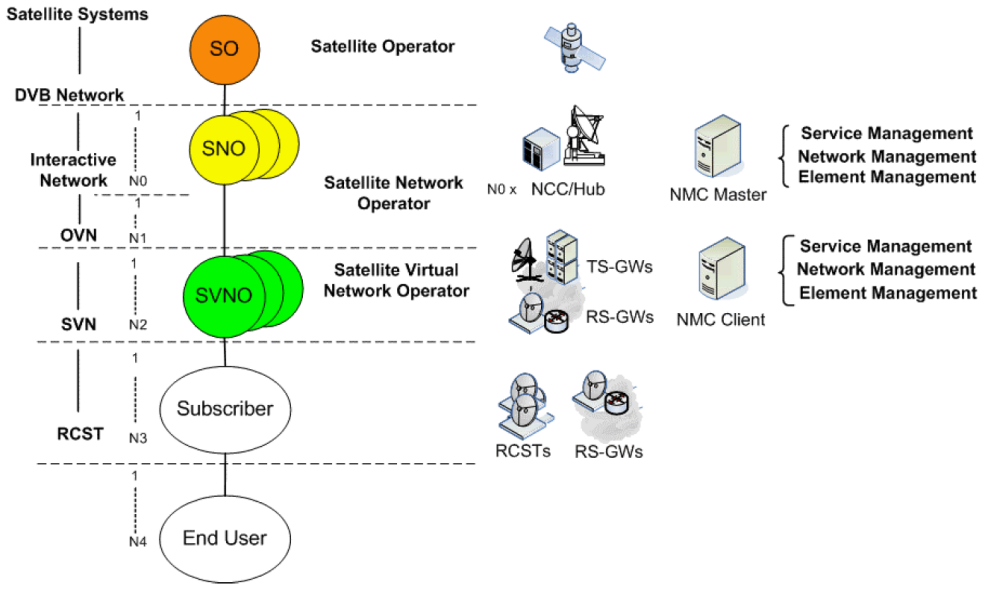

Figure 23. Satellite Virtual Network in DVB-RCS2 [DVBRCS2]

In this model, several Satellite Network Operators share the overall physical frequency/time resource on their own, operate their system with their GWs, and provide the management of the network services. Satellite Virtual Network Operators (SVNOs) are possible resellers of the capacity at the network/service level. SVNOs (when present) are the network providers with which customers subscribe. From a more technical perspective, SNO/SVNOs are identified according to specific labels inserted in the protocol stack at the DVB-RCS2 sublayers (network_ID and SVN-MAC).

Other tunnelling protocols

In the frame of general networking, other tunnelling protocols are already available for several years (e.g. Generic Routing Encapsulation [GRE]). And recently, falling into the Virtualisation networking technologies, solutions like VXLAN, already discussed in Chapter 2, are proposed by network vendors.

Note that the applicability of those tunnelling protocols is not only limited to service isolation. Indeed, those solutions also support 'connectivity extension' meaning that any L2/L2.5 island or domain (e.g. MPLS, etc…) could be interconnected over e.g. a satellite segment supporting only IP (or Ethernet) connectivity service. However, the price to pay in the case non-native support is related to performance degradation/capacity reduction (overhead), compared to native transport.

Expectations from Cloud Networking

As explained in the previous paragraphs, Service Isolation is clearly not a new feature in satcom – even if actual implementation in systems is much probably limited. However, Cloud Networking could bring good opportunities with respect to some issues:

- Scalability: fixed header fields used to discriminate services are not always adapted to provide service/provider isolation. Note that the issue of overhead shall not be forgotten here – an important requirement of the satcom community is still to keep overhead as low as possible; for example by avoiding tunnelling and using stateful header rewriting (address translation) instead.
- Flexible and smooth Service reconfiguration: change of service for a given flow or group or flows shall be supported dynamically (ultimately: on the fly). A reconfiguration only supported and managed via software means shall favour the development of such feature. This feature also involves the ability to deploy easily any Tunnelling technology in the system to offer any type of connectivity (possibly on-demand) – but with the limits mentioned above. More functionalities can therefore be provided but with performance reduction. However, this appears as a suitable solution if used for a low number of users/service in the system (e.g. High value added services) and could need overprovisioning of resources.

We finally assess the criticality level of this dimension as Low.

3.1.1.3. Reconfigurability and programmability

Strong expectations on reconfigurability and programmability[5] are expressed by operators in order to adapt the (satcom) system configurations to actual conditions of operations. Reconfigurability and programmability are two essential characteristics in SDN technologies, including OpenFlow, but could be also supported with NFV in the

[5] Network programmability refers to the ability to migrate the control logic of network appliances to software entities i.e. have software programs to arbitrarily define the switching/routing/forwarding policies rather than rely on "hard-wired" protocols.

way configuration changes/upgraded are easy (no need to change Harwarde and interrupt services, possibility to test alternative configurations, etc…)

Short-term reconfigurability/programmability may address diverse parameters, such as number of active users in the system, their repartition in the different spot beams, interference levels (power control), type/quantity of resource allocated, and of course all the network configuration.

Longer-term reconfigurability/programmability capabilities are more in-line with system planning activities. Different aspects can be highlighted:

1. The lifetime of a satcom system (e.g. 10-15 years) necessarily introduces assumptions or models that include some uncertainties, that may require upgrades and/or changes after some years. In broadband access systems with High Throughput Satellite (HTS) and several GWs, the ground segment can be more affected by such changes, sometimes with a joint impact on the space segment.
2. New opportunities related to external technological improvements would also arise. Examples can include the evolution or change of a protocol (new versions or new protocols adopted) that may add functionalities to the satcom components. Also, the satellite backbone connecting the GW can be heavily affected with the evolution of backbone transport technologies. For example, around 2000, ATM was a very widespread standard in transport networking-now progressively replaced by other technologies (MPLS, IP/MPLS, MPLS-TE and other MPLS extensions, Carrier-grade Ethernet). Such evolutions may require major replacements and restructuring in the network infrastructure.
3. Due to novel reconfigurability capabilities, regulatory and commercial frameworks may also be affected – either relaxed or constrained further. For example, operated frequency bands represent a usual and permanent concern for satcom; an impact on spectrum usage regulations might be crucial. In addition, affected regulations could also address the network level (in the domain of security and encryption, lawful interception, right or not to distribute or modify Contents, etc…)

Additionally, a remarkable quality of Cloud Networking is their ability to smoothly support successive life cycle phases of products and systems. This may also include research, experimentation, and temporary test and validation phases that aim to mitigate risks for upgrading systems. This finally contributes to reduce the Time to Market of initial deployments, and of further upgrades with possibly significant impact on costs. For example at some point a temporary/alternative configuration (including change of some protocols, such as new tunnelling/connectivity service) might be tested in a deployed system without risk. Temporary deployment for the entire system, or with restriction to some services or users, are both possible.

We assess the criticality level of this dimension as High.

3.1.1.4. Mobility

In the satcom context, mobility may refer to different scenarios, considering the various categories of satcom systems.

Mobility in MSS systems is by nature fully supported from the beginning of their design. At lower layers and at system level, hardware (antennas, payload, platforms, on board and ground processing…) and procedures (e.g. handovers) must be extensively specified, tested and optimized. On the other hand, in case new mobility algorithms, policies, criteria, etc. are found and developed, it can be very difficult to introduce them after the system deployment. Also, different mobility solutions per user terminal could hardly be supported by the satcom GW/NCC(s) that cannot easily run distinct algorithms or use distinct interfaces.

For other types of systems, including hybrid systems, and extension of FSS system for Mobile terminals, mobility can be supported at different levels. Compared to MSS, no guarantee at the networking and service level is given. Mobility shall remain transparent in (hybrid) FSS system (e.g. DVB-RCS+M). This is possible at the physical layer (with signal processing techniques) but not completely at higher layer (issue of resource re-allocation, QoS, and in the maintenance of connectivity and reachability) if a spotbeam/GW/satellite handover happens. Inter-system "handover" (or reception of service) is also a possible event. So far, the implication of the space segment was practically excluded, motivated by the low occurrence of such events (for example in system where spotbeams are in the order of few hundreds of kilometers). Finally, for purely network-level mobility solutions (such as Mobile IP, Multiple Path TCP etc…) used over the satcom network, relevant services indeed exist, but still with performance issues and limited network stack compatibility. Relying on those solutions might be an issue for providers that cannot always guarantee end systems will support those protocols. Transparent Proxy-based solution (e.g. Proxy Mobile IP, etc…) can alternatively be implemented but imply additional costs/complexity to be supported (just like TCP PEP).

Considering that this last type of systems shall be generalized in the future (e.g. due to the advent of 5G federated architectures), software/Cloud Networking-based solutions would be ideal to facilitate handover management, by implementing open policies, possibly distinct per terminal, group of terminal, provider, taking into account the different connectivity options and associated service parameters (cost, resource, latency, QoS etc). In this context, Cloud Networking is therefore not expected really to introduce a new feature but could be helpful in the way mobile support is facilitated. For example, for proxy-mobility solution, any NFV-compatible implementation would be ideal to support multiple configuration and/or smooth configuration transition.

We assess a criticality level of this dimension as Low.

3.1.1.5. Resources elasticity

At system perspective, resource allocation is an important topic for satcom networking. It is of prime importance to optimize resources allocation at system level,

so as to accept the largest number of users in the system, without downgrading the individual performance beyond an acceptance threshold.

It shall be recognized, however, that limitations do exist for such scalability. Available hardware (at terminal side), often related to link budget issues or processing capacity) sometimes sets the hard limit; further improvement of performance is possible only at the cost of hardware upgrade, with variable impact at system level (might be fully transparent, or not). But even more powerful hardware may sometimes not be sufficient in case of system constraints. For example, a given terminal is allowed to access a limited number of carriers in the system according to its position, even if it could support to access other carriers from adjacent spot beams.

At customer level, Subscriber Level Agreement (SLA) constraints also limit the ability to acquire new resources. Other methods, such as Dynamic QoS [DynQoS] can sometimes be a partial answer to add instantaneously resources (CBR for instance) – but it is not suitable for all services and requires adequate billing models in order to support it. Flat rates (sometimes with restrictions such as Fair Access Policy (FAP) that introduces Volume thresholds after which service limitations apply) are the common case in broadband commercial systems. On the other hand, fully "On Demand" charging remains very limited (e.g. access to specific services, such as VoD, and in that case Digital Right Management fees may apply in addition). However, some interesting cases could still apply in this area; for example, the ability to scale down the resource for customer that consume low capacity and that could desire to decrease costs as much as possible. The scaling response time should be here at the level of a minute at most.

At system level, resource elasticity can affect the amount of resources leased by the capacity reseller or wholesale provider) to e.g. enterprise customers or Virtual operators for specific needs. This elasticity may refer to the flexible allocation of resources to virtualized services. Such operations usually fall in the domain of the network management for which more or less satisfactory solutions may exist, but can be highly dependent on the initial system design. Other domains such as Institutional, Governmental, Emergency and Military services are known to be demanding of the so-called "Flexibility". Increasing demand has been observed during the last years for the development of flexible payload and systems. For such systems, automating the network management and operations, increasing the number of possible settings, and at the same time reducing the cost of re-configuring the system, is a very attractive perspective. In this case, the scaling response time might not be so important (few hours should generally be acceptable).

The criticality level of this dimension is assessed as High.

3.1.1.6. Availability and resiliency

Availability and resiliency are crucial in satcom services and any introduction of new technologies (such as Virtualisation and Cloud Networking) should not add – and shall even reduce – any risk of increasing faults or outages (occurrence). At the same time each occurrence should have equal or less effect on the system performance.

Intrinsically, software-based networks can appear less reliable than topologies based on traditional hardware appliances. However, replacement of faulty hardware is also complex and acquisition of new hardware might be costly. Continuity of service and operations can be also complex with pure hardware solutions. Hence, for a given level of targeted availability and resiliency, software might be cheaper to acquire and operate, and present sounding advantages in terms of scalability.

Without arguing at this stage that Cloud Networking can enhance or not availability and resiliency, the criticality level of this dimension is assessed as Low.

3.1.1.7. Security and privacy

In the domain of security and privacy, Cloud networking introduces at the same time threats as well as opportunities. The threats refer to security issues related to faults (e.g. software bugs) or authorization issues (e.g. remote configuration). On the other hand, the development of new security methods (e.g. security protocol) and rapid reconfiguration shall be globally easier via pure software means and therefore appears as an interesting opportunity. In addition, software-based security can be seen as an additional level of protection, provided that lower-layer solutions are not removed (e.g. link-level encryption, etc.)

The criticality level of this dimension is assessed as Moderate.

3.1.1.8. Accounting, billing and SLAs

Billing, accounting and SLAs are closely related to the impact to the revenues model of providers when resources are re-allocated (at least for the commercial sector). “On-demand” billing models shall clearly be the norm in future telecommunications systems.

As a consequence, the adoption of Virtualisation and Cloud Networking could require important changes of the Management planes. It cannot be definitely concluded whether such models can be economically feasible for satcom, where flat-fee billing is generally applied. In any case, current billing and accounting tools will need to be adapted (level of adaptation can vary) but this adaptation should come at a minimum cost. Efficient billing and usage tracking of elastic resources still remains a big challenge.

The criticality level of this dimension is assessed as High.

3.1.1.9. Performance

It has been highlighted that performance in networking environments could be affected by Cloud Networking technologies. Satcom are systems were resources are generally much constrained, addressing coverage areas where possibly no other connection alternatives can be found. Performance degradation (such as latency increase or bandwidth decrease) could negatively affect the perceived quality of service. However, virtualization technologies are not considered to have significant

impact in satcom performance. Also, in some rare occasions (e.g. during/after reconfiguration phase), transient and punctual reduction of performances could be accepted with very moderate impact for service.

Therefore the criticality level of this dimension is assessed as Low, although it is vital that introducing Cloud Networking shall not (or only minimally) affect network performance.

3.1.1.10. Satellite specific capabilities

At long-term horizon, future satellite systems will gain in flexibility, with some reconfiguration capabilities already identified for example for the payload (e.g. flexible power distribution for different RF input/outputs and spotbeams). Optimization will be possible in reducing more and more the applicable configuration time – up to smooth and fully automatic processes (in the goal of reducing the Telemetry commands to only cope with very specific and unpredictable changes). For example a large-coverage satellite payload could adapt power/frequency resources according to the actual traffic – compensating (maybe partly) traffic imbalance between the spots (e.g. different time zones covered). In that case, "software and virtualized Payload" concepts should be of help. This could be seen as an extension of the satellite OS and processor virtualization (see below).

At the same time, new technologies related to on-board switching/routing shall also be developed – at least provided that systems with multiple satellites interconnected with ISLs use this capability (this is still uncertain for GEO-based platforms). For the reasons cited previously, internal routing/switching in the satellite segment may also benefit from these Cloud networking Technologies.

Finally, other satellite-specific capabilities with impact on networking, such as large unicast multicast/broadcasting; relatively high degree of independence from terrestrial infrastructures, ability to reach remote and isolated areas, etc. are considered more or less neutral with respect to Cloud Networking.

Hypervisors for satellite payloads

Virtualisation concepts have only recently (2009-2010) started to be investigated for adapting the implementation Virtual Machine (also called "Partition") and the so-called Hypervisors (in charge of managing/interfacing the VM with Hardware) for satellite payloads that runs with specific processors and OSs. Hypervisor allows for virtual processes and payloads to run on one or more processors simultaneously. In the space environment, the advantage of a hypervisor is that if one of the processes "crashes", experiences SEUs, or experiences malicious code attacks, it will not impact the performance and operation of the other processes. The degree upon which Spatial isolation is possible may therefore be different from one solution to another one and therefore may constitute an important factor of merit or constraint. According to their design, Hypervisors also imply different characteristics in terms of access performances to different hardware resource types (CPU, memory, timers, bus). Finally, different level of affinity/heritage to generic OSs (e.g. Linux, RTLinux)

may be found in each solution, seen as advantages (e.g. good maturity levels) but also sometimes with side flaws (such as: patents or licensed technologies) .

Consequently, different approaches have been proposed to develop or to adapt generic/terrestrial solutions for the specific space environment and its constraints. We briefly describe this trough two examples found in research projects: QuickSAT/Xen [QSAT] in the frame of university research programs for CubeSAT mission, supported by the US Navy, and XTRATUM for LEON processor, supported by the space industry in Europe (ESA, EADS-Astrium and CNES) [XTRAT].

- QuickSAT/Xen (extracts from ([QSAT])

The QuickSAT/Xen Space Hypervisor, is an open source space hypervisor that supports the virtualization of satellite payloads, systems and software modules on a range of satellites including CubeSATs and MicroSATs to satellites over 1000 kg.

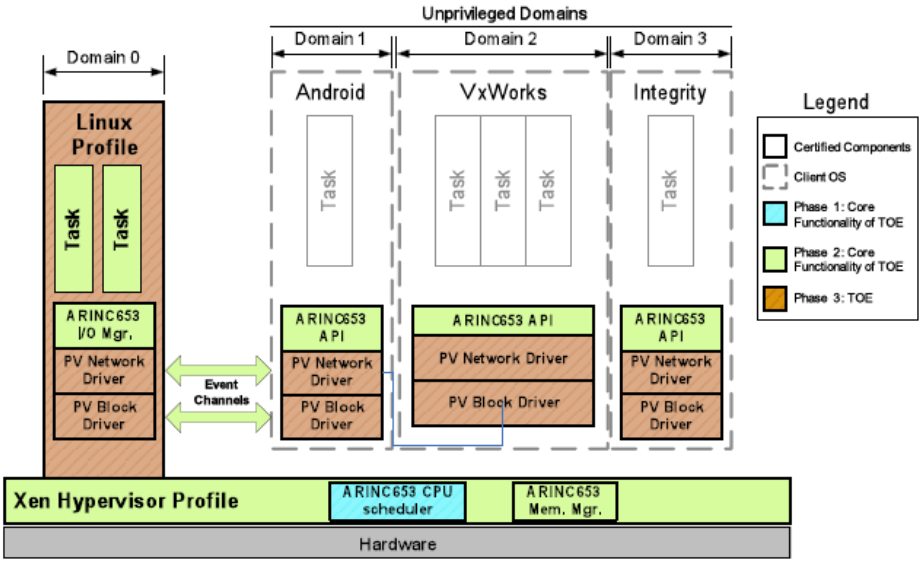

Figure 24. XEN/Arinc Hypervisor used for QuickSAT

- XTRATUM (extracts from ([XSTRAT])

XTRATUM is another (IT) Hypervisor, open-source, for the space environment targeting the LEON processors family. Once again, the objectives of Virtualization is the safe and secure partitioning of resources when multiple applications have to be executed in parallel on the same hardware platform. XTRATUM has been supported in the CNES LVCUGEN project, then in the ESA Securely Partitioning Spacecraft Computing Resources project.

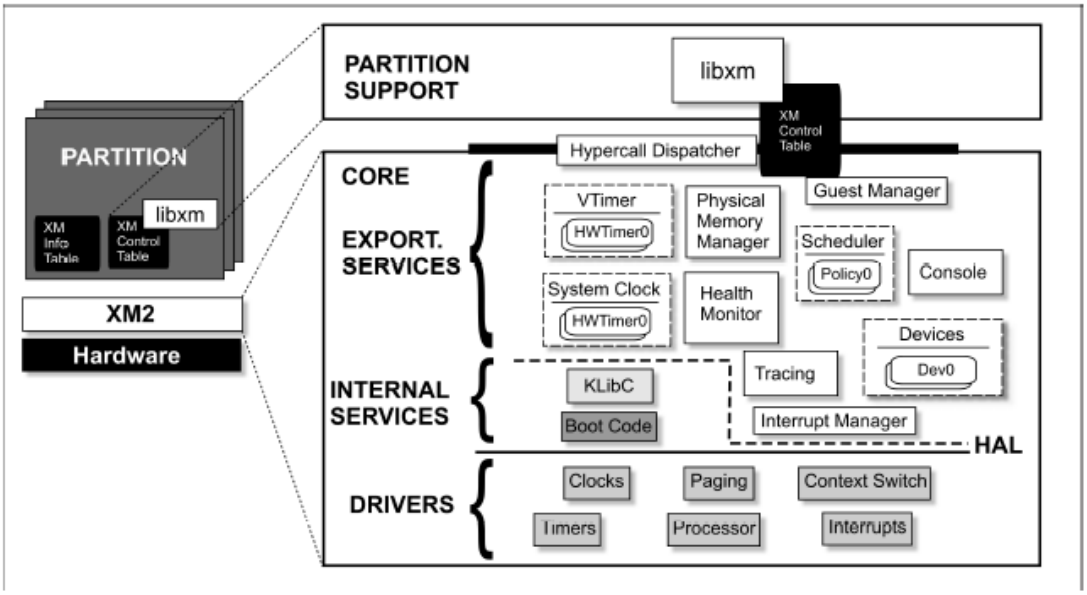

Figure 25. XtratuM architecture

The initial XTRATUM, design on Linux Real-Time OS as a Linux Kernel module, paved the way to payload virtualisation, on platforms using LEON2 processor. The work has been developed further with XTRATUM 2.0 for LEON3 processors in order to overcome important limitations, implemented full spatial isolation between partitions, and implement inter-partition shared memory when needed.

Conclusion

Virtualized payload for satellite is therefore not a completely new idea. The benefits and interests of those approaches are clearly different from the scope of Cloud Networking. However, this proves feasibility of the concept. Hypervisors may introduce some overhead (i.e. dedicated resources) and logical point of failure but in general they will improve isolation between processes - as they are expected to. The question of extending this model to offer Cloud Networking services in the Sky (e.g. implementing SDN-compatible switch on-board) seems not only relevant but fully realistic in few years.

Therefore the criticality level of this dimension is assessed as Low (GEO) or Moderate (MEO/LEO systems with ISLs) – according to the different requirements and missions associated to each type of satellite.

3.1.2. Cloud Networking Integration dimensions for satcoms

3.1.2.1. Satellite integration/coexistence with terrestrial cloud networking

Generic Integration with terrestrial networks

Satcom and terrestrial network integration (e.g. fixed access) have already been envisaged. In a recent example, the European BATS project (Broadband Access Terrestrial Satellite) [BATS] (see Annex 1) has defined a possible integration between a HTS satellite system with fixed and mobile terrestrial networks. At network level, the integration is focused on two key components, the Intelligent Network Gateway and Intelligent User Gateway. The questions of operability and deployment/configuration arise, as two (or more) providers shall have access to shared equipment (physically or per management). To this respect, Virtualisation could realistically pave the way to (remote) management assured by multiple entities – and avoiding chained delegation of operations (security issues).

What shall be noted also is that the introduction of virtualization technologies has already begun in terrestrial networks – and their adoption in satcom would become an additional factor of acceptance by the terrestrial community.

Therefore the criticality level of this dimension is assessed as High.

5G integration

The same argument as above applies for integration with future 5G networks. The marketing perspectives are also more than tremendous here.

Therefore the criticality level of this dimension is assessed as High.

Standards and technologies

Standardization perspectives for Cloud Networking are at the crossroad of satcom, networking, and mobile 5G networking standardization areas. The adoption of existing technologies (e.g. SDN/OpenFlow) shall remain transparent to satcom although some satellite-specific recommendations could be formulated as generic guidelines. Since Cloud Networking mostly focuses on operational aspects, any possible work would have marginal impact on their adoption. This assertion shall call for some nuance, as the global Network Management System (NMS) or OSS/BSS component could offer new interfaces and/or features with Cloud Networking, and where specific information in the domain of Virtualization would be exchanged between centralized controllers or managers with the lower-level entities (e.g. “low-level” NMS) of the Transport satcom networks. In the next diagram (extracted from BATS) the envisaged interfaces between NMS and centralized OSS/BSS are shown for an hybrid system. Those specific interfaces shall therefore be impacted with the adoption of a Cloud Networking model.

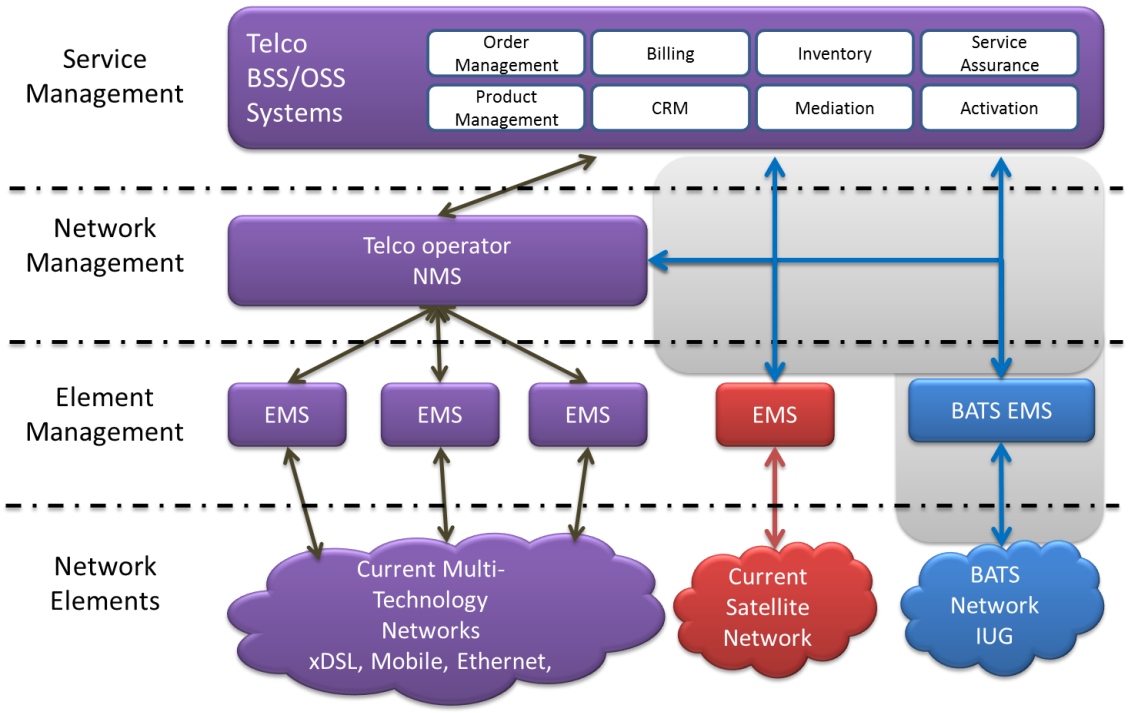

Figure 26. Network Management model in hybrid systems (BATS source)

Therefore the criticality level of this dimension is assessed as Medium.

3.1.2.2. Ground segment impact

On the ground segment, the Virtualisation of support network functions and appliances (e.g. PEP, firewall, etc…) will have moderate impact on the global system and operations, regardless the actual benefits (costs and functionality provided). Some constraints on the management side will be relaxed.

The criticality level for network support functions is assessed as High.

A much higher degree of impact will be found when essential network functions or appliances (e.g. SDN-enabled switches) will be introduced (Gateway/NCC). The resource management model will be possibly completely redesigned, also affecting aspects such as on-demand SLA, QoS, etc. And consequently, the Network Management System would be affected, requiring e.g. new interfaces and operational models to be defined.

The criticality level at GW/NCC components is assessed as High.

Finally for user terminal side, mostly the "highest" customer profiles (professional/corporate customers or trunking stations for terrestrial operators) present the highest interests for Cloud Networking (e.g. when integrated/interfaces to large networks with sustained requirement services). For low cost end-user terminals (mass market) the impact shall remain limited.

The criticality level at terminal components is assessed as Moderate.

3.1.2.3. Impacted phases in lifecycle

Another dimension related to cloud networking integration, which is interesting to observe, is the impact on lifecycle phases. Lifecycle may refer to a given component (e.g. a Virtualisation networking appliance) or to the entire system. The main stages are:

- Research, experimentation, prototyping, for which the main benefits should be to accelerate the Time to Market (TTM) **(High impact)**
- Tests to be conducted during validation (Assembly, Integration and Test phases), where specific (and costly) hardware used for test purpose can be replaced. Tests can be parallelized easily and are expected to be more efficiently and rapidly performed **(Moderate impact)**
- Operational phase: benefits are related to the capabilities to support smooth and quick re-planification (e.g. resource reallocations). Also whether simultaneous configurations can be executed in parallel, when needed **(High impact)**
- Continuous product/system upgrades: progressive upgrading could be supported at any time, while limiting the system unavailability. In the meantime, it is important for satcom that the deployment of new technologies/implementations is less risky and can be easily tested in well-established and deployed systems (e.g. Military systems). This could even favour the desire and ability for operators and providers to constantly improve their deployed systems. Most of the time, reconfiguration in a satellite network requires costly changes in order to maintain compatibility **(Moderate impact)**

3.1.2.4. Technology Readiness

Technology Readiness Level (TRL) is an important indicator to keep in mind for assessing a technology in the quite monolithic domain of satcom. Usually, development cycles of satcom technologies are very long compared to terrestrial systems and frequently only high-graded TRL technologies from terrestrial domain are deployed and used (sometimes with adaptations).

Technology Readiness is assessed to be of a Moderate criticality level.

3.1.3. Cloud Networking Business dimensions for satcom

3.1.3.1. Satcom Market attractiveness

Broadcast Services

A main interest of Broadcast Service providers in Cloud Networking would be related to upgradability. Broadcast Services still represent an important share of the satellite

revenues so that some marketing opportunities can be identified for Cloud Networking technologies.

Current broadcasting systems are based on rather fixed, known, and predictable service and characteristics, and very limited requirements in terms of flexibility. Technology improvements are introduced at very limited rate in such systems, except when considering hybridization with broadband. However in the domain of service delivery and applications, a virtualized STB/terminal could make sense to ease the introduction of new services (support/evolutions of new codecs such as MPEG-4/H264 towards H.265 (HEVC), 3DTV, etc…). Further, the large scale (number of units) of this market calls for sounding business opportunities.

Attractiveness is assessed as High for these services.

Commercial Telecom / Broadband Services

For Broadband Services, the identified benefits related to flexibility (in the deployment, and during operations) are especially applicable.

Professional Services and (Mobile) Backhauling, two other important sectors for satcom, shall also find innovative and attractive business models to support "on-demand" provision (and reconfiguration) of services with Cloud Networking potential.

Attractiveness is assessed as High for these services.

Institutional, Governmental and Defence Telecom services with FSS

Flexible deployment and (re)-configurability aspects are essential in those systems where for example, steerable beams can be activated/deactivated to cover a variable number of theaters of operations. The limited number of such systems per country usually involves high budgets (billion(s) of euros), and the replacement rate of the installed systems is quite limited. In addition, the actual location of theatres and spots is not predictable. This plainly justifies efforts to be put on optimizing the system before and after its deployment, via flexible and dynamic resource management.

Attractiveness is assessed as High for these services.

MSS

In MSS, the need for scalability and resource reallocation (case of unequal load per spot) may be even more exacerbated given the scarcity of resources that often induce system load (case of global constellations), and the high cost of access.

Attractiveness is assessed as High for those services.

Vertical Markets

Vertical markets are a fast developing sector for satcom, related to very high value-added services, in industrial domains of utilities (energy, water, gas…), transportation,

automotive, agriculture and safety/security. All these sectors may include sensor monitoring, Machine Type Communications (MTC) or IoT/M2M, Vehicle to Vehicle communications (V2V) etc.

Vertical markets can share most of the concerns and characteristics of MSS, even though service usage commonly involves lower bit rate applications. Thus, when it comes to the applicability of Cloud Networking, it seems rather limited from a network point of view. However, the possibility to deploy added-value services e.g. for data processing within the network in virtualized IT assets seems very attractive.

Attractiveness is assessed as High for those services.

3.1.3.2. Stakeholders' interests

Satellite manufacturers/integrators and manufacturers of ground segment equipment

For manufacturers and integrators, Cloud Networking could offer the potential to rapidly test and deploy the system in different and any possible configurations. Functional tests may require developments of features that may be much easily supported via software rather than hardware (e.g. chipsets may not be available), even if the level of performance is degraded. Time to Market can also be reduced.

Interest is assessed as High.

Satellite operators

Satcom operators shall firstly be interested in flexibility (such as power and frequency allocation with multi spotbeams) aspects. Rapid reconfiguration and capacity re-planning are also two key advantages.

Interest is assessed as High.

Satcom (virtual) network operators

Satcom network providers can dynamically re-allocate capacity to their different customers according to requested traffic for instance. In some cases, this request can be anticipated but sudden unpredictable events may also occur. The response time is also an important aspect to make such feature realistic.

Also, the ability to test (smoothly) new technologies or to work in parallel with non-interoperable technologies is especially attractive. As an example, a given satcom provider could support with the same Gateway IP connectivity for private residential end-users; VLAN for the need of interconnecting distant sites of a company, and 802.3ad interconnection for Provider Bridging services (in the case of terrestrial network customers).

Interest is assessed as High.

Professional customers

Similar situations/benefits can be expected from professional customers (also including network operators which use satcom for backhauling), where resources can be requested (or freed) on-demand, without affecting other services.

Interest is assessed as High.

Individual customers

Individual customers could benefit from instantaneous SLA renegotiation to increase the capacity of their connection – or to save money when they don't use the connection service. This possibility is available only if their hardware equipment (modem, ODU part, etc.) can support the additional requirements.

Interest is assessed as Moderate.

3.1.3.3. High-level costs

When it comes to costs associated with deploying, integrating and operating virtualisation technologies, different elements will have to be assessed (at high-level):

- Costs for developing the technologies, and/or for adapting them for the satcom context
- Additional CAPEX for buying Cloud Networking / Virtualisation equipment and to train operators and users. At least the same level of performance in any area shall be targeted.
- Revenue gains provided from the virtual additional capacity: for example gains may be expressed as the number of additional customers a satcom system could accept with Cloud Networking compared to a traditional system (note: may be in combination with other technologies).
- Cost savings related to test configurations of devices and interactions in the system. As an order of magnitude, tests of any sorts "can represent up to 1/3 of the CAPEX of a satellite, and indirectly - due to its characteristics to lengthen the time before the satellite is in service - it gives a cumulative time-value-of-money cost of generally about a year" [VIPS].
- Cost saving for deployments of a new releases in an operational system (e.g. related to reduction of service outage, etc.)
- Enhancement of billing models oriented towards "on-demand" vs. "flat rates"

For any tenants identified above, costs will represent of course a main driving factor in the development and adoption of this technology for satcoms. Moderate economy and impacts can be expected at the system level (space segment + ground segment).

The criticality level of this dimension is assessed as Moderate.

3.1.4. Synthesis

The tables below present an overview of the dimensions reviewed as well as their criticality.

Table 4. Criticality assessment for Functional dimensions

Functional dimensions	Criticality Assessment
Networks federation and coupling	High
Isolation of services	Low
Reconfigurability and Programmability	High
Mobility	Low
Resource elasticity	High
Availability and resiliency	Low
Security and privacy	Moderate
Accounting, billing, SLA	High
Performances	Low
Satellite Specific Capabilities	Low to Moderate (long term)

Table 5. Criticality assessment for Integration dimensions

Integration dimensions	Criticality Assessment
Convergence / standards	Generic Integration with terrestrial: High 5G integration : High Standards: Medium
Ground Segment impact	Network support: High Gateway/NCC: High Ground Segment (satcom Terminal): Moderate
Impacted lifecycles	Experimentation/Prototyping: High Validation/Tests: Moderate Operational phases: High System upgrading capabilities: Moderate
Technology Readiness	Moderate

Table 6. Criticality assessment for business dimensions

Business dimensions	Criticality Assessment
Market attractiveness	Broadcast Services: High Commercial Telecom / Broadband Services: High Institutional, Governmental and Defence Telecom services with FSS: High MSS: High Vertical Markets: High
Stakeholders interests	Manufacturers and system integrators: High[6] Satellite operators: High Satcom network operators: Moderate Professional/Terrestrial Network customers: High Individual customers: Moderate
High-Level costs	Moderate

3.2. Assessment of enabling technologies

Following the identification of the dimensions of suitability for satcom/cloud networking integration, we proceed by recalling the enabling techniques and technologies surveyed in Chapter 2 and assessing their suitability for integration into satcom infrastructures.

The methodology adopted is as follows:

- We examine separately each of the four technological domains
- For each domain, we consider the possible integration with satcom taking into account the dimensions of suitability identified in Chapter 0.
- For the "Functional" dimensions, we identify explicitly:
    - the added-value brought to satcom from the integration of the technology
    - the added-value brought to terrestrial cloud networks from the integration of satcoms
    - potential disadvantages and/or side effects which may be associated with from this integration

[6] Given the components considered

(Please note that dimensions which are considered as not relevant to a specific technology domain have been omitted)

- For the "Integration" and "Business" dimensions, we discuss in high-level the aspects and opportunities associated with the integration, the business aspects involved, as well as the impact on the timescale and roll-out of terrestrial and satcom cloud network services.

3.2.1. Infrastructure Virtualisation

As overviewed in Chapter 2, the Infrastructure Virtualisation technology domain includes all the techniques and technologies which enable the virtualisation and abstraction of physical resources, either networking or IT.

3.2.1.1. Functional dimensions

Table 7 discusses the applicability of infrastructure virtualisation technologies to satcom, with regard to the Functional dimensions identified in Section 3.1.1.

Table 7. Functional dimensions for the integration of infrastructure virtualisation technologies in satcom

Dimension of Suitability for Integration	Added-value to satcom	Added-value to terrestrial cloud networking	Disadvantages
Networks federation and coupling	Also, IT virtualisation enables the enhancement of satcom services with IT resources for application hosting and/or traffic processing	The statistical multiplexing within the satellite forward link can accommodate the high dynamicity of the network traffic for IT cloud services without significant management overhead. Generic traffic tunneling (such as GRE may be implemented to provide any type of interconnection , whatever the service interface supported at the satellite segment (e.g. IP only)	Satellite and terrestrial domains are usually administered by different business entities, which restricts federated management.
Isolation of services	Cloud technologies enable the logical isolation of IaaS IT	Multi-spot transmission and on-board switching (when implemented)	Due to the broadcast nature of satcom, user traffic is physically

	services provided over satcom.	each facilitate service and resource isolation in communication among remote sites	received across the entire spot (reduced area in case of multispot beam systems)
Reconfigurability and Programmability	IaaS IT services over satcom can be easily reconfigured upon user and operators request.	-	End-to-end network slices cannot be easily reconfigured in case of highly dynamic environments e.g. LEO constellations or satellite relays. This is a challenge for the long term.
Resource elasticity	Virtualised IT services over satcom can be up- and down-scaled on-demand.	-	Signal quality as well as other limits, such as hardware processing capacity must be taken into account during resource up/down scaling.
Security and privacy	-	-	Due to the broadcast nature of satcom, network virtualisation via tunneling by itself is not adequate for privacy, unless encryption is also employed. For the same reason, privacy issues associated with IaaS cloud services are even more important in a satcom environment.
Performance	-	-	In network virtualisation, adding an extra encapsulation layer further reduces satellite available useful capacity. In IT virtualisation, the performance of cloud services can be greatly affected by poor satcom performance (e.g. delay, temporary link outages). For the same reason, the applicability of network virtualisation for Data Centre

			interconnection (e.g. with VXLAN or NVGRE) over satcom is limited.
Satellite-specific capabilities	-	The broadcast nature of satellite can inherently support virtualized point-to-multipoint services, which usually pose significant overhead in terrestrial networks. Also, the PID or equivalent field in DVB-based networks (and/or equivalent signaling at GSE/RLE level) are often used statically in current system but higher dynamicity would be possible.	Existing satcom encapsulation mechanisms such as GSE already provide sufficient signaling overhead, so the use of an additional encapsulation technology might have no clear added-value.

3.2.1.2. Integration dimensions

Virtualisation technologies, especially network virtualisation, while often transparent to satcom, can bring added value in federated scenarios. In particular, Network slices can be extended to reach locations not covered by terrestrial networks. Also, isolated terrestrial backhauls can use satellite to seamlessly interconnect over tunnels.

With regard to virtualisation of IT infrastructures, its adoption across a hybrid domain of both terrestrial and satellite data centers, creates the need for a new approach for designing and building IT assets into the satcom infrastructure, taking into account aspects such as elasticity, migration and multi-tenancy.

Concerning elasticity, it should be noted that virtualisation creates the need for elastic provision plans, according to which computing resources will be scaled up or down, on demand, by adjusting the provided allocated resources. This is an important paradigm shift for satcom provision plans, which are in their majority static.

Concerning migration, virtual machines can be migrated on the fly while in service from one physical infrastructure to another one, with scope to reassure high availability provision and minimizing maintenance impact. This implies the decentralisation of the satellite operator infrastructure, from a single infrastructure site (satellite gateway) to multiple sites with failover domains.

Concerning multi-tenancy capabilities that are achieved by IT infrastructure virtualisation, the opening of satcom IT assets to multiple tenants achieves economies of scale by aggregating resources across applications, business units, and even separate corporations to a common infrastructure.

In addition, elasticity, migration and multi-tenancy implies that the satellite network itself will be able to respect: a. the elasticity of the resources (and therefore the elasticity of the bandwidth), b. the geographical portability of the resource, c. the provision of the appropriate and agreed QoS per tenant end-to-end and d. reports/interaction to billing and network management. Thus, consistent network-supported and virtualisation-driven policy and controls are necessary.

In terms of technology readiness, infrastructure virtualisation is considered a mature technology on which reliable solutions can be defined, designed and developed, especially for private infrastructures.

The introduction of IT infrastructure virtualisation in a satcom platform, although it is not expected to directly affect the product lifecycle for core satcom components (unless it is used in conjunction with NFV, as discussed), it has in any case the potential to increase the responsiveness of satcom stakeholders to new trends and services.

3.2.1.3. Business dimensions

Till today, telecom (including satcom) investments were focused on non-virtualizeable hardware infrastructure, involving high cost, limiting business strategic decisions, slowing down the product lifecycle and hampering investments in novel technologies.

However, recent developments show that the telecom stakeholders are gradually shifting their focus from a hardware centric growth model (i.e. specialized IT infrastructure) to a software centric (IT infrastructure agnostic) business model. The primary reason behind this shift is the noticeable progress that has been observed in the fields of cloud computing and virtualisation, leading to commercially mature technologies and platforms that allow the virtualisation of the infrastructure and the utilization of the virtualized IT resources for the implementation of a wide range of software-based systems. Satcom industries are expected to be favored by the virtualisation trend, since will give them the opportunity to rebrand their specialized IT infrastructure into a flexible, general-purpose software centric (IT infrastructure agnostic) product.

Therefore, from a business perspective, the satellite operator who follows the IT virtualisation paradigm, becomes more flexible to market changes and can adapt faster to new technologies and to the new challenges of the international environment and competition.

Moreover, for a satcom operator, beyond CAPEX/OPEX reduction, infrastructure virtualisation greatly simplifies IT management, which was previously being distributed among various management frameworks of specialized products/systems Centralized and integrated management reduces the budget, time and resources devoted to application coordination.

Other suitability perspectives that make IT virtualisation attractive for satcom could be considered in a variety of factors, such as:

- minimization of the need for rack space in data centers (physical advantage)
- reduced power consumption (environmental advantage)

- improved scalability (technical advantage)
- optimized workloads (performance advantage)
- simplified cable infrastructure (maintenance advantage)
- reduced complexity (design advantage)
- centralized administrative tools (administrative advantage)

With regard to service portfolio enhancement, IT virtualisation allows the satcom service provider to provide "native" (hosted) cloud services to the customer. However, in this field, the competition with large scale cloud providers (e.g. Amazon, Microsoft etc.) is hard to confront, unless significant discounts are offered in service bundles (e.g. satellite connectivity + cloud hosting) or there exist critical privacy constraints which do not allow the use of public clouds. In this case, the "satcom+cloud" bundle might match the needs of specific customer group.

Another dimension regarding customer-facing services lies in the fact that IT virtualisation is essentially an enabler technology for NFV, thus allowing the satcom operator to offer innovative VNFaaS offerings, as will be discussed in Sec. §3.2.3.3. .

3.2.2. Programmable and Software-Defined Networking

As overviewed in Chapter 2, Software-Defined Networking –and, in general, network programmability- technologies enable the decoupling of the forwarding and control planes in network nodes, allowing the control logic to be off-loaded to a centralized software-based entity.

In a satellite platform, SDN can be adopted as a management paradigm to control the Satellite Gateway (SG), (in the longer term) the on-board processor (OBP), or even the Satellite terminal/CPE (Figure 27). This implies that either the SG, the OBP and/or the satellite terminal are properly upgraded in order to expose an Openflow-enabled control interface, allowing to control the entire traffic which is traversing (or part of it).

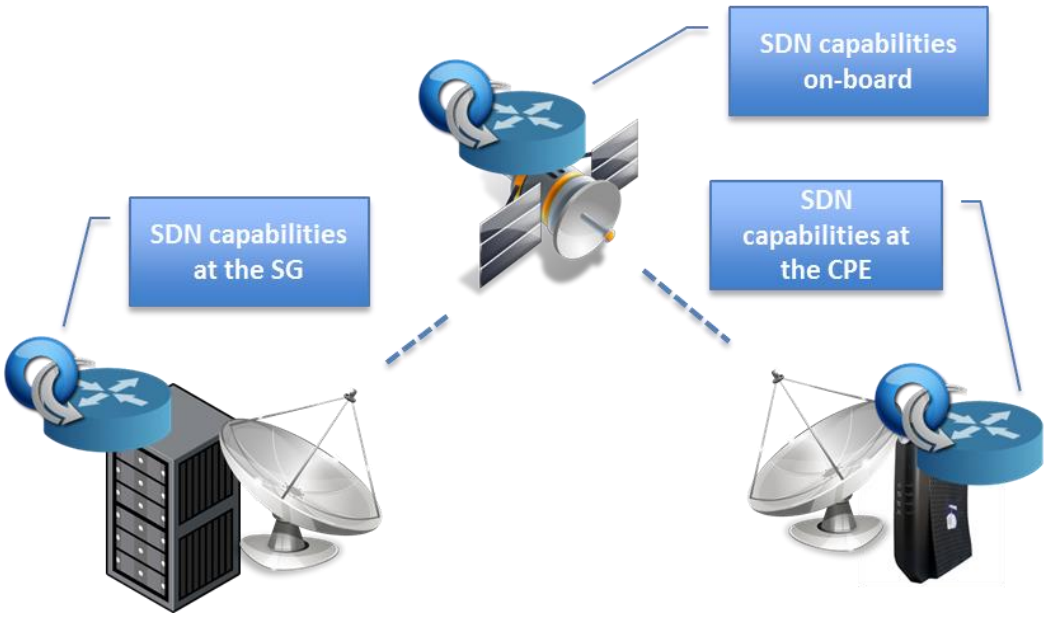

Figure 27. SDN support at various segments within the satcom infrastructure

SDN capabilities on-board shall also be envisaged for systems with ISL requiring internal routing (e.g. case of global LEO constellations for MSS services), not shown in this figure.

The scope of interests for adopting and integrating SDN capabilities in those components will of course largely depend on the scenario and applicable network architectures. Such scenarios will be addressed in the next chapter. At first sight, any combination could be envisaged. SDN could be independently supported only at SG, satellite, or terminal side, or in any 2 or even the 3 components.

3.2.2.1. Functional dimensions

Table 8 discusses the applicability of network programmability and SDN technologies to satcom, with regard to the Functional dimensions identified in Section 3.1.1.

Table 8. Functional dimensions for the integration of network programmability and SDN in satcom

Dimension of Suitability for Integration	Added-value to satcom	Added-value to terrestrial cloud networking	Disadvantages
Networks federation and coupling	By embracing SDN, the satellite network can expose a vendor-neutral, universally supported interface, enabling unified management with terrestrial networks. The Satcom platform can be managed via any SDN controller.	-	Per-flow management in a SG handling traffic for thousands of users may be unscalable.
Isolation of services	State-of-the-art SDN controllers such as OpenDaylight/VTN can be used for the provision of logically isolated NaaS services to multiple tenants over satcom.	In the case of multi-spot transmission, services across different beam footprints are physically isolated.	In single-beam configurations, the inclusion of all the traffic into a common satellite forward link significantly hampers the SDN switching capabilities.
Reconfigurability and Programmability	SDN allows the application of arbitrary per-flow logic at the SG, beyond existing, "hard-wired" network protocols. Furthermore, the perspective of SDN support at the OBP results in a very flexible payload, allowing custom per-flow operations on-board.	-	The Openflow protocol is rapidly evolving, thus not yet stable enough for OBP deployment. Increase of complexity, resources, weight, occupied space and (launching) costs for the satellite payload with a SDN-compatible OBP needs to be assessed

Mobility	SDN allows rapid flow switching and/or address rewriting, thus significantly facilitating mobility at L2/L3 level across beams or carriers or even LEO handovers.	Mobility scenarios across a large geographical area in terrestrial networks require management of a distributed and extended network of access points, which is quite complex. Instead, SDN-based control of a single satellite gateway can allow multi-beam handovers in a much simpler manner.	-
Resource elasticity	SDN facilitates rapid capacity up/down scale within the pool of satellite customers, per-flow granularity, in a much more flexible manner than current SG management allows.	Forward link services can reach high degree of bandwidth scalability as compared to wired terrestrial network (Point-to-Point links) as resource separation is mostly based on Time Division Multiplexing with very high throughput links. Resource reallocation would have no impact on MAC/physical layers, if limited on the same TDM multiplex.	Return link services cannot provide so important upscaling capabilities, as resource are frequently shared according to MF-TDMA access with fixed allocation to carriers, and given that terminals cannot generally transmit on simultaneous carriers
Resiliency and availability	Within the ground segment, SDN mechanisms can rapidly divert the traffic to failover units in case of failures.	Centralised SDN management can divert the traffic via satellite in case of terrestrial network outage or failure.	The network SDN-controller (and associated interconnection means) appears as the single point of failure and needs redundancy
Security and privacy	Openflow metrics can be leveraged for detection of attacks and/or anomalies. Targeted blocking of specific flows can be applied to eliminate the incident.	-	An SDN-enabled OBP may pose high security risks, if the control channel link is compromised. Generally, SDN security issues are much more critical in a satcom environment. Likewise, access control

			for the centralized SDN-switch/router at SGs is also critical to protect from DoS attacks
Performance	User QoE can be enhanced thanks to SDN-driven hybrid delivery (per-flow switching from satellite to terrestrial and vice versa)	High-bitrate popular content can be switched via SDN management to satellite, so as to relieve terrestrial access.	Intensive Openflow signaling may degrade management performance, especially when a large number of users/flows is involved.
Satellite-specific capabilities	-SDN can be exploited in OBPs and/or LEO/relay scenarios to allow intelligent on-board switching based on arbitrary rules. -Smooth and non-risked introduction of new forwarding/switching/routing technologies or standards in deployed systems is fully supported with SDN. -Distinct switching technologies can be supported simultaneously if needed (case of a system that would support radical different customer profiles)	Per-customer and per-flow QoS guarantee is simpler in a satcom environment, since the radio resources are known, shared and directly controllable at the SG. Also, SDN management can be enriched with satellite-specific features e.g. PID or DVB label management and association.	Higher complexity with higher resource (processing/memory) is expected for SDNenabled payloads

3.2.2.2. Integration dimensions

Programmable and software-defined networking is considered a key driver towards satcom/terrestrial coexistence and federation. By becoming SDN-enabled, the satellite network becomes an integral part of the 5G landscape, in which SDN is expected to play a dominant role in network control and management. From a terrestrial point of view, typical terrestrial SDN use cases are significantly widened thanks to satellite access and satellite capabilities, thus further promoting SDN and widening its applicability.

The most important feature of SDN, which is expected to greatly facilitate satellite/terrestrial integration is the unification of the network programming. Currently, network engineers have to face the diversity of proprietary device programming interfaces, which are vendor-specific and require managers to program each vendor's device using vendor-specific commands (e.g. Cisco IOS vs. Juniper JunOS). SDN introduces a single programmable layer, merging the current vertical

operational isolation and independency of different vendor devices, and enables horizontal communication to the network devices through a single common protocol.

This capability can be a major driver to network federation, since it hides the management specificities of the satellite infrastructure and abstracts it into a standard SDN domain, which can be managed in the same manner as other (terrestrial) SDN domains, by means of high-level controllers. The latter can have a single view of the entire satellite/terrestrial network in comparison to the vendor silos that currently require administrators to administer the network through vendor-specific tools/protocols[7]. The same simplicity also applies to network monitoring. SDN-managed satellite networks provide opportunities to better leverage advanced data analytics and real-time monitoring, down to per-flow granularity, in order to simplify administration and provide automation opportunities that do not exist in traditional networks.

From a technological readiness and maturity perspective, although SDN is widely popular with cloud providers, carriers, and universities, today's SDN solutions are not mature yet for enterprise deployment for the following reasons:

- Major technology vendors, while acknowledging that SDN is a future direction, have yet to agree on a common set of interoperability standards for all their network products, despite SDN's open heritage.
- Key satcom vendors don't offer SDN functionality as part of their mainstream product lineup
- Most enterprise teams lack the experience, the training and the maturity to manage such an SDN environment.

However, SDN is still recognized as the main driver towards infrastructure openness and network federation, and that is why it is adopted as a key enabling technology in most future network architectures, including the 5G landscape.

3.2.2.3. Business dimensions

Currently, the business model of the satellite operators is limited to rigid, slow and hard to change schemes, which do not allow fast adaptations. Network architecture today is closed, complicated, and still focused on network elements that only allow individual (and not federated) management. In other words, the current business model of the network provision is closed, protocol-centric and vertically fragmented with management software being embedded within each network node/device, lacking a horizontal approach with coherent and unified management per tenant (i.e. per network slice).

In this context, SDN can be seen by satcom stakeholders as the enabling technology towards a new business model, which moves the current situation from rigidness to

[7]It must be noted, however, that vendor-specific enhancements on OpenFlow, as well as differences among Openflow versions may create interoperability challenges.

elasticity, from slow adaption to fast adaption and from vertical to horizontal provision of the network resources, enabling new network capabilities.

From a business perspective, SDN-enabled satcom extends the service programmability model to the satellite domain, allowing to dynamically allocate resources in a horizontal (i.e. slicing) way and apply arbitrary control logic to user flows. In this new business model, the satcom offering is augmented, from a monolithic provision of a single, autonomous and independent connectivity domain to a universal virtual network infrastructure (slice), which is not only logically isolated, but also programmable. Moreover, SDN creates the basis to shift from flat-fee billing to a pay-as-you go business models, which are greatly facilitated by SDN. Instead of requesting an upfront payment for fixed vertical and autonomous service provision, satellite customers become able to subscribe to services only when they need them on an elastic basis.

In addition to service enhancement, SDN has the potential to transform the satellite equipment vendor market; SDN technology exports the embedded network functions from the devices and offloads the control logic to SDN management software. So, SDN technology can rebrand the satellite network industry from a costly status quo based on closed, monolithic hardware, individually managed, to a cost-effective abstracted network domain, where the network controllers are integrated into a single control unit and decoupled from the data plane. By using SDN, it will be possible that the control plane (software) and the data plane (hardware) of satellite network elements are developed by different vendors, yet remain fully interoperable. Control components for terrestrial networks can be re-used, achieving economies of scale. The satcom market can thus be significantly widened, as new market entrants gain the potential to develop control code which will be used to manage high-performance data plane hardware. This market openness can be further promoted by NFV, as will be explained in Sec. 3.2.3.3. .

On the other hand, however, a controversial issue has been raised on SDN technology adoption. While the concept of SDN is well-defined and clear, its implementations are still evolving; the opponents of the SDN argue that the current versions of SDN APIs (such as OpenFlow) are being released too quickly, without allowing for stable product development. This quick and forced market penetration is faced critically from traditional networking vendors, whose products are a result of several years of research, development and testing. This consideration can be even more critical in the case of satcom, where well-established (and well-tested) product series from a few vendors dominate the market. Moreover, the lack of training, expertise and experience of the new networking architecture will cause internally in the organizations a reaction to the technology change and the SDN adoption.

3.2.3. Network Functions Virtualisation

As explained in Chapter 2, Network Functions Virtualisation refers to the migration/support of network functions traditionally hosted in monolithic hardware-based network elements, to software entities hosted in generic commodity servers.

As a concept, NFV is applicable to functionalities across all layers of the OSI model, from the physical up to the application layer.

In a satellite platform, NFV could be used[8] to virtualize:

- *core SG functions*, such as firewalling, PEP/acceleration, scheduling, media transcoding etc.
- *radio front-end functions*, such as modulation and coding (according the Cloud-RAN concept) (much probably, in the long term for satcoms)
- *on-board functions,* such as switching, replication, or other kind of traffic processing (in the long term)
- *customer premises equipment (vCPE) functions,* such as firewalling, traffic inspection, intrusion detection, etc. vCPEs can be instantiated either at SG side (i.e. before the satellite segment) or at the customer side -depending on the function being virtualized- in specifically configured satellite modems, equipped with additional computing resources in order to be able to accommodate software VNFs.

NFV enables all the aforementioned functionalities to be instantiated and offered "as-a-Service" to tenants/customers.

These scenarios imply that NFV can be applied at either the satellite gateway and/or the customer side, without excluding the long-term perspective of NFV-enabled satellite payloads, as shown in Figure 28.

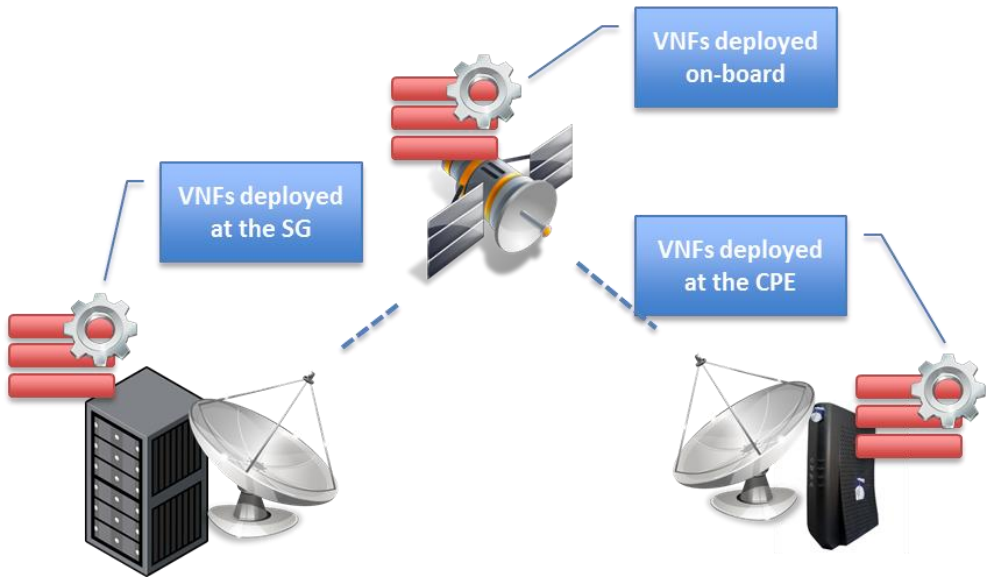

Figure 28. NFV support at various segments within the satcom infrastructure

8 The actual opportunities chance and realistic timeframes to implement those different "flavors" of NFV greatly vary. As previously explained in section 3.1.2.2. , ground segment functions then pure satellite segment shall progressively be able to host and integrate Cloud Networking technologies in general, and VNFs, in particular.

3.2.3.1. Functional dimensions

Table 9 discusses the applicability of Network Functions Virtualisation to satcom, with regard to the Functional dimensions identified in Section 3.1.1.

Table 9. Functional dimensions for the integration of NFV in satcom

Dimension of Suitability for Integration	Added-value to satcom	Added-value to terrestrial cloud networking	Disadvantages
Networks federation and coupling	The NFV paradigm simplifies the provision of added-value networking services in satcom systems, by re-using the NFV management platforms developed for terrestrial networks. Federated NFV management is further promoted if the satellite network is SDN-enabled.	In satellite architectures, VNF deployment will mainly take place at the Gateway side in a centralized manner, which simplifies VNF mapping and placement decisions, as opposed to distributed terrestrial infrastructures.	-
Isolation of services	NFV achieves the logical isolation of the traffic processing entities, in addition to capacity. Thus, dedicated isolated VNFs such as virtual firewalls and virtual PEPs/caches can be offered to (and managed by) the customers.	Remote management of VNF (case of network provider as a satcom customer) is highly facilitated	-
Reconfigurability and Programmability	Satcom network functionalities are no longer restricted in monolithic “boxes”, but become totally reconfigurable. Introduction of new protocols and services is greatly facilitated. The perspective of NFV on-board allow arbitrary processing within the OBP – in the limits of the underlying hardware reconfigurability	-	Higher complexity with higher resource (processing/memory) is expected for NFV-compatible payloads
Mobility	Mobility protocols can be implemented and deployed as VNFs within the satellite infrastructure. Moreover, live VNF migration will allow the VNFs “follow” the customer upon switching from one network to the	Due to the large satellite footprint, user terminal mobility across an extended geographical area can be supported without having to continuously migrate the user VNFs from one NFVI-	VNF mobility can be an issue in highly dynamic environments (LEOs, satellite relays etc.)

	other.	PoP to another.	
Resource elasticity	As opposed to hardware network elements of satellite infrastructures, VNFs can be up/down scaled on-demand.	The aggregation of most VNFs in a single data centre (hosted at the SG) facilitates resource scaling, as opposed to performing adjustments in several distributed terrestrial NFVI-PoPs	-
Resiliency and availability	In case of infrastructure faults, the software nature of VNFs allows new instances can be rapidly deployed in failover servers. Costs to target an equivalent level of resiliency/availability may be reduced using generic hardware platform	After a switch to satellite access due to a failure in the terrestrial infrastructure (e.g. after a disaster or an outage), NFV-enabled satcom enables also the proper re-deployment and recovery of the customer VNFs. End-to-end fault monitoring may be easier thanks to unification of interfaces	NFV availability issues are critical in the satcom context, since e.g. a failure in a core VNF at the SG (or on-board) may affect a large number of users.
Security and privacy	Security-oriented VNFs such as virtual security appliances can be deployed per-customer, adding an extra layer or specific/customized feature or configuration for security.	-	NFV security issues (due to malicious or malfunctioning VNFs) are much more critical in a satcom environment, since VNFs are commonly installed in the SG (or on-board) and possible security incidents may affect a large number of users.
Performance	Acceleration, media transcoding and caching functions can be deployed on-demand as VNFs, boosting satcom performance and improving user QoE. Also, CloudRAN aspects in satellite transmission enable the faster uptake of new more spectrum-efficient technologies.	Dedicated VNFs for the acceleration of hybrid satellite-terrestrial access, based on context-aware network delivery (e.g. intelligent load balancers and multipath proxies) can leverage satellite capabilities to boost end-user QoE.	Performance issues of core VNFs deployed at the SG (or on-board) may affect a large number of users. In some situations where common hardware platforms (e.g. behind SG) would be used by different customers (e.g. terrestrial operators) VNF shall be coupled with IT Virtualisation to manage and allocate properly the

			hardware resource (processing power, memory and storage) between customers, in order to maintain acceptable network performances
Satellite-specific capabilities	NFV is considered to be a crucial driver to overcome the "ossification" of satellite platforms and enable the faster uptake of future network technologies. Also, NFV support at the customer side (via NFV-enabled CPEs) enables rapid deployment and upgrade of customer operations, such as encryption and PEP, QoS, or access control enforcement. Components of customer VNFs can be shared between the SG and the CPE.	Due to the broadcast nature of satcom, VNFs for customer-side installation can be "pushed" simultaneously to thousands of terminals, relieving the terrestrial network from individual VNF transfers. NFV support at the centralized sides (SGs) can be deployed and (re)configured remotely by terrestrial providers	-

3.2.3.2. Integration dimensions

Being a virtualisation technology, NFV aims at abstracting infrastructure-specific details, exposing a common virtualised environment across the satellite and terrestrial domains, on which generic VNFs can be deployed and run. In this context, Satcom customers and operators gain access to (and can benefit from) a wide variety of software VNFs currently being developed for the terrestrial domain. From the terrestrial point of view, the capabilities of satellite access significantly widen and expand the envisaged use cases for NFV, as identified in [NFVUC] for terrestrial networks.

Focusing on the technical issues of integration, the deployment of a virtualised IT environment in the satellite gateway is not considered to involve significant technical complexity and should be relatively straightforward. At the customer side, NFV-enabled CPEs require that IT virtualisation capabilities are embedded in the satellite terminal. This is becoming increasingly possible with the advent of virtualisation capabilities for lightweight, low-power computing platforms such as ARM architectures, given of course that the VNFs deployed are not too resource-demanding. Finally, NFV capabilities on-board can only be considered a long-term perspective, since virtualisation-enabled platforms do not yet qualify (in terms of size, power and required resources) for inclusion in the satellite payload.

Concerning the technology readiness of NFV, it should be noted that already commercial exploitations of NFV solutions have been applied by a number of top-tier communications service providers as a standards-based approach to virtualizing a range of telecom applications. However, while the technological foundation of NFV (including IT virtualisation technologies) is well-established, the NFV management frameworks and unified APIs are currently at early stage. In any case, it can be quite safely predicted that the NFV industry will in the short term converge to a specific toolbox of management technologies, most of them de facto standardized via open-source initiatives such as OPNFV.

3.2.3.3. Business dimensions

From a business point of view, NFV offers the potential to radically redefine the architectural logic of satellite networks, by softwarising key network functions at the core and the edge. The business value introduced by NFV to satcom could be identified in the following aspects:

- Consolidation of satcom hardware resources, leading to reduced equipment investment and maintenance costs (reduction of both CAPEX and OPEX).
- Power consumption reduction and enhanced green fingerprint via consolidation of several functions into a small number of servers.
- Enhanced resource utilization due to sharing/reallocation of resources among different network functions and tenants.
- Enhanced elasticity of satcom resources assigned to each network function, further optimizing resource usage and reducing costs.
- Opportunity to expand the satcom service portfolio with VNF-as-a-Service (VNFaaS) offerings; NFV enables satellite service providers to offer new types of services, creating virtualized service instances specifically for each customer/tenant and customizing them accordingly.
- Accelerated deployment of novel/upgraded satcom functions, leading to significant decrease of Time-To-Market (TTM). Such a reduced TTM, gives to the satcom industry unprecedented flexibility and adaptability to market needs and changes. For this purpose, concerning the impacted product lifecycle, it should be claimed that NFV has the potential to radically speed up the satcom product lifecycle, since it provides a solution towards the fast turn-around of the production evolution by low cost software updates and deployment.
- Promotion of innovation and competition, by opening a part of the satcom market and transforming it to a novel virtual appliance market, facilitating the involvement of software stakeholders. However, this opening could also be seen as a major threat by the most important actors of the sectors (big satcom operators, for instance). For some markets and/or geographical areas where a given provider or operator are in deep competition, open/virtualisation should be seen as valuable differentiators. But, on the other hand, in some well-established businesses, reluctance shall be expected in general on opening

and virtualising products and solutions. Arguments are mostly similar to the ones related to the effectiveness of integration between satellite and terrestrial networks.

From a service point of view, the automation of service deployment reduces significantly the time required for delivering new services according to the customer needs. This is a significant improvement in the market performance of satellite operators, for whom currently even a small change or modification in the service delivery chain, requires a corresponding change in the related hardware device, usually done manually. NFV is able to leverage the programmability of the network functions using appropriate APIs, shortening the time provision significantly.

Moreover, the ability offered by NFV to compose various network services rapidly and efficiently by chaining together virtual network appliances, is very important for satcom operators, who currently deal with a very specific hardware-based service chain, which does not allow any service bundling and assembling. The ability to blend services dynamically will promote a marketing mix, which opens new opportunities to the niche markets of the satellite industry.

Another significant stakeholder interest on NFV and VNF deployment should be spotted on the transformation of the fault resilience and availability cases. For example, in a conventional network, the failure of a hardware system is typically service-affecting and it requires urgent action to replace the failed unit in order to restore the proper level of fault tolerance. With NFV this business process is modified since the failure of a piece of hardware has no more impact on the service delivery except than a temporary reduction in the QoS and performance of a given service, which is accommodated by the rest available hardware resources (e.g. by migrating the VNF from a faulty server to a failover one). With the NFV reform, it is not anymore necessary to replace the failed hardware immediately and the service is continued even at a reduced QoS level. This modified business process introduces significant changes in the operational demands of the satcom provider, providing the business opportunities for new QoS constraints, high availability services and fault resilience schemes.

3.2.4. Federated Management and Orchestration

As highlighted in Chapter 2, the Federated Management and Orchestration technology domain embraces all the architectures and techniques which enable the joint management of network and IT assets, possibly across heterogeneous domains, for the provision of integrated cloud network services.

In the satcom context, Orchestrators are envisaged to be deployed on top of heterogeneous terrestrial and satellite infrastructures expected to develop in several contexts, in order to jointly manage their resources, establish end-to-end connectivity services and instantiate VNFs across the traffic path. For this purpose, the orchestrator needs also to perform actions such as path computation and service chaining (interconnection of service elements and VNFs across the traffic path).

3.2.4.1. Functional dimensions

Table 10 discusses the applicability of federated management/orchestration technologies to satcom, with regard to the Functional dimensions identified in Section 3.1.1.

Table 10. Functional dimensions for the integration of federated management/orchestration technologies in satcom

Dimension of Suitability for Integration	Added-value to satcom	Added-value to terrestrial cloud networking	Disadvantages
Networks federation and coupling	Orchestrator architectures significantly simplify network management, allowing high-level view of the satcom infrastructure and enabling optimal allocation of resources in an automated manner.	Extending the orchestration to satcom is quite efficient and scalable, since by only managing a single entity (the SG) it is possible to control service delivery to thousands of customers (as opposed to distributed terrestrial networks, in which efficient Orchestration is a challenge)	Link quality and satellite footprint are additional parameters which need to be taken into account during service mapping, making the resource calculation and allocation procedure even more complicated. Important efforts needed to develop novel Network Management and SLA/billing models (but this exceeds the single scope of Virtualization)
Isolation of services	Orchestration mechanisms abstract a full subset of the satcom functionalities and capacities, providing a cross-domain logically isolated slice, as an independently managed virtual network infrastructure.	-	Services isolation is hard to achieve when the service crosses several satellite/terrestrial administrative domains.
Reconfigurability and Programmability	Orchestrators commonly expose a northbound API for programmability, allowing the management of the satellite resources by arbitrary customer-defined network applications.	-	Current service programmability architectures do not scale well and may be unsuitable for a satellite network with a large number of users.
Mobility	Federated management allows centralized path establishment and	Federated management allows centralized path establishment and switching, providing	In federated environments, service setup time may be

	switching, providing alternative paths over terrestrial access where needed (i.e. when more capacity is required)	alternative paths over satellite access where needed (i.e. when the user moves in a region outside terrestrial coverage)	significant and may not allow fast handovers from satellite to terrestrial and vice versa.
Resource elasticity	Federated management involves a global view of the infrastructure capacity and thus allows the close-to-optimal re-allocation of satellite resources.	Including satcom in federated management allows for more dynamicity, since it provides the capabilities to offload traffic and functions to the satellite domain, where appropriate.	In federated environments, service rescaling time may be significant and may not be able to accommodate abrupt demand fluctuations.
Resiliency and availability	The joint management of satellite and terrestrial resources widens the failover options for the satellite infrastructure (usage of multiple terrestrial backhauls, allocation of capacity in terrestrial data centers for VNF migration etc.)	Federated management can achieve automatic migration of the entire cloud network service to the satellite domain in case of failure or unavailability of the terrestrial infrastructure.	-
Satellite-specific capabilities	Federated management allows satcom to fully expose its inherent capabilities (native multicast, wide coverage, uniform QoS) to a much wider customer group, acting complementarily to terrestrial infrastructures.	The uniform QoS offered by satellite, independent of user location facilitates SLA fulfillment in diverse user contexts in a federated satellite/terrestrial landscape.	-

3.2.4.2. Integration dimensions

With regard to satcom/terrestrial coexistence, federated management allows satellite networks to be an integral component in future network/5G infrastructures. The management and orchestrating solutions are expected to have significant impact on the current satcom landscape. The satellite community is exposed to a wide ecosystem, which is composed of heterogeneous players (terrestrial wired and wireless/cellular providers, network function vendors, infrastructure providers etc.) and which would allow multi-vendor schemes and End-to-End QoS orchestration.

Within an SDN- and NFV-enabled Infrastructure ecosystem comprised of multiple vendors with a single orchestrator and management system, VNFs can be provided across different infrastructure segments, as components of a network service over heterogeneous network domains (satellite, terrestrial). A fully software-defined, open, and virtualized network platform will offer to satellite service providers –and

also to their customers- the flexibility to select multi-vendor solutions across a wide range of virtualized networking functions.

The terrestrial community is also expected to benefit from such interplay, since cloud network services can be extended to locations/situations only covered by satcom. However, the considerations commonly associated with federated management still apply here; usually, satellite and terrestrial networks are managed by different business actors, limiting federated management capabilities. The conditions under which a satcom provider could be willing to expose part of the management interface to a third party, need to be carefully examined.

When it comes to technology readiness, the TRL of the orchestrating and management solutions is low, since no standards exist, apart from some open-source initiatives with significant momentum. In the field of NFV-enabled orchestration, the ETSI MANO framework provides just an architectural blueprint [NFVMAN], accompanied by high-level recommendations, while the current proof-of-concept implementations are far from being compatible with each other.

Finally, concerning the impacted system and the product lifecycle, it should be noted that till today multi-tenant networks have been operated using per tenant and vendor specific management systems, without these systems being aware of adjacent tenants or resources. This fragmented, vertical and distributed management approach forces network operators to manage different vendor's equipment with multiple management systems, which results in high complexity and poor network utilization, along with configuration errors, slow response to resolution of issues and higher operational costs. Orchestrators should impact established but also private OSS systems and applications, without replacing them but interfacing with them and other proprietary individual management systems, achieving interoperability between the vendor community and the carrier IT organizations.

3.2.4.3. Business dimensions

The new ecosystem that is formulated by the federated management paradigm has a major business impact on all involved stakeholders.

First of all, the satellite community gains access to an expanded portfolio of services, including both terrestrial and satellite resources that could dynamically join and leave federation partnerships. This novel elastic provision of resources will create the opportunity, beyond political, economic and geographical concerns, for the definition of novel value chains, the definition of new strategic joint ventures and the creation of novel services and applications, creating new opportunities for revenues. This federation paradigm augments the role of satellite operators as indispensable nodes of the value chain of the new virtualized market.

Moreover, novel management and orchestration schemes promise to efficiently cope with the dynamic network changes and reconfigurations. This ability greatly promotes the business value of satellite offerings, since it reduces the setup/configuration delay for federated terrestrial and satellite network services, while on the other hand it extends the manageability of the satellite domain.

Furthermore, the novel management and orchestrator systems will greatly facilitate customer-side service management through service status visualization and configuration via intuitive dashboards and GUIs. This novel paradigm of network management will affect respectively the current satellite service value chain, allowing user-centric service management, following the paradigm of virtualized service offerings provided by terrestrial players.

Last but not least, an interesting dimension is the ability of the satellite vendors to expand to the Orchestrator market. The interest of the stakeholders and vendors towards designing and developing orchestrating solutions is high, because the anticipated standardized framework is expected to provide a stable future market for the sales and commercial success of the orchestrating products. In this framework, SatCom vendors may be interested to get involved in the provision of off-the-shelf orchestrating solutions, which will give them the opportunity to gain a share of the network management market, which currently they are poorly involved. Moreover, orchestrating solutions developed by satcom vendors are expected to have a competitive advantage in comparison to "purely" terrestrial ones, since they will be tailored to also deal with the specificities and the special requirements of the satellite domain.

3.3. Technology Selection and justification

Following the analysis which took place in the previous sections, there is enough input so as to conclude to specific recommendations on the applicability and the perspectives of the integration of the different cloud networking technologies and techniques into satellite networks.

With regard to *infrastructure virtualisation:* In the field of Network virtualisation, there seems to be limited added-value in the application of network encapsulation/tunneling protocols such as VXLAN to satellite access networks, since satellite networks already provide Service isolation mechanisms - at least for systems using DVB-based standard technologies (DVB-S2/RCS2). Moreover, satellite links are not considered suitable, in the general case, for Data Centre interconnection due to the high delay involved and bandwidth limitation. However, the applicability of transparent tunneling mechanisms at the ground part for terrestrial domains over satellite link offering limited service interfaces (e.g. most frequently only IP, although sometimes VLAN Ethernet connectivity may also be available natively) appears as an enabler for network federation.

In the field of IT/Cloud virtualisation, it seems that the integration of traditional IaaS cloud computing services into the satcom platform does not present any significant added-value, since these services can be hosted in a public cloud provider and accessed via satellite without noticeable QoE degradation (however, strong bandwidth limitations may apply with respect to usual performances in terrestrial networks). However, the usage of IaaS cloud platforms as NFV enablers seems quite promising and thus need to be studied. In particular, the Openstack platform [Openstack] appears as the most candidate for this purpose, since currently it exhibits the highest momentum among all cloud management systems (e.g. Cloudstack,

Eucalyptus etc., as listed in Chapter 2). Moreover, Openstack is the most likely candidate for the reference Open Platform for NFV (OPNFV), currently under elaboration.

With regard to *Programmable and Software-Defined Networking,* network programmability is considered as a major driver for the inclusion of satellite platforms into 5G future networks. Satellite Gateways and –in the longer term- on-board processors need to become SDN-enabled in order to allow vendor-agnostic, unified management with terrestrial networks. SDN-enabled satcom allows the use of SDN controllers to be used for the provision of logically isolated NaaS services to multiple tenants. Also, SDN can radically alleviate the ossification of satellite platforms, facilitating the application of arbitrary per-flow logic at the SG, beyond existing, "hard-wired" network protocols.

When it comes to specific technology selection, Openflow is by far the SDN control protocol with the highest momentum for the time being, so its adoption is mandated, although its quick evolution and the absence of a stable, long-term version raises issues about its deployment on-board. With regard to Openflow controllers, OpenDaylight [ODL] seems the most likely candidate, being a complete network management solution with a wide spectrum of services and –most important- huge industrial and community support; although not yet mature enough for deployment in a "production" environment.

With regard to *Network Functions Virtualisation,* the flexibility and agility introduced by softwarising network functions can be seen so as to significantly contribute to alleviating the current ossification of satellite platforms. By migrating network functionalities to software entities, new network protocols can be deployed rapidly, tested and fine-tuned without affecting the operation of the satellite network. Ultimately, in the long term, this concept could also be applied on-board, adding unprecedented flexibility and reconfigurability to OBPs. Routing/switching paradigm in future constellation-based systems could also be heavily affected.

In any case, network functions can be offered on-demand to satellite customers in the VNF-as-a-Service context, significantly expanding the service portfolio of satcom service providers. For this purpose, VNFs initially developed for terrestrial networks can be slightly adjusted to match the specificities of satcom and introduced to the satellite domain, thus vastly enriching the functionalities and capabilities of the latter. Finally, expanding the NFV paradigm to also embrace the radio front-end (long-term), allows the full virtualisation of the satellite delivery chain, its slicing and its provision "as-a-Service" to multiple tenants, greatly facilitating the concept of Satellite Virtual Network Operators (SVNOs).

Coming to specific technology selection, NFV management platforms are yet in an early stage, so no mature technology exists in place in order to consider its applicability. However, in any case, the compatibility with the ETSI NFV reference architecture [NFVArch] is considered a mandate, as well as the alignment with the current architectural concepts and terminology introduced by the ETSI NFV ISG recommendations. Also, emerging NFV platforms with strong industrial and community support, such as the Open Platform for NFV (OPNFV) [OPN] are likely candidates for integration with satcom.

With regard to *Federated Management and Orchestration,* the perspective of federating and jointly managing satellite and terrestrial infrastructures is indeed the primary focus of this study and thus appropriate federation/orchestration mechanisms are more than relevant. Unfortunately, many of the network/IT orchestration architectures (some of them surveyed in Chapter 2) mostly aim at optimizing cloud computing IaaS services over multi-domain networks. Since incorporating IaaS capabilities in satcom platforms solely for computing services, as aforementioned, does not seem to be a promising perspective, most of these cloud-centric federation mechanisms are of little relevance. Instead, network-centric orchestrators should be targeted, whose primary offering is the network service itself, rather than the computing resources. Such orchestrators should make extensive use of the SDN and NFV technologies, which were already identified as key drivers for satellite/terrestrial integration. The orchestrating platforms developed in T-NOVA [TNOVA] and Mobile Cloud Networking [MCN] projects are examples of such SDN/NFV-enabled solutions. However, no specific strong recommendation can be made in this field, since -to date- there exist neither standardized orchestrator architectures, nor community/industry-led network management projects with sufficient momentum.

Summarizing the above conclusions, Table 11 overviews the recommendations on the applicability and possible integration of terrestrial cloud networking techniques and technologies with satcom.

Table 11. Summary of recommendations on technology applicability

Dimension of Suitability for Integration	Major applicability and interests for satcom	Specific candidate technologies
Infrastructure virtualisation	Partial; limited to NFV and federation enablers.	Openstack
Programmable and Software-Defined Networking	Yes	Openflow (protocol), OpenDaylight (controller)
Network Functions Virtualisation	Yes	ETSI NFV architecture; OPNFV
Federated Management and Orchestration	Partial; limited to orchestrating NFV and connectivity services	T-NOVA or MCN project architectures

4. INTEGRATION SCENARIOS

The aim of this chapter is to elaborate candidate scenarios for the integration of cloud networking technologies with satcom. These scenarios should present a clear technical added-value as well as considerable market potential, as well as facilitate the seamless integration of satellite and terrestrial infrastructures in future networks.

As a starting point, we survey and identify the most prominent use cases identified for SDN and NFV in terrestrial networks. For each of these use cases, we also identify the relevance to the satcom domain.

The next step is to combine the most relevant of the use cases with the major satcom services and technical challenges in order to derive a set of candidate integration scenarios. These integration scenarios, are essentially use cases of cloud networking techniques, especially tailored to the satcom context. These scenarios may either correspond to novel services, or provide significant techno-economic benefits for existing services.

Last, the identified scenarios are assessed in terms of technical added-value, market potential and technology readiness, and a specific subset is selected for further elaboration and experimentation.

4.1. Terrestrial SDN/NFV Use Cases

4.1.1. SDN Use Cases

Thanks to the logical separation of control and data planes and the ability to deploy arbitrary control logic into the network driven by software applications, SDN greatly facilitates network and resource management. Therefore, most of the SDN use cases listed in this section reflect this management flexibility introduced by SDN.

4.1.1.1. Experimentation and innovation

Testing and experimenting of new protocols has been one of the primary use cases of SDN. Especially the Openflow protocol was initially introduced as an enabler for innovation in university networks and other experimental infrastructure [McKeown08]. The decoupling of the control logic from the switch fabric allows researchers to deploy and test new services and protocols not only in software lab emulators, but also in operational networks. Even clean-slate technologies and protocols (e.g. non-IP) can be easily deployed, and also coexist with legacy services in the same infrastructure. In this sense, large-scale experimentation facilities such as the OFELIA infrastructure [OFELIA] or the GENI experimental testbed have been deployed, allowing researchers to run experiments using SDN and especially Openflow.

In a satcom context, however, satellite capacity is always a valuable resource and thus large-scale experimentation of new services using Openflow would incur significant costs, which inhibit use by the wider research community. That is why we consider this use case of partial relevance to the satcom domain.

Nevertheless, the ability of SDN to accommodate new (tested) protocols and services with minimal upgrade overhead is still an important added value for satcom, as explained in Chapter 3.

4.1.1.2. Network slicing and multi-tenancy

Thanks to real-time per-flow management capabilities, SDN greatly facilitates network slicing i.e. the partitioning of the network infrastructure into logically isolated virtual networks (slices).

When compared to traditional network virtualization via static encapsulation, SDN-controlled virtualisation can achieve much higher resource efficiency, due to its dynamicity, and also much faster service setup times. State-of-the-art SDN controllers such as FlowVisor or OpenDaylight act as virtualization middleware ("network hypervisor") and enable the abstraction of network slices as well as their use by applications on multi-tenant basis via appropriate Northbound Interfaces (NBIs). SDN-based network slicing is especially valuable for data centre networks (see Sec. 4.1.1.7. below).

Given that network slicing has been identified as a functionality of prime importance for satellite networking, also driving satellite/terrestrial federation, we consider this use case as of significant relevance to satcom.

4.1.1.3. Service automation

SDN greatly promotes the automation of the network service provisioning lifecycle. This is achieved by adopting SDN-based management in Operational Support Systems (OSSs). Rapid service creation and service changes can be thus achieved through e.g. software architectures such as REST APIs in a uniform and easy-to-use manner.

In addition, this capability allows the development of self-service portals, enabling customers to tailor the network to accommodate their application or service needs. This increases service innovation, flexibility and responsiveness and drives increased service profitability.

Although the long service setup time, caused by the manual procedures which are involved, is a significant drawback of satcom, SDN by itself cannot speed up this procedure significantly. This is because satcom service setup and reconfiguration also involves actions such as terminal registration at the gateway, radio resource adjustment, antenna pointing and calibration etc, which cannot be accelerated by SDN. This is why SDN-driven service automation is considered as of partial relevance to satcom.

However, customer-driven resource management, involving resource-on-demand and QoS offerings, are quite appealing, as discussed in Sec. 4.1.1.4. .

4.1.1.4. Resource management and QoS

Although QoS has not been the initial focus for SDN (e.g. early Openflow versions do not support any QoS-related actions), the agility and reconfigurability introduced by SDN control is seen to greatly facilitate resource management and QoS enforcement within the network. Specific flows can be marked according to dynamic policies and differentiated using e.g. a DiffServ model. QoS policies can be changed on-demand with only minimal setup delay, allowing a great degree of elasticity. This agility greatly facilitates bandwidth-on-demand service models, in which the acquired resources (mainly the guaranteed capacity) can be up and down scaled according to customer requirements and SLAs.

Moreover, in small- and medium- sized networks, centralized SDN control can easily manage all the network elements across the data path, thus making end-to-end QoS assurance feasible.

Similarly, in carrier networks, programmatic controls can be applied on carrier links to request extra bandwidth when needed (for e.g. disaster recovery or backups).

As an example, Figure 29 shows an architecture for QoS provisioning for unified communications and collaboration (UC&C) services, where the resources needed to maintain an acceptable quality of experience for the telepresence service are reserved across the network by a network service application via an SDN controller.

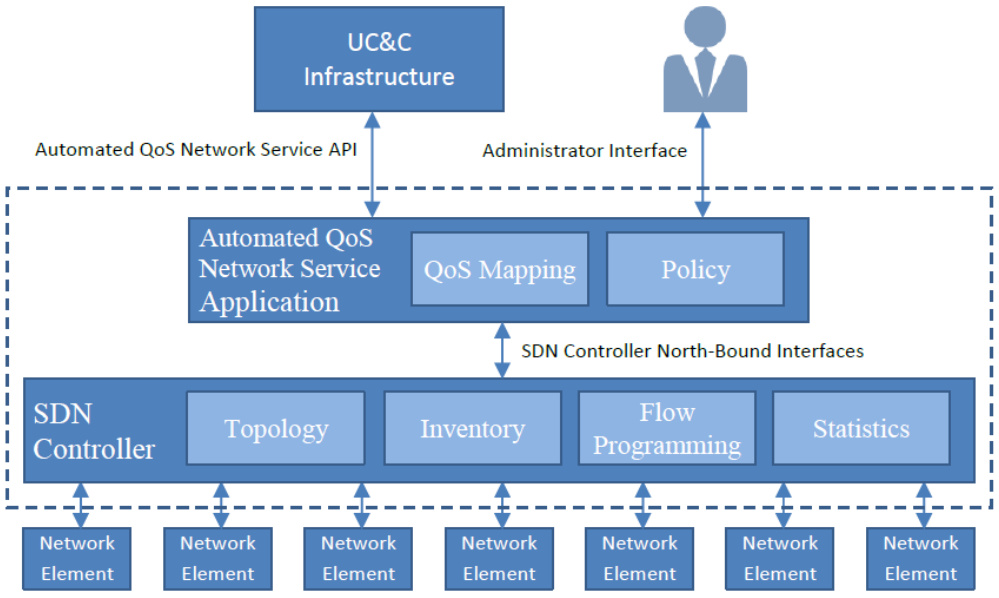

Figure 29. An architecture for SDN-driven QoS management [AQOS14]

This use case is considered of significant relevance to satcom, where QoS provisioning has always been an issue. However, SDN mechanisms need to be coupled with RRM (radio resource management) techniques to ensure that QoS decisions also propagate to (and also take into account the state of) the lower layers (L2/L1).

4.1.1.5. Network monitoring and analytics

SDN greatly facilitates the collection of statistics, such as packet/bit rate or packet loss down to per-flow granularity, from all network elements. SDN controllers can collect these metrics and give network operators the traffic statistics at the exact granularity they need – be it aggregate IP statistics, per-MAC-address statistics, or even per-application statistics. These metrics not only facilitate network supervision and management decisions, but also allow flexible accounting and billing, especially when elastic resource plans (see Sec. 4.1.1.3.) are in effect.

In addition to monitoring based on observable flow metrics, SDN and especially Openflow simplifies the task of network "tapping" i.e. the process of replicating a specific portion of the switched traffic and redirecting it to e.g. a DPI (Deep Packet Inspection) module for on-line or off-line analysis. These inspection modules can be either isolated or even organized in a cluster ("analytics network") [AGC13], as shown in Figure 30. In this context, per-flow network monitoring can go beyond metrics collection and extend to deep inspection of selected flows.

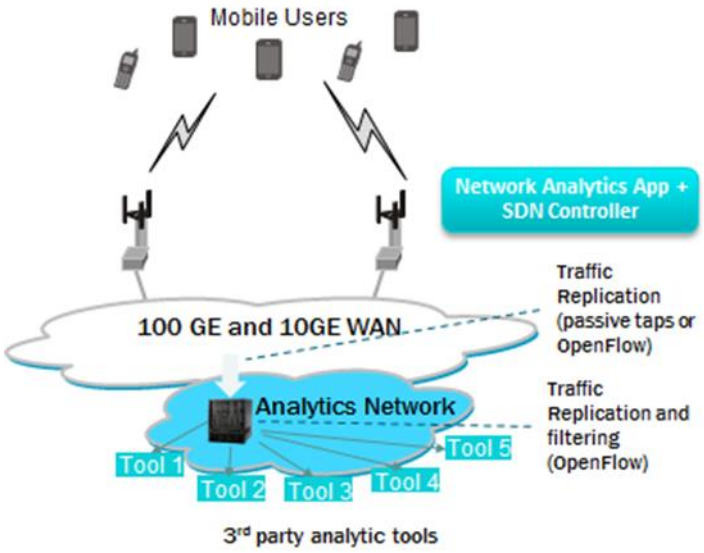

Figure 30. Tapping and analyzing network traffic with SDN/Openflow [AGC13]

All this rich network information exposed thanks to SDN offers a better insight not only into the status of the network but also into the nature of services conveyed. This insight can be used as input to Business Intelligence (BI) processes, allowing the network operator to develop enhanced customer quality systems and offer better user experience.

We consider SDN network analytics as an interesting use case for satcom, especially when combined e.g. with other services such as QoS (See Sec. 4.1.1.4.).

4.1.1.6. Network security

The results of network monitoring –either via metrics processing or flow inspection, see Sec. 4.1.1.5. - can be used to derive specific decisions about how to handle each flow. These decisions can be easily applied via SDN and dynamically changed on-the-fly, resulting in the selected flows being dropped, rerouted and/or QoS shaped.

In this context, (D)DoS mitigation solutions can use traffic statistics provided by SDN switches to detect traffic anomalies and engage traffic redirection/steering capabilities (see Sec. 4.1.1.9.) to divert suspicious traffic to a DoS detection appliance (physical or virtual). Based on the results of this analysis, specific flow entries can be rapidly enforced via SDN into ingress switches, in order to block the offending traffic.

Network security is a critical issue for all network infrastructures, including satellite. However, in a satcom setup typically all the traffic traverses the satellite gateway, and security policies can be applied locally; in this way, the capabilities of SDN to monitor and mitigate distributed threats and attacks across the network cannot be exploited. In this context, the network security UC is only of partial relevance as the added-value is limited.

4.1.1.7. Datacentre network management and interconnection

Management of Inter- or Intra-data centre networks –including private and public cloud infrastructures- via SDN is probably one of the most mature use cases today. Within a data centre hosting a virtualized IT infrastructure comprising hundreds or thousands of Virtual Machines (VMs), the traffic load can be huge, often making the network the bottleneck of the entire infrastructure. Thus, the proper management of the datacentre network, including the proper routing of flows across the hierarchy of switches, is more than crucial.

For this purpose, SDN is used to create location-agnostic virtual networks, across racks or across datacentres. These virtual networks can be rapidly reconfigured, e.g. in case of the migration of a VM across compute nodes (or even across datacentres), also involving the dynamic reallocation of resources. SDN is also used to manage the tunnels established across datacentres towards their federation.

Since it mostly refers to network management to support hosted IT services, this UC is of limited interest to satcom, unless the satcom operator hosts one or more datacentres to offer IT cloud services.

4.1.1.8. WAN management

Wide-area network (WAN) management via SDN is still at its early stage, mostly due to the complexity and high availability requirements for the WAN (which inhibits the adoption of new management paradigms), as well as the scalability issues associated

with centralised SDN management. Yet, SDN brings specific benefits when it comes to WAN management.

For example, SDN can be used to create dynamic interconnects at Internet interchanges between enterprise links or between service providers using cost-effective high-performance switches. The ability to instantly connect reduces the operational expense in creating cross-organization interconnects, providing ability to enable self-service.

Moreover, SDN provides network operators an accurate depiction of network topology and usage and the control to eliminate unnecessary capacity increases and, thus, result CapEx savings. Arbitrary routing and traffic engineering mechanisms can be applied, further improving resource usage.

In the same context, SDN can divert specific flows to specific paths via traffic steering (see Sec. 4.1.1.9.), thus overriding the default behaviour of routing protocols. This can achieve a more fine-grained handling of specific portions of traffic, towards more efficient resource usage and fulfilment of service- and customer-specific requirements.

In order not to interfere with existing services, in production networks, wide-area SDN is often used as an overlay, i.e. using tunnelled traffic and co-existing with traditional network elements.

Another issue in wide-area SDN is associated with the scalability of centralised management; per-flow signalling of globally distributed network elements to and from a centralised controlled induces considerable delays. That is why, in large-scale SDN deployments, a swarm of distributed controllers are used, each of which controls a specific segment of the infrastructure.

Controller distribution commonly follows two strategies; a) either the controllers form a mesh and communicate with peer controllers so as to maintain a distributed view of the infrastructure and the flow rules which are applied, or b) they are organised hierarchically, in multiple levels/tiers. In this case, the lower tiers interact directly with the network elements, while the higher ones achieve the overall coordination, maintaining only a high-level view of the network policies.

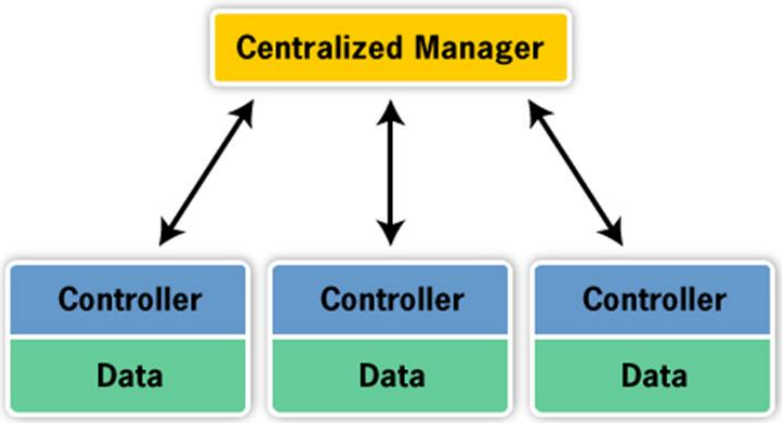

Figure 31. Distributed SDN model

One of the most known examples of SDN application in WAN management is Google's global B4 network [Jain13]. Google uses SDN control in order to cope with some unique characteristics of its internal traffic: i) the massive bandwidth requirements deployed to a modest number of sites, ii) the elastic traffic demand that seeks to maximize average bandwidth, and iii) the need to have full control over the edge servers and network, which enables rate limiting and demand measurement at the edge. B4 (Figure 32) uses a set of SDN controllers to manage each physical site, which in turn control a cluster of Openflow switches.

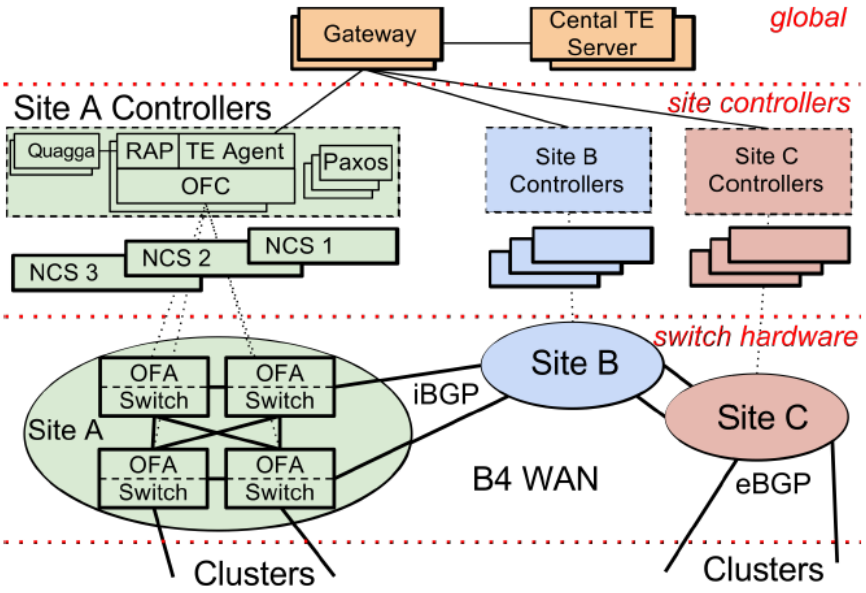

Figure 32. Components of Google's B4 network [Jain13]

WAN management via SDN is considered of particular interest to satcom, especially when used to achieve federation across large-scale terrestrial and satellite networks.

4.1.1.9. Service chaining and traffic steering

In the network context, service chaining is the task of interconnecting several network functions together, forcing the traffic to traverse them in a specific sequence. This is achieved via traffic steering i.e. forcing selected flows to follow a specific forwarding path into the network, even diverting from their "natural" (default) path. This capability allows to create dynamic chains of (commonly L4-L7) services on a per-tenant basis. These chains can be set up with minimum delay and can also be reconfigured on the fly.

Service chaining is closely associated with NFV, in the sense that VNFs deployed at various points into the network need to be interconnected in order to form an integrated end-to-end network service. SDN can help in this, and achieve service chaining with per-flow granularity via rules dynamically installed at critical nodes into

the network. Steering can be implemented either via routing/address rewriting or encapsulation/tunnelling.

Figure 33 depicts a simple example of this concept, where mobile value-added services (VAS) are enabled via redirecting mobile traffic to a data centre hosting virtualised network appliances (VNFs) for video optimisation, caching and parental control.

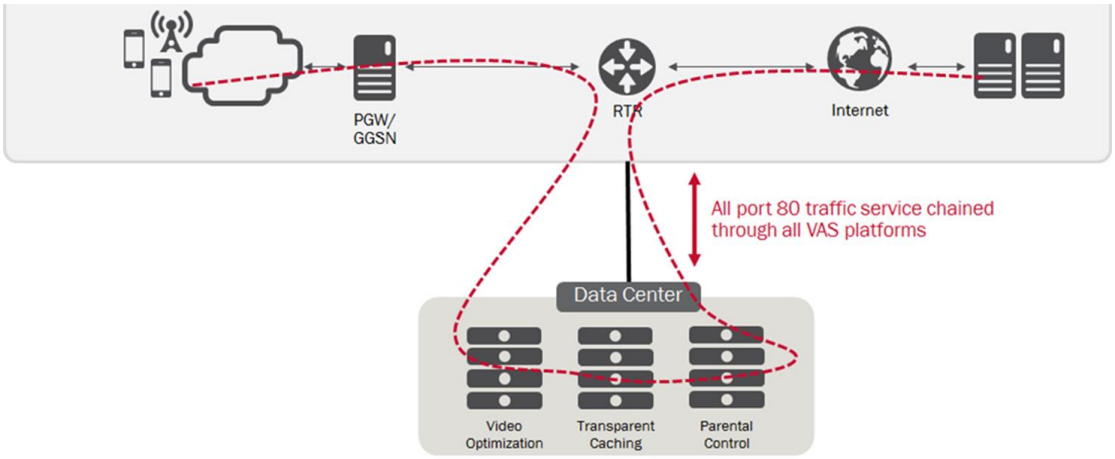

Figure 33. Service chaining/traffic steering for mobile VAS services

Service chaining is considered of significant importance for satcom, since it is a fundamental enabling technology for most of the NFV use cases (discussed in the next section). Service chaining allows per-customer value-added services to be inserted on demand, thus creating novel service bundles, beyond plain connectivity.

4.1.2. NFV Use Cases

NFV provides the capability to virtualise various functionalities within the network infrastructure and thus presents various interesting use cases, some of which are associated with new service offerings, while others only present internal benefits to the network operator in terms of management flexibility and resource efficiency.

The use cases listed below originate from an adapted subset of the ones identified by ETSI ISG NFV [NFVUC]. We have also added a couple of additional use cases, derived from the emerging trends of end-to-end infrastructure virtualisation and Mobile Edge Computing (MEC), which could be of particular interest to satcom.

4.1.2.1. Virtual Network Functions as a Service (VNFaaS)

Virtual Network Functions as-a-Service is probably one of the most promising UCs for NFV, it could be seen as similar to the software-as-a-service (SaaS) offering in traditional IT cloud environments. It targets at customers which want to migrate functionalities from hardware network appliances to their virtualised counterparts.

The customer selects specific VNFs which are used to be included in his/her network service and process his/her own network traffic.

To realise this service, the NFV service provider initiates VNF instances (virtual firewalls, service classifiers, load balancers etc.) per-customer within a multi-tenant virtualisation environment. The resources of the VNFs are allocated so as to match the requirements/SLA of the customer. Then, using traffic steering mechanisms, the customer traffic is redirected to the data centres (NFVI-PoPs) which host the VNFs.

Figure 34 depicts this use case, showing as an example an enterprise network where specific network functions (Access Router/AR, WAN Optimisation Controller/WOC, DPI and VPN appliances) are off-loaded from hardware appliances to virtualised components using the VNFaaS paradigm.

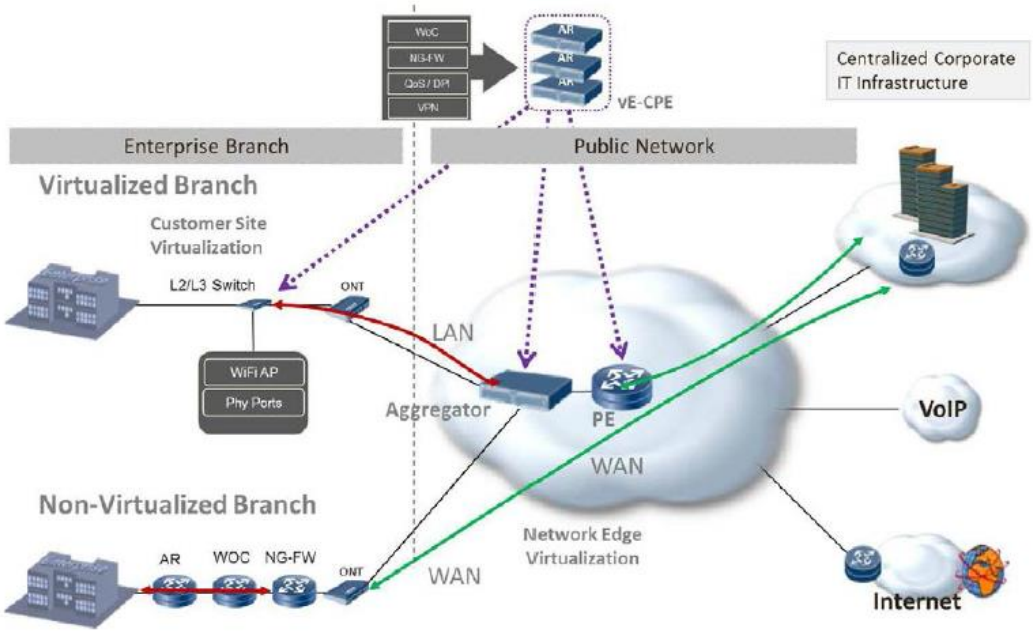

Figure 34. VNFaaS in an enterprise network [NFVUC]

The VNFaaS UC is seen as significantly relevant to satcom, since it provides the capability to enhance the satcom network service offering with added-value VNFs.

4.1.2.2. Virtualisation of core network functions

As opposed to the VNFaaS UC, which is mostly targeted to the customers, this use case refers to the virtualisation of core network functions, to be exploited by the network operator per se. Core network functions candidate for virtualisation include carrier-grade Network Address Translation (NAT), DPI, proxies/caches, as well as mobile network components, such as IMS or EPC modules.

The benefits, from the side of the network operator, are reduced total cost of ownership (TCO) for networking appliances -since the functions are hosted in commodity hardware-, more flexible resource utilisation and elasticity, unified management, as well as the capability to easily maintain and upgrade.

In the satellite context, although the satcom operator could also receive similar benefits from virtualising its core functions, this benefit would be somehow limited, given that these functions are aggregated in few points of presence (normally integrated in the satellite gateway), without many requirements for resource elasticity and frequent upgrades. Therefore, the relevance of this UC to satcom seems somehow limited.

4.1.2.3. Virtual CDN (vCDN) as-a-Service

In most deployments of Content Delivery Networks (CDNs), CDN nodes are hardware appliances integrated into the network infrastructure. This creates some business implications, since the CDN nodes and the network might be managed by different parties (CDN provider and network operator respectively). Moreover, CDN nodes are normally overprovisioned i.e. dimensioned for peak content demand, which results in low resource utilisation.

In the vCDN use case, the network operator offers IT computing resources at several points into the network, where the vCDN providers can deploy their own nodes as virtualised appliances. This simplifies the management and the maintenance of the vCDN overlay and allows for resource elasticity and easy upgrades/modifications. In this context, a vCDN can be deployed at almost zero CAPEX over a virtualised infrastructure.

Given that media content delivery is considered a core business for satcom, the vCDN use case might appear as an interesting perspective, especially when combined with a hybrid access scheme (satellite + terrestrial).

4.1.2.4. RAN virtualisation

As explained in Chapter 2, Radio Access Network (RAN) virtualisation refers to the migration of L2/L1 operations carried out in mobile base stations/access points to software entities (VNFs). This means that functionalities typically performed by hardware, such as coding/decoding, modulation/demodulation, MIMO processing etc. are now off-loaded to VNFs. The concept is similar to Software-defined Radio (SDR), with the addition of multi-tenant features (several virtual base stations hosted in a physical infrastructure), resource elasticity and reallocation among virtual BSs, as well as unified management using common interfaces.

RAN virtualisation is expected to provide advantages such as lower footprint and energy consumption thanks to dynamic resource allocation and traffic load balancing, easier management and operation, as well as faster time-to-market for new RAN protocols.

In the satcom context, RAN virtualisation should be considered only for the long term. In addition, radio technologies, especially in the Forward Link (FL) are already quite spectrum-efficient and operate close to the Shannon limit (e.g. DVB-S2), so no major advancements in the short-term should be expected to justify the use of RAN virtualisation for rapid technology upgrade. However, the concept of the sharing of

the satellite capacity across several virtual operators using virtual radio head-ends seems an interesting approach, especially when used for end-to-end infrastructure virtualisation (see Sec. 4.1.2.7.)

4.1.2.5. Virtualisation of customer premises equipment (vCPE)

This UC is quite similar to the more generic VNFaaS scenario (see Sec. 4.1.2.1.), in the sense that it involves VNFs instantiated in the network infrastructure to be offered to customers. However, while the VNFaaS scenario mostly refers to enterprise customers, the vCPE case is more oriented to residential or SOHO use.

In this case, the virtualisation target is the functionalities commonly hosted in home gateway, which, apart from traditional networking functions (firewall, NAT etc.), also include content management (such as parental control) as well as media handling operations (media recording/PVR, transcoding, indexing, content creation and sharing etc.). All these operations are quite resource-intensive and also rely on rapidly advancing technologies, which justifies their migration to VNFs.

In the vCPE scenario, all vCPE functionalities are deployed as VNFs in an NFVI-PoP as close to the customer as possible. The equipment which is finally installed at the user premises only provides basic L2/L3 connectivity, since most of the intelligence has been off-loaded to the virtualised infrastructure.

This UC is considered as significantly relevant to the satcom domain, since it provides the capability to augment the basic networking functionalities commonly provided by satellite access terminals with several added-value services.

4.1.2.6. Edge traffic processing

This UC is inherited from the more generic concept of Mobile Edge Computing (MEC), which is increasingly gaining support by the mobile, network and IT communities. As outlined in [ETSIMEC], MEC assumes the deployment of virtualised IT assets (“edge cloud”) into the access part of the network -either wired or wireless-, hence the name “edge”. These IT assets are used to deploy virtualised components (virtual machines) which can either process traffic or even host user applications very close to the user. The benefit is that the service latency is minimised, and also the network backhaul is relieved from excess traffic to remote servers. Furthermore, edge applications and network services have direct awareness of the network context (e.g. radio conditions, network statistics), so that they can easily self-adapt to maintain a consistent Quality of Experience.

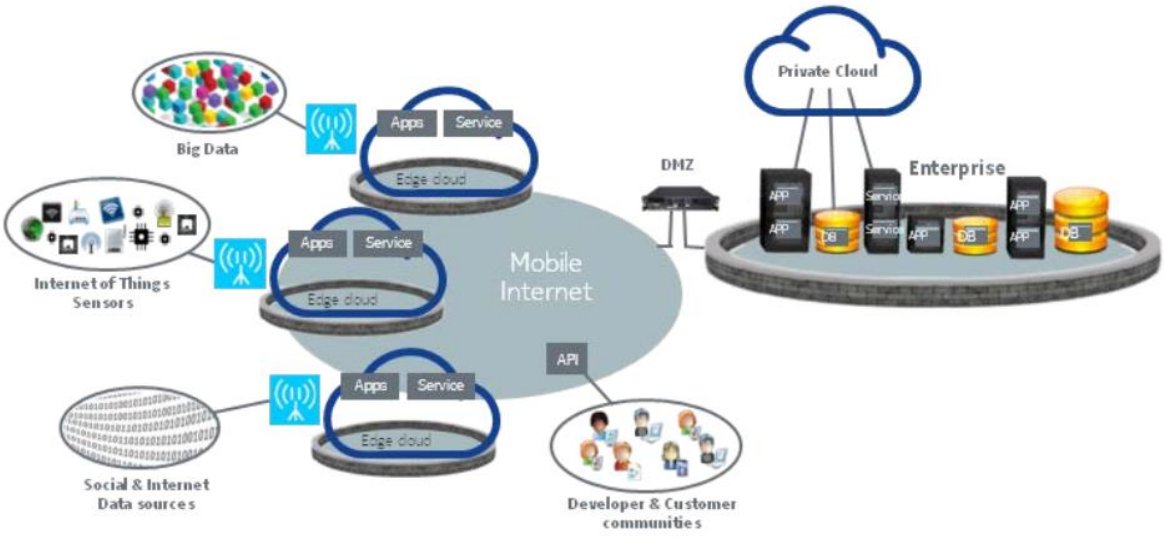

Figure 35. The "edge cloud" concept, as an enabler for edge VNFs [ETSIMEC]

It is obvious that MEC has a strong overlap with NFV, in the sense that it enables VNFs to be deployed at the edge of the network, processing the traffic shortly before it is eventually delivered to the customer.

In this sense, some promising scenarios of edge traffic processing are: content caching, network-aware content optimisation (e.g. video transcoding), M2M traffic processing, DNS caching etc.

Edge traffic processing could be of significant benefit for satcom, especially if the "edge cloud" is deployed at the satellite terminal, given that the latter is used as a gateway which redistributes the service in a local network – or used for backhauling. In this sense, edge processing capabilities might significantly mitigate the satellite service delay and also relieve the satellite communication channel, thus saving valuable capacity[9].

4.1.2.7. End-to-end network infrastructure virtualisation

This UC is actually a composition of the core functions virtualisation (Sec. 4.1.2.2.) and RAN virtualisation (Sec. 4.1.2.4.), glued together under unified management and supported by SDN-driven network virtualisation. The aim is to support the Virtual Network Operator (VNO) scenario and to further augment it. While VNOs are typically offered just capacity slices over the network infrastructure, the end-to-end NFV-based virtualisation offers to VNOs a complete virtualised instance of the entire infrastructure as-a-Service, including core functions as well as virtual RAN head-ends. In this way, the VNO has almost full management capabilities over the virtual infrastructure, from the core to the access, also allowed to fine-tune radio parameters.

[9] This can be achieved e.g. in a MTC scenario, when the edge processor locally processes bulky sensor data and sends back over the satellite only a small set of extracted features.

Such a concept is promoted by e.g. the Mobile Cloud Networking project [MCN] which elaborates a management and virtualisation framework able to completely slice a mobile network and offer virtualised instances with enhanced management capabilities. All the components of the mobile network, including EPC core, monitoring, registries, network links, base stations (eNodeBs) are virtualised, deployed as VMs and offered as a service (Figure 36), even across different physical infrastructure domains.

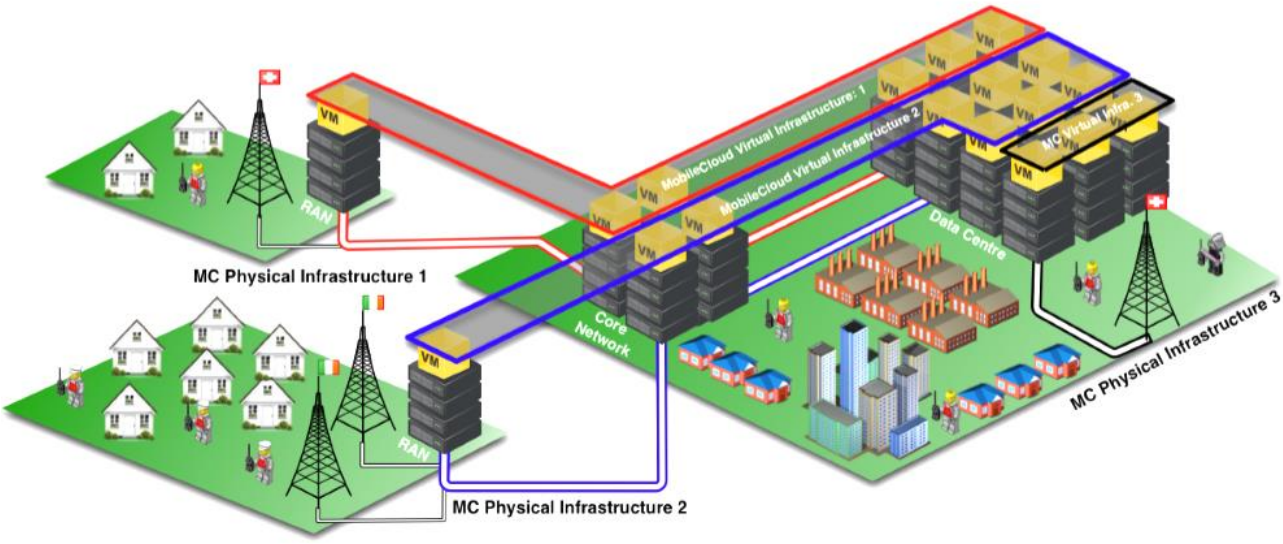

Figure 36. Virtual mobile networks as a Service [MCN]

This UC is considered as of significant interest to satcom, since it augments the already emerging Satellite Virtual Network Operator (SVNO) concept and allows the deployment of virtual satcom infrastructures with very low setup delay and under minimal CAPEX.

4.2. Integration Scenarios

The scenarios elaborated in this chapter correspond to promising cases for the integration of cloud networking techniques into satellite networks. These integration scenarios have been derived using concepts from the terrestrial SDN/NFV use cases identified in the previous chapter, and adapting them to the satcom context. We have also taken into account the services with currently the highest market share for satcom, i.e. content delivery, broadband access and M2M, and oriented the integration scenarios to correspond to these services.

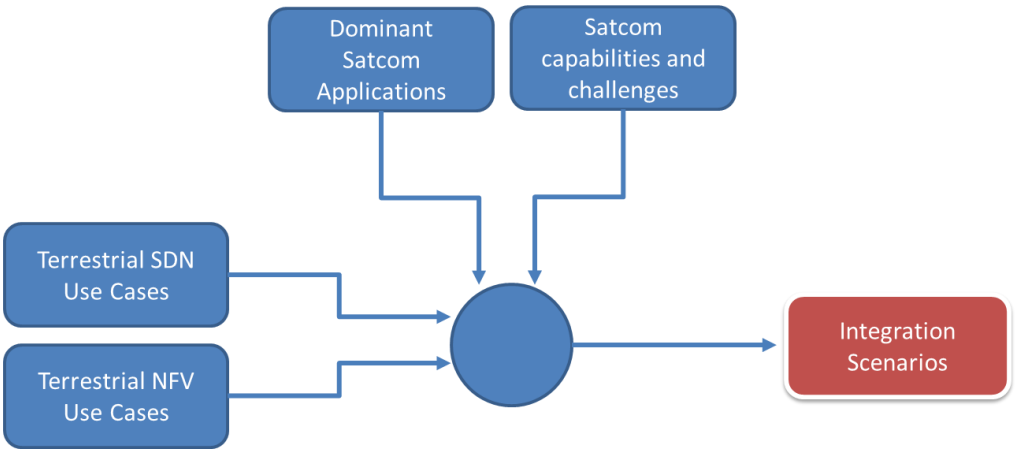

Figure 37. Methodology for deriving the integration scenarios

For each integration scenario, we identify:

- the actors/roles involved;
- the high-level description[10];
- the technical added-value for satcom with regard to existing services and technologies;
- the aspects and challenges associated with the implementation of the scenario, also including an assessment of the readiness of the required technological framework;
- the market potential.

The technical added-value and market potential of each scenario is only described in high-level; a more precise estimation of the techno-economic efficiency for the selected use cases will be attempted in Chapter 7 (Analysis of costs effectiveness, economic gains vs. constraints)

Regarding the value chain and the business roles involved, Figure 38 below depicts a generic model including most of the roles which are associated with satellite/terrestrial cloud network service offerings.

[10] The technical details of the implementation of the identified scenarios are not included in this deliverable; instead, they will be discussed in TN3.1 “Definition of Integrated Cloud Networking Architectures” for the selected use cases.

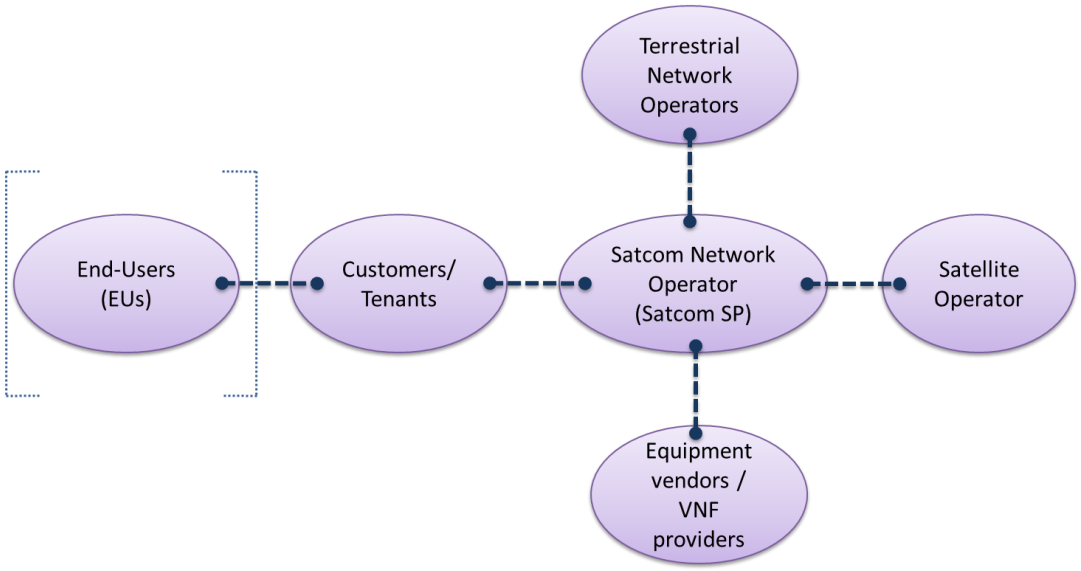

Figure 38. Generic value chain for satellite/terrestrial cloud network services

Satellite operators offer the satellite platform as well as the raw capacity to be used for the establishment of the satellite network. In most cases, the application of cloud networking techniques is transparent to them.

Satcom and Terrestrial Network Operators/Service Providers (SPs) possess a virtualization-capable networking infrastructure, able to offer cloud network services. SPs fulfil the customers' service requests by allocating and orchestrating infrastructure resources in order to compose the virtualized service.

Customers or Tenants are the "operators" of the virtual tenant service. Commonly, Customers establish SLAs with the SPs for the desired service level and have specific management, control and monitoring rights on the provisioned slices. In case of federated satellite/terrestrial services, Customers maintain a unified view of the provisioned slice, regardless of the multiple infrastructure domains on which it may be built.

The Customers may exploit the network slice for own internal use (e.g., in the case of an enterprise user establishing a corporate VPN). Moreover, Customers may also in turn act as Service Providers themselves and exploit the slice for offering a service to their customers (e.g., in the case of a content provider leasing the slice to distribute an IPTV service). In this case, the model also includes *End-Users (EUs),* who receive the application/content over the slice. The existence of the slice is totally transparent to the EUs, who interact only with the offered application/content.

Finally, in a cloud network model, the role of the *Equipment vendor* is expanded in order to encompass also the *VNF providers*, i.e. the developers of virtual network functions, which constitute crucial components of the network service, along with their hardware counterparts.

Chapter 7 will further elaborate on the precise value chain configurations, as well as the business interfaces, corresponding to selected integration scenarios.

4.2.1. Scenario #1: Elastic bandwidth-on-demand

The aim of the elastic bandwidth-on-demand scenario is to augment the typical satellite broadband access service with the capability offered to the customer to dynamically request and acquire bandwidth and QoS, in order to fulfill the requirements of specific time-critical applications.

4.2.1.1. Actors and roles

Being probably the simpler of the scenarios described in this chapter, the elastic BoD service does not involve any modification on the typical satcom value chain; all interactions still take place between the Customer and the satcom Service Provider (Network Operator). The satellite operator, if a separate entity from the SP, does not actually participate in the scenario, since the elastic BoD capabilities are transparent to him/her, given that the allocated satellite capacity remains constant.

4.2.1.2. Description and added-value

Bandwidth-on-demand is not a new term in satcom; since many years, service providers have provided the option to acquire satellite network capacity for a specific time window. This bandwidth is commonly been provided on-demand, without the need to schedule or book in advance, while the service is charged according to the traffic conveyed (e.g. per MB).

Typically, the BoD service has been targeted at customers who require satellite-based delivery of mission-critical services on demand, from fixed or frequently moving sites. This service has been proved satisfactory for certain segments of the market, particularly military, homeland security and broadcast customers, in establishing initial communications facilities from remote locations.

The elastic BoD (eBoD) scenario aims to radically simplify service provisioning and elasticity and thus enhance and expand the BoD offering, beyond institutional users, also to home customers. In this sense, a typical user with a broadband “always-on” satellite connection can easily acquire network resources on-demand at minimal cost.

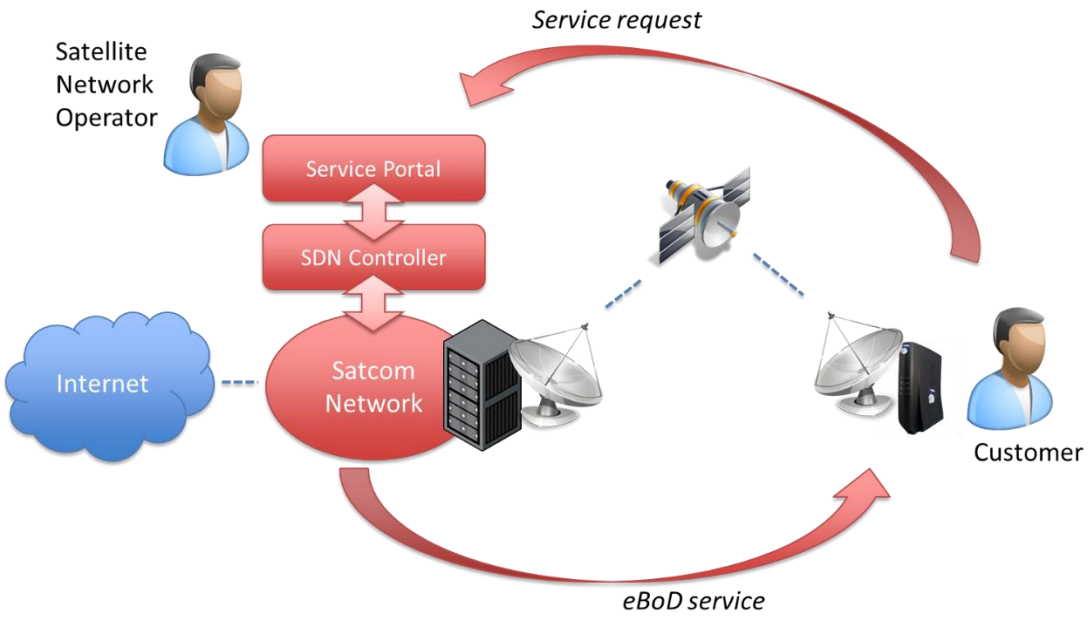

Figure 39. Elastic Bandwidth-on-demand scenario

Depending on his/her needs, the customer may require one of the following:

- **Dedicated bandwidth (typical BoD offering)** in this case, he/she is guaranteed a specific throughput (e.g. 500 kbps), which can be used as desired, for a specific time window.
- **Dedicated bandwidth for a specific application**; in this case, the guaranteed bandwidth and also class of service is applied only to a specific application. For example, whenever a VoIP session is initiated, it should be allocated a guaranteed 64 kbps of bandwidth and much lower delay than other applications. This differentiation may be requested for a specific time window, or alternatively can be statically configured and be constantly active throughout the customer subscription period.
- **Relative traffic precedence**; in this case, the customer is offered a higher precedence compared with the other customers in the multiplex. The differentiation of service classes into e.g. "Gold/Silver/Bronze" or other similar classes, can be applied. Nevertheless, no specific throughput or latency guarantee is given; only a relevant priority is granted. Again, this precedence can be requested for a specific time window, or for the entire customer subscription period.

The customer requests can be submitted in a self-service web portal, via which the customers log in and adjust their service configuration, during service operation. This kind of automation facilitates the roll-out of self-configurable services and responds to time-sensitive changes in bandwidth requirements.

Accounting and billing follows the same concept; depending on the pricing models supported by the service provider, the customers may be charged a flat rate plus an extra rate whenever they request some sort of QoS, or, alternatively, may completely

move to the pay-as-you-go model, in which they are billed only for the resources they consume (e.g. cost per MB, according to the service class). In all these cases, pricing models from IaaS computing services can be exploited, which have been proven adequately efficient in trading elastic resources.

In this sense, the elastic offering gives the satcom operator the potential to escape from the flat-fee model and actually monetize on the offered resources. In addition, it offers a direct competitive advantage since the elastic plans might be especially attractive for certain customers.

The technical added-value of this scenario is assessed as Medium, providing measurable added-value compared to current BoD offerings.

4.2.1.3. Implementation aspects and challenges

The elastic BoD scenario requires that traffic control, inspection, prioritization and also metering capabilities are present at both the satellite gateway as well as the terminal. Although these operations can be achieved via traditional networking technologies, their implementation via SDN offers much more flexibility and agility. For this purpose, SDN capabilities are assumed to exist at both the gateway and the terminal, both being managed by a centralized SDN controller, normally placed at the Gateway. The controller has a global view of the satellite network resources at L2/L3 (yet is unaware of the radio resources) and dictates the appropriate flow rules, configuring the precedence and the policing as necessary.

Deploying eBoD in an SDN architecture with a programmatic northbound API allows the service provider to have centralized, granular control over the networking infrastructure. It also enables customers to automatically request dynamic changes to bandwidth allocation and other Quality of Service parameters, either immediately or scheduled in the future. The SDN controller can leverage per-flow management to cost-effectively provide guaranteed performance on a per-connection or flow basis to meet SLA requirements.

Although SDN-driven traffic engineering should be feasible and quite effective to manage services statistically multiplexed in the satellite forward link (FL), where capacity among services and customers can be easily repartitioned, for the return link (RL), the approach is not so straightforward. The reason is that per-customer RL throughput also depends on the radio resources assigned. Therefore, in order to provide also guarantees on the RL, it is necessary to couple SDN control with radio resource management. In this case, a vertical cross-layer integrated management architecture should be considered as most appropriate.

When it comes to per-application granularity (for assigning precedence e.g. to VoIP services etc.), the application identification needs to be based on simple rules, in order to be carried out through SDN only. That is, applications should be differentiated based on source/destination IP address/domain name, protocol or port. That is, it should be possible to give precedence to e.g. UDP VoIP traffic using a specific port, or to HTTP traffic from a specific site. In case that multiple applications running over the same protocol and to the same destination need to be differentiated

(e.g. HTTP video streaming vs. plain HTTP web page download), then deep packet inspection methods need to be applied, which, although feasible, significantly complicate the implementation. The DPI-based service classification scenario is included in Scenario #8 – Customer functions virtualization (See Sec. 4.2.8)

Last but not least, another challenge is associated with accounting and billing. The satellite Network Management Center (NMC) and OSS/BSS system should be able to accommodate the elastic resource provisioning and charge the customer accordingly, as also highlighted in [Bertaux15]. Although service metering data are typically directly offered by the SDN controller, the elastic billing paradigm should also be supported by all organizational processes of the satcom provider.

The technology readiness of this scenario is assumed as **High**, assuming the most simple deployment with FL resource control only and minimal intervention of the Gateway.

4.2.1.4. Market potential

The market potential for the eBoD scenario is mainly driven by the ever-increasing demand for real-time services, for which eBoD is particularly applicable. Indeed, the continuously rising demand for efficient and affordable communication services across the globe keeps up fuelling the growth of real-time services such as VoIP and video conferences. The real-time services market has experienced continuous growth over the last years due to its unique characteristics such as cost effectiveness, quality of service and rising demand for product and service differentiation by end consumers. This market penetration can be further enhanced by reinforcing the provision of real-time services with elastic provision of resources through the use of satellite networks.

Nowadays, VoIP is considered to be the most mature example of such services having gained a remarkable market share over the last decade. VoIP market has grown in terms of subscribers, revenues and traffic, and keeps up restructuring voice revenues worldwide. According to Point Topic (2012), there were 135.4 million subscribers of VoIP services worldwide in the end of 2011 – 12.6% more versus the 4th quarter of 2010 and 2.8% more compared to the 3rd quarter of 2011 (Figure 40).

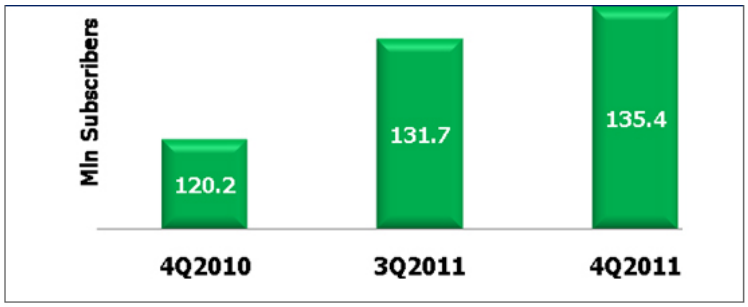

Source: Point Topic, 2012

Figure 40. VoIP Service Subscribers across the world, 2010-2011

According to Infonetics Research (2012), similar rates of growth exist in the mobile VoIP (mVoIP) market mainly due to the increase in sales of smartphones and the active penetration of mobile Internet. In 2011, there were 91 million active mVoIP subscribers in the world (compared to 47 million in 2010) compared to 5.6 billion mobile subscribers across the globe. Thus, the global penetration of mobile VoIP may be estimated at 1.6% of the total number of mobile communication users at year-end 2011. However, taking into consideration the high actual and expected growth rate this indicator should grow significantly. Infonetics Research (2012), forecasts that the number of mVoIP users will grow up to nearly 410 million by the end of 2015, an increase of 4.5 times versus 2011 (Figure 41).

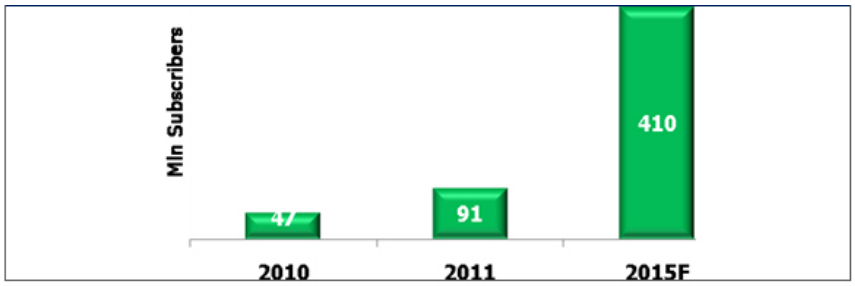

Source: Infonetics Research, 2012

Figure 41. Number of Mobile VoIP Subscribers, 2010-2015

Similarly to VoIP, video conferencing has made its own technological breakthrough as a real-time service in the business world. Nowadays, video conferencing can be used without the technological and hardware limitations of older video-call systems, and can be provided over diverse networks (including satellite) as either over-the-top (OTT) or managed service.

To that end, the eBoD scenario can achieve sustainable quality of experience (QoE) for real-time services over satcom and thus promote the role of satellite as carrier for such services. The exploitation of satellite network characteristics along with the dynamicity and elasticity in resource allocation can achieve a significant technological, financial, business and social impact by allowing high-quality and at the same time cost-effective telepresence applications over satellite.

Concluding, the market potential for this scenario is considered **High**.

4.2.2. Scenario #2: Hybrid media distribution network as-a-Service

This scenario focuses on the federation of satellite and terrestrial domains and the provision of a hybrid satellite/terrestrial access network slice to a media service provider for content distribution.

4.2.2.1. Actors and roles

In this scenario, we consider one or more network infrastructure providers (**Network Operators**) who operate the terrestrial and satellite access network segments[11]. The network operator(s) leverage virtualization mechanisms to partition the network and lease slices to a **Media Service/Content Provider** who is as the NO Customer.. The latter acquires the slices in order to offer media services to **End-users**, who are equipped with either single- or hybrid-access terminals, i.e. attached simultaneously to both satellite and terrestrial access network. Normally, the Customers are assumed to contract only with the media service provider and not with the network operators.

4.2.2.2. Description and added-value

Hybrid distribution of digital media, combining satellite broadcast and terrestrial IP, is a scenario which is gaining increasing attention during the last years, due to the fact that it brings together the best of both worlds: high-bitrate and high-quality 2D/3D broadcast content, coupled with interactive personalized services.

Currently, media streaming, especially over the terrestrial segment, is seen as an over-the-top (OTT) service for media providers. Therefore, the service is distributed in a best-effort manner, without QoS guarantees. Even more, network operators operate just as bit carriers, without being able to actually claim a part of the revenue from the media services which their networks are combined.

Fortunately, contemporary virtualization technologies allow network operators to partition their networks into virtual slices, with specific capacity and QoS, and to offer these slices to content service providers. This capability is promoted more and more via EU and global research efforts, such as the EU project ALICANTE [ALICANTE] as well as novel network management architectures and even close-to-market products.

This scenario extends this concept to also embrace the satellite segment. Virtualisation technologies can abstract the satellite and terrestrial access network and also federate them, so they can be offered to the Media Service Provider as a single logically isolated virtual infrastructure, as-a-Service (Figure 42). In this manner, the MSP can extend his/her customer coverage area, almost without any requirement for upfront investment.

[11] These can be combined into a single entity; i.e. a single network operator owning both the terrestrial and satellite network segments.

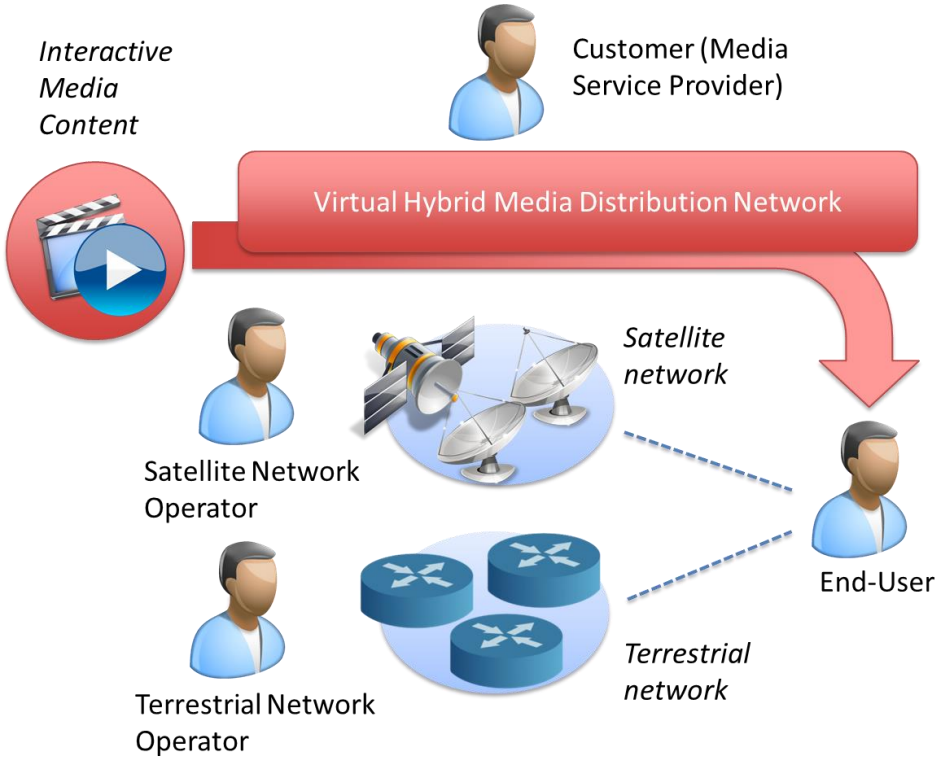

Figure 42. Hybrid media distribution network as-a-Service scenario

In the simplest approach, the MSP (Media Service Provider) just uses the hybrid virtual network as a "dumb pipe" (yet with specific SLA) to convey media streams. However, a significant added-value of the use of virtualization and programmability technologies would be to offer to the MSP elevated management and control capabilities on the hybrid virtual network. This means that the MSP may develop his/her own network control logic in order to dynamically configure the network at runtime, allocate resources and also influence routing/forwarding decisions as desired (i.e. divert streams from the terrestrial to the satellite channel and vice versa on-the-fly or adjust the load balancing between the two networks)

Furthermore, thanks to resource elasticity, the capacity and QoS offered to the MSP virtual network may fluctuate over time, enabling the MSP service to be up and down scaled on-demand or automatically, to react to the customers' demand. This means that the MSP may dynamically request more capacity if needed (e.g. in case of highly popular content)

From the application point of view, three different approaches are considered for hybrid media access:

- *Hybrid broadcast/interactive services*, where the broadcast content is received via satellite, while the terrestrial channel is used for auxiliary and/or interactive content, such as e.g. second screen applications. Standardised technologies such as Hybrid Broadcast-Broadband TV (HbbTV) can be used for the composition of the service at application layer.

- *Scalable media delivery services,* where the media content is transmitted in scalable format. In the most common approach, the satellite network would be considered the primary distribution channel. Thus, the base layer of the scalable media would be transmitted over satellite (and received by all customers, even with satellite-only access). The enhancement layers could be transmitted over terrestrial, so that customers with hybrid access can consume the media content with enhanced quality.
- *Multipath media access*, where unicast media streams are load-balanced between satellite and terrestrial, according to the access capacity of each customer and available resources. In the unicast scenario, the primary distribution channel should be the terrestrial one; the customer would receive the media content over terrestrial and, in case of insufficient terrestrial capacity, a part of the traffic could be diverted over satellite. This could be done via a load balancing service if the media consists of multiple flows, or via a multipath mechanism if a single flow is to be split. In a more advanced scenario, unicast realtime streams (e.g. TV broadcasts), if selected by a considerable number of users, could be decided to be diverted from terrestrial unicast to satellite broadcast, in order to maximize the efficiency and overall utilization of the hybrid network.

The technical added-value of this scenario is assessed as Medium, although this depends heavily of the service to be offered, as well as the hybrid access approach.

4.2.2.3. Implementation aspects and challenges

Software-Defined Networking is a key enabler for network virtualization, partitioning and dynamic control, as required for this scenario. By means of a single or distributed SDN controllers, as described in Sec. 4.1.1.8. , the network operator can manage the partitioning of the network and offer the slices to several MSPs. On top of that, these slices can also be programmable; the MSP can develop an arbitrary SDN application which will control the hybrid virtual network and manipulate/divert the media streams across multiple paths as desired. This is the so-called “SDN as-a-Service” (SDNaaS) service paradigm, which is a significant added-value compared to static, non-programmable virtualization, in which the MSP just uses the offered capacity, without any control capabilities. However, in this case, SDN security aspects need to be taken into account, since the network control applications should neither affect the stability of the infrastructure resources, nor interfere with (possible) other MSPs using the same infrastructure.

The use of SDN brings benefits not only to network control, but also to network monitoring. Sets of network metrics, such as per-flow latency, loss etc. can be provided to the MSP in real time, so that the latter can dynamically decide the balancing of the load between the terrestrial and satellite segment. In any case, the exposure of these metrics should be done in a controlled manner, so as not to affect the privacy of the network operator and to avoid exposure of sensitive data about the status of the infrastructure.

Synchronisation between streams is always challenge in hybrid/multipath media delivery. That is, the delay difference between the satellite and terrestrial path should not exceed specific thresholds, so that the receiver's buffering mechanisms can compensate it. Again, SDN monitoring can help observe the delay of the flows and (in some cases) mitigate the difference in delay by means of network control (e.g. by adjusting relevant flow priorities).

Another challenge is the administrative federation of the satellite and network infrastructure segment, particularly when they are owned by different business entities. In this case, federation could be achieved via a third party, who should establish a "federation umbrella" (super-controller) on top of both infrastructures. However, this might have some implications on the management freedom delegated to such a federation architecture and would for sure restrict the control capabilities which could be offered to the MSP.

The technology readiness for this scenario is considered **High**, requiring federated SDN control with minimal intervention to the satellite network.

4.2.2.4. Market potential

Hybrid media delivery services have currently a significant market momentum, due to the fact that they offer feature-rich broadcast and Internet services either simultaneously on the TV screen or as "second screen" applications in parallel with the typical broadcasting program. HBB technologies and particularly HbbTV enable such application scenarios. In terms of market size, especially for the second screen applications and services, the market continues to grow at a rapid pace (Figure 43) also because of the massive proliferation of smartphones and tablets and the ever present consumer behavior to turn to them in any lull of activity during entertainment (TV or otherwise). The opportunity presented by the second screen phenomena is to capture this opportunity with an engaging consumer experience related to their video experience (companion or viewing).

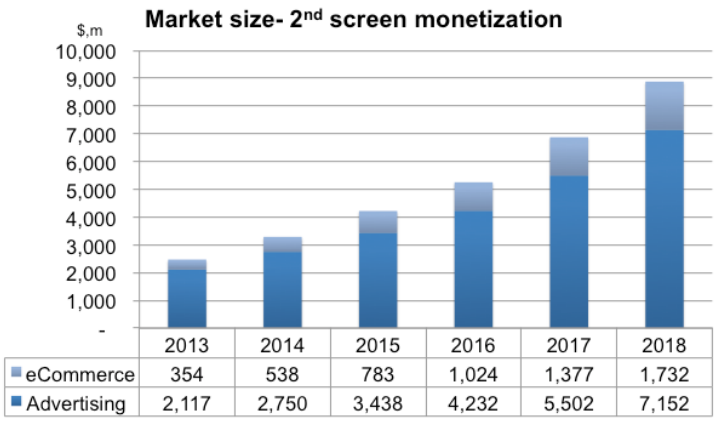

Figure 43. Second screen market size

(Source: NPD DisplaySearch Quarterly Smart TV Shipment and Forecast Report, 2014)

ProSieben channel in Germany reports 1.4M connections of unique TV sets to its HbbTV service per month with a growth of 20%.

Nearly 9.5 million consumer-controlled (open internet access) smart TVs shipped in Q2 2012, in total of 43 million that shipped in 2012. This figure is forecast to grow to 95 million in 2016 (Figure 44)

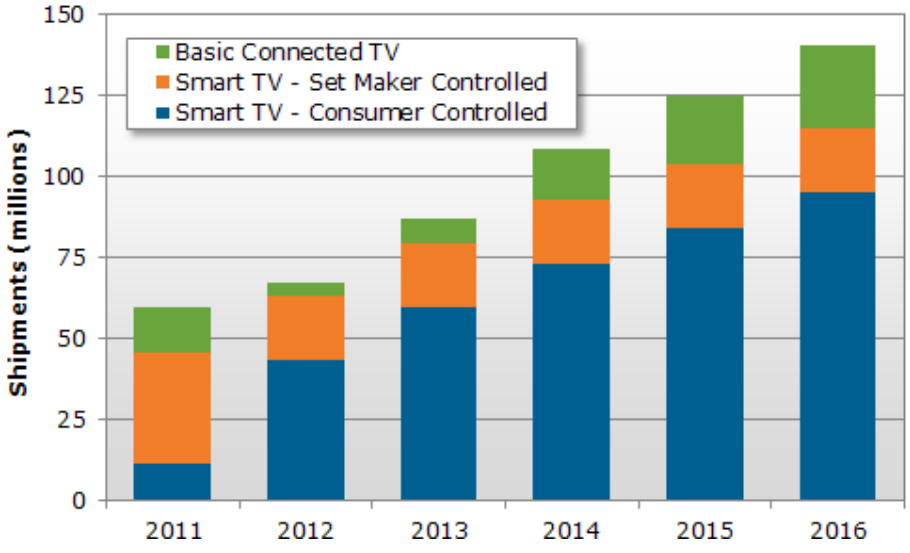

Figure 44. Smart TV shipment forecast

(Source: NPD DisplaySearch Quarterly Smart TV Shipment and Forecast Report)

At the same time, basic sets that link to the HbbTV services will enter at the lowest price points in Europe and Latin America. The challenge for TV manufacturers will be to bring enough value to their sets with extra functions. New open standards such as HTML-5 will help solve the problem of software updates and obsolescence in smart TVs, which should enable such sets to compete with cheap streaming boxes. Towards this, the "Tizen" Operating System is expected to boost the hybrid broadcasting and second screen applications. Samsung is expected to release Tizen-compatible devices in additional Asian markets later this year, including Indonesia, Sri Lanka, and Nepal.

Concluding, the market potential of this scenario can be assessed as **Medium**, depending primarily on the eventual penetration of hybrid broadcast/broadband services.

4.2.3. Scenario #3: Virtual CDN as-a-Service

The Virtual Content Delivery Network (CDN) as-a-Service (vCDNaaS) scenario involves the virtualization, abstraction and offering of slices of the satellite network -enhanced with in-network functionalities such as content caching and transcoding- as a virtual CDN infrastructure, to be used for efficient content distribution over satellite.

4.2.3.1. Actors and roles

The scenario involves a **satcom Network Operator** employing virtualization mechanisms to facilitate the deployment of a virtual CDN (vCDN) service over its infrastructure. The latter is offered as a service to one or more **vCDN Providers**. The **Customers** are the eventual consumers of the content; they commonly have contract with the satcom SP - in this case the CDN service is transparent to them.

Additionally, vCDN provider(s) may in turn offer the content handling service to one or more **Content Providers**. However, the latter are not expected to actually interact with the satcom infrastructure, so their participation in the scenario is rather limited.

4.2.3.2. Description and added-value

Content Delivery Networks are widely used to improve the distribution of content (mostly Web and media) over the Internet, allowing content providers to provide high-quality live and on-demand content to end users with quality similar than –and often superior to- end users. Integrating CDN nodes into networks has been an effective and cost-efficient way to boost customers' Quality of Experience (QoE), mostly by caching content close to the consumers, thus relieving core and backhaul links from unnecessary retransmissions of highly popular content. CDN providers either exploit the CDN infrastructure to deliver their own content, or offer these capabilities as a wholesale service to third parties (e.g. content providers).

Currently, a CDN provider who seeks to extend their coverage using satellite access would have to physically install CDN nodes i.e. dedicated physical appliances into the satellite infrastructure. This installation would require an agreement with the satcom network operator, who would also (optionally) offer some dedicated capacity for the delivery of the content, if network QoS is desired. This traditional approach, besides requiring significant CAPEX from the CDN provider to acquire and install equipment, would be quite inflexible, mainly because:

- Physical devices would need to be over-provisioned to match peak demand requirements
- Upgrades and modifications on the CDN node operations (e.g. updates on video formats, installation of new protocols etc.) would be costly and resource-demanding.

Another very important limitation specifically associated with satellite CDN is that, in the traditional approach, CDN nodes could only be installed in the satellite gateway side (i.e. before the satellite access segment). This limitation would significantly hamper the efficiency of caching, since there would be no saving on the valuable satellite link capacity; cached content would still be served over satellite every time it is consumed. Instead, it would be desirable that caching be also possible after the satellite access, by the satellite terminal. This deployment could also exploit the broadcast satellite capabilities for content distribution in a "push" manner, as described in [VillasenorDC]. However, with the traditional hardware-based approach,

this is particularly complex, inflexible and costly, especially when many CDN providers share the same satellite infrastructure.

Virtualisation technologies promise to alleviate most of the aforementioned limitations by completely virtualizing the CDN infrastructure, as proposed by [NFVUC] and already described in Sec. 4.1.2.3. . The application of the vCDN as-a-Service paradigm to satcom would mean that:

- the vCDN nodes are instantiated as software entities within the satcom infrastructure, while still fully managed by the vCDN Provider like physical devices.
- the vCDN nodes would be able to scale up/down on-demand, rather than rely on statically allocated resources
- the vCDN nodes would be able to be instantiated also at the terminal equipment, thus allowing content caching mechanisms to partially relieve the satellite network from multiple transmissions of the same content – as well as radically reducing access latency for popular content. This approach would make sense when multiple customers are served by a single terminal and would greatly benefit from the inherently broadcast nature of satellite, since popular content could be simultaneously pushed to hundreds or thousands of remote caches and served locally.
- the vCDN provider would very easily deploy (and offer to content providers) additional added-value services, such as media transcoding, content pushing or DRM (Digital Rights Management), in addition to passive caching.
- the vCDN provider would be able to acquire network resources on-demand for content delivery (e.g. bandwidth and QoS on-demand, see Sec. 4.2.1), rather than operating on a best-effort basis. This capability would be particularly useful for maintaining an acceptable customer QoE level during peak hours.

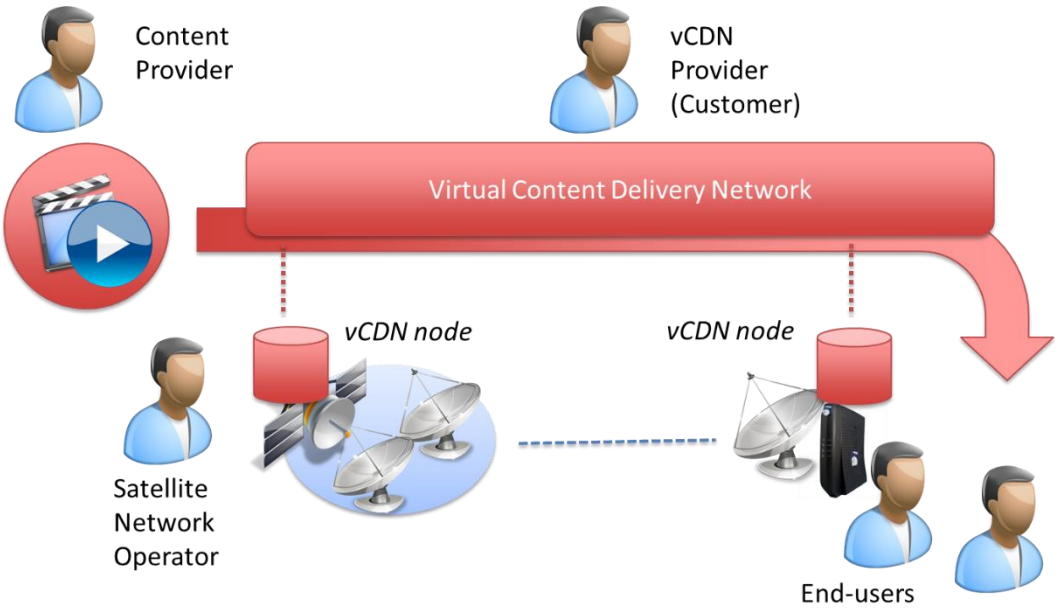

Figure 45. vCDN as-a-Service over satellite scenario

Apart from the vCDN nodes, the centralized CDN controller could also be a target for virtualization. In this approach, the entire vCDN service would be completely virtual and could be deployed with minimal upfront investment.

Last but not least, the vCDN scenario, which was described to apply to a single satellite infrastructure, could be expanded to address multi-domain deployments. In a federated concept (see Scenario #4, Sec. 4.2.4), the vCDN service could span across multiple satellite and terrestrial domains, in order to reach a wide range of customers.

However, since elasticity and on-demand deployment seem not to introduce considerable added value for the vCDN provider, given the relatively limited number of nodes which will be deployed in the satellite network, the technical added-value of this is assessed as Low for the scenario which foresees vCDN only at the GTW side, and as Medium for the scenario which assumes vCDN nodes at the terminal side.

4.2.3.3. Implementation aspects and challenges

Since the virtualization of CDN functions is the core concept of this scenario, NFV appears as the most prominent enabling technology. In order that vCDN functionalities (not only caches, but also transcoders, security appliances etc.) be deployed as VNFs, the satellite network infrastructure needs to be NFV-enabled. That is, the satellite gateway must also feature private cloud infrastructures for VNF hosting and management. Moreover, an NFV management mechanism must be in place, supporting among others multi-tenancy, i.e. allowing each vCDN provider to manage his/her own vCDN nodes.

Additionally, if network resource management is also desired i.e. elastic Bandwidth-on-demand and QoS for content delivery, then SDN-based network control would greatly assist, as described in Sec. 4.2.1.

If the virtual resources (both computing and network) allocated to each vCDN providers are not fixed but dynamically resized, then appropriate metering/accounting/billing mechanisms should be established in order to properly bill the resources used, so that the vCDN provider can be charged in a pay-as-you-go model.

Last but not least, while the instantiation of virtual CDN nodes at the satellite gateway seems quite straightforward, the deployment of vCDN functions at the satellite terminal poses some technical as well as business challenges. Technical challenges are associated with the potentially limited computing resources at the terminal, which need to be carefully managed, especially when shared among various vCDN providers. Business challenges arise when the satellite terminals are not owned by the vCDN provider or the satellite network operator, but by the customer. In these case, the business model must elaborate specific benefits for the customer as a compensation for borrowing local resources in order to support the vCDN service.

For this purpose also, the technology readiness for this scenario is considered Medium.

4.2.3.4. Market potential

Today, multimedia content streaming occupies more than half of Internet traffic. In addition, the growth of cloud networks dramatically increases the amount of content being stored over the web, and forcing virtual CDNs to deal with the constraint of delivering such content as if it was locally stored. Complementary to the growth of today's multimedia traffic, the requirements on quality are also evolving. According to Frost & Sullivan, as video content continues to be consumed more often, for longer periods of time, at higher quality and on more devices, there is an increasing demand for continuously improved video content delivery services. In 2013, a CDN survey by ATLANTIC-ACM unveils the different levels of content type demand (Figure 46)

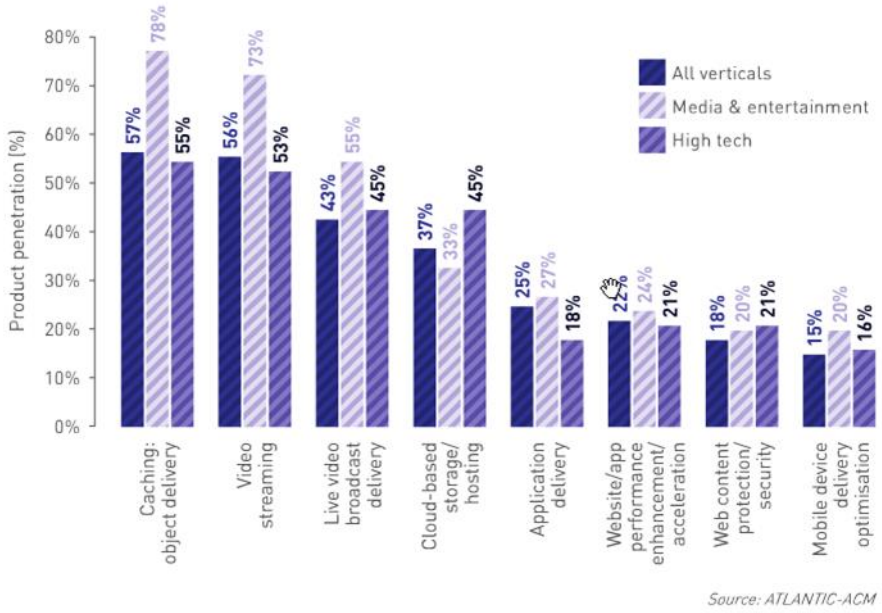

Figure 46. CDN product penetration, 2013

Source: Capacity Journal/ATLANTIC-ACM: CDN Survey (Nov 2013)

Taking into consideration the above evolution and diversity of multimedia services along with the ever increasing user demands, service environments provided by virtual CDN offer many opportunities, from the point of view of service cost and quality, as well as differentiated service provision. To achieve these targets, features such as dynamic resource allocation, on-demand instantiation of content handling and scheduling for the provisioning of virtual CDNs through cloud infrastructures are quite critical.

CDN market and financial dynamics are far more than promising for the years to come and satellite networks can further support such a growth. AccuStream Research (2014), estimates that the CDN revenues touched $3.36 billion in 2014, an increase of 19% from last year. The report, CDN 2014 – 2017: Operations and Analytics, also projects a steady growth in the CDN market from now to 2017, reaching almost $5.5 billion by then (Figure 47)

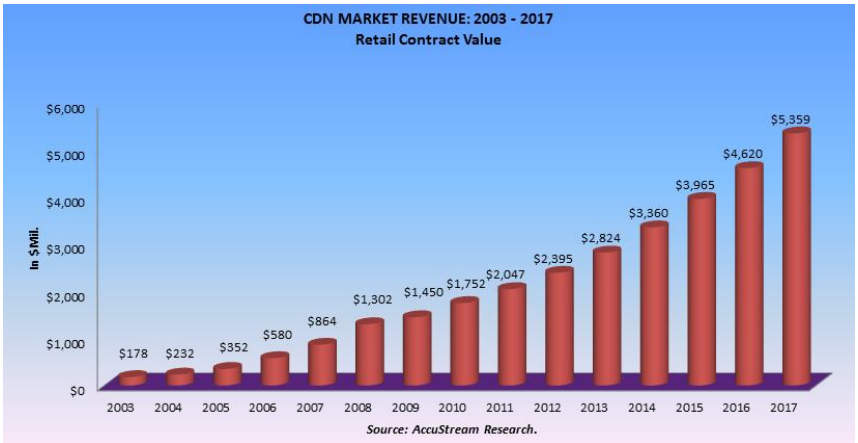

Figure 47. CDN Market Revenue, 2003-2017

Source: AccuStream Research (2014)

According to AccuStream research, the total 2013 commercial value of media and entertainment video (views and advertising), movie/TV files, music listening and downloads (including self-hosting entities such as Google and Amazon) stood at $3.35 billion, of which $1.05 billion (31.3%) was delivered through CDN contracts. Video viewing and advertising (combining self-hosted networks), TV/movies and music alone accounted for 2.4+ billion gigabytes of data transfer, worth $1.6 billion in commercial market value.

Optimization of content delivery over virtual CDNs is certainly a competitive advantage that every CDN provider would strive to acquire in order to strengthen their market position, increase their market share and consequently their profitability. However, in the satellite case, the technical added-value seems somehow limited and thus the market potential is also considered Medium.

4.2.4. Scenario #4: Federated terrestrial-satellite VPN

This scenario is mostly oriented to enterprise or institutional use and assumes a customer with several distributed Points of Presence (PoPs), e.g. headquarters, branches, remote offices, mobile units etc which need to be interconnected into an integrated corporate Virtual Private Network (VPN). We also assume that some of the PoPs are outside terrestrial network coverage, for example in isolated areas or in long-haul routes (ships, airplanes etc.). For this reason, the VPN needs to encompass both the terrestrial and the satellite infrastructures in order to cover all PoPs.

Satellite connectivity is not required only when terrestrial coverage is missing, but also in cases when a backup link is required for redundancy, when the availability requirements are strict (e.g. in mission critical applications etc.)

4.2.4.1. Actors and roles

In this scenario, the Customer is the consumer of the VPN service e.g. the enterprise. The infrastructure for the VPN is offered by the Satellite Network Operator(s) and the Terrestrial Network Operator(s) – although in some cases a single entity might operate both. The VPN service may either be provided by one of the participating operators, or a third party (Virtual Network Provider), who undertakes the federation and integration of the virtualized infrastructures.

4.2.4.2. Description and added-value

VPNs are commonly implemented as logically isolated overlays over the public Internet (or, less commonly, over private networks) and realized via tunneling mechanisms. All VPN endpoints have private IP addresses assigned to virtual interfaces, and appear as if they were interconnected in the same physical network.

The easiest option is to establish a VPN “over-the-top” (OTT), e.g. establish a tunneled communication with one or more remote hosts over the network, without any intervention of the network operator. This approach, although fairly simple, does not provide any service guarantees (QoS, availability etc.) and is not suitable for stable VPN networks, especially interconnecting corporate branches, which have more stringent SLA requirements.

Instead, this scenario assumes that the terrestrial and satellite operators employ virtualization and programmability technologies to offer end-to-end managed VPN services, as shown in Figure 48.

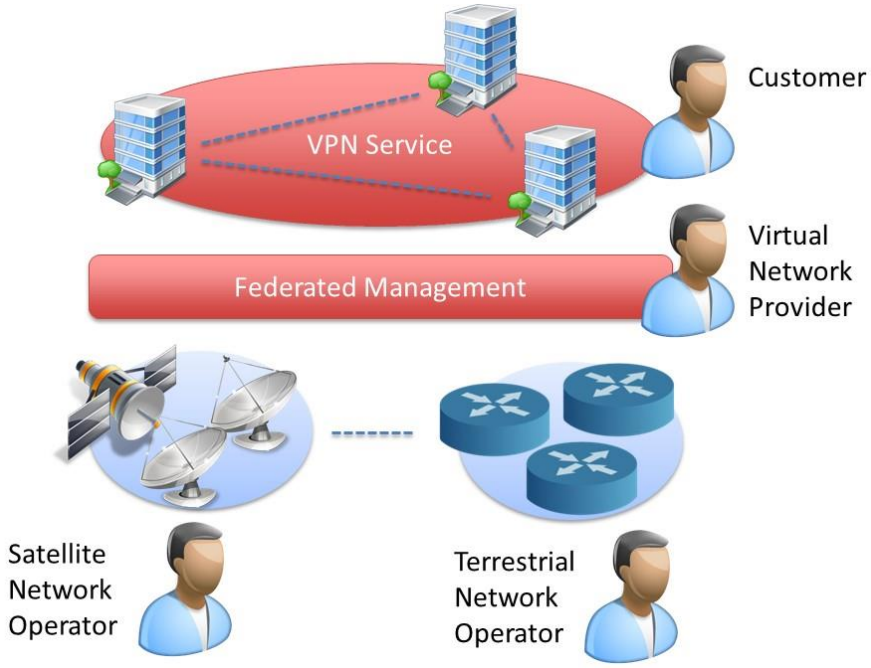

Figure 48. Federated satellite/terrestrial VPN as-a-service scenario

It is true that the establishment of a L2 or L3 VPN via static tunneling, using e.g. VLXAN, GRE or MPLS encapsulation, is feasible without the use of cloud networking technologies and is already offered by most operators. However, the use of programmability and virtualization technologies for the establishment and control of the VPN would bring several valuable benefits such as:

- Rapid setup as well as reconfiguration of the VPN service, with a delay of minutes or even seconds. This capability is especially useful in cases where the service needs to be quickly deployed and/or reconfigured i.e. disaster recovery or high mobility
- Unified control of the satellite and terrestrial domains via standardized protocols, enabling vendor-agnostic setup of VPNs flexibly using also plain IP encapsulation. This capability is discussed in [Gall13].
- Direct mesh routing without the need of a VPN concentrator. Via centralized control, traffic can be encapsulated close to the user and re-routed through the network directly to the peer. This means that the typical “hub and spoke” VPN topology can be avoided, allowing the traffic to be directly diverted to the peer node, increasing network efficiency and minimizing latency.
- Better support of user mobility i.e. dynamic reconfiguration of the VPN network as one of the end nodes roams across networks and its attachment address changes
- More efficient monitoring of the entire VPN service, providing detailed insight of the traffic in all branches of the VPN topology.
- Elastic resource scaling, on-demand allocation and flexible billing (see Sec. 4.2.1)

- Exposure of advanced control capabilities to the customer over the VPN. That is, the customer may (in some cases) be offered the capability to apply some arbitrary flow handling logic over the VPN e.g. block/prioritise/reroute flows etc. This is in line with the "SDN-as-a-Service" paradigm.

Ultimately, using SDN and assuming a very closely coupled satellite and terrestrial network (i.e. assuming that all elements are directly managed by the same controller), it would be possible to realize a VPN even without tunneling, by rewriting the packet headers at the endpoints and properly routing them into the network, keeping the state of the flows. It would be also possible to convey and manage heterogeneous traffic (e.g. IP, MPLS, carrier Ethernet etc.) natively over the same satellite forward link. This approach would avoid the VPN tunneling overhead –and thus save valuable satellite capacity- however it would require SDN capabilities at every node of the network.

This scenario can be combined with Scenario #8 (Customer functions virtualization), in order to enhance the VPN service with added-value in-network functionalities, implemented as VNFs (such as encryption, VoIP PBX etc.)

Overall, the technical added-value of this scenario is assessed as Medium, compared to existing VPNs with static allocated capacity.

4.2.4.3. Implementation aspects and challenges

As aforementioned, even though the VPN service per se can be implemented without any cloud networking technologies via the traditional static manner, the employment of SDN brings significant benefits with regard to service setup, topology management, enhanced control and resource elasticity.

In order to achieve these benefits, several nodes of the network, at least the ones in the Provider Edge (PE) and Customer Edge (CE) domains, need to be programmable and SDN-compliant in order to be able to manipulate the traffic at the service entry points. The VPN service is established and managed by establishing the tunnel virtual interfaces and configuring the appropriate flow rules in order to divert selected flows into the tunnel interfaces.

The end-to-end management of the VPN service requires centralized SDN management by a controller with global view of the network. We assume that each administrative network domain is managed by such a controller (or a cluster of distributed controllers for large-scale networks) and exposes a northbound SDN control interface. This control interface can be leveraged e.g. by a web self-service portal, in which customers can login to setup, configure and monitor their VPN service in an automated manner.

In the case that the satellite and terrestrial segments belong to different business entities, it is required that the different segments are jointly managed by a federating module (federated manager), operated by either operators or a third party. That would be the role of the Virtual Network Provider (VNP). However, the federated scenario raises the common concerns of federated management, associated with the general reluctance of infrastructure operators to expose management and control

capabilities to third parties. However, the evolution of SDN as an enabler for secure multi-tenancy and infrastructure sharing tends to alleviate this reluctance.

Finally, with regard to traffic manipulation, e.g. routing, filtering etc., if the capabilities inherently provided by SDN are not sufficient to fulfill the customers' needs, then NFV would be employed to deploy specialized virtual middleboxes (VNFs) into the VPN service (cf. Scenario #8)

The technology readiness for this scenario is considered **Medium**, although single-domain (non-federated) management can be also implemented in the short term.

4.2.4.4. Market potential

To date, the greatest drawbacks of using VPN over a satellite connection have been the limited bandwidth, increased overhead and the poor QoS due to the high-latency of satcom. The scenario described in this section promises to allow satcom operators to take advantage of the high market penetration of VPN services and boost their growth even more. Specific industries, like the shipping industry, would have a remarkable gain from the use of terrestrial and satellite VPNs since everyday business processes such as vessels positioning and office-vessel data exchange would be improved.

As an indicator of the VPN market potential, a Frost and Sullivan survey (2011) reports that the mobile VPN market has been grown over the last years providing even more revenues to the corresponding service providers (Figure 49)

Figure 49. Mobile VPN Market Revenue Forecasts, 2007-2015

Source: Frost & Sullivan (2011)

According to the same study, a sales break down of VPN Revenues prove the diversity of the market-business activities being involved (Figure 50). Most of these sectors, especially healthcare, utilities and telecommunications, involve very important use

cases for satcom, in the sense that satcom is essential for reaching remote assets (e.g. health centers, power plants and network access points/base stations respectively), but also for providing a failover solution.

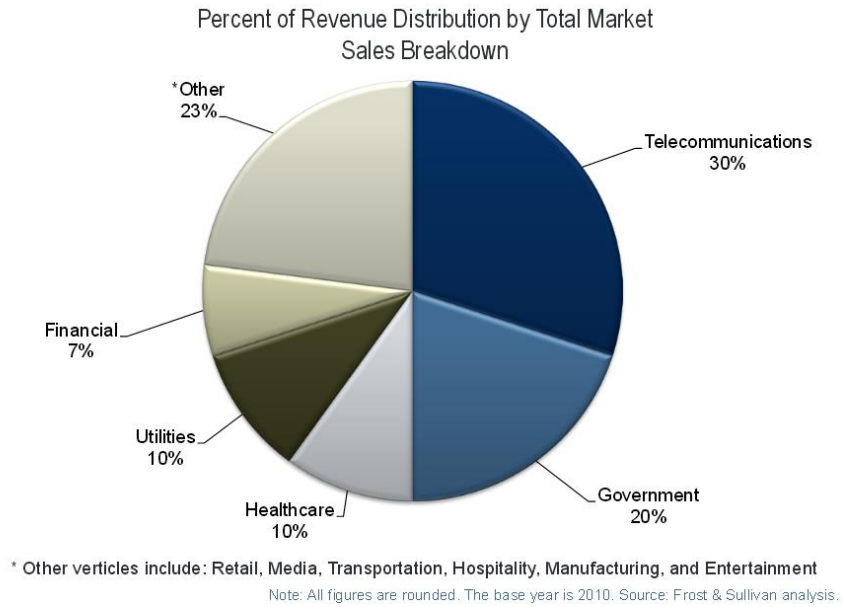

Figure 50. VPN Market Revenues Distribution

Source: Frost & Sullivan (2010)

Consequently, market trends for VPNs are more than positive and markets along with technology seem to be mature enough so as to welcome on-demand provisioning of end-to-end virtual network slices across satellite and terrestrial segments. As an example, Points-of-Presence interconnection for a shipping company, while being served in combination by either satellite or terrestrial networks is certainly a feasible, reliable and promising option which over time may become more robust and more efficient both in terms of cost, performance and QoS.

Despite the demand on the VPN service bundles, the market potential of this scenario is assessed as Medium, mostly targeted to customers who have specific requirements with regard to elasticity, management, monitoring and programmability.

4.2.5. Scenario #5: Satellite Virtual Network Operator (SVNO)

This scenario is inherited from the concepts of virtual network operators (VNOs) in terrestrial wired infrastructures and Mobile VNOs (MVNOs) in cellular networks. The SVNO scenario involves the partitioning of the satcom infrastructure into logically isolated end-to-end slices with dedicated network, IT and radio resources. These slices, in the form of “virtual hubs” are leased as-a-Service to several SVNOs, who are offered full control of the virtual infrastructure, as if it were a physical network.

4.2.5.1. Actors and roles

The main interactions of this scenario take place between the satcom nework operator, who will be called **Satcom Infrastructure Provider** in this scenario to be clearly distinct from the virtual operator, and the **Satellite Virtual Network Operator**, who corresponds to the Customer in this case, leases the slice and consumes the SVNO service. In this scenario, the **End-users** are assumed to maintain relationships only with the SVNO.

Terrestrial network virtualization value chains often also include the role of the Virtual Network Provider (VNP). The VNP uses the resources of the infrastructure provider (InP) to provide the virtualized service to the VNO. However, in a single-domain satcom context, it would make sense to assume that this role is also undertaken by the InP.

4.2.5.2. Description and added-value

With the advent of virtualization technologies and enablers, the concept of Virtual Network Operators (VNOs) and especially Mobile VNOs (MVNOs) is gaining ground, and the VNO business case is becoming more and more attractive.

During the last years, the VNO concept has extended to encompass the satellite segment, and Satellite Virtual Network Operator (SVNO) offerings have emerged [SVNOTT] [iSATVNO]. The DVB-RCS2 technology [DVBRCS2] supports SVNO by dividing the capacity into several logical and independent networks – Operator Virtual Networks (OVN). Each OVN is assigned a set of customer terminals and dedicated capacity, staying logically isolated from the rest OVNs.

By exploiting the virtualization paradigm, the scenario described herein extends the SVNO concept from the plain slicing of capacity, to the full virtualization of the entire hub – i.e. the core gateway and front-end functions, including traffic control (caching, firewalling, PEP etc.), multiplexing, multiple-access and also radio (coding and modulation). Each of these functions are implemented in logically isolated virtualized appliances (VNFs) and are chained together to become components of a “virtual hub”- and eventually of an end-to-end SVNO service.

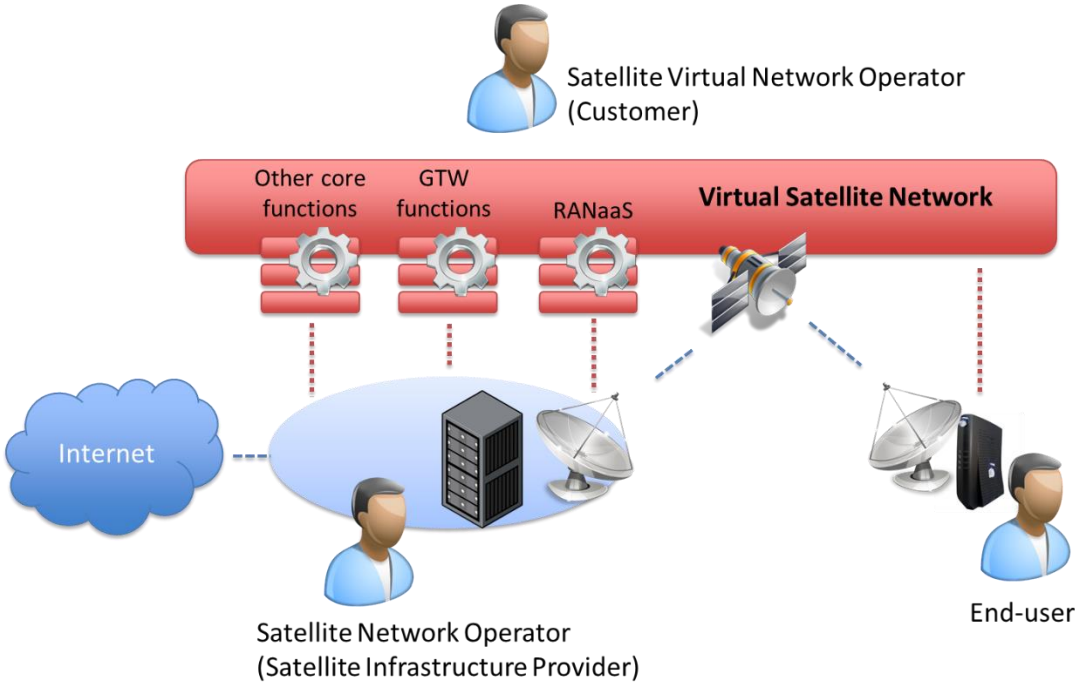

Figure 51. SVNO service scenario

A key added-value stemming from this approach, compared to current SVNO offerings, are the full administrative privileges which are offered to the SVNO, who is able to manage all the virtual appliances involved in the service independently, as if he/she was managing physical devices. For example, he/she could configure the PEP, change scheduler priorities, manage the multiplexing process and even fine-tune the modulation/coding parameters – respecting of course the satellite power and link budget constraints. That is he/she can enjoy (almost) the same administrative freedom as a physical satcom network operator. However, depending on the operating model and also on the technical competence of the SVNO, the latter might decide to outsource some management functionalities to the Infrastructure Provider.

Another benefit, which can be potentially offered to the SVNO under this scenario, is the capability to choose among multiple virtual appliances and combine (chain) them as desired. For instance, the SVNO service could combine the virtual firewall of vendor A with the virtual multiplexer of vendor B and the virtual modulator of vendor C. In this mix-and-match case, it would make sense to extend the value chain to also include the role of the Virtual Appliance Vendors (VNF developers), since they play a more active role in the scenario.

The fast setup time as well as the resource elasticity are also advantages to be considered. According to the traffic served and the customer density and demand, the SVNO might request to scale up or down the resources assigned to the virtual network, however this scaling would not be considered highly dynamic, as explained in the next section.

Last but not least, it would be also possible (although with several technical and business considerations) that a SVNO combines resources from several satcom infrastructures to form a federated virtual infrastructure. In this case, the virtual

network service would span across several administrative domains. This approach would achieve e.g. increased capacity (via bandwidth aggregation from multiple satellites) and/or extended footprints (via exploiting multiple satellites covering diverse areas).

For all these purposes, under several business and operational models, the SVNO paradigm could be suitable for a wide variety of actors, including but not restricted to:

- small data service providers who wish to enter the market with low CAPEX investment,
- terrestrial ISPs who wish to add a satellite "branch" to reach certain customers – or to offer hybrid access,
- M2M service providers who also own M2M application platforms and wish to offer turn-key, and-to-end M2M solutions via satellite,
- large enterprise users who want the virtual network for internal use and seek a service more "owned" and self-managed than the VPN service described in Scenario #4.

Overall, the technical added-value of this scenario should be considered **High**.

4.2.5.3. Implementation aspects and challenges

SDN and NFV appear as key enabling technologies for the SVNO scenario. In order to fully support the SVNO offering, with the capabilities described, the satcom infrastructure needs to be fully SDN- and NFV-enabled.

As described in the previous scenarios, SDN can be used to i) reserve SVNO capacity within the infrastructure ii) establish network tunnels where necessary and iii) implement the service chaining, interconnecting the various virtual appliances of the "virtual hub".

In addition, while current SVNO offerings provide specific –often limited- management capabilities based on protocols such as SNMP or even on proprietary protocols, an SDN-driven SVNO may (optionally) expose an SDN northbound interface for network control; in this sense, the virtual operator can control the service by any standard SDN controller, even developing his/her own control applications, as also highlighted in [Bertaux15]. This capability paves the way towards fully programmable satellite virtual networks.

SDN-based control also means that SVNOs can make the provisioning process of the services delivered to their customers fully automated. Indeed, a provisioning engine can be used to orchestrate and perform all the required configurations via SDN. In other words, services such as the elastic BoD (Scenario #1) can now be offered over the virtual network, rather than the physical one.

In turn, NFV is needed for the virtualization and unified management of the virtual appliances which are the components of the "virtual hub", assuming that all VNFs will expose a common, standards-compliant interface for management

Although the technological enablers are in place, the SDN/NFV-driven SVNO remains a highly challenging scenario. As with any infrastructure virtualization approach, two main considerations are security and resilience. Since the virtual service has the same availability requirements as the physical one, any malfunctions (accidental or deliberate) should be rapidly mitigated –by means of e.g. live migration of virtualized appliances- and should not affect the SVNO services of other tenants using the same infrastructure.

Another challenge concerns the dynamicity of the SVNO resources. Although, thanks to SDN, the resources among the customers within the virtual network can be rapidly reallocated, the scaling of the SVNO service as a whole would be rather limited and would not be assumed to take place often. Especially -in realistic conditions- the RF bandwidth offered to the virtual radio front-end would not be considered a dynamically scalable resource.

Concluding, although L2/L3 logical network partitioning mechanisms are already well-established, the application of the radio access virtualisation concept can only be considered for the long term, and that is why the technology readiness for this scenario is Low.

4.2.5.4. Market potential

The business model of Satellite Virtual Network Operator (SVNO) is based on the concept that the satellite resources of a SNO are leased to one or multiple SVNOs, thus allowing the SNO provider to partition its satellite resources between multiple SVNOs efficiently by delivering dedicated satellite capacities with different levels of QoS guarantees.

The SVNO business opportunities can be deduced from the market share of the Mobile Virtual Network Operator (MVNO), an important role in the telecommunication value chain, considering that the number of MVNOs has significantly increased and is still steadily rising during the last years.

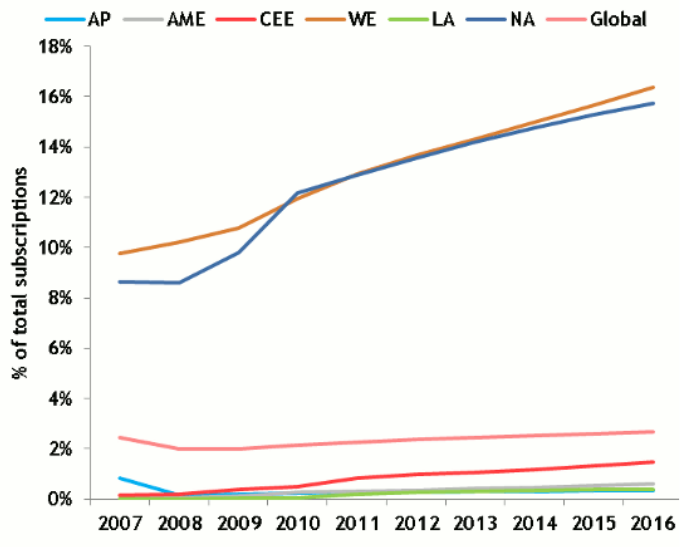

Figure 52. MVNO Market Forecast (2007-2016) as a percentage of total mobile subscribers per region

(Source: Pyramid Research, 2012)

As shown in Figure 52, the number of subscribers per MVNO is expected to keep increasing in the coming years, in particular in the Western European (line WE in the graph above) and North American (line NA in the graph above) markets. MVNOs, in order to deal with this increasing demand, enrich the bouquet of the available services by integrating different functions that can either be handled in-house by the MVNO itself or outsourced to a Mobile Network Operator (MNO), meaning that MVNOs can adopt different operating models, including various operators. Figure 53shows an example of the various ranges of MVNO operating models, which can be also applied to the SVNO approach with minor adaptations.

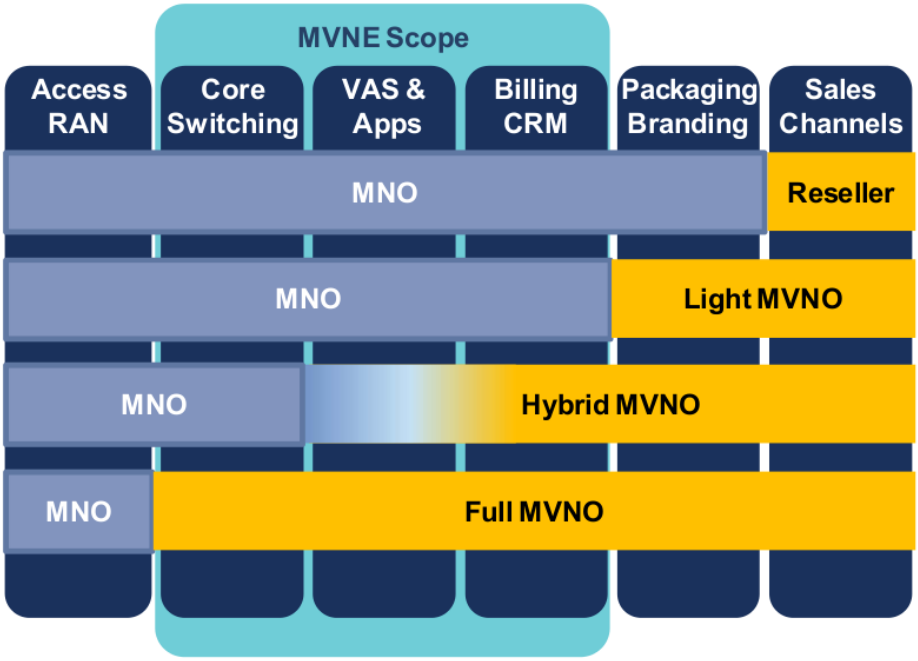

Figure 53. MVNO business models

(Source: Copeland, 2011)

In this context, further expansion of MVNO to the satellite industry may be considered, since it increases further the business opportunities of the VNO and creates novel value propositions by offering enhanced value added service to the end-customers, such as dynamic backhauling.

According to the Satellite Industry Association (SIA), the satellite industry growth brought in $195 billion in revenue for 2013, making it a significant niche market for the business expansion of the VNO model. More specifically, fixed broadband satellite traffic is expected to grow due to a combination of increases in both the number of subscribers and the traffic per subscriber, while mobile satellite services are also in favor by the significant increase of mobile data use.

The market potential of this scenario is considered Medium, expected to further rise in the mid term.

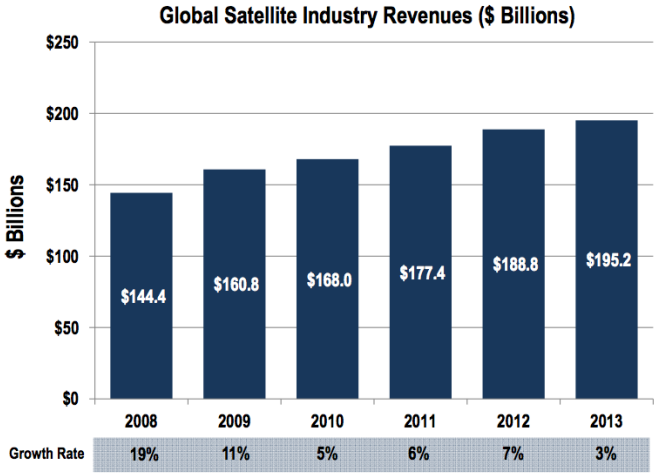

Figure 54. Satellite industry revenues by year from 2008 to 2013

(Source: SIA, 2014.This report is published Sept 2014)

Overall the satellite industry grew by 3 percent, a slight edge over U.S. economic growth of 2.8 percent and slightly more than the global average of 2.4 percent. With close to 1200 satellites operating today, SIA reported that global industry revenues have nearly tripled since 2004, averaging an annual growth rate of 11 percent, which denotes a highly developing market, which together with the already successful MVNO may create significant business opportunities and value propositions.

4.2.6. Scenario #6: Programmable payloads and flexible inter-satellite links

This scenario involves the exploitation of programmable, software-defined mechanisms on-board, allowing to deploy arbitrary routing/switching logic per service, per customer or per flow. These mechanisms can be exploited either in regenerative GEO deployments or in LEO/MEO constellations, paving the way towards truly flexible inter-satellite links. Another step towards payload flexibility will be the inclusion of virtualization-capable computing architectures on-board, which, although a long-term vision, seems to be gaining momentum [DSI15].

4.2.6.1. Actors and roles

The capabilities introduced in this scenario are mostly internal to the Satellite Network Operator, who in this case will most likely be the same entity as the Satellite Operator. However, this is not the case where (in the long term) regenerative payloads are shared among many operators in a “virtual hosted payload” approach. In any case, payload flexibility is expected to offer to the End-users enhanced service QoS/QoE, improved bandwidth and reduced latency and eventually lower service fees due to improved resource usage.

In the future, and if the programmable payload vision becomes a reality, it might be possible to lease slices of the programmable infrastructure to third parties (e.g. SVNOs), who would have full programmability privileges on the allocated slices,

according to the “SDN as-a-Service” (SDNaaS) paradigm. However, this should only be considered as a very long-term perspective, especially for the LEO/MEO scenario.

4.2.6.2. Description and added-value

Although the typical “bent-pipe” configuration, based on transparent payloads, is still the dominant scenario for most satellite networks, On-Board Processing (OBP) capabilities are still attractive in many cases. In GEO configurations, OBP improves resource utilization and reduces latency by allowing on-board switching among terminals in the same or different beams. The value of OBP is further exploited in LEO/MEO scenarios, where inter-satellite links are established to relay the traffic across the satellite constellation.

In most OBP deployments, routing and switching strategies and algorithms are statically built into the payload. This approach significantly hampers the introduction of new routing protocols -and new network technologies in general- into the satellite network. This issue is partially addressed by emerging LEO constellations such as Iridium NEXT [IrNEXT], whose payloads are software-upgradeable. However, and even more important, this static, monolithic approach limits the operator’s flexibility to differentiate the behavior of the network against different services or different subscribers. For example, the constraints of a specific service for low latency or higher bandwidth might require the selection of a different routing strategy, which would fulfill such requirements.

In order to alleviate these limitations, this scenario involves the introduction of network programmability capabilities into the satellite payload, i.e. keeping only the data plane operations in the payload and offloading the control plane (decision logic) into a centralized controller.

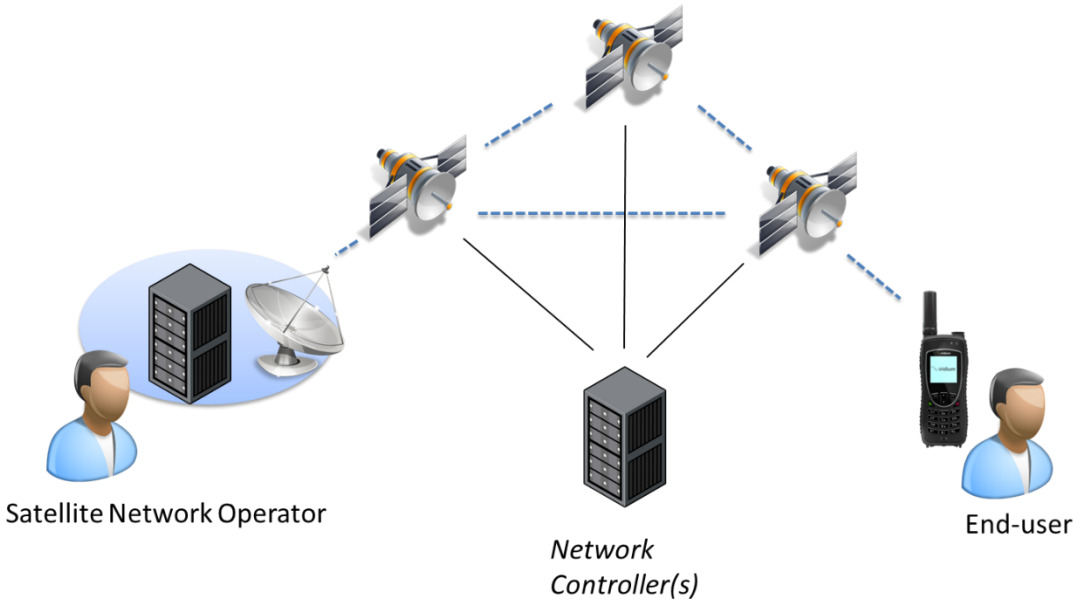

Figure 55. Programmable payloads / flexible ISLs scenario

In this approach, the controller receives notifications from the OBP about new flows and network events, and dictates flow rules i.e. blocking/QoS-shaping flows and redirecting them to a specific next hop, according to customer requirements and application policies.

The SDN controller can be either single (in the GEO case) or distributed (in the LEO/MEO case) and may be deployed:

- within the payloads themselves; this achieves fast control plane-data plane communication, but does not allow frequent reconfigurations
- in the ground gateway network; this assumes that each gateway controls the satellites which are under coverage and the SDN payloads hand over from one controller to the other as they move. In the GEO scenario, a single controller in the single GTW is assumed.
- in an overlay GEO network; this is an even more ambitious approach, proposed by [Bao14]. It assumes that the SDN controllers are distributed in a cluster of GEO satellites, which maintain better visibility to the LEO/MEO constellation, thus minimizing the controller handovers. In turn, the GEO controller cluster communicate with a federated manager on the ground in order to be coordinated.

Whatever the deployment scenario to be followed, the programmable payload has the potential to further augment the already unique capabilities of LEO/MEO constellations, which currently provide the only means for truly global coverage to fixed and mobile users. Programmable payloads are expected to add considerable flexibility and allowing fine-grained QoS-aware management with dynamic resource allocation, thus improving the quality and lowering the cost of the satellite service.

Concluding, the technical added-value of this scenario is considered Medium.

4.2.6.3. Implementation aspects and challenges

Openflow would be the most candidate SDN protocol for the flexible payload scenario, as also proposed by [Bao14]. However, if the SDN controller is not located on-board, then the frequent communication between controller and data plane would induce a considerable overhead in the satellite links. That is why the network control applications should be designed in a way to keep this communication down to the minimum.

Another challenge is associated with the dynamicity of the LEO/MEO network, whose topology changes frequently (yet predictably). Moving user terminals with variable network conditions add even more complexity. Given this dynamicity, it is essential that the SDN controller(s) i) adopt an efficient coordination scheme among them, so as to appropriately accommodate controller handovers and ii) closely interact with the constellation management system, in order to be aware of at least the current location of the satellites, so as to optimize routes.

In addition to technical challenges, there are also certain considerations associated with SDN deployed onboard. Apart from the hardware resources which will be needed (memory, processing, power requirements etc.), a probably even more important issue refers to SDN maturity and stability; some argue that current versions of Openflow are being released too quickly, without allowing for stable product development. This quick market penetration is faced critically from traditional payload manufacturers, whose products are a result of several years of research, development and testing. What is more, OBP payloads are expected to operate for several years without any possibility of manual intervention and upgrade, apart probably from remote software updates. Therefore, it will probably take some years until SDN technology stabilizes, so that SDN payloads can successfully undergo the qualification procedures foreseen for space equipment.

For these purposes, the technology readiness level for this scenario is Low.

4.2.6.4. Market potential

SDN onboard is seen as a mid/long-term approach, yet with clear market potential, targeting at various application scenarios such as aim tracking, dense earth-observation, communication relays etc.

This market potential is emphasized by the growth in the satellite manufacturing market, combined with the exponential demand for SDN equipment. In the first domain, the average of 122 satellites to be launched per year after 2010 is up significantly from the annual average of 77 satellites launched in the previous decade, a sign that government and commercial operators require more satellite capabilities. According to Euroconsult's "Satellites to be Built & Launched by 2019, World Market Survey," the company projects that revenues from the manufacturing and launch of about 1,220 satellites will reach $194 billion worldwide for the decade (Figure 56).

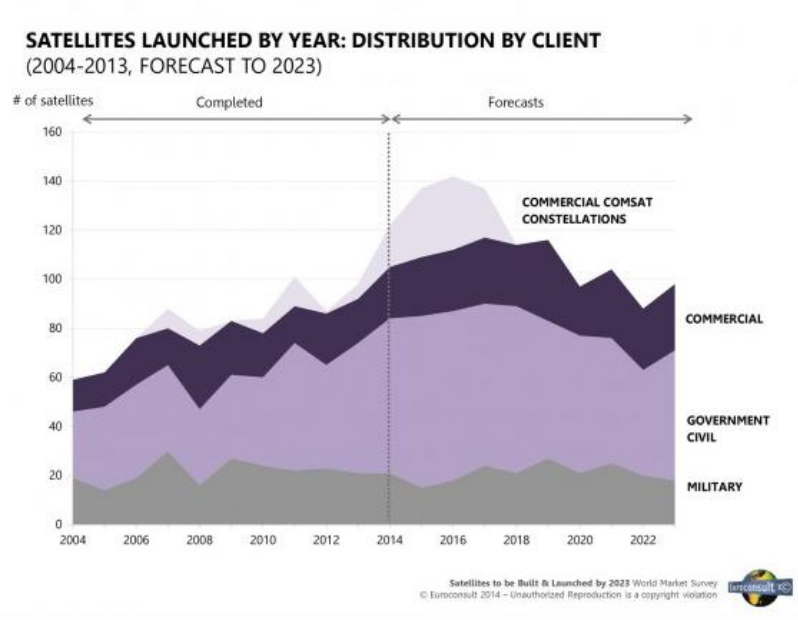

Figure 56. Satellite Market growth per client type

(Source: Satellites to be Built & Launched by 2023, by Euroconsult, 17th edition, July 2014)

Commercial satellite services outside the geostationary orbit will get a boost with a total of 200 satellites to be built and launched into medium and low Earth orbits (MEO and LEO) during the period. Most of them (80 percent) will be communications satellites to replace the first LEO generation operated by Iridium, Globalstar and Orbcomm and to create the first generation constellation of O3b, an innovative system to be launched into MEO. Additionally almost 40 satellites will be launched into low Earth orbit for commercial optical and radar imagery (e.g., Infoterra, GeoEye). Thus, the LEO/MEO platforms maintain a considerable share of the market (Figure 57).

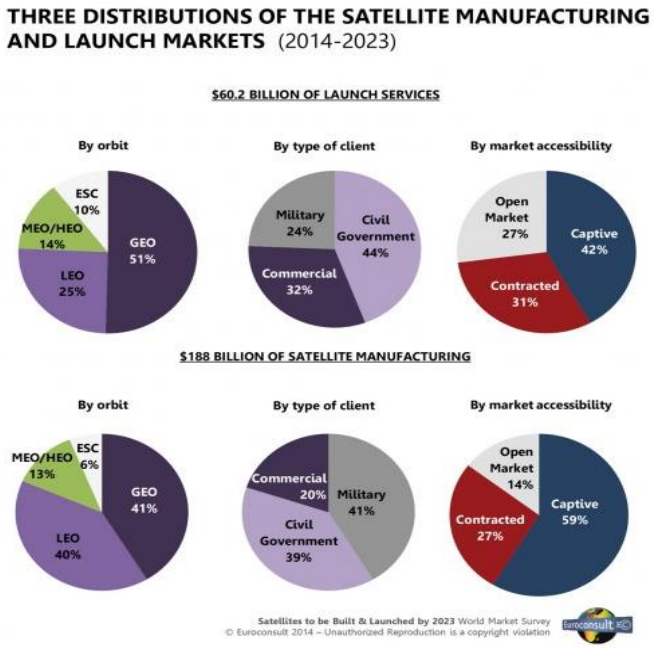

Figure 57. Satellite Market growth of manufacturing and launch markets

(Source: Satellites to be Built & Launched by 2023, by Euroconsult, 17th edition, July 2014)

Insert also the potential impact of the new announced LEO/MEO mega constellation.

Similarly promising is the market growth of SDN services. According to 2013 SDN market size forecast by Plexxi and Lightspeed Venture Partners, the SDN market is growing much larger and much faster than anyone had anticipated. The numbers cited show the SDN market size close to $3.5B in 2014 and it is anticipated that the market will grow to nearly 3x that size by 2015 and to a full 10x that size by 2018.In specific, the impact of software-defined networking (SDN) will exceed $25B per annum by 2018, and could grow as high as $35B annually (Figure 58).

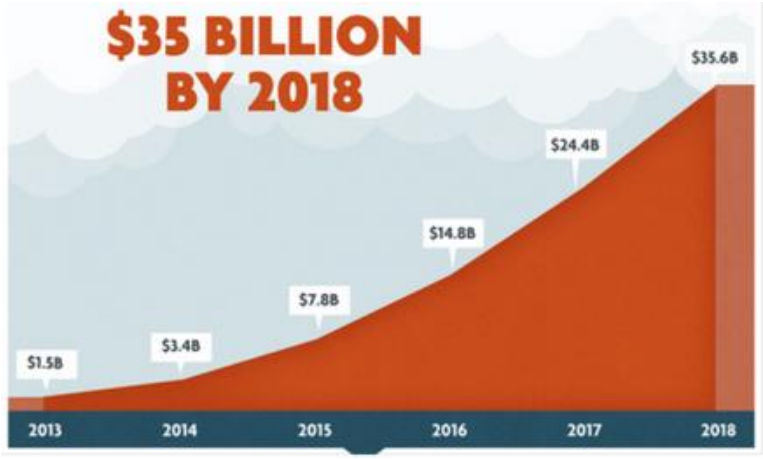

Figure 58. SDN Market size

(Source: Plexxi and Lightspeed Venture Partners)

It is therefore obvious that the significant momentum of both the LEO/MEO satellite and the SDN markets can be exploited to create new business opportunities based on reprogrammable payloads and flexible ISLs. We thus consider the market potential of this scenario as **Medium**, given that LEO/MEO services will always occupy a specific (yet limited) share in the satellite market.

4.2.7. Scenario #7: Dynamic backhauling with edge processing

The dynamic backhauling with edge processing as-a-Service scenario investigates the dynamic extension of terrestrial networks via satellite links, in cases where terrestrial coverage is inadequate. Beyond allocating capacity on-demand and providing the necessary QoS per service, it becomes possible to also deploy instances of specific services of the terrestrial network, such as LTE EPC components as VNFs on the satellite access segment. This is the concept of satellite edge processing, which is inline with the emerging paradigm of Mobile Ede Computing (MEC).

Apart from backhauling support for mobile networks, this scenario also aims to augment the typical satellite M2M service by dynamically deploying data processing components as VNFs at the satellite access segment i.e. at the gateways providing satellite connectivity to the local M2M network. This capability allows local preprocessing of the M2M traffic at the aggregation point (e.g. data aggregation, statistical processing, video feature extraction etc.) in a reprogrammable/ reconfigurable manner.

4.2.7.1. Actors and roles

Although this scenario has considerable technical implications, the value chain is simple. The **Satellite Network Operator** offers the dynamic backhauling service, also providing satellite terminals with edge computing/processing capabilities. The **Customers** are expected to be e.g. mobile operators (using the satellite segment to

extend network coverage), M2M platform operators, institutional users etc. Generally, this type of service is not targeted to retail/residential customers.

4.2.7.2. Description and added-value

Mobile backhauling (e.g. for 2G/3G/4G networks) has been one of the typical use cases for satcom. Integrating satellite in the cellular infrastructure by feeding remote base stations via satellite allows mobile network operators to extend their services to areas and cases not covered by terrestrial backhauls (e.g. fiber or microwave). These cases include remote, isolated locations, where the extension of terrestrial backhauls is not technically feasible or economically advisable. Satellite backhauling is also used where the terrestrial infrastructure has suffered considerable damage (e.g. after a natural disaster).

As also explained in the previous scenarios, the use of network programmability technologies greatly facilitates the allocation, management and optimization of the backhaul capacity. Thus, short service setup time and resource elasticity are key benefits to be introduced.

However, in a virtualization-enabled world, backhauling can mean much more than capacity. Specifically, one of the envisaged key elements of the 5G technological framework is the capability to deliver intelligence directly to network's edge, in the form of virtual network appliances, jointly exploiting the emerging paradigms of Network Functions Virtualisation (NFV) and Edge Cloud Computing. Novel edge infrastructures promise to offer dynamic processing capabilities on-demand, optimally deployed close to the user. Following this direction, novel business cases will produce added value from any kind of infrastructure or application that has the potential to be offered 'as a Service'.

The satellite edge processing scenario assumes the extension of this paradigm to the satellite domain; specifically, it foresees that the backhauling service is coupled with virtualization capabilities at the satellite terminal, able to host virtual traffic processors close to the end users (Figure 59). Such local traffic processing can achieve significant savings in satellite capacity.

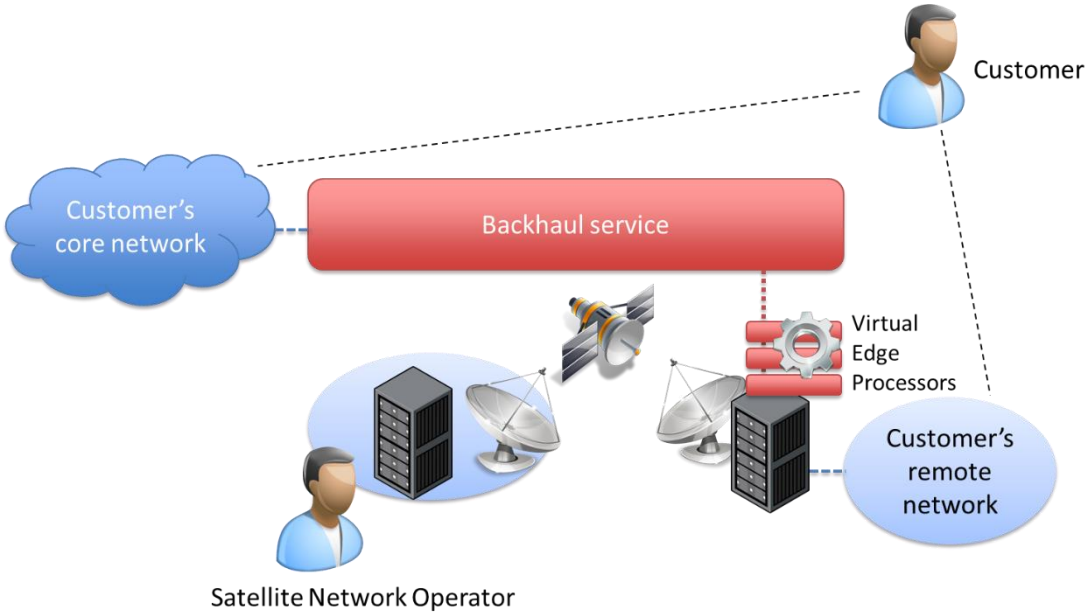

Figure 59. Dynamic backhauling with edge processing scenario

Two examples of this scenario are:

- in 3G/4G mobile backhaul services, some IMS or EPC components could be deployed at the edge, so that user traffic is processed and rerouted locally, without the need to traverse the satellite segment
- in M2M services, sensor data can be aggregated and processed locally at the virtual processor(s) of the terminal, for example-
    - measurements from multiple sensors can be aggregated, and only aggregates and possibly detected events are transmitted back over satellite
    - video streams can be dynamically transcoded, features can be extracted and only the features/processing results are transmitted back over satellite

The NFV agility allows customers to deploy such traffic processing functionalities on-demand in professional satellite terminals, upgrade them and configure/manage them in a unified manner. Resources of virtual appliances can be scaled up and down on-demand, matching the traffic characteristics and customer requirements.

This concept eventually results in a totally new service mix, in which traditional backhauling is coupled with edge processing resources, offered on-demand, as-a-Service. The terminal is essentially transformed to a virtualization-capable remote head-end, able to serve a wide range of use cases.

Last but not least, although the scenario, as described, assumes the use of the satellite terminal by a single customer, virtualization technology allows also multi-tenancy at the edge segment; this means that the professional terminal itself may be partitioned into multiple "virtual terminals", offered to different customers. This capability can be exploited in scenarios where the satcom operator has already deployed a network of terminals and leases portions of the terminals to different

customers. For example, a set of terminals covering a remote village can be leased and shared among two or more mobile operators. This interesting and novel approach demonstrates the power of virtualization technology to introduce new market opportunities and to transform the typical telco value chains.

For all these reasons, the technical added-value of the solution is considered **High**.

4.2.7.3. Implementation aspects and challenges

With regard to reserving bandwidth capacity for the backhaul service, the use of SDN greatly simplifies network control and facilitates QoS assurance with per-flow or per-application granularity. As explained in Scenario #1 (Sec. 4.2.1), it is advisable that SDN capabilities are integrated in both the satellite gateway and the remote terminal, which are centrally controlled by the satellite operator who uses SDN management to allocate bandwidth on-demand. Although, as explained, BoD is already feasible with legacy technologies, SDN allows elasticity, per-application differentiation and flexible SLAs and pricing – specifically suited to more dynamic use cases. However, as highlighted in Sec. 4.2.1.3. , it is essential to couple SDN with radio resource management in order to efficiently control and share the satellite capacity, especially for the return link.

When it comes to edge processing, then NFV, coupled with emerging Mobile Edge Computing (MEC) concepts for deployment of cloud resources at the network edge, are the key enabling technologies. The satellite terminal needs to encompass virtualized IT resources in order to host the traffic processors, as virtual network functions (VNFs). When it comes to management, since it is not advisable to deploy an entire cloud system (e.g. Openstack) on the terminal, it could be assumed that the cloud controller is located centrally at the satellite gateway, controlling remote compute nodes at the terminals. In a more lightweight approach, the terminals can encompass plain IT virtualization (e.g. via a KVM hypervisor or even via Docker containers), without any cloud framework. This approach has the cost of reduced elasticity and management features. However, it saves IT resources and also relieves the satellite segment from excessive signaling, thus it would be more appropriate for edge VNFs (rather than for VNFs hosted at the Gateway, where Openstack-based management is still advisable).

The technology readiness level is considered **Medium**, mostly due to low maturity of edge computing mechanisms.

4.2.7.4. Market potential

Global coverage and dependability are and will remain the main added value of space-based communication services. Integrated in the 5G network infrastructure, SatCom solutions are well positioned to target the 4 main types of use cases identified in Figure 60.

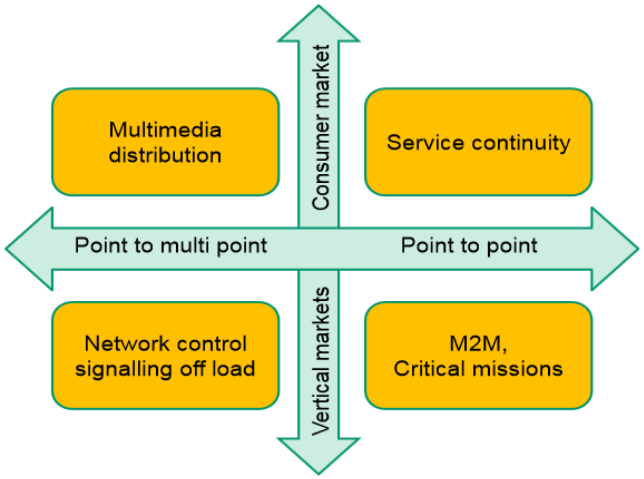

Figure 60 . Satellite use cases in 5G

The advent of satellite platform technology, new players in the launch services market, and new technologies such as High Throughput Satellites (HTS), is allowing the satellite industry to improve its value proposition to wholesale customers. These developments are opening up and expanding market opportunities particularly in cellular backhaul, which is seeing explosive demand on the back of high growth rates in 3G and 4G. ABI Research expects the small cell backhaul equipment market to exceed US$5 billion in 2019, which aligns with operators' deliberate approach to small cell deployments. Fuelled by the increased penetration of smartphones globally, which is leading a growth in Internet applications and data demands from mobile users across the world, emerging regions are becoming a gold mine for satellite operators to offer backhaul services.

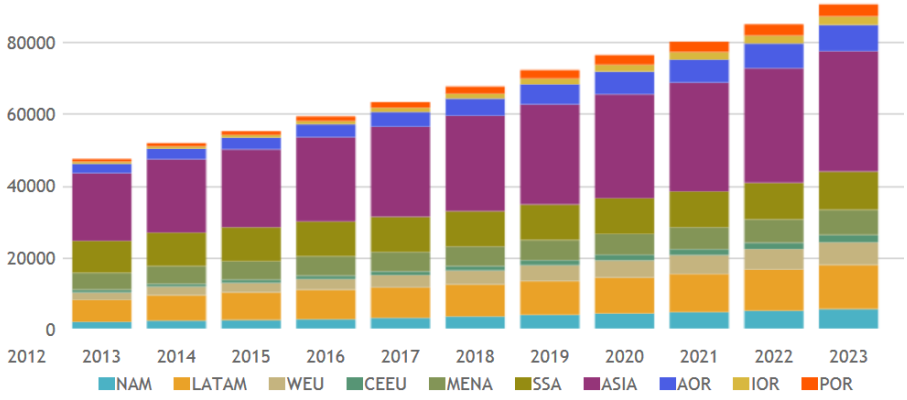

Figure 61. Global wireless backhaul, trunking and video offload capacity demands

(Source: Northern Sky Research, 2014)

According to Northern Sky Research (NSR), Asia is expected to dominate the demand given the large population base and geographic coverage that requires satellite reach. The study forecasts Latin America to come in at a distant second, given the high levels of urbanization where terrestrial technologies can address the population base. Given

this increasing demand for backhaul services, the scenario presented in this section presents an attractive business case for satellite operators.

The market potential for this scenario is considered **Medium**, targeting to a specific customer group (enterprise users with specific needs)

4.2.8. Scenario #8: Customer functions virtualization

This scenario is based on the VNF-as-a-Service (VNFaaS) paradigm and assumes the dynamic offering of virtual network appliances to satcom customers in the form of VNFs (e.g. firewalls, traffic filters, home gateway functionalities, media storage and processing etc.). According to their nature, these VNFs can be instantiated either at the satellite gateway or at VNF-enabled satellite terminals.

It must be noted that this scenario focuses on consumer use, as opposed to scenarios #7 (edge processing) and #5 (SVNO) which, although also employing NFV, serve purely professional use cases and are targeted to enterprise Customers who leverage virtualization technologies to resell services to end-users.

4.2.8.1. Actors and roles

This scenario assumes that the **Satellite Network Operator** also undertakes the role of the NFV service provider and offers VNFaaS as added-value services along with satellite connectivity to **Customers**. In a more pluralistic scenario, the **VNF Providers** (developers) play a more active role, advertising and dynamically pricing their services which are published in a catalogue. The Customers may select the services that best suit their needs. In some business models, the VNF Providers may receive direct profit from the customers, either indirectly as a share of the satcom service fee or directly, as a license fee for using the VNF.

4.2.8.2. Description and added-value

In the most common scenario of satellite broadband access, the satellite terminal itself exposes some basic network functionalities to the customer, such e.g. as firewalling, NAT, port forwarding etc. If more capabilities are needed, the customer has to acquire and install additional physical appliances.

Furthermore, there are some capabilities that would be advisable to be present before the satellite segment, for the sake of saving satellite capacity. For example, firewalling should be conducted at the satellite gateway to avoid transmitting over satellite traffic which will be eventually blocked at the terminal. Same with media transcoding; it would be advisable to transcode streams before they are transmitted, so they occupy less satellite capacity. However, such capabilities cannot be currently provided per customer; the network functionalities at the Gateway apply to the entire traffic and of course cannot be managed by the customer.

The VNFaaS scenario promises to alleviate these limitations by allowing network functions in the form of virtual network appliances to be acquired on demand by the Customer and instantiated either at the satellite terminal or in a shared resource space (mostly private cloud infrastructure) at the Gateway (Figure 62). Some functions, such as PEP and application classification, could be installed at both ends.

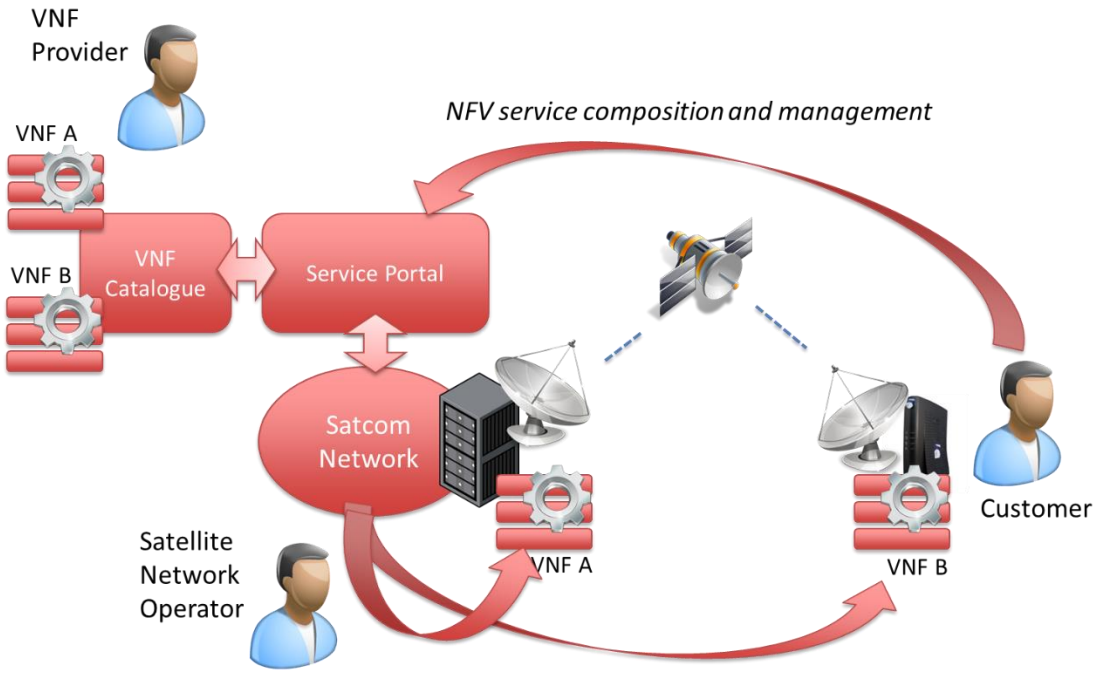

Figure 62. Customer functions virtualization scenario

In a more static scenario, the Satellite Network Operator manually deploys the VNFs and interconnects them, following a customer request. In a more interactive and dynamic approach, the customer composes the NFV service in a completely automated manner by accessing a service portal, browsing the VNF catalogue, selecting the VNFs which best match his/her needs and integrating them into a satcom service package. This is e.g. the concept of the "NFV Marketplace" developed by the T-NOVA project.

The same service portal could then be used for the monitoring and the management of the service. VNFs may be managed either via the portal or via individual management interfaces.

Examples of VNFs which would bring added-value when offered as-a-service in a satcom context would be:

- Firewalling and content filtering (GTW side)
- Application classification (GTW and Terminal sides)
- Caching (Terminal side to cache traffic from external networks; GTW side to cache traffic stemming from the terminal)
- Media transcoding (GTW side for media streams consumed by the customer)
- Performance Enhancement Proxy (GTW and Terminal sides)

Due to the new services introduced, coupled with flexibility and resource efficiency, the technical added-value of this scenario is considered **High**.

4.2.8.3. Implementation aspects and challenges

For the implementation of the use case, a VNFaaS platform such as the one being developed in the T-NOVA project [TNOVA] needs to be integrated into the satellite infrastructure. Commonly, the NFV management entities are deployed at the Gateway side, controlling NFV resources both local (at the Gateway) and remote (at the terminals). In order not to pose excessive capacity overhead in the satellite segment, the remote management of NFV resources at the terminals should involve as little signaling as possible.

The NFV management is expected to carry out procedures for controlling the entire NFV lifecycle, including:

- NFV service mapping, i.e. allocating the resources which match the service requirements and characteristics
- VNF instantiation, i.e. launching of the VNF images in the host machines
- Service chaining, i.e. controlling the network to interconnect the various VNFs of the service and directing the customers' traffic through the VNFs
- Service monitoring, i.e. collecting and aggregating metrics from VNFs and virtual networks
- Service rescaling, including rescaling of VNF resources and network resources
- Service starting/stopping and teardown

Apart from the aforementioned management procedures, the NFV platform also needs to accommodate interactions with the customers, allowing them to select, deploy, manage and monitor VNFs. An NFV service catalogue is essential in order to allow customers to customize the services according to their needs. Proper SLA and billing mechanisms must also be in place.

In order allow deployment of VNFs in the satellite terminals, the latter need to offer generic computing resources, as well as the proper management interfaces, in order to accommodate VNFs. Given that terminals have generally constrained hardware resources, it is of particular interest to exploit novel virtualization techniques for non-x86 processors (e.g. suitable for ARM processors) as well as lightweight virtualization schemes (e.g. Linux containers or Docker containers), rather than full virtualization based on Virtual Machines. This approach would allow the deployment of multiple VNFs chained together in a single terminal with minimal resource overhead.

Furthermore, SDN support within the satellite network (at least in the Gateway local network) is considered essential, since not only the T-NOVA platform, but also other emerging NFV architectures, are based on SDN for network management.

The technology maturity for this scenario is considered Medium, closely associated with the foreseen progress of the NFV architectures in the years to come.

4.2.8.4. Market potential

According to Infonetics Research Report (2014), the global service provider software-defined networking and network function virtualization market is expected to grow from less than $500 million in sales last year to $11 billion in sales by 2018. For this forecast, the report broke down the market into three categories, with revenues from new SDN and NFV software to make up 20% of the market by 2018, displaced revenue from NFV-related products that a company buys instead of buying network hardware to make up 12% of the market; and revenue from newly identified segments of existing markets, including virtualized network functions, ports on routers, switches and SDN-capable optical gear making up 68% of revenues (Figure 63)

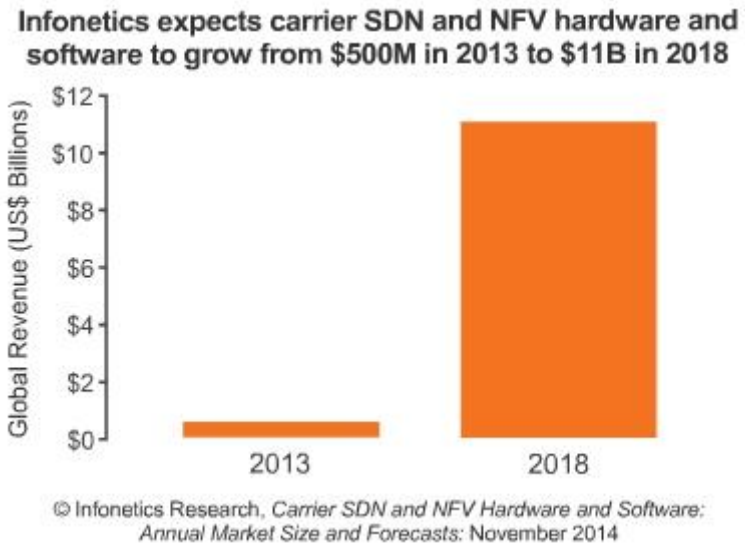

Figure 63. SDN and NFV Revenues Forecast

(Source: Infonetics Research Report (2014))

According to Infonetics, NFV is expected to represent the lion's share of the market between 2014 and 2018, with the value of NFV coming mostly from VNF software rather than orchestration and control. It is noted that VNF makes up more than 90% of the NFV software segment, and that it is expected SDN and NFV software to comprise three-fourths of the total revenues in 2018.

Driven by the thriving ecosystem Software Defined Networking (SDN), Network Functions Virtualization (NFV) and network virtualization market is expected to account for nearly $4 Billion in revenue 2014, reveals The SDN, NFV & Network Virtualization Bible: 2014 - 2020 report. A further growth at a CAGR of nearly 60% over the next 6 years is expected despite barriers relating to standardization and co-existence with legacy networks. This report also estimates that by 2020, SDN and NFV can enable service providers (both wireline and wireless) to save up to $32 Billion in annual CapEx investments.

According to Analysys Mason's new report Cloud computing, NFV and SDN: worldwide market sizing and forecast 2014–2018, CSPs will be slow to spend on NFV and SDN technology in the next 5 years (see Figure 64), but the forecast is that this trend dramatically changes towards NFV deployment till 2023.

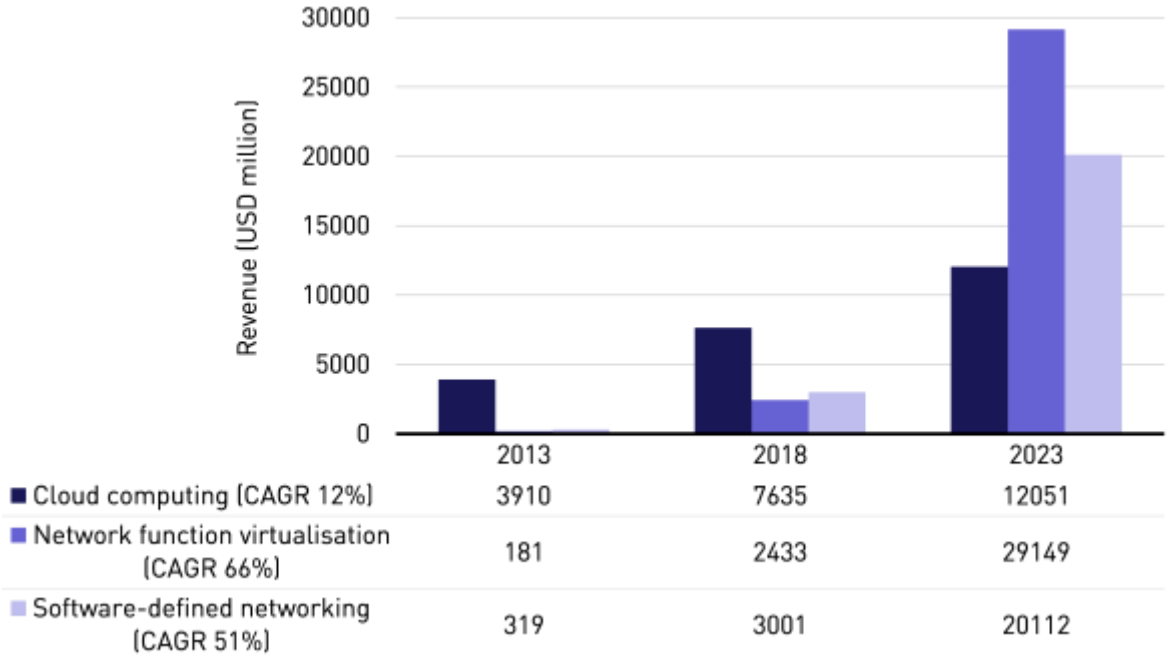

Figure 64. Software-controlled networking revenue by market segment, worldwide, 2013, 2018 and 2023

(Source: Analysys Mason, 2014)

The Customer functions virtualization scenario has the potential to expose the satcom community to the NFV market perspectives and contribute to the significant growth of the satcom-oriented virtual appliance market. Therefore, the market potential of this scenario is High.

4.3. Scenarios Consolidation and Selection

The scenarios presented not only address a wide range of use cases, but also involve diverse system characteristics. Table 12 overviews the scenarios described in terms of system characteristics, and especially:

- *Preferred satellite network configuration*: whether the scenario is more appropriate to MEO/LEO or GEO configuration, single- or multi-spot (SS/MS)
- *SDN*: whether the scenario requires SDN capabilities in the satellite network
- *NFV*: whether the scenario requires NFV capabilities in the satellite network
- *Hybrid access*: whether dual-access terminals are involved (featuring both terrestrial and satellite access). In the hybrid access scenario, network traffic to and from customers can be routed either via the satellite or terrestrial access links. Multipath techniques also belong to this category.

- *Satellite-terrestrial federation*: whether federated management of the terrestrial and satellite segments is required.

Table 12. System characteristics for the integration scenarios

Sce-nario no.	Scenario name	Preferred satellite network configuration	SDN	NFV	Hybrid access	Satellite-Terrestrial federation
1	Elastic Bandwidth-on-Demand	GEO SS	✓	-	-	-
2	Hybrid media distribution network as-a-Service	GEO SS	✓	-	✓	✓
3	Virtual CDN as-a-Service	GEO SS	-	✓	-	-
4	Federated Terrestrial-Satellite VPN	GEO MS	✓	-	✓	✓
5	Satellite Virtual Network Operator (SVNO)	GEO SS	✓	✓	-	-
6	Programmable payloads and flexible ISLs	LEO/MEO	✓	-	-	-
7	Dynamic backhauling with edge processing	GEO SS	✓	✓	-	✓
8	Customer functions virtualization	GEO SS	✓	✓	-	-

As a general overview of the scenarios, Table 13 recalls the assessment of each scenario in terms of technical added-value, technology readiness and market potential.

Table 13. Scenarios characterization and selection

Sce-nario no.	Scenario name	Technical added-value	Technology readiness	Market potential	Selected for further study
1	Elastic Bandwidth-on-Demand	M	H	H	
2	Hybrid media distribution network as-a-Service	M	H	M	✓
3	Virtual CDN as-a-Service	L	M	M	
4	Federated Terrestrial-Satellite VPN	L	M	M	
5	Satellite Virtual Network Operator (SVNO)	H	L	M	
6	Programmable payloads and flexible ISLs	M	L	M	
7	Dynamic backhauling with edge processing	H	M	M	✓
8	Customer functions virtualization	H	M	H	✓

As seen in the table, the following scenarios are selected:

- *Scenario #2 (Hybrid Media distribution network as-a-Service)* because it can be applied in the short term, it is based on media distribution which is a key satcom market segment and also it encompasses the hybrid access approach.
- *Scenario #7 (Dynamic backhauling with edge processing)* because it addresses backhauling, which is a key satcom use case, and aligns with the MEC concept which currently has significant momentum
- *Scenario #8 (Customer functions virtualization)* because it presents significant technical added-value and inherits most of the technical and market benefits of NFV.

Scenario #1 is not selected, although it has both high market potential and technology readiness, because this feature (dynamic bandwidth assignment and elasticity) is also demonstrated in Scenarios #2 and #7.

Scenarios #3 and #4 are not selected due to the relatively limited added-value brought by cloud networking technologies, compared to the current service offerings.

Scenarios #5 and #6, although attractive, are not selected due to their low TRL, which among others inhibits their implementation in the CloudSat testbed (absence of a Cloud-RAN platform and an SDN-enabled satellite constellation emulator, respectively).

5. INTEGRATED CLOUD NETWORKING ARCHITECTURES

This chapter presents a concrete and sound architectural proposal for the provision of federated satellite/terrestrial network services. This architecture should be in line with recent SDN/NFV technologies and architectural trends, align with the specific capabilities and constraints of the satellite networks and should also be able to be realized via several deployment scenarios and evolutionary paths. Most important, it should be able to accommodate the integration scenarios/use cases identified as of particular interest for satellite/terrestrial cloud networks.

The methodology for deriving this architectural proposal first includes the survey of state-of-the-art architectural frameworks, mostly in the field of SDN and NFV and the identification of common concepts which can be exploited. In parallel, specific high-level architectural requirements are derived from the identified integration scenarios. These two inputs are used to derive a proposal for the integrated CloudSat architecture.

5.1. Overview of Relevant Architectural Proposals

The following sections aim to survey a number of integrated SDN/NFV enabling architectures, as proposed by R&D projects currently running, industry frameworks and solutions as well as efforts from standardisation bodies related to NFV. This survey, which focuses mostly on architectural aspects rather than specific enabling technologies, provides a useful insight on how virtualization-capable network infrastructures are commonly structured.

5.1.1. ETSI ISG NFV

The scope, work and current status of ETSI ISG NFV was described in Chapter 2. Here we review the architectural model which has been proposed by ETSI in [ETSINFV]. The high-level NFV architectural framework includes three main working domains which can be identified (Figure 65)

- The Virtualised Network Function (VNF), as the software implementation of a network function which is capable of running over the NFVI.
- The NFV Infrastructure (NFVI), which includes the diversity of physical resources and how these can be virtualised. NFVI supports the execution of the VNFs.
- The NFV Management and Orchestration (NFV MANO), which covers the orchestration and lifecycle management of physical and/or software resources that support the infrastructure virtualisation, and the lifecycle management of VNFs. NFV MANO focuses on all virtualisation-specific management tasks necessary in the NFV framework.

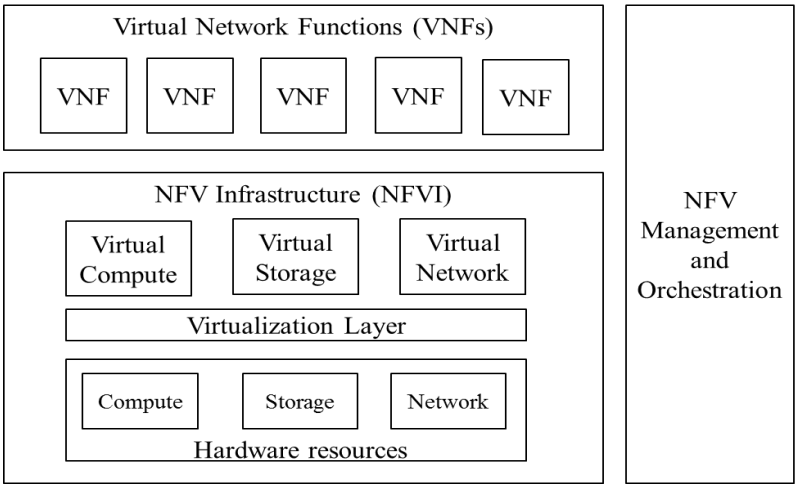

Figure 65. ETSI High-level NFV domains (Source: [ETSINFV])

The NFV architectural framework handles the expected changes that will probably occur in an operator's network due to the network function virtualisation process. Figure 66 shows this global architecture, depicting the functional blocks and reference points in the NFV framework.

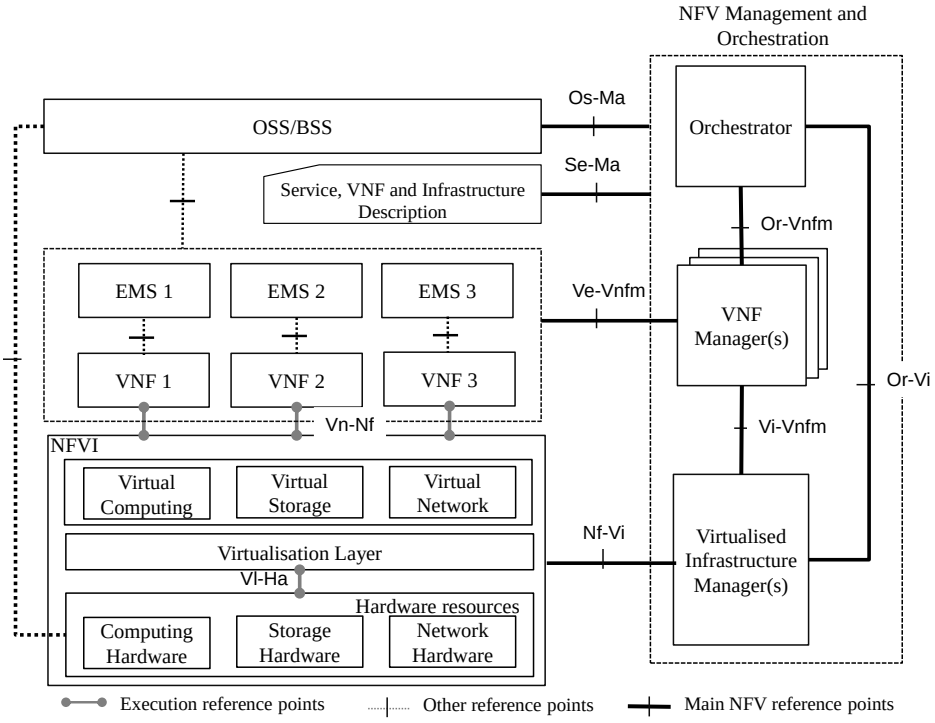

Figure 66. ETSI NFV reference architectural framework[12]

The architectural framework shown in Figure 66 focuses on the functionalities necessary for the virtualisation and the consequent operation of an NFV-enabled operator network. It does not specify which network functions should be virtualised, as that is solely a decision of the owner of the network.

[12] Source: gs_NFV002v010101p - NFV - Architectural Framework

5.1.2. CloudNFV

CloudNFV [CloudNFV] is an open platform for implementing NFV based on cloud computing and Software Defined Networking (SDN) technologies in a multi-vendor environment. The involved companies are: 6WIND, CIMI Corporation, Dell, EnterpriseWeb, Overture Networks, and Qosmos. CloudNFV, still preserving the alignment with the NFV ISG of ETSI, deploys a mixture of virtual network functions, cloud application components, real network devices and services, and multi-operator federated services.

CloudNFV architecture is based on management and orchestration applications built around an agile data/process model called Active Virtualization, which provides for order/contract and policy storage ("Active Contract") and resource state information ("Active Resource") provided by EnterpriseWeb. Service orders are optimized through Active Virtualization then provisioned on cloud infrastructure using Overture Network's Ensemble Service Orchestrator, which instantiates the virtual network functions through OpenStack Nova and connects them using OpenStack Neutron. The overall CloudNFV architecture is depicted in Figure 67.

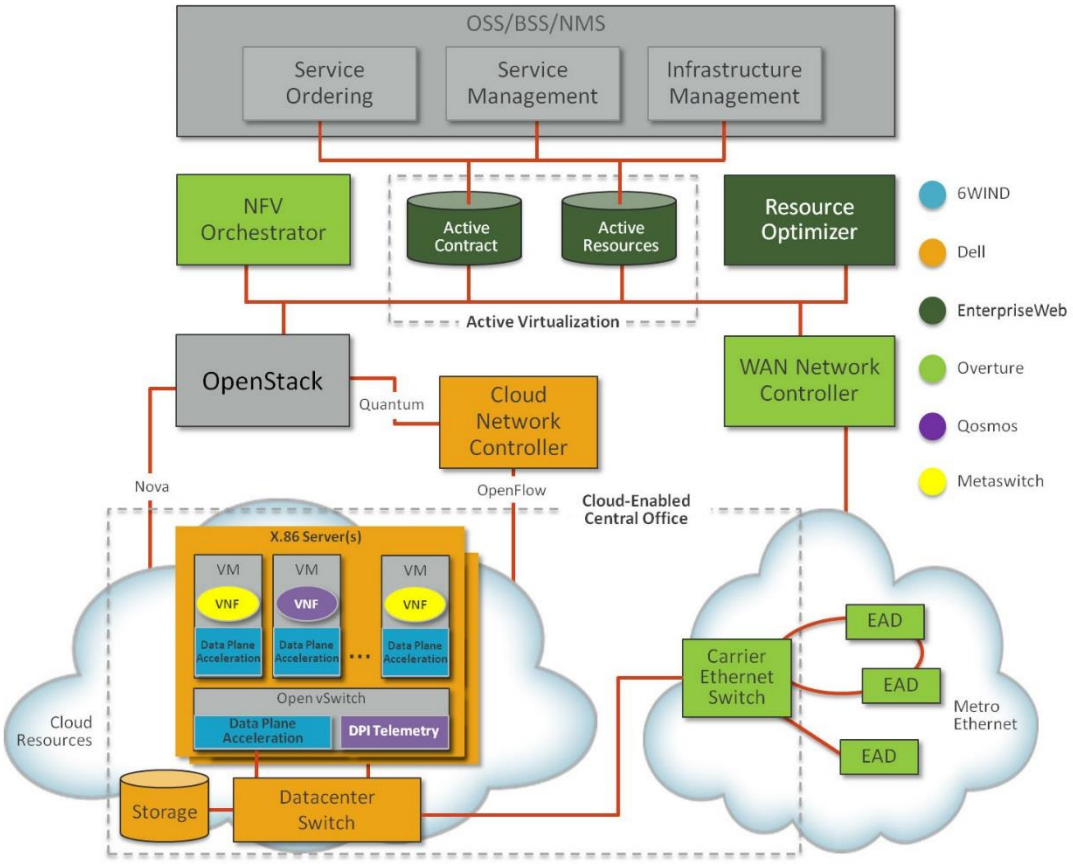

Figure 67. CloudNFV Architecture

The architecture is designed to support open interfaces for carrier federation at both infrastructure and orchestration levels. This provides the capability to CloudNFV to deploy assets on a per-customer basis and also as Infrastructure Services (shared between users). Moreover mixed physical devices and virtual functions configurations are supported.

5.1.3. HP OpenNFV

OpenNFV [OpenNFV] is a comprehensive project launched by HP, built around a proposed open reference architecture, encompassing a service portfolio, and enforced by an ecosystem of ISVs, NEPs and application developers.

HP architecture is aligned with the ETSI model, and HP has a number of active contributors in the NFV ISG. OpenNFV main components are a NFV Infrastructure and a NFV Orchestrator module, in turn based on HP Converged Infrastructure and HP Converged Cloud propositions. It also capitalizes on the SDN role, and on HP's SDN technology assets. It is a modular architecture, basically vendor agnostic and allowing a modularized approach to NFV take-up. A high level picture of OpenNFV architecture is presented in Figure 68.

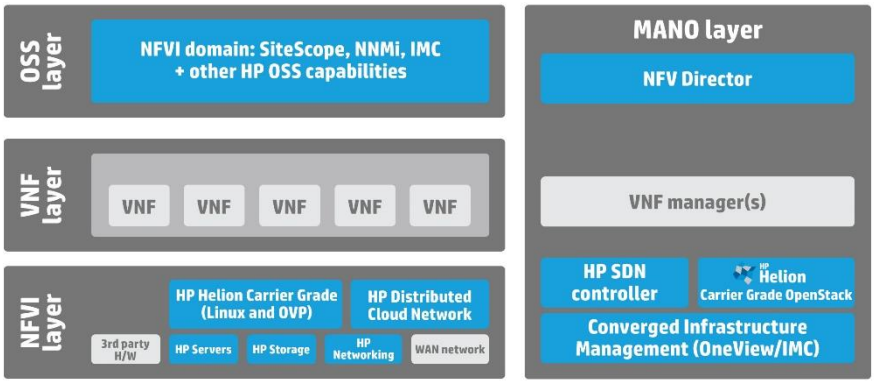

Figure 68. HP OpenNFV architecture

The NFV Orchestrator module (implemented by HP in the *NFV Director* component) implements the functions prescribed by the ETSI model MANO layer specification. The NFV Orchestrator is hypervisor-agnostic, so it can support different solutions both proprietary and open source.

5.1.4. Qosmos/Intel/Tieto

Intel has been an active player in supporting the development and evolution of SDN and NFV through industry and vendor specific initiatives. The goal of the program is to make it easier to build, enhance, and operate SDN/NFV-based infrastructure, while lowering capital and operating expenditures. The program publishes function specific architectures such as vEPC, vBRAS, vCPE etc. [INTEL].

The key components of the architecture are shown in Figure 69. This architecture was used to form the basis for a proof of concept (POC) that demonstrated a number of ETSI defined VNF use cases.

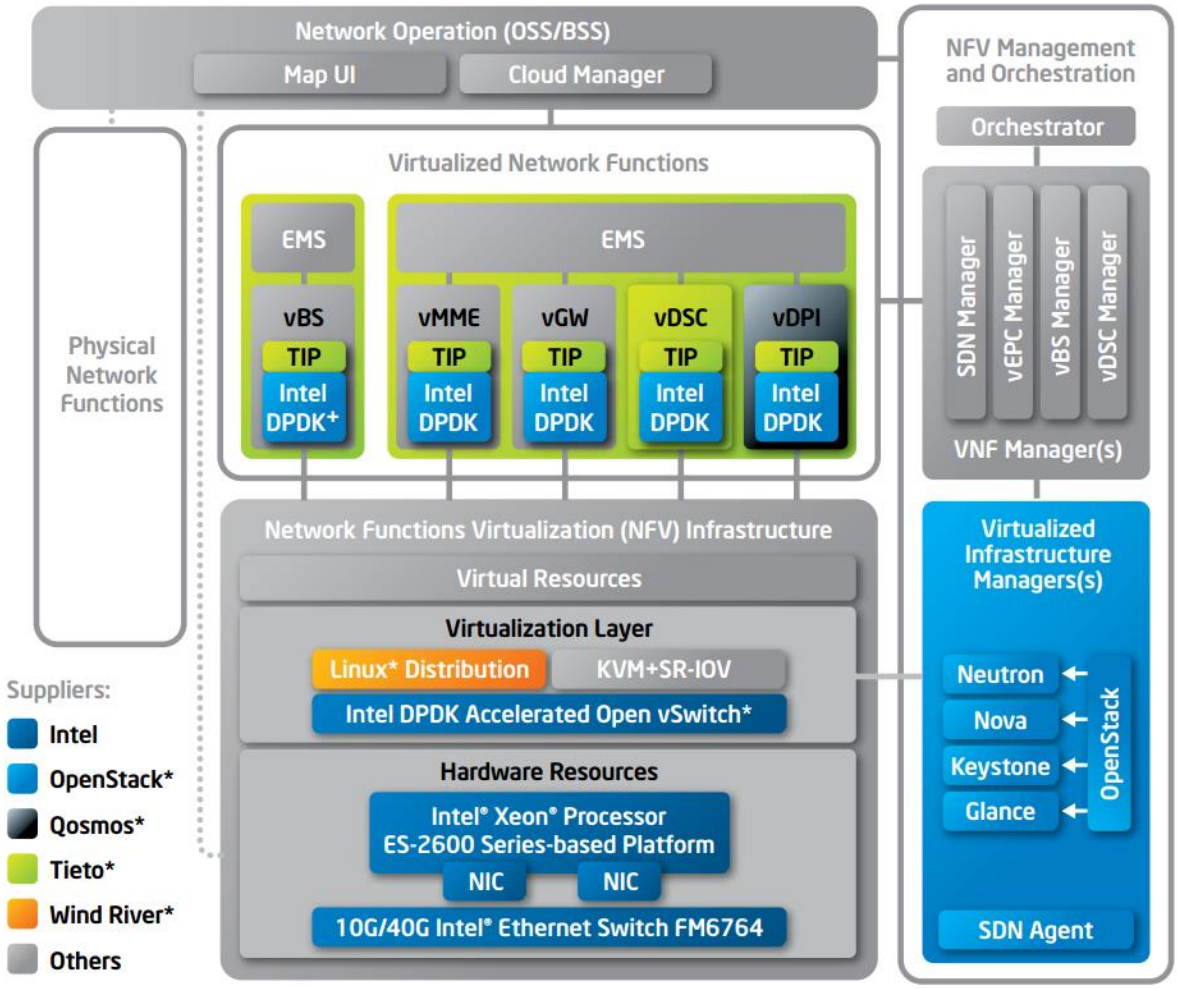

Figure 69. Intel NFV High Level Architecture

The key components of the architecture are as follows:

Virtualised Network Functions – The PoC implemented reference software for LTE eNodeB and EPC (MME, SGW, and PDN GW), along with Tieto's Diameter Signalling Controller (DSC), which was deployed as a VNF. Qosmos' intelligent Deep Packet Inspector (DPI) was included and could be deployed either within a VNF or as a standalone virtual networking function component (VNFC). The associated Element Management System (EMS) for each VNF is integrated within the VNF subsystem, which monitors the operational condition of the VNF's as part of the overall Telecommunications Management Network (TMN).

NFV Management and Orchestration – This sub-system is responsible for the management of VNF deployments and lifecycle. The SDN controller is also contained within this sub-system with responsibility for flow control to enable intelligent networking. The architecture is orchestrator-agnostic however the reference implementation is based on OpenStack. The network management solution interfaces to OpenStack via Heat are used for automation and deployment.

Network Operations (OSS/BSS) – The architecture and reference implementation supports a NETCONF interface for OSS and a Cloud Infrastructure Management Interface (CIMI). The intent is to demonstrate that new VNFs can be deployed and managed from an end-to-end perspective within a telco cloud environment.

Cloud Platform – This sub-system is based on OpenStack, which has been modified by Tieto to include telco-grade supervision, statistics, diagnostics, fault and performance management capabilities.

SDN Networking – The SDN controller is compatible with OpenStack Neutron and supports SDN networking and legacy network management system (NMS) integration, and provides supervision, statistics and performance management. OpenFlow is used to communicate between the SDN controller and the Open vSwitch, which is managed by the OpenFlow Controller.

5.1.5. Alcatel-Lucent CloudBand

The CloudBand NFV platform [ALC] aims at transforming carrier-grade service provider (B2B proposal) networks with distributed footprint into a single, manageable, virtual cloud. The overall architecture of CloudBand is illustrated in Figure 70.

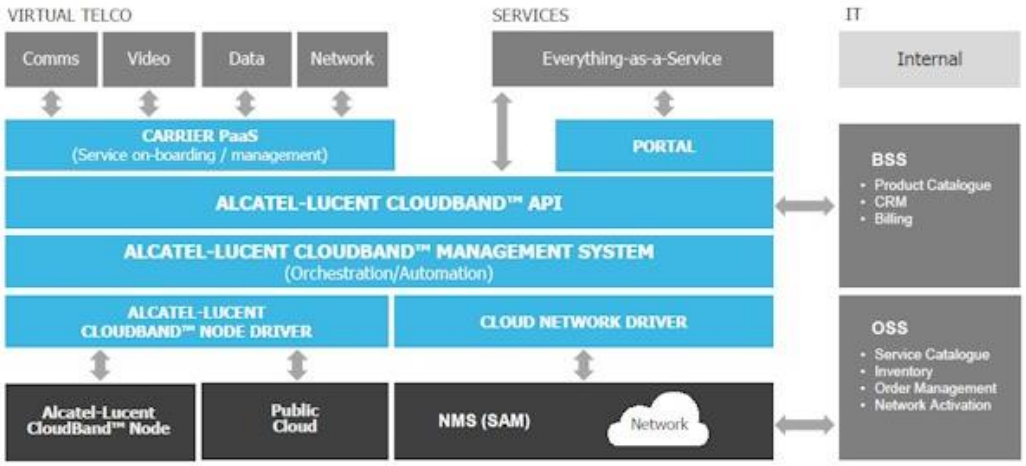

Figure 70. CloudBand Overall Architecture

The major characteristics of the architecture in Figure 7 are:

- Mono-vendor solution NFV / Cloud management
- Possible expansion to Everything as a Service offer[13]
- Cloud Infrastructure accesses via API from NFV and BSS
- Interface to public cloud
- NMS includes management of the CloudBand Network
- Multi-vendor support for cloud nodes

[13] The platform is supposed to be flexible to accommodate other services in the paradigm XaaS beside the ETSI VNF.

In addition, concerning the deployment, a distributed cloud infrastructure can accommodate a large number of small and medium datacentres, placed in different sites to spread services across multiple locations.

CloudBand also comes with a carrier PaaS (cPaaS) management tool that manages application lifecycle automatically and on-boards VNFs in the cloud. It hides the complexities of infrastructures and OS. It can automate and optimise application services like IMS or any other carrier grade service and facilitates operators to concentrate on other aspects of the application lifecycle namely provisioning, monitoring, healing and scaling.

The management of the PaaS allows defining rules for the placement, SLA, placement Zones, monitoring, Cloud resources, High availability (HA) and redundancy, tracking the full lifecycle of the deployed applications.

5.1.6. Telefonica OpenMANO

Driven by Telefonica NFVlabs, OpenMANO [OpenMANO] is an open source project that provides a practical implementation of the reference architecture for Management & Orchestration under standardization at ETSI's NFV ISG.

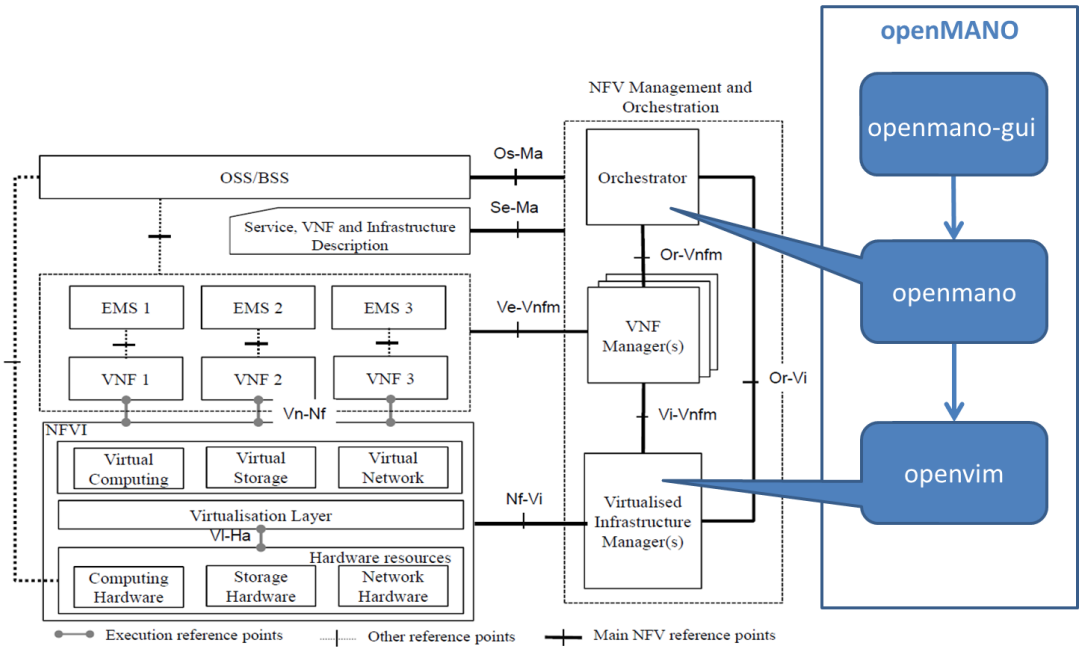

Figure 71. OpenMANO mapping to ETSI NFV architecture

The OpenMANO platform consists of three main functional components:

- **openvim** is an implementation of an NFV VIM (Virtualised Infrastructure Manager). Openvim interfaces with the compute nodes in the NFV Infrastructure and an Openflow controller in order to provide computing and networking capabilities and to deploy virtual machines. It offers a northbound

interface, based on REST (openvim API), which allows the creation, deletion and management of images, flavors, instances and networks.

- **openmano** is an implementation of an NFV-O (Network Functions Virtualisation Orchestrator). It interfaces with an NFV VIM (openvim) through its API and offers a northbound interface, based on REST (openmano API), where NFV services are offered including the creation and deletion of VNF templates, VNF instances, network service templates and network service instances.
- **openmano-gui** is a web GUI to interact with the core openmano service, through its northbound API in order to facilitate human interaction.

5.1.7. Project MCN

The scope, work and current status of EU-funded MCN project was described in Chapter 2. Here we review the MCN architectural approach for mobile service virtualization.

Mobile Cloud Networking overall architecture design [MCN] is mainly governed by service oriented design principles. Every service in MCN has the same provisioning and lifecycle management pattern and architecturally follows the global MCN reference architecture.

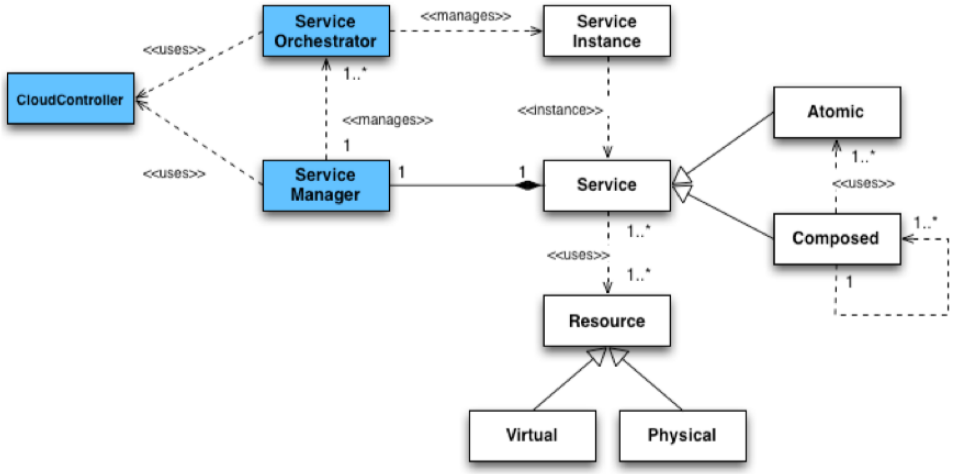

Figure 72. MCN Architectural Entities Relationship

A brief description of key components follows:

Service Manager (SM): provides an external interface to the Enterprise End User (EEU) and is responsible for managing service orchestrators, it has business and technical management functions.

Service Orchestrator (SO): it oversees the end-to-end orchestration of a service instance. It is responsible for managing the Service Instance and in particular its components (SIC), once it is created and running.

Cloud Controller (CC): provide the signaling and management interfaces to enable the common (northbound), and technology-specific (southbound) control planes. It provides both atomic and support services required for realizing SO needs. The main MCN architectural entities that interact most with the Cloud Controller are the SM and SO.

Service Orchestrator implementation in MCN project is service specific as it depends on the domain knowledge of the respective VNF it is implementing. Some of the prominent VNFs being virtualized and managed as a service in MCN are: EPC, IMS, DNS, OSS/BSS (RCB), AAA, CDN, etc.

An EEU (Enterprise End User) can request a service, which is in turn realized through composition of many services including atomic and other composed service. The MCN architecture through SM and SOs support such service composition, provisioning and runtime management in a standard manner.

5.1.8. Project CONTENT

The EU-funded CONTENT (Convergence of Wireless Optical Network and IT Resources in Support of Cloud Services) project [CONTENT] aims at offering a network architecture and overall infrastructure solution to facilitate the deployment of conventional Cloud computing as well as mobile Cloud computing introducing new business models and facilitating new opportunities for a variety of business sectors.

CONTENT proposes a layered architecture with the aim to facilitate the main principles of its novel proposition i.e. cross-technology virtualization in support of optimised, seamless and coordinated cloud and mobile cloud service provisioning across heterogeneous network domains. The overall CONTENT architectural structure is illustrated in the figure below and includes the following layers:

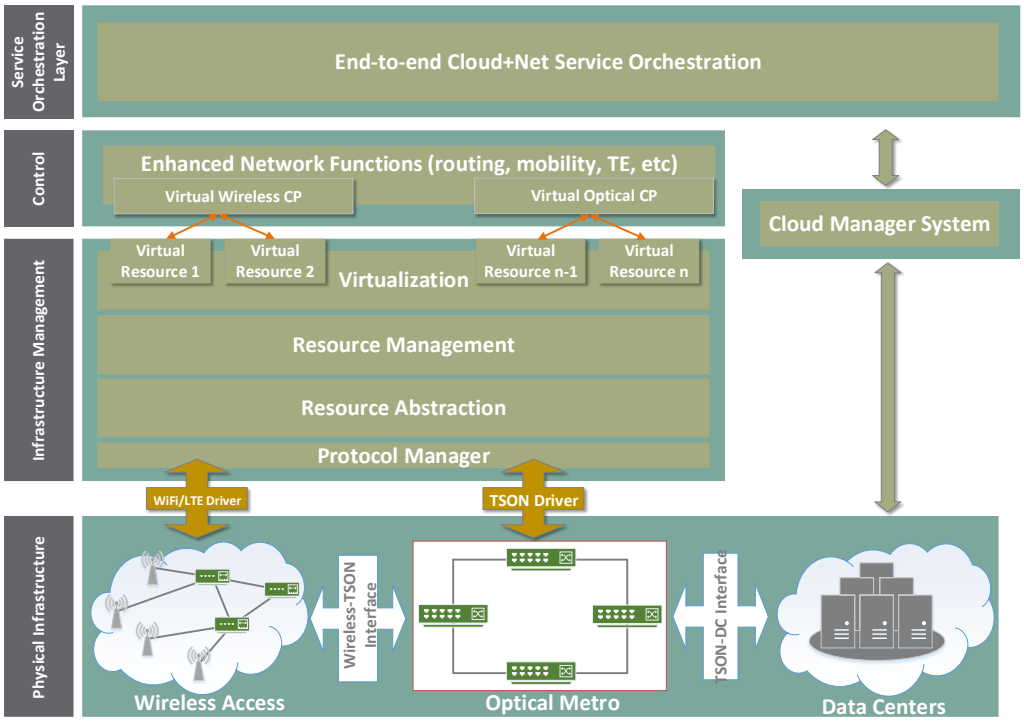

Figure 73. Overall CONTENT layered architecture

Heterogeneous Physical Infrastructure Layer: including a hybrid wireless access network (LTE/Wi-Fi) domain, and an optical metro network domain (TSON) interconnecting geographically distributed data centres, supporting frame-based sub-wavelength switching granularity.

Infrastructure Management Layer: is responsible for the management of the network infrastructure and the creation of virtual network infrastructures over the underlying physical resources. This involves functions including resource representation, abstraction, management and virtualization across the heterogeneous network domains. An important feature of the functionalities supported, is orchestrated abstraction of resources across domains, involving information exchange and coordination across domains.

Control Layer: responsible to provision IT and (mobile) connectivity services in the cloud and network domains respectively. The focus of the project is on the network side, where the control layer establishes seamless connectivity across heterogeneous technology domains (wireless access and optical metro) through a coordinated, end-to-end approach to support optimized performance, QoS guarantees as well as resource efficiency and sustainability.

Service Orchestration Layer: responsible for efficient coordination of the cloud and network resources, in order to enable the end-to-end composition and delivery of integrated cloud, mobile cloud and network services in mobile environments supporting the required QoE.

5.1.9. Project UNIFY

The EU UNIFY (Unifying Cloud and Carrier Networks) project [UNIFY] has the goal of increasing the potential of virtualization and automation across the whole networking and cloud infrastructure. The project is focused on enablers of a unified production environment and will develop an automated, dynamic service creation platform, leveraging a fine-granular service chaining architecture.

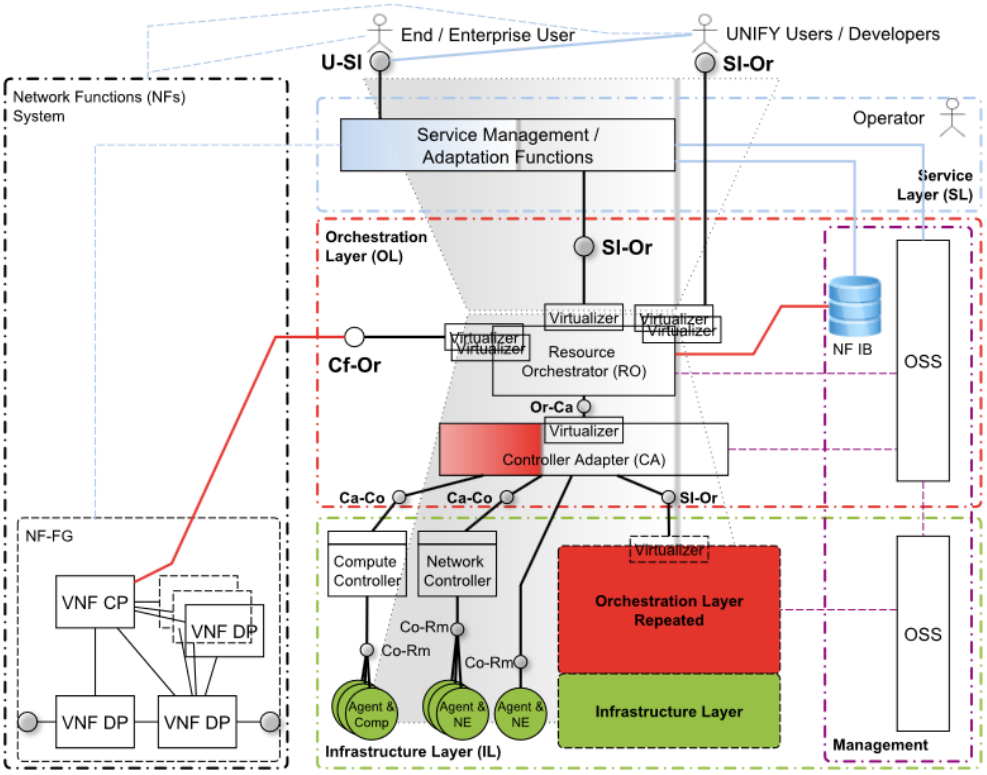

Figure 74. Overall UNIFY architecture

The overarching view of the UNIFY architecture comprises three layers, namely, the Service Layer (SL), the Orchestration Layer (OL) and the Infrastructure Layer (IL). The architecture also includes management components, a Network Functions System (NFS) and reference points between the major components. UNIFY proposes a service abstraction model and a service creation language to enable dynamic and automatic placement of networking, computing and storage components across the infrastructure.

5.1.10. Project T-NOVA

The scope, work and current status of the EU-funded T-NOVA project was described in Chapter 2. Here we briefly recap the T-NOVA architectural approach which is based on four logical layers (Figure 75).

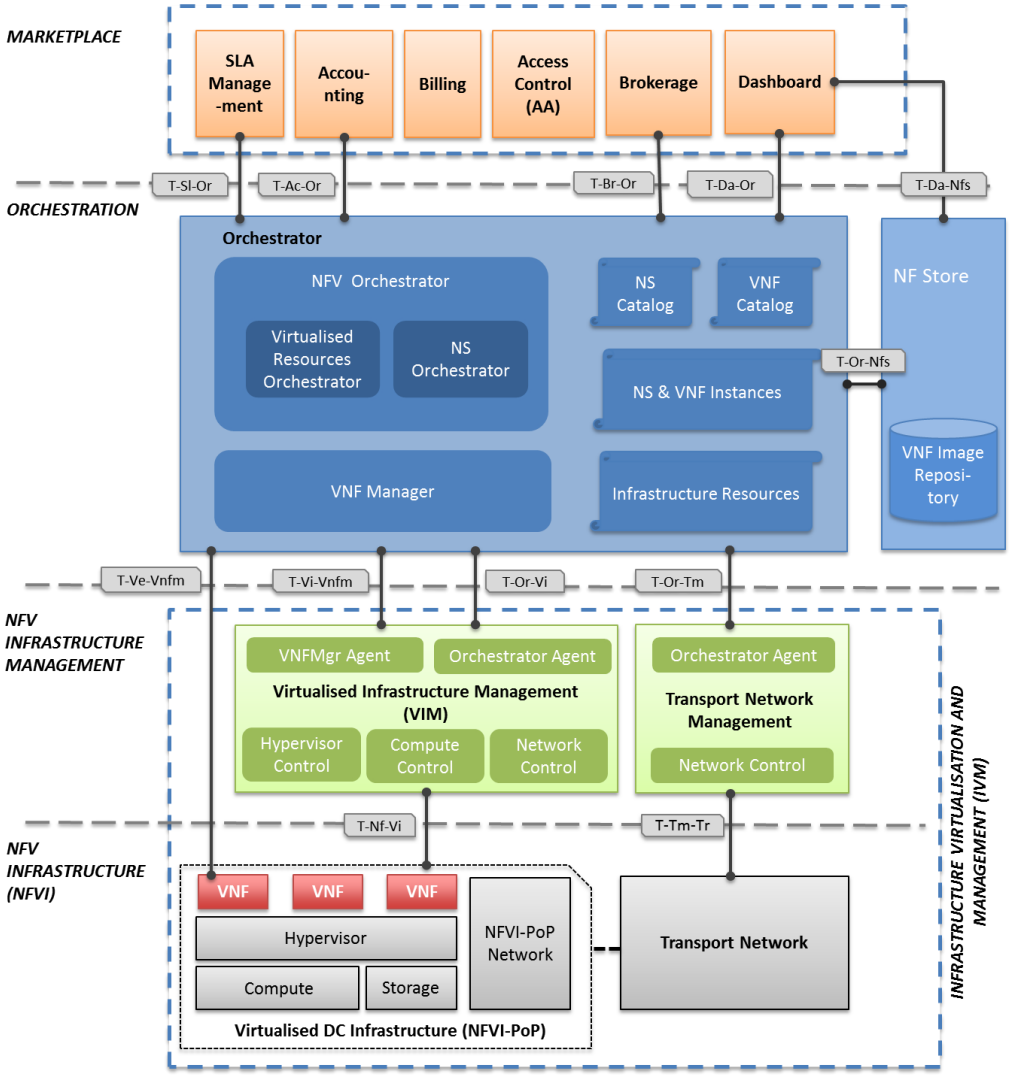

Figure 75. High-level view of overall T-NOVA System Architecture

- The *NFV Infrastructure (NFVI)* layer – comprising the physical equipment
- The *NFVI Infrastructure (NFVI) Management* layer comprises the management entities (Virtualised Infrastructure Manager - VIM, Transport Network Manager - TNM) which control physical and virtual resources
- The *Orchestration* layer which coordinates end-to-end services across several domains
- Finally, the *Marketplace* layer contains all the customer-facing modules which facilitate multi-actor involvement and implement business-related functionalities.

5.2. High-Level Architectural Requirements

This section discusses the system-level requirements for the CloudSat architecture, as driven by the integration scenarios already defined in Chapter 4. The ETSI NFV requirements document [ETSIREQ] has also been taken into account.

In the following list, which summarises the main high level architectural requirements, we refer to the "CloudSat system" as the complete end-to-end implementation of the proposed architecture, encompassing all infrastructure, management and orchestration entities in a federated satellite-terrestrial context.

Overall Requirements

- Interoperability. The CloudSat system shall support interoperability between the satellite and the terrestrial network domains.
- Virtualisation, infrastructure sharing and multi-tenancy. The CloudSat system shall be able to abstract infrastructure resources and allocate them between different tenants.
- NFV support. The CloudSat system shall be able to accommodate and dynamically manage virtualised network appliances (VNFs) offered as-a-Service, either as single instances or as sets (service function chains).
- Federated multi-domain management. The CloudSat system shall be able to manage the terrestrial and satellite domains in a federated manner, in order to achieve e.g. load balancing, mobility management etc.
- Resiliency. The CloudSat system shall be resilient to changes and faults, performing the necessary actions to minimise their impact in service continuity and performance.
- Modularity. The CloudSat system shall be able to be deployed in a modular manner, following an evolutionary path, also exploiting already deployed infrastructure as well as legacy technologies.

Service-related Requirements

- Network service request. The CloudSat system shall be able to handle and process network service requests in an automated manner.
- Network service mapping. The CloudSat system shall be able to optimally map network service requests received from customers to the infrastructure, such that all service requirements are met (e.g. computational requirements, network topology and QoS).
- Network service deployment. The CloudSat system shall be able to deploy the requested service by automatically assigning the necessary computing, network and storage resources needed.
- Network service monitoring. The CloudSat system shall be able to monitor the deployed services and expose monitoring information to the Customer.
- Network service elasticity and reconfiguration. The CloudSat system shall be able to dynamically scale-up/down the allocated resources to existing network services to deal with traffic variation and SLA contracts.
- Network service programmability. The CloudSat system shall be able to offer programmatic interfaces for the interaction with customers' high-level network control applications.

Resource Handling Requirements

- **Resource awareness.** The CloudSat system shall be aware of the status of network, computing and storage resources across the infrastructure.
- **Resource allocation on-demand.** The CloudSat system shall be able to allocate physical resources on-demand according to the requested network services.
- **Resource isolation.** The CloudSat system shall be able to logically isolate resources dedicated to collocated network services.
- **Resource efficiency.** The CloudSat system shall optimally allocate resources among services, yielding satisfactory resource utilisation.

Other Requirements

- **SLA monitoring.** The CloudSat system shall be able to support Service Level Agreement (SLA) handling. Any violations in SLAs should be promptly reported in order to trigger a service reconfiguration.
- **Billing.** The CloudSat system shall be able to support diverse billing models, such as flat rate and pay-as-you-go billing.

5.3. CloudSat Reference Architecture

5.3.1. Overview

The previous chapter presented an overview of the high-level system requirements of the generic CloudSat architecture, which need to be fulfilled in order to accommodate all the use cases / scenarios described in Chapter 4. It is possible to distill them down to a basic set of features/functionalities of the system and use them to derive architectural entities which must be defined. These driving features are:

- **Support of both satellite and terrestrial domains**; the architecture needs to include multiple discrete satellite and terrestrial administrative network domains, which maintain their independency.
- **Federation of satellite and terrestrial domains**; the different satellite and terrestrial domains need to adopt similar management technologies and expose similar northbound abstractions, so that they can be managed in a uniform manner. A federated management entity also needs to be in place for this purpose.
- **Accommodation of Virtual Network Functions (VNFs);** for this purpose, the infrastructure needs to include NFVI-PoPs (points of presence), based on clusters of commodity servers.
- **Flexible and scalable infrastructure management;** for this purpose, it is deemed appropriate to employ a two-tier hierarchical management structure, where each infrastructure segment is managed by a virtualised infrastructure manager, able to control both SDN and non-SDN network elements. At higher

level, an orchestrating entity needs to be in place in order to enable end-to-end service management.

- **User-friendly interface to customers for service deployment and management**; a dedicated front-end portal needs to be considered to facilitate interaction with the customers.

Taking into account the aforementioned features, it becomes possible to draft a high-level integrated satellite/terrestrial architecture which fulfils the requirements and which is in principle able to accommodate all the use cases / scenarios described in Chapter 4. The proposed architecture is depicted in Figure 76.

It must be noted that the presented architecture is technology-agnostic; i.e. it does not mandate a specific technology for e.g. SDN or cloud implementation. However, some technology recommendations are given and justified in the sections to follow, which describe in more detail the various architectural entities.

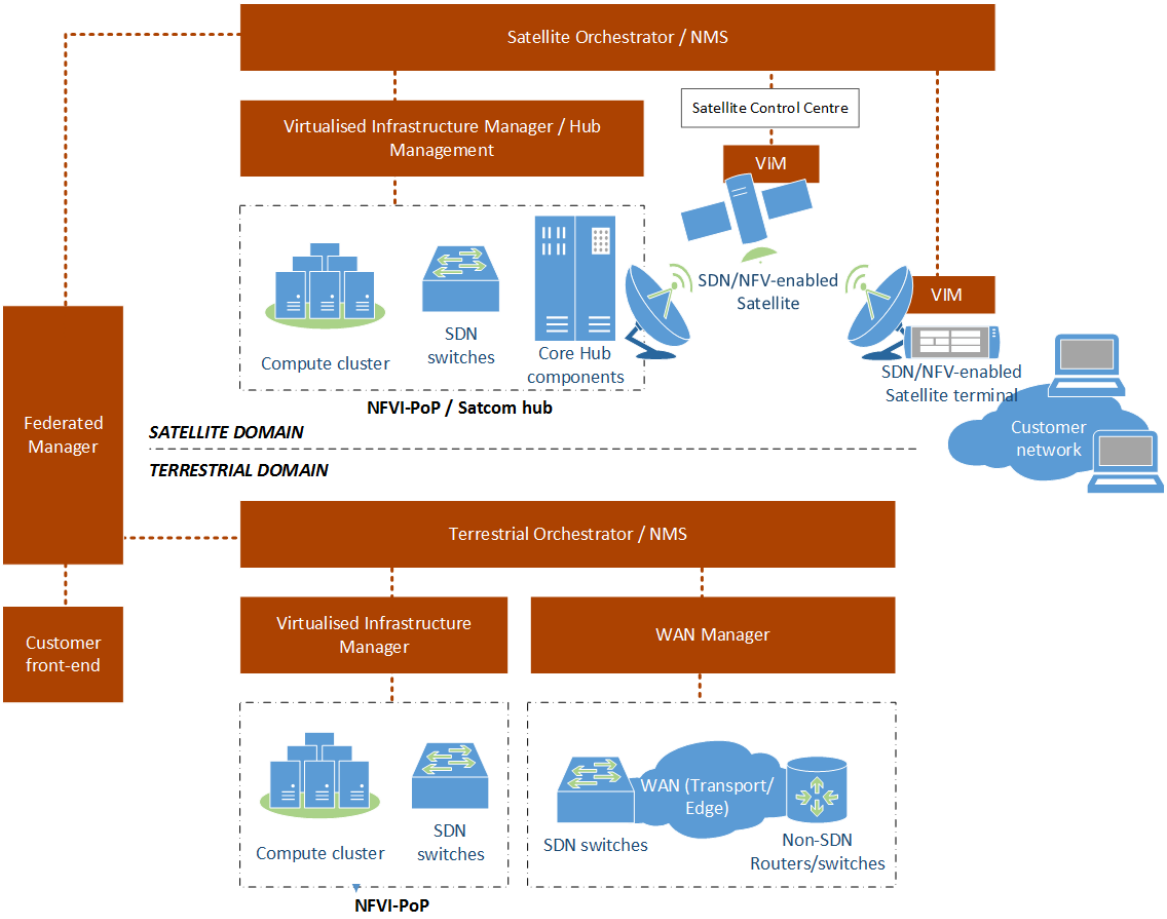

Figure 76. CloudSat reference architecture

It can be seen that the architecture is split into three main parts: the satellite segment, the terrestrial segment and the federated management/front-end. Each segment corresponds to a different administrative network domain; although the

architectural diagram includes only one satellite and terrestrial domain, it can be expanded to include as many domains as appropriate.

Also, although the figure illustrates a single satcom hub and satellite in the satellite operator domain, the architecture can support multiple hubs and satellites e.g. in a multi-spot or multi-satellite scenario involving High Throughput Satellites (HTSs) and constellation-based systems.

In order to derive a better insight of the hierarchical logic of the reference architecture, the architectural figure can be transposed from a domain-centric to a layered view (Figure 77). As it can be seen, each satellite and terrestrial domain are split into three logical layers:

- The Infrastructure (lower) layer includes the virtualization-capable equipment on which the network service is deployed. This layer includes:
    - the SDN and non-SDN network elements of the terrestrial Wide Area Network (transport and edge)
    - the distributed NFVI-PoPs (data centres with compute clusters with the supporting SDN network). In the satellite segment, the hubs themselves are seen as NFVI-PoPs comprising of cloud assets, SDN network plus the “traditional” core hub components
    - in the satellite segment case, the satellite itself as well as the customer terminal also belong to the Infrastructure layer, both assumed to be SDN/NFV capable.
- The Infrastructure Management (middle) layer includes distributed management entities for the various parts of the infrastructure. The SDN/NFV enabled segments (NFVI-PoPs) are managed by a Virtualised Infrastructure Management (VIM) entity. In the satellite segment, a tailored (lightweight) VIM is also assumed to control the SDN/NFV-enabled satellite (via the satellite control centre) as well as the terminal. If these two components are not SDN/NFV capable, then the VIM can be omitted and replaced by the respective traditional management modules, leaving the rest of the architecture unaffected. In addition to the VIM, a legacy management entity needs to be foreseen for the management of the non-SDN/NFV components of the satcom hub. Last but not least, for the core/edge terrestrial network part, a Wide Area Network (WAN) Manager is assumed, supporting both SDN and non-SDN elements; while it is beneficiary to employ SDN at several parts of the WAN network, support of legacy non-SDN elements is deemed necessary to ensure interoperability.
- The Orchestration (top) layer is responsible for the coordination of the entire administrative domain, the infrastructure as well as the services which run on it. It orchestrates virtualised resources to compose end-to-end services and optimizes them dynamically. This role is undertaken by an Orchestrator entity, which normally closely interacts with (or ideally is integrated in) the operator’s overall Network Management System (NMS).

Furthermore, to achieve federation among different satellite and terrestrial domains, a fourth, top-most architectural layer (**Federation** layer) is foreseen, which acts as an umbrella and coordinates inter-domain services. This layer includes a federated manager entity, which interacts with the Orchestrators of the cooperating domains. It also includes a customer front-end, which facilitates user-friendly interaction with the customers for service deployment, management, monitoring and teardown. This front-end also exposes a programmatic interface (API), which enables customers to deploy and operate arbitrary applications, controlling their own tenant networks services.

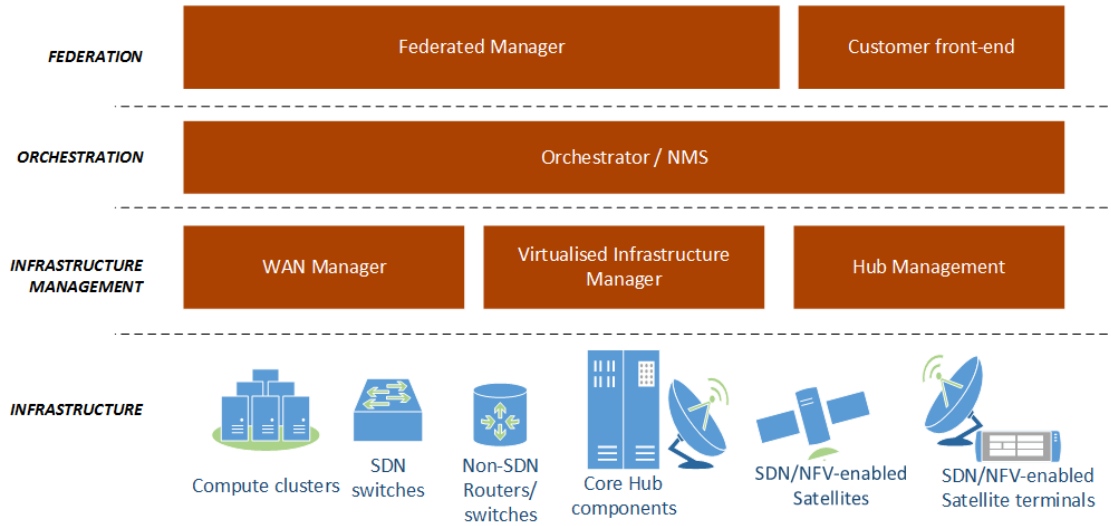

Figure 77. Layered view of the CloudSat reference architecture

All the aforementioned architectural entities are described in more detail in the sections to follow.

Finally, with regard to the cardinality of the various architectural entities within the integrated architecture, since it is not explicitly illustrated in the above figures, it is clarified in the table below.

Table 14. Cardinality of entities in the overall CloudSat architecture

Architectural entity	Cardinality
Customer front-end	One
Federated manager	One
Satellite domains	One or more - typically one per operator
Terrestrial domains	One or more - typically one per operator
Satellite Orchestrator	One per satellite network operator domain
Terrestrial Orchestrator	One per terrestrial network operator domain

VIM	One per NFVI-PoP
NFVI-PoPs	One or more per domain
Satcom hub	One or more per domain
Satellite	One or more per domain
Wide Area Network (Transport/Edge)	Typically One per domain

It must be noted that the four-layer approach has been adopted to achieve an optimal trade-off between the federated management capabilities with the key requirement to preserve the administrative independence of each network operator; this requirement mandates that the top-most management entity of each domain (Orchestrator) is fully controlled by the network operator and only exposes certain service (and not infrastructure) management services via its northbound interface to the Federated manager, which may be operated by a third party. Alternatively, the Federated manager could be operated by one of the network operators involved in the federation (satellite or terrestrial), who would play the role of the federating actor in this case.

If we assume that the satellite and terrestrial domains are owned and managed by the same business entity, then the architecture of Figure 76 can be simplified to omit the Federated manager. In this case, the entire satellite/terrestrial infrastructure can be managed by a single Orchestrator platform (Figure 78).

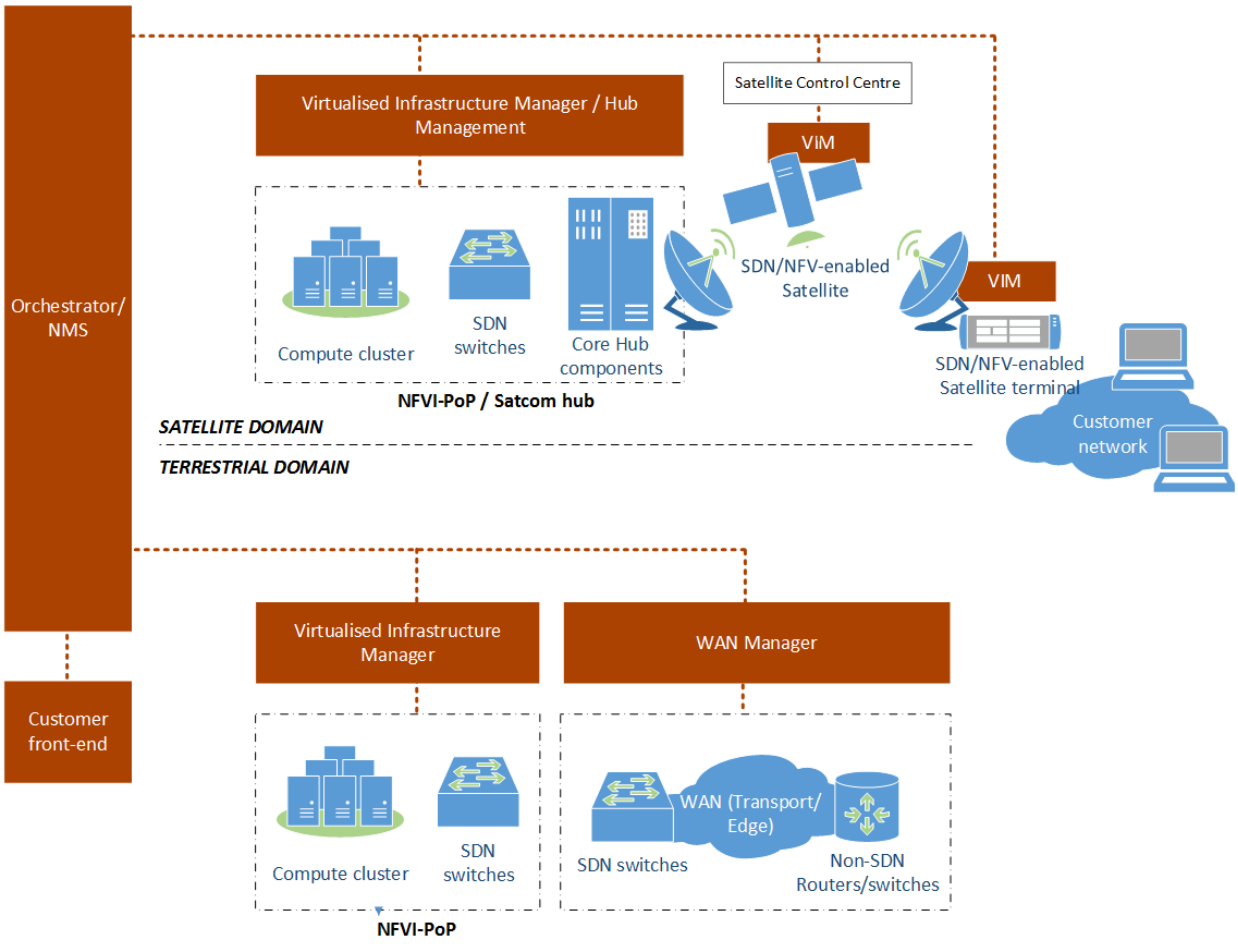

Figure 78. Single-operator variant of the CloudSat reference architecture

This variant (i.e. with the Orchestrator as the top-most management entity) is applicable also to terrestrial-only and satellite-only configurations (i.e. when no satellite/terrestrial federation is foreseen).

5.3.2. Main architectural entities

5.3.2.1. Infrastructure layer

Compute clusters

Compute clusters (clusters of commodity servers) are included in the architecture in order to be able to accommodate Virtual Network Functions (VNFs) as workloads. These clusters are organized in islands (data centres), the NFVI-PoPs (NFV Infrastructure Point-of-Presence). We also assume the integration of compute infrastructure with the satcom data hubs in order to enhance the latter with NFV capabilities.

A NFVI compute cluster is inherently virtualization-capable and consists of:

- The *Compute* domain, which represents the lowest (physical) level, comprising the computing and storage equipment (standard high-volume servers with or

without specialized hardware accelerations and storage infrastructure). For standard-scale data centre implementations, servers based on the x86 architecture are a common choice. The adoption of features for hardware-assisted virtualization, such as DPDK (Data Plane Development Kit) support [DPDK] and SR-IOV (Single-Root I/O VIrtualisation) [SRIOV], seems quite promising for the enhancement of VNF performance and is thus recommended.

- The *Hypervisor* domain, which is responsible for the abstraction of the physical compute and storage resources (possibly aggregated across multiple physical elements) and their assignment/allocation to VNFs. The hypervisor commonly exposes a northbound interface for the interaction with the Management layer. Several choices are available for the hypervisor technology (Hyper-V, VMware, KVM, Xen etc.), heavily depending on the compatibility with the VIM and also with the physical infrastructure. Kernel-based Virtual Machine (KVM) [KVM] would be a safe recommendation, given its openness, wide compatibility and full-featured integration with Openstack.

SDN switches

SDN-enabled switches are considered throughout the CloudSat architecture in order to facilitate SDN-based network management. Both physical and virtual switches are foreseen:

- Within the NFVI-PoP: both physical and virtual switches supporting the communication within the NFVI-PoP and interconnecting the compute clusters as well as the virtual machines/VNFs.
- As a part of the terrestrial Wide Area Network (WAN): physical switches
- As a part of the SDN-enabled satcom hub: physical and virtual switches enabling SDN control of the hub network

With regard to the SDN technology, Openflow would be the prevailing option, due to its flexibility and significant momentum, preferably at its latest version (1.4) [OF14] although other SDN alternatives such as NETCONF/YANG [NETCONF] could be considered.

For the physical switches, no specific vendor recommendation can be made, since Openflow-compatible switches are already available from several vendors. However, if compatibility with the latest Openflow version is an absolute requirement as well as upgradeability, more “open” switch architectures in terms of hardware and firmware should be preferred.

For the virtual switches, Open vSwitch [OVS] is by far the prevailing option, used in many production infrastructures, including large-scale cloud deployments.

Non-SDN routers/switches

In addition to SDN elements, non-SDN (legacy) components are also considered at various parts of the network infrastructure substrate, especially at the WAN core and also at the data centre edges. These elements are managed in a traditional, non-SDN manner and are assumed to be semi-statically configured; i.e. they do not actively participate in the lifecycle of a cloud network service.

For example, non-SDN elements could be core and edge routers, optical switches etc. Their role is to maintain core network connectivity and also to interconnect NFVI-PoPs and network domains. Although, in theory, the CloudSat architecture could solely rely on SDN networking, it is considered realistic to assume non-SDN elements also in order to yield a feasible architecture, able to be deployed in an evolutionary manner, also taking advantage of legacy hardware.

Core hub components

The typical ("traditional") components of the satcom hub are also considered part of the infrastructure. These include e.g. the FL encapsulator/multiplexer, FL modulator, RL demodulator, as well as network-layer functions such as firewall, QoS shaper, PEP, etc.

Depending on the "depth" of SDN/NFV adoption in the satcom hub, several configuration variants could be considered.

In its simplest form, SDN-managed NFV-PoPs could be connected "back-to-back" with existing hubs, leaving the satellite hubs unaffected. This allows NFV deployment and SDN control at the edge of the satcom network, yet deprives the core of the satellite hubs of the benefits of softwarisation and virtualization.

In a more tight integration, many of the L3/L2 network components, such as the firewall, entry router, QoS shaper and the multiplexer could be also SDN-enabled, enabling granular SDN control at several stages within the transmission/reception chain.

Finally, in a long-term scenario, we could consider most hub components to be virtualised and deployed as VNFs, allocated to either a single or multiple tenants. Of course, this virtualization/multi-tenancy should be applied only to operations/functions which the hub operator is willing to expose to customers.

Satellite payload

Depending on the use case/scenario to be realized, the CloudSat architecture may assume or not an active role of the space segment in the lifecycle of the cloud network service.

The simpler approach would consider a typical transparent satellite providing bent-pipe connectivity and independent from any cloud networking mechanisms and services. In this case, the satellite is considered as a passive network link.

In the long term, an SDN/NFV enabled regenerative satellite would be SDN-controlled and an NFV capable.

SDN capability implies that the switch fabric of the regenerative satellite can operate at L2 or higher and comply to SDN flow rules and commands. Regenerative multi-spot satellites would be especially candidate for such an evolution, since the satellite, from an SDN point of view, could be seen as a switch with multiple input/output ports. As discussed, a software switch such as Open vSwitch could be ported to a satellite platform to provide SDN features – although an FPGA SDN implementation such as NetFPGA [Naous08] would be more suitable for on-board use.

NFV capabilities would imply that the satellite payload would be virtualization-capable and able to host VNFs as virtual machines on-board. This means that the payload should play the role of a small NFVI-PoP – albeit rather constrained in terms of resources and capabilities. For this purpose, a lightweight computing architecture (e.g. based on ARM processors) would be more appropriate, rather than the x86 platforms mentioned for terrestrial NFVI-PoPs. For the hypervisor domain, it would be more appropriate to move away from the "full" virtualization realized through traditional hypervisors such as KVM and adopt a more lightweight approach based on containers, such as the popular Docker platform [Docker].

In any case, since, in the normal satellite value chain, the payload is expected to be administered by the satellite operator (and thus it cannot be controlled by the satcom SP), any configuration/management commands/requests should go through the satellite control centre. The latter is responsible to handle multiple requests by various SPs, validate them and then forward them to the payload.

Last but not least, whatever the technical approach adopted, there does not seem to be any dependency/influence among the SDN/NFV mechanisms and the RF parameters of the satellite transponders, such as band of operation (Ku, Ka, S etc.), EIRP or bandwidth.

Satellite terminal

As with the satellite case, the CloudSat architecture also proposes the extension of the SDN/NFV capabilities to the satellite terminal, although it can also accommodate legacy terminals.

SDN capabilities at the satellite terminal allow the SDN-based control of flows i) to and from the satellite network and ii) among different segments of the local network, served via different physical ports of the terminal.

In order to introduce SDN capabilities at the terminal, the solution of integrating a software SDN implementation, such as Open vSwitch, as discussed, in the terminal firmware, seems quite feasible.

NFV capabilities at the satellite terminal allow the local instantiation and hosting of VNFs, mostly serving as middleboxes, processing the traffic to and from the satellite network, as well as the local traffic. Given the resource constraints at the terminal, a container-based solution such as Docker, deployed on a lightweight ARM-based architecture for NFV at the payload, seems suitable for the terminal case also.

5.3.2.2. Infrastructure Management

Virtualised Infrastructure Manager (VIM) for terrestrial NFVI-PoPs

Conforming to ETSI ISG NFV terminology, the Virtualised Infrastructure Manager (VIM) is the functional entity that is responsible for controlling and managing the infrastructure (compute, storage and network) resources. The management scope of the VIM is generally restricted within a single NFVI Point of Presence (NFVI-PoP). This means that a NFVI-PoP is assumed to be managed by a single VIM platform (1:1 cardinality)

While a VIM, in general, can potentially offer specialisation in handling certain NFVI resources, in the CloudSat context (but also in other architectures), the VIM is seen to encompass all management and control functionalities needed for the proper administration of the infrastructure, as well as the virtualised services running on top of it.

In specific, the following key tasks are performed by the VIM:

- Maintenance of a resource, capability and topology repositories/inventories, thus establishing a comprehensive “map” of the underlying hardware;
- Joint management of the infrastructure (compute, storage, networking) resources;
- Association/mapping of the virtualised services to the infrastructure resources;
- Basic network control services, including topology management and path computation;
- Management (create, query, update, delete) of service function chains i.e. the interconnections among VNFs, by creating and maintaining Virtual Links, virtual networks, sub-nets, and ports;
- Management of VNF software instances (add, delete, update, query, copy);
- Management of virtual networks, tunnels and QoS, where applicable;
- Collection and communication of measurements and faults/events information relative to physical and virtual resources.

In order to realise these tasks, the VIM needs to comprise the following components, as shown in Figure 79:

- A Resource repository database – for maintaining a comprehensive landscape of the underlying infrastructure, the exposed capabilities and the available resources;
- A Topology and service function chain management module – this component undertakes most network management tasks, including virtual network management, interconnection of virtualised components and tunnel establishment;

- A Compute/hypervisor management module – this component undertakes the management of VMs/VNFs;
- An integrated monitoring and event/alarm management framework – for efficient and effective collection of metrics and production of events/alarms;
- A set of southbound interfaces for managing the infrastructure (compute nodes, hypervisors, network elements). These commonly come in form of "plug-ins" in order to accommodate multiple infrastructure technologies (such as several hypervisors, network management protocols etc.);
- A set of northbound APIs (commonly REST-based) for communication with the upper layer (Orchestrator).

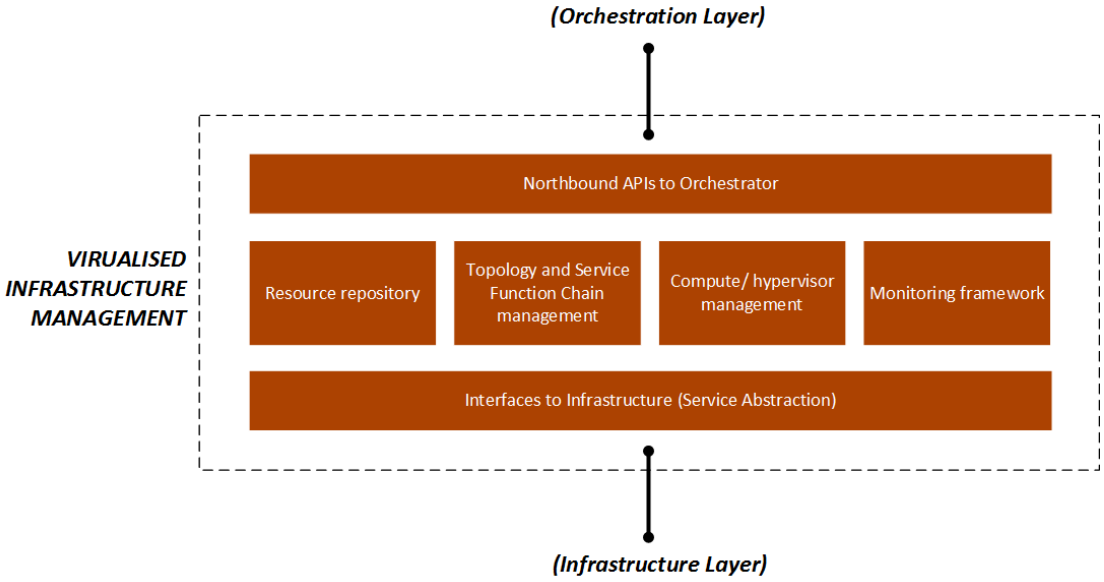

Figure 79. VIM and core components

From the implementation point of view, a VIM is commonly realized by coupling a network controller and a cloud controller platform. As an example, most VIM platforms under development, including the one to be released via the OPNFV initiative [OPNFV] are based on a customized combination of the Openstack [Openstack] cloud controller with the OpenDaylight [ODL] network controller platform. Indeed, Openstack and OpenDaylight constitute the two most popular candidate technologies for VIM implementation to date.

Virtualised Infrastructure Manager (VIM) for satellite payloads and terminals

In case of SDN/NFV-enabled satellite payloads and terminals, it would make sense to consider these architectural entities as a special case of NFVI-PoPs, although with considerably reduced resources and capabilities.

Therefore, a dedicated VIM component should be assigned to manage the SDN/NFV resources of the payload and/or terminal. This assumption would on the one hand

preserve the layered logic of the overall architecture and on the other hand (most important) promote self-management capabilities on the payload and the terminal, which would greatly reduce the interchange of management-related signaling over the satellite link. In other words, the SDN/NFV-enabled satcom component, "bundled" with the local VIM would constitute an independent, virtualization-capable infrastructure islet, which would require only minimum communication with the rest of the platform for management/administration purposes.

As opposed to what was presented in the previous section for terrestrial NFVI-PoPs, the adoption of a cloud and network controller, such as Openstack/Opendaylight «as-is» for managing satcom payloads and terminals would be not only an overkill but also inappropriate, since these controllers i) are quite "heavy" and resource-demanding to be deployed, ii) assume multiple physical infrastructure elements, iii) involve heavy management signaling and iv) employ multiple southbound plugins for interfacing with the infrastructure (which are not necessary in this case).

Thus, it is considered essential to develop tailored VIM components for payload/terminal use. These VIM components need to have a minimal footprint, being specific for the underlying infrastructure and tailored for it. They should implement only a basic essential subset of the general VIM functionalities, in specific:

- Management of VNFs (as VMs or containers)
- Installation of flow rules on the local SDN switch
- Management of network interfaces and physical compute resources
- Infrastructure and service monitoring

In additional, this lightweight VIM should expose a northbound API to the Orchestrator. For the sake of uniform management, this API should be compliant with the API exposed by the terrestrial "full-scale" VIM.

As stated, a payload VIM should only be directly controlled by the satellite control centre. Any commands or queries to the payload VIM by the Orchestrators of the satcom SPs using the satellite should go though (and possibly checked/validated/policed by) the CC.

WAN Manager

In order to realise federated satellite/terrestrial end-to-end network services, the management entity for the Terrestrial Wide Area Network (WAN) also needs to be considered as an active element of the reference architecture. The main function of the WAN manager is to carry out all administrative tasks of the WAN (i.e. setup and configuration of network elements, supervision, monitoring etc.). Moreover, it has to manage the virtual network services installed on top of the WAN infrastructure, and especially the interconnection between the NFVI-PoPs, mostly in the form of inter-data-centre tunnels; in this sense, the WAN Manager implements VIM functionalities and can be seen as a specific VIM entity. The Orchestration layer indeed treats the WAN Manager as a VIM and uses its services to setup, configure and teardown virtual networks across service endpoints.

Overall, the functionalities to be provided by the WAN Manager, are:

- Management and control of WAN physical network elements
- Monitoring of WAN resources and virtual links in order to provide the Orchestrator with useful statistics (such as jitter, RTT, delay, bandwidth, etc.) to make decisions about allocation of network resources;
- Management of virtual/physical links between NFVI-PoPs via configuration of:
    - SDN-enabled network elements, that enable network slicing techniques,
    - Legacy network elements, relying on tunnelling protocols (e.g., VXLAN, NVGRE, STT, see relevant overview in Chapter 2) in case of L3 network elements or on native trunking/aggregation protocols in case of L2 elements (e.g. VLAN, Q-in-Q, etc.),
- Interfacing to the Orchestrator in order to accept provisioning requests and to submit monitoring information.

Regarding the southbound communication with the infrastructure, as highlighted in Sec. 5.3.2.1. , the WAN segment is assumed to comprise both SDN and non-SDN physical network elements. For this purpose, the WAN needs to interface with the network using several network configuration protocols, including Openflow, Netconf, SNMP, etc.. In some cases, there will also be the need to interact with the infrastructure with vendor-specific management protocols also, in order to maintain compatibility with installed large-scale infrastructures based on existing hardware. For this purpose, the southbound part of the WAN manager should employ a Service Abstraction layer, as also discussed for the VIM, within which support for various network configuration protocols should be enabled in the form of plug-ins.

Other components which should be considered a part of the WAN Manager, are very much related to the VIM components for network control i.e. topology manager, resource repository and monitoring framework.

Coming to the implementation technologies, although a multitude of well-established network management platforms exist, which can scale-up in a satisfactory manner to manage a complex WAN, it should be highly advisable to employ an emerging open network control platform, such as OpenDaylight also for this purpose (in addition to NFVI-PoP network control, as suggested). The motivation behind such a choice is the openness, dynamicity and wide community support of ODL which make it a promising solution for managing virtualised services over future heterogeneous networking infrastructures. Nevertheless, it may take some years until ODL reaches the stability and scalability which is required to manage a “production” large-scale operator network.

Hub Management

With the satcom hub being an integral part of the satellite communications infrastructure, the management components of the hub need to be included in the Infrastructure Management layer of the reference architecture.

Traditional hub management functionalities include the management of terminals and subscribers, QoS and applications configuration, PEP parameters configuration and FL multiplexer / modulation settings as well as RL parameters.

Since, as aforementioned, the satcom hub is assumed to be enriched with SDN/NFV capabilities, the hub management is expected to co-exist with a VIM platform, "horizontally" placed in the same architectural layer.

As discussed, three configuration variants are considered with regard to the "depth" of SDN/NFV adoption in the satcom hub. We review here these variants, discussing their impact on the configuration of the infrastructure management layer.

- Simplest form: "back-to-back" connection of an SDN-managed NFVI-PoP with an existing hub. In this case, the hub management and VIM platforms can be totally independent, probably including some sort of horizontal communication to ensure that e.g. a set of flows processed by the NFVI-PoP is admitted and assigned an appropriate level of QoS within the satcom gateway.
- More tight integration: introduction of SDN capabilities into several hub L3/L2 network components (such as the firewall, QoS shaper and the multiplexer). In this case, the management scope of the VIM expands to embrace the SDN-enabled satcom components, to achieve a more uniform and integrated management.
- Long-term scenario: Full virtualization of most hub components as VNFs, allocated to either a single or multiple tenants. In this case, we should consider the practical elimination of the traditional hub management platform and the undertaking of most management tasks by the VIM, enriched with the appropriate plug-ins to manage the hub infrastructure.

Nevertheless, even in the long term scenario, a standard VIM platform and especially the network controller part should be considerably expanded in order to encompass all the specific requirements of satcom network management. This indeed is very challenging and goes far beyond the current controller paradigm. For example, taking into account an SDN-oriented generic network controller such as OpenDaylight, it is currently not possible to administer any L1 aspects, such as e.g. modulation and power control settings, even via a specialized plugin. Nor is it possible to remotely monitor and control satellite terminals. In order to achieve total integration in the management plane, the current vision of "network controller" should not also expand horizontally (to support satcom vendor-specific capabilities and features) but also vertically (to encompass aspects such as radio access management and terminal control, which are currently out of SDN management scope). Such an evolution could pave the way towards a fully SDN/NFV-enabled satellite Network Management System, capable to effectively deploy and manage virtualised services over a satellite infrastructure.

5.3.2.3. Orchestration

The term "orchestration" is a heavily overloaded term, in the sense that is used very often during the last years with diverse meanings. In the CloudSat architecture, we

consider an Orchestrator platform as the single top-level management entity of the domain. This is commonly the role played by a traditional Network Management System (NMS), that's why we refer to an integrated "Orchestrator / NMS" system.

The Orchestrator communicates with the underlying localized managers of the different network segments (see Sec. 5.3.2.2.) in order to "orchestrate" resources across the entire domain and thus realise the deployment and management of end-to-end services. The Orchestrator is the entity which maintains a complete view of the whole infrastructure of the domain; it keeps a record of installed and available resources, as well as of the infrastructure topology. For the sake of scalability, the Orchestrator maintains only a high-level view of the resources and the services, while the detailed mapping of services to resources is undertaken by the local managers (Infrastructure Management layer). This two-tier approach is compliant with many architectural approaches, as surveyed in Chapter 5.1.

Coming to the architectural details of the Orchestrator, we consider appropriate to inherit the functional structure of the Orchestrator platform developed in the T-NOVA project, as described in [TNOVAD231], which in turn is fully compliant with the ETSI proposal for network service orchestration, as defined in [ETSIMAN]. Most of the functional blocks described below are inherited from the ETSI specification document.

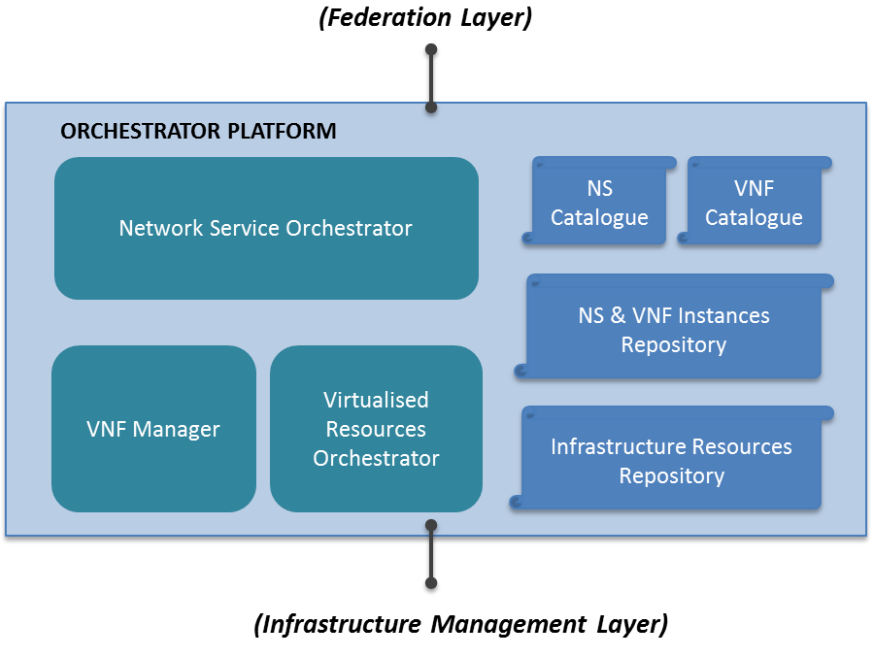

Figure 80. Orchestrator Platform functional components

As seen in Figure 80, the Orchestrator platform comprises both catalogs/repositories as well as execution components:

- The **Network Service Catalogue** contains a description of all available ("on-boarded") Network Service[14] templates, including descriptors such as Network

[14] It is reminded that the term "Network Service" refers to a connectivity service plus the associated VNFs, organized in a service graph.

Service Descriptor (NSD), Virtual Link Descriptor (VLD), and VNF Forwarding Graph Descriptor (VNFFGD). ETSI has already defined the data to be included in these descriptors. It is clarified that the NS Catalogue only contains a list of the available service templates which can be selected by the customer, not the deployed services themselves.

- The **VNF Catalogue** represents the repository of all of available on-boarded VNF Packages, supporting the creation and management of the VNF Package (VNF Descriptor (VNFD), software images, manifest files, etc.). The information contained in the VNFD is defined by ETSI. Again, it is clarified that the VNF Catalogue contains a list of the available VNFs which can be included in an NS, not the deployed VNFs themselves.
- The **NS&VNF Instances Repository** contains information of all service instances which have been actually deployed. The repository is frequently updated, to reflect the status and the lifecycle of the deployed virtualised services.
- The **Infrastructure Resources Repository** holds information about available/reserved/allocated NFVI resources as abstracted by the VIM across operator's Infrastructure Domains, thus supporting information useful for resources reservation, allocation and monitoring purposes.
- The **VNF Manager (VNFM)** is responsible for the lifecycle management of VNF instances. Each VNF instance is assumed to have an associated VNF Manager. A VNF manager may be assigned the management of a single VNF instance, or the management of multiple VNF instances of the same type or of different types. Operations carried out by the VNF Manager are VNF instantiation and feasibility checking; integrity management; VNF instance modification/scaling/healing/termination.
- The **Resource Orchestrator (RO)** is the Orchestrator component which mainly interacts with the VIM for resource discovery, allocation and management. The RO allows the Orchestrator platform to manage and control distributed resources across multiple NFVI-PoPs.
- Finally, the **Network Service Orchestrator (NSO)** is the core decision-making component, actually the “kernel” of the Orchestrator platform. The NSO instantiates network services (using the NS and VNF templates in the corresponding catalogues) and manages the whole NS lifecycle. For this purpose, it communicates with the VNFM and the RO for the control of VNF instances and the (re-)allocation of virtualised resources. This task includes the control of the network assets, for virtual network establishment and QoS provision. Moreover, the NSO monitors the deployed NSs, deciding on auto-scaling or healing functions, if necessary.

Regarding southbound communication, the Orchestrator interfaces with the Virtualized Infrastructure Management (VIM), for managing the data center network/IT infrastructure resources, as well as with the WAN and satcom hub management for controlling the wide-area network resources (connectivity management). Finally, the Orchestrator interacts directly with the deployed VNFs in order to ensure their lifecycle management.

For the northbound communication, the Orchestrator exposes a set of API services to the federation layer, allowing the latter to request, deploy and configure services. In the single-domain scenario, where the Federation layer is missing, a customer front-end GUI is directly attached to the Orchestrator, enabling customers to request, deploy and configure their services. In the single-domain case, the Orchestrator also needs to encompass customer-oriented BSS (Business Support System) operations, such as accounting and billing, while in the multi-domain operation these functions are transferred to the Federation layer (see next section).

Last but not least, an administrative GUI needs to be foreseen, enabling the domain administrators to manage the infrastructure and the deployed network services, thus allowing the Orchestrator platform to serve similar functionalities as a traditional NMS or Operations Support Systems (OSS). If a "legacy" NMS/OSS needs to be in place, for any reason, then a horizontal interface between the NMS/OSS and the Orchestrator needs to be established, allowing these two entities to exchange information about the status of the infrastructure and the deployed services.

5.3.2.4. Federation

Federated Manager

In the multi-domain scenario, the Federated Manager is the highest-level management entity of the CloudSat infrastructure. A Federated Manager is deemed essential so as to enable coordinated managemend of inter-domain services (e.g. services spanning across the satellite and terrestrial segments), yet without violating the administrative independence of the involved domains, which may belong to different business entities.

For this purpose, the Federated Manager only maintains a very high-level view of the underlying infrastructure, without being aware of implementation and topology details, as well as monitoring data, which are internal to each domain operator. The Federated Manager does not actually control the underlying infrastructure; its operation is restricted in high-level functionalities, such as:

- Processing customer requests and mapping them to the domain capabilities;
- Issuing service creation requests to the underlying Orchestrators and coordinating inter-domain connectivity;
- Managing the interactions among several business actors;
- Implementing business-related functionalities, such as SLA management accounting and billing.

For this purpose, the following functional components are considered as part of the Federated Manager:

- Access control: this component provides authentication and authorization functionalities to manage the access to the federated system by the different stakeholders. In other words, it regulates who is allowed to access the federated system and what is it allowed to do.

- Accounting module: this module is in charge of registering all business relationships (subscriptions, SLA evaluations and usage) and making the related information available for the billing system.
- Brokerage module: it is the component that receives the customer requests, and will present at the customer the most suitable offerings that matches his/her requirements; depending on the applicable trading-policies the necessary actions to get the best price for each service+SLA when creating a new service starting from the VNFs will be carried out.
- SLA management module: it is the component that will register all the SLA agreements among the involved parties, checks if the SLAs have been fulfilled or not, and informs the accounting system for the pertinent billable items.
- Billing system: It is the component that produces the bill for a customer on behalf of the Service Provider.

As also stated in the previous section, in the single-domain scenario where the Federation layer is omitted, then all the aforementioned functionalities should be integrated in the Orchestrator platform, in order to fulfil all the requirements of a traditional OSS/BSS.

Customer front-end

The customer front-end enables the interaction between customers and the federated CloudSat architecture. The basic operations exposed via the customer front-end are:

- Service advertisement (i.e. presentation of catalogs of available network service templates and VNFs)
- Service request
- Service deployment, management and teardown
- Service monitoring
- SLA and billing management.

In the multi-domain scenario, the customer front-end interacts with the Federated Manager for implementing the aforementioned functionalities, exploiting the corresponding functional blocks of the FM, as discussed, In the single-domain scenario, the interaction is performed directly with the Orchestrator, exploiting the Orchestrator northbound API for network service management.

The interaction between the customer and the front-end can be realized in two ways:

- Either in a graphical, user-friendly manner, via a web-based Graphical User Interface (GUI). The graphical interface allows direct human interaction and helps to visualize service parameters, such as network service topology and service monitoring metrics, which can be consolidated and presented altogether in a comprehensive “Dashboard”, also including billing information and SLA status.

- Or in a programmatic manner, via an API (commonly REST-based) to be exposed to custom user applications, especially developed and tailored to automate service deployment and management processes, according to the customer's special needs.

Overview of technology recommendations

Although, as previously explained, the CloudSat reference architecture is proposed to be technology-agnostic, the former sections also indicated some specific technology recommendations for some of the architectural entities. These recommendations are overviewed in the table below and are aligned with the overall conclusions on technology suitability presented in Chapter 3. The technologies referred to have been already identified in Chapter 2.

Table 15. Overview of specific technology recommendations

Architectural layer	Architectural entity	Recommended technology
Infrastructure	Compute cluster (compute domain)	Multi-core x86-based architecture Hardware-assisted virtualization (DPDK support [DPDK], SR-IOV [SRIOV] etc.)
Infrastructure	Compute cluster (hypervisor domain)	Kernel-based Virtual Machine (KVM) [KVM]
Infrastructure	SDN physical and virtual switches (SDN protocol)	Openflow 1.4 [OF14]
Infrastructure	SDN virtual switches (Openflow switch implementation)	Open vSwitch [OVS]
Infrastructure	Satellite payload (NFV enabler)	Docker [Docker] on ARM-based architecture
Infrastructure	Satellite payload (Openflow switch implemenetation)	FPGA-based
Infrastructure	Satellite terminal (NFV enabler)	Docker [Docker] on ARM-based architecture
Infrastructure	Satellite terminal (Openflow switch implemenetation)	Open vSwitch [OVS]
Infrastructure Management	Virtualised Infrastructure Management (VIM) – Cloud Controller	Openstack [Openstack]
Infrastructure	Virtualised Infrastructure	OpenDaylight [ODL]

Management	Management (VIM) – Network Controller	
Infrastructure Management	WAN Manager	OpenDaylight [ODL]

For the architectural entities not included in the above table, mostly from the Orchestration and Federation layers, no well-established industry-standard technology exists to date, which would justify a recommendation.

5.3.3. Reference points

This section presents the interaction between the architecture subsystems, as described in Sec. 5.3.2., by identifying specific reference points at the subsystems' boundaries, as shown in Figure 81.

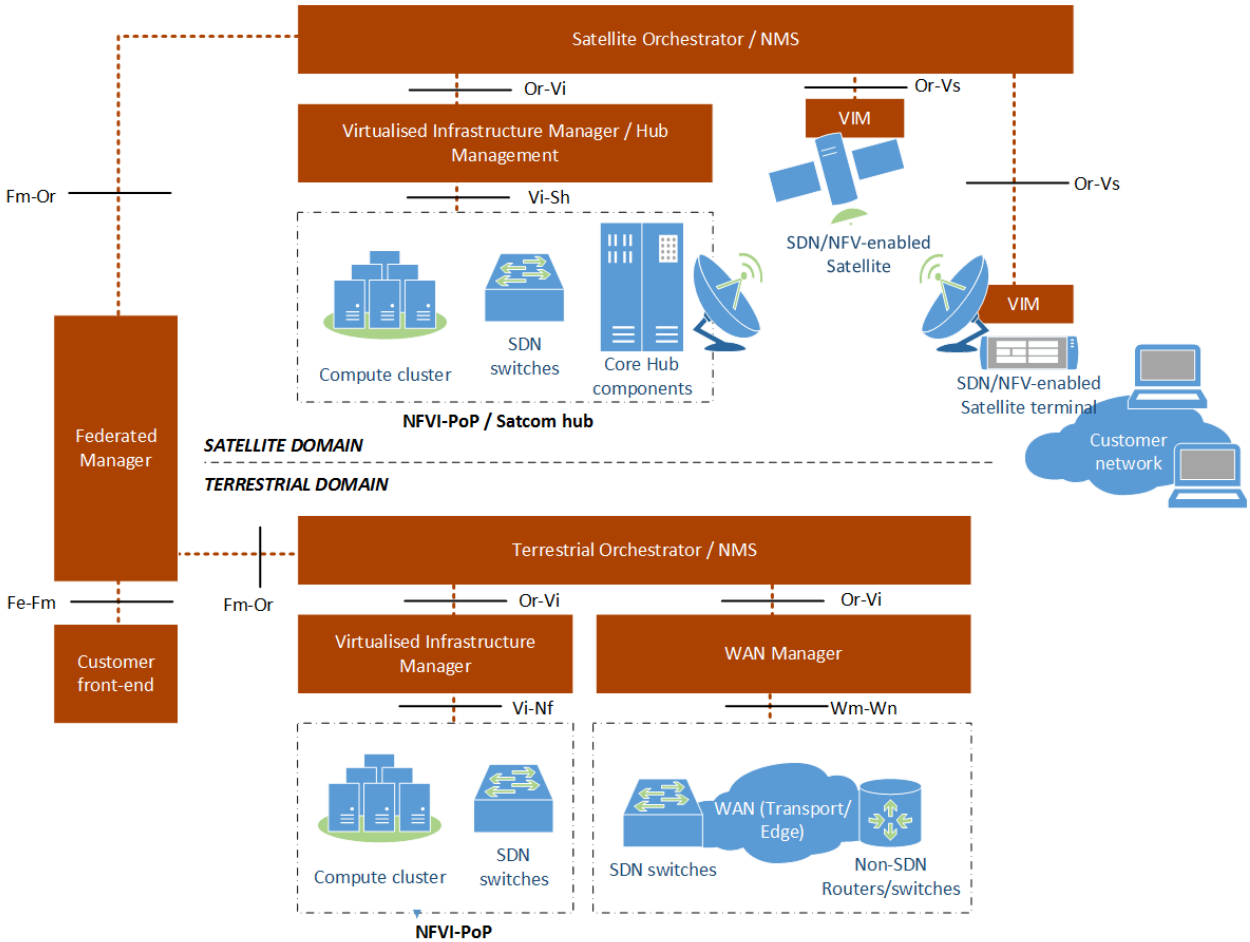

Figure 81. CloudSat architecture with reference points

Some of the reference points and in particular Or-Vi and Vi-Nf have been inherited from the ETSI NFV architectural model, while the rest ones have been inserted/adapted in order to match the specificities and the requirements of the

federated satellite/terrestrial cloud network. The tables to follow present a high-level description of the information exchanged via each reference point and suggest a candidate implementation technology for the respective interface.

Reference point	Fe-Fm
Entities connected	Customer front-end, Federated Manager
Information exchanged	CFE → FM: Service requests (deployment, reconfiguration, teardown) FM → CFE: Service offerings, Service status (request acknowledgements, monitoring metrics, SLA status, billing data)
Candidate implementation technology(ies)	HTTP REST (XML/JSON format)

Reference point	Fm-Or
Entities connected	Federated Manager, Orchestrator (Satellite or Terrestrial)
Information exchanged	FM → OR: Service requests (deployment, reconfiguration, teardown) OR → FM: Service status (request acknowledgements, monitoring metrics)
Candidate implementation technology(ies)	HTTP REST (XML/JSON format)

Reference point	Or-Vi
Entities connected	Orchestrator (Satellite or Terrestrial), Virtualised Infrastructure Manager (VIM)
Information exchanged	OR → VIM: Virtual Machine/Virtual Network Function lifecycle management (instantiate, rescale, start/stop/suspend, terminate), virtual network instantiation and management, service function chaining (SFC) management, network configuration, flow rules VIM → OR: Status and monitoring information of virtualised entities (VNFs, vNets), status and resource information of physical infrastructure, events and alarms
Candidate implementation technology(ies)	HTTP REST (XML/JSON format)

Reference point	Or-Vs
Entities connected	Orchestrator (Satellite or Terrestrial), Virtualised Infrastructure Manager (VIM) for satellite components (payload, terminal) *(NOTE 1: Or-Vs is meant to be a "scaled-down" version of the Or-Vi reference point)* *(NOTE 2: Or-Vs interactions are assumed to take place via the satellite control centre – not shown in* Figure 81*)*
Information exchanged	OR → VIMs: Virtual Machine/Virtual Network Function lifecycle management (instantiate, rescale, start/stop/suspend, terminate), network configuration, flow rules VIMs → OR: Status and monitoring information of VNFs, status and resource information of physical infrastructure
Candidate implementation technology(ies)	HTTP REST (XML/JSON format) or CoAP (Constrained Application Protocol) or raw TCP socket

Reference point	Vi-Nf
Entities connected	Vistrualised Infrastructure Manager (VIM), NFV Infrastructure (NFVI-PoP)
Information exchanged	VIM → NFVI: VM/VNF management requests, SDN flow rules, physical infrastructure management NFVI → VIM: VM/VNF status, flow information and events, status of physical infrastructure
Candidate implementation technology(ies)	Hypervisor API (e.g. libvirt API) Openflow

Reference point	Vi-Sh
Entities connected	Vistrualised Infrastructure Manager (VIM), NFV Infrastructure (NFVI-PoP) & Satcom hub *(NOTE: Vi-Sh functionality is essentially a superset of Vi-Nf)*
Information exchanged	VIM → NFVI/SH: VM/VNF management requests, SDN flow rules, satellite network management, terminal management NFVI/SH → VIM: VM/VNF status, flow information and events, satellite network information, terminal information
Candidate implementation technology(ies)	Hypervisor API (e.g. libvirt API) Openflow Hub-specific management (e.g. SNMP)

Reference point	Wm-Wn
Entities connected	WAN Manager, WAN
Information exchanged	WANM → WAN: Network service configuration (network management, virtual networks/tunnels establishment) WAN → WANM: Status/information from WAN Network elements
Candidate implementation technology(ies)	SNMP Netconf/YANG Openflow

5.4. Architecture Refinement for Specific Scenarios

This chapter discusses the specific "instantiation" of the overall architecture illustrated in Chap. 5.3., as shaped to suit the needs of the integration scenarios which were selected in Chapter 4. It also describes the sequence of interactions between the architectural entities which are required in order to realise each scenario.

5.4.1. Scenario #1: Hybrid media distribution network as-a-Service

5.4.1.1. Overview and scenario-specific requirements

Scenario #1 involves the provision of a virtual networking infrastructure as-a-Service, spanning across the satellite and terrestrial domains. The potential customer is a media content provider, who uses this virtualised infrastructure to deliver the content along diverse paths. The actors involved in this scenario, as well as its high-level description and the added-value of the use of virtualization technologies, are described in Chapter 4.

The main feature to be demonstrated in this scenario is the federation of satellite and terrestrial domains as well as the virtualization and programmability aspects enabled via the use of software-defined networking aspects. Therefore, of the high-level requirements listed in Sec. 5.2, the ones which are most applicable for this scenario are multi-domain management and interoperability, infrastructure sharing and multi-tenancy, service programmability as well as business-related requirements (SLA monitoring and billing).

5.4.1.2. Architecture refinement

Taking into account the specific requirements and technical capabilities required for this scenario, the generic architecture laid out in Chap. 5.3 can be adapted as shown in Figure 82.

The architecture needs to encompass both satellite and terrestrial domains in order to realise hybrid media delivery. The Federated Manager is in place to allow coordination between satellite and terrestrial domains. In this scenario, the management of the network service is realized via SDN, so all SDN-enabled components across the infrastructure are retained. On the other hand, NFV aspects are not included, so the presence of NFVI-PoPs powered by compute clusters are not required per se.

It must be noted here that, although the term "VIM" is mostly associated with NFV architectures (for purely network infrastructures a "Network Controller" would be more appropriate), we retain this term also in this scenario to maintain consistency with the overall architectural terminology.

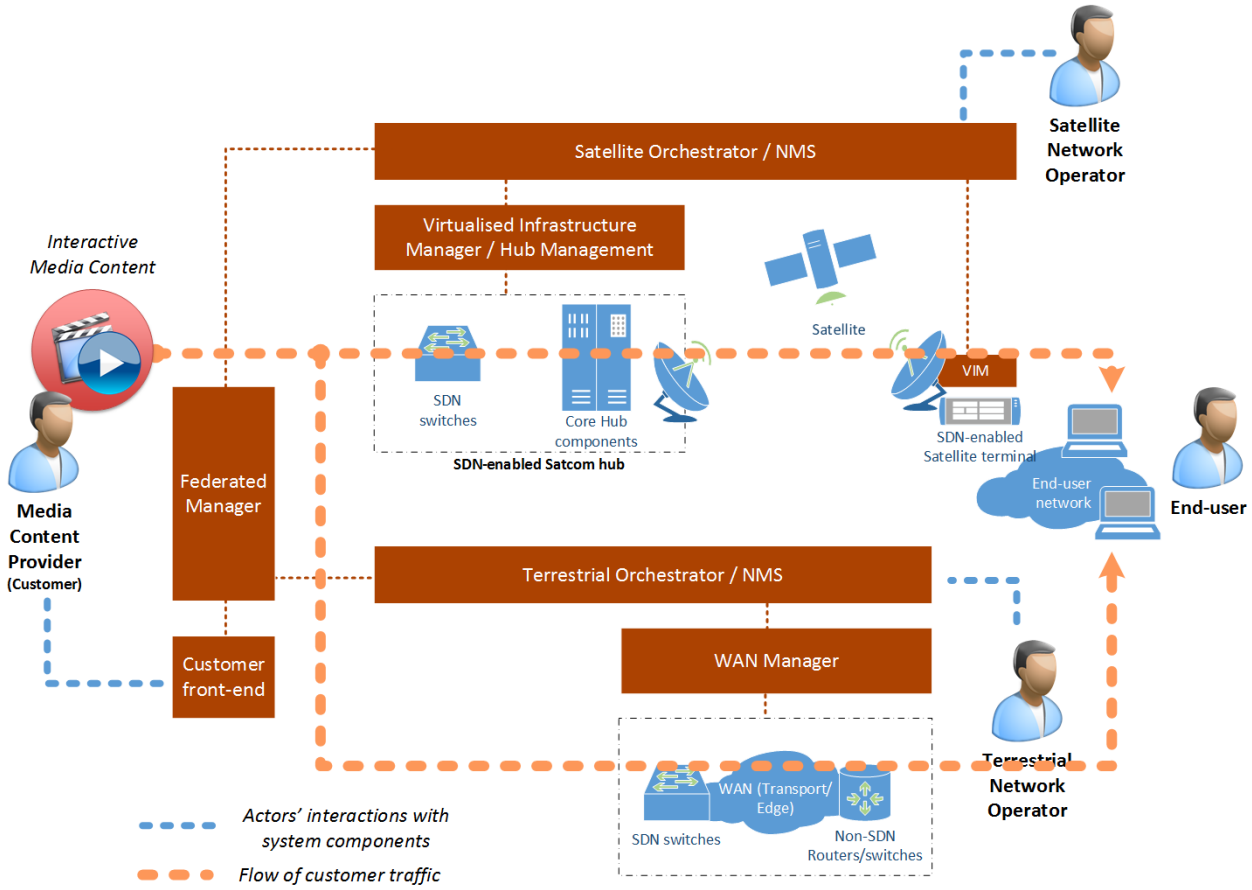

Figure 82. Architecture refinement for the "Hybrid metria distribution network as-a-Service" scenario

The actors of the scenario interact, for management-related actions, with the corresponding appropriate management entities; the Media Content Provider is the customer for this scenario and accesses the Customer front-end GUI for service request and configuration. Furthermore, the Customer's service control applications

(if any) for dynamic service management, interact with the programmatic interface (API) exposed by the front-end for this purpose. The network operators (satellite and terrestrial) typically manage their infrastructures via the corresponding Orchestrators.

5.4.1.3. Scenario realization

Figure 83 displays a high-level sequence for the realization of the scenario. Three stages of the service lifecycle are shown: service setup, monitoring and adaptation/reconfiguration.

For the sake of brevity, the interactions between the VIM and the infrastructure (SDN elements) are not included in the sequence diagram. Also, it can be understood that the successful execution of a service deployment or reconfiguration request is verified by a respective acknlowledgement message, originating from the infrastructure and propagating back to the customer; this ACK transmission is also not shown. These assumptions also apply to the sequence diagrams of the other two scenarios.

Service *setup* involves the following steps:

1. The Customer accesses the GUI/front-end and issues a network service request with specific requirements and SLA constraints (topology, endpoints, capacity, QoS etc.)
2. The front-end dispatches the request to the Federated Manager
3. The FM processes the request and maps it to the available resources
4. The FM decomposes the service request to the satellite and terrestrial part and dispatches the corresponding request to the satellite and terrestrial Orchestrators
5. The satellite and terrestrial Orchestrators perform a second-level, more fine-grained service mapping (not shown in the picture) in order to decide about the actual resources to be reserved
6. The Orchestrators issue virtual network (vNet) setup requests to the corresponding VIMs. For the satellite domain, the VIMs controlling the Hub network as well as the satellite terminal (in case of SDN-enabled terminals) are notified. For the terrestrial domain, the WAN Manager is involved.
7. The VIMs execute the request by installing the appropriate flow rules at the SDN-enabled network elements and configuring the traffic queues for QoS support.

Service *monitoring* is realized as follows:

1. SDN elements periodically report flow statistics to the VIMs
2. The VIMs aggregate flow statistics to vNet monitoring information and communicate them to the terrestrial and satellite Orchestrators.
3. The Orchestrators in turn aggregate this information to produce service-level information and communicate it to the FM.

4. The FM composes the overall status of the network service, checks the conformance to the SLA and performs billing operations.
5. Service, SLA and billing information are communicated to the Customer (via the GUI) and the customer's control application (if any, via the API).

Finally, service *reconfiguration* can include any modification in the service parameters, e.g. topology change, QoS parameters adjustment as well as per-flow rules (flow switching from terrestrial to satellite etc.). This reconfiguration can take place either manually (e.g. requested by the Customer via the GUI) or automatically (by the Customer's control application over the API). In any case, the service reconfiguration request is propagated from the FM to the Orchestrators, where it is translated to specific changes in the installed flow rules. These changes are communicated to the VIMs and enforced in the SDN network elements.

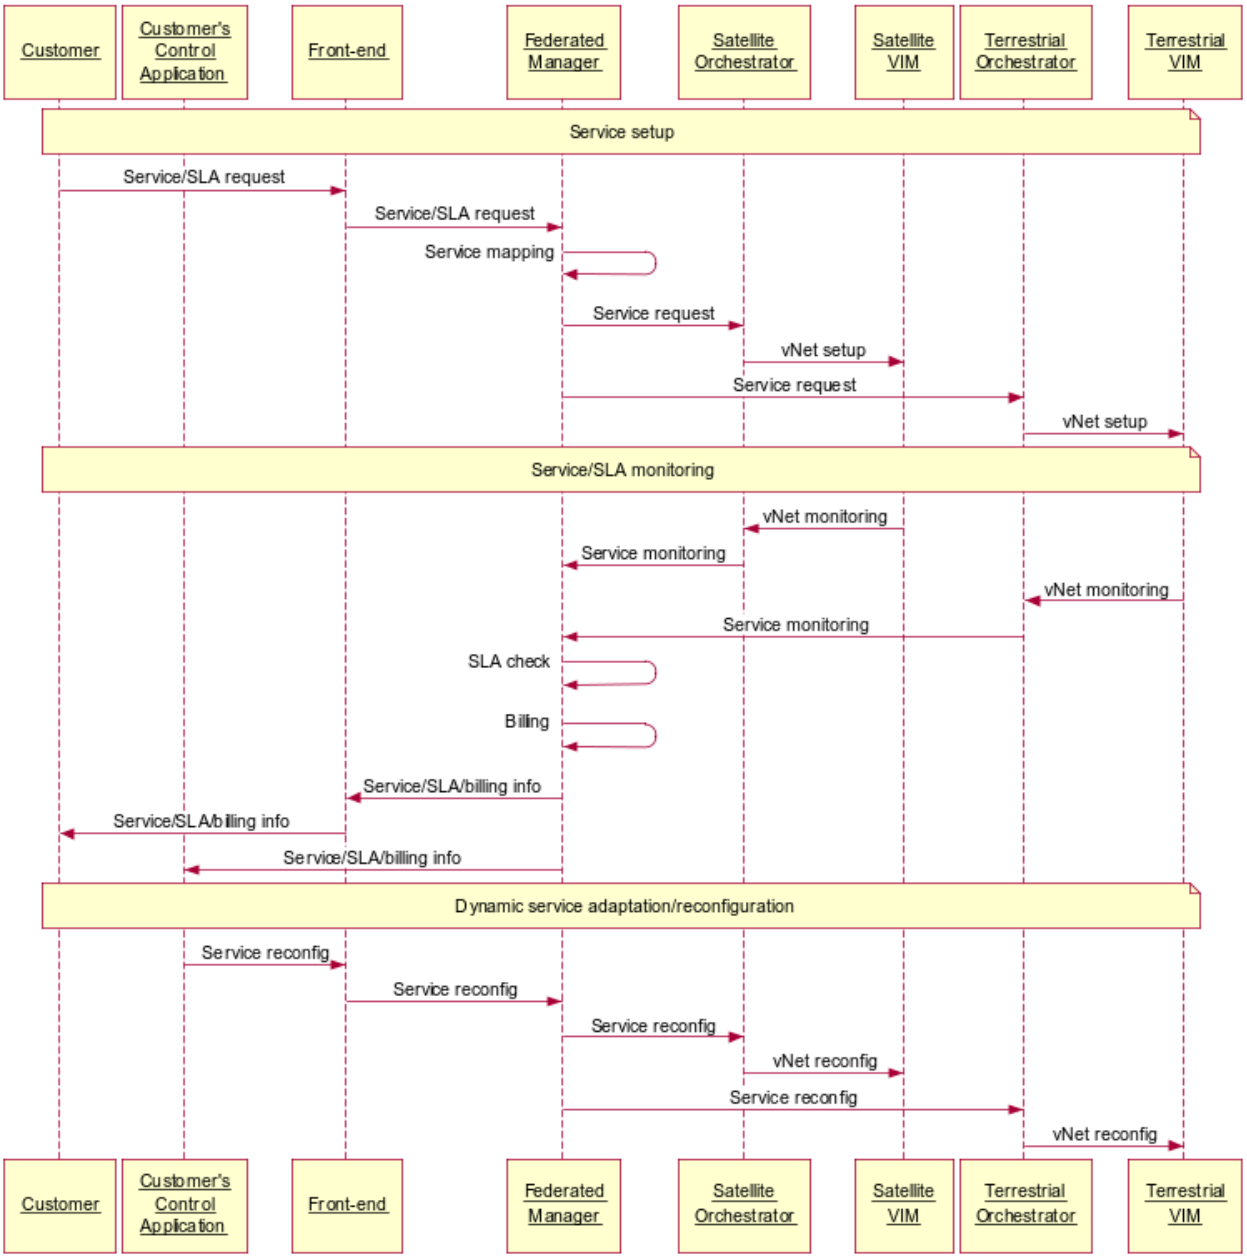

Figure 83. High-level sequence for Scenario #1 realisation

Service *teardown/termination* is not shown, but it follows a similar logic with service setup; a service termination message is originated by the customer and propagated to the infrastructure, where the traffic flow rules are removed and the corresponding resources are freed.

5.4.2. Scenario #2: Dynamic backhauling with edge processing

5.4.2.1. Overview and scenario-specific requirements

Scenario #2 involves the on-demand extension of a terrestrial customer network via satellite links to reach remote locations. This backhauling service is set up in an automated manner and its resources can be dynamically scaled. In addition, virtualised components as edge VNFs can be deployed in the professional satellite terminal to enable satellite edge traffic processing. This scenario has been extensively described, including the involved actors, the aspects/challenges and the market potential, in Chapter 4.

The main features to be demonstrated in this scenario is the dynamic allocation and up/down scaling of satellite network capacity, with specific QoS constraints, as well the on-demand deployment of VNFs in the customer's remote network (i.e. in the satellite terminal). Thus, the main requirements associated with this scenario refer to NFV support as well as service elasticity and reconfiguration.

5.4.2.2. Architecture refinement

Given the specific features and capabilities required for this scenario, the generic architecture laid out in Chap. 4 is refined as shown in Figure 84.

As with Scenario #1, federated management achieves the integrated control of the federated satellite/terrestrial service, although in this scenario specific emphasis is put in the management of the satellite network. SDN capabilities at both the satellite and terrestrial segment enable the uniform management of the multi-domain connectivity service. NFV capabilities are assumed only at the (professional) network terminal, which feature a lightweight VIM for the remote management of locally deployed VNFs.

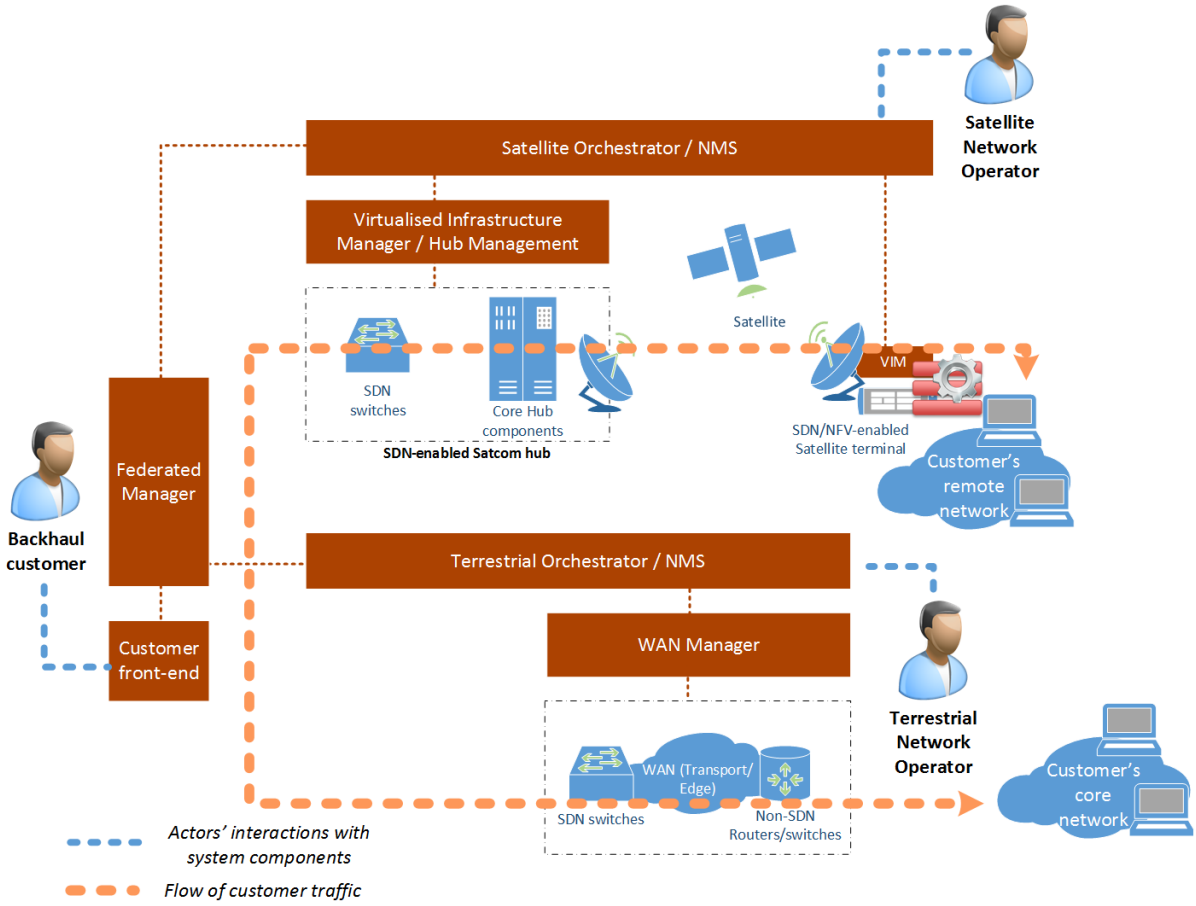

Figure 84. Architecture refinement for the "Dynamic backhauling with edge processing" scenario

The Customer (who requests the backhaul service) controls the service lifecycle via the front-end by means of human interactions; no automated programmatic control is foreseen, as in Scenario #1. The network operators (satellite and terrestrial) manage the respective network domains via the orchestrators/NMSs.

5.4.2.3. Scenario realization

It can be seen that the implementation of the scenario is quite similar with that of Scenario #1. The main difference is the inclusion of the edge VNF(s) as part of the network service. This inclusion requires the addition of the corresponding interactions with the satellite terminal VIM, via which the edge VNFs are deployed and monitored. SLA and billing aspects operations are not displayed in the diagram for the sake of brevity, but can be included in the same manner as Scenario #1.

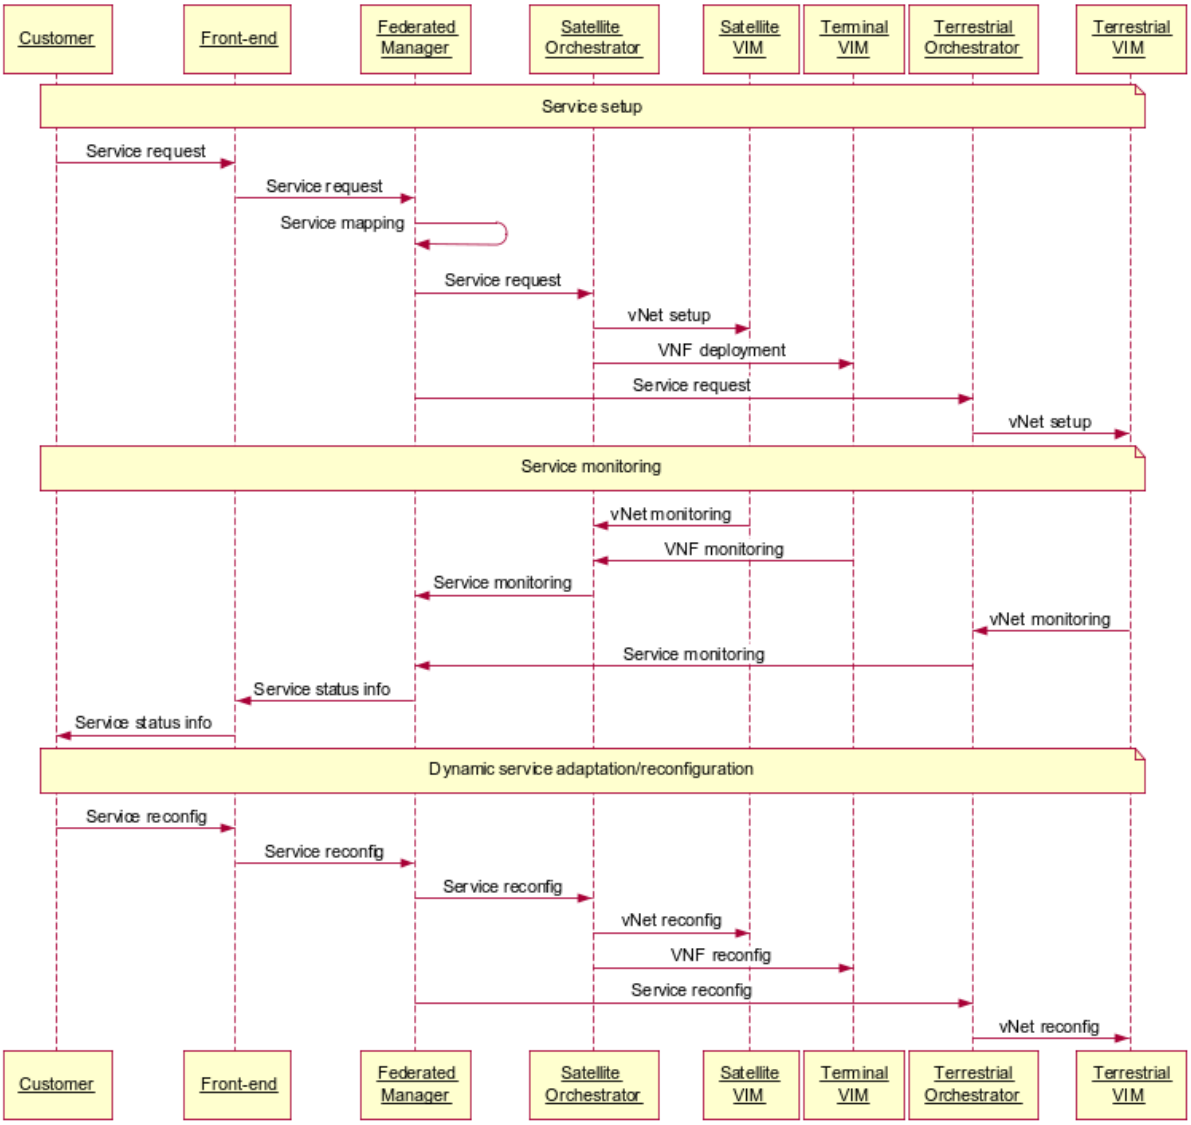

Figure 85. High-level sequence for Scenario #2 realisation

The service *setup* phase includes the following steps:

1. The Customer accesses the GUI/front-end and issues the network service request including the network endpoints, the bandwidth/QoS and the edge VNF(s) to be deployed. These VNFs can be selected from a catalogue of generally available VNFs, or be customer-specific.
2. The front-end dispatches the request to the Federated Manager
3. The FM processes the request and maps it to the available resources
4. The Federated Manager processes the request, maps the service to infrastructure resources and forwards the request to the satellite and terrestrial Orchestrators, which perform a second internal service mapping to the exact infrastructure resources.

5. The satellite and terrestrial Orchestrators install the network connectivity service (vNet) via the satellite VIM and the terrestrial WAN manager, eventually applying the relevant policies to SDN-enabled network elements.
6. The satellite Orchestrator interacts with the VIM at the terminal to deploy (instantiate) the selected VNF(s).

As a last step, depending on the nature of the VNF, additional configuration of the VNF may be needed in order to operate properly according to the customer's context and requirements. In principle, initial configuration of a VNF is performed automatically by the VNF Manager (an Orchestrator module). This initial configuration typically includes basic parameters, such as e.g. configuration of virtual network interfaces, execution of startup configuration scripts etc. However, the VNF will most likely need additional configuration, which needs to be carried out by the Customer him/herself, as if the VNF was a physical network element. For example, if the VNF is a local VoIP softswitch, although after automatic deployment it will be fully functional, the Customer will need to manually configure VoIP users, stream formats, SIP parameters etc. This VNF-specific configuration is assumed to take place via direct interaction with the VNF administrative interface, which can be either graphical or console-based.

Service *monitoring* is carried out pretty much like as described in Scenario #1 – network service metrics are reported from the infrastructure to the VIMs and then filtered/forwarded to the Orchestrator and eventually to the Federated Manager in order to construct the overall status of the service. The main addition in this scenario is that service metrics do not include only network parameters, but also monitoring information from the VNF(s) deployed. This information includes generic VM metrics, such as e.g. CPU utilization, memory usage, network interface load etc. but also VNF-specific metrics, which represent the operation of the VNF application per se. For example, in the case of a VoIP softswitch VNF, such metrics may include number of calls established or dropped, users' statistics etc.

Service *reconfiguration* is also similar to Scenario #1, with the addition that the VNF may need to be reconfigured. Since the VNF is constrained within the satellite terminal and no migration process is foreseen (from e.g. one physical host to another), VNF reconfiguration mostly includes rescaling i.e. reallocation of resources assigned to it. This corresponds to a "resizing" of the VM hosting the VNF.

Finally, service *teardown/termination* (not shown) involves the propagation of the termination request across the involved infrastructure elements, removal of flow rules and termination of the VNF, with the consequent release of allocated resources.

5.4.3. Scenario #3: Customer functions virtualisation

5.4.3.1. Overview and scenario-specific requirements

This scenario involves the enhancement of a typical satellite connectivity service with several Virtual Network Functions (VNFs) offered to the satcom customer as-a-Service, in the form of virtualised network appliances (VNFaaS). According to their nature, these VNFs can be instantiated either at the satellite hub or at VNF-enabled satellite terminals. More than one VNFs can be interconnected to form a custom service chain. The actors involved in this scenario, as well as its high-level description and the added-value of the use of virtualization technologies, are described in Chapter 4.

This scenario is mostly focused on NFV aspects; the main capabilities involved refer to the handling of heterogeneous customer VNFs at various parts of the satellite network and the management of the entire NFV lifecycle.

5.4.3.2. Architecture refinement

Taking into account the specific requirements and technical capabilities required for this scenario, the generic architecture laid out in Chap. 5.3 can be adapted as shown in Figure 86.

The focus here is in the satellite domain, since the scenario is applied to a satcom-only service. Since this is a single-domain scenario, the single-domain variant of the generic architecture is used (see Figure 78), which omits the Federated Manager, while the service lifecycle is managed solely by the domain Orchestrator and the customer front-end is directly attached to the latter.

Across the infrastructure, all NFV enabling entities are retained; NFV support is assumed not only at the satellite terminal but also in the hub side, thus assigning to the hub all the functionalities of an NFVI Point of Presence (NFVI-PoP). SDN is also assumed to be supported throughout the infrastructure, not only for establishing the connectivity service, but also for interconnecting the various VNFs of the customer and redirecting the traffic through them, thus forming a service function chain.

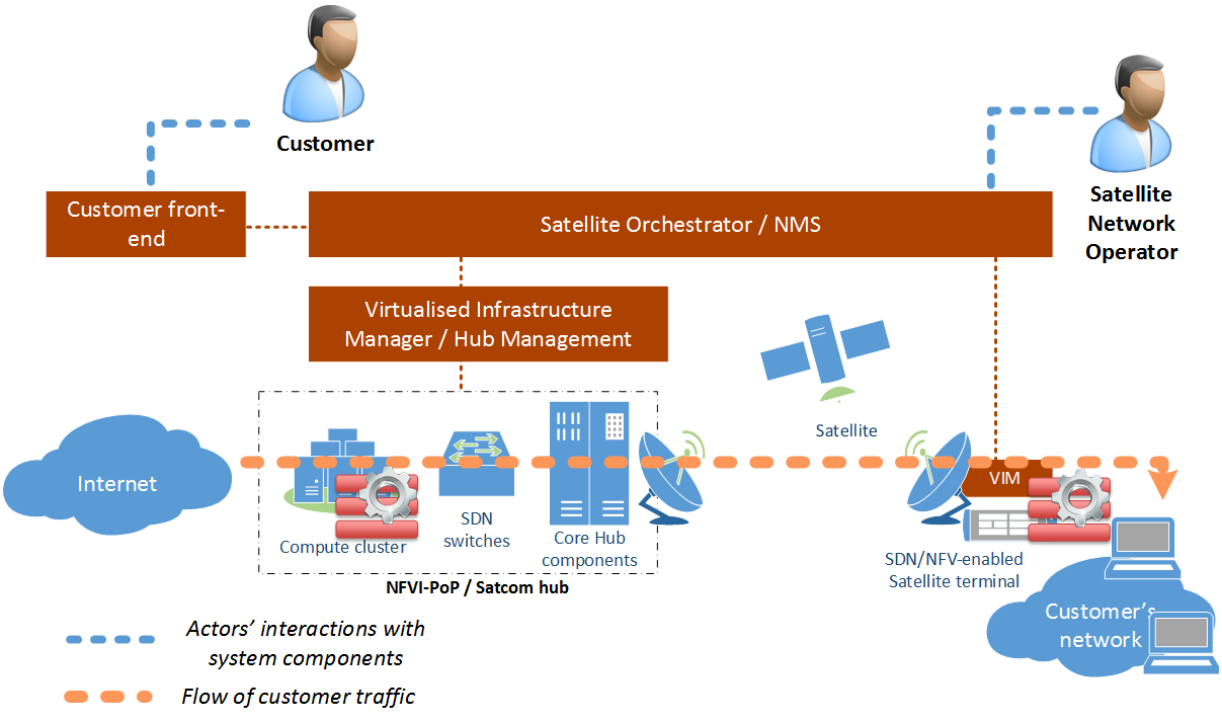

Figure 86. Architecture refinement for the "Customer Functions Virtualisation" scenario

In this scenario, the basic network service is assumed to correspond to a typical connectivity service, where the customer traffic is subject to specific QoS constraints and traverses the instantiated VNFs. In this context, the Customer has to interact with the front-end in order to request the service, choosing among various connectivity bundles and selecting available VNFs from a catalogue. After deployment, the service is also managed via the FE; in this scenario, this procedure should be quite simplified and adapted since the service is also targeted to non-professional users with limited technical knowledge (as opposed to the previous two scenarios).

The deployed services as well as the infrastructure are supervised and managed by the satellite network operator via the Orchestrator/NMS.

5.4.3.3. Scenario realization

Figure 87 illustrates the proposed high-level sequence for the realization of the various phases of Scenario #3, namely service setup, monitoring and reconfiguration.

Service *setup* involves the following steps:

1. The Customer accesses the GUI/front-end and browses the catalogue of the available services. This catalogue contains connectivity bundles with a set of VNFs. For example, a service bundle would include:
   - 1Mbps guaranteed download rate / 8 Mbps max DL rate
   - 256 kbps guaranteed upload rate / 1 Mbps max UL rate
   - a virtual security appliance
   - and a virtual content filtering appliance

The computing resources applied to the VNFs should be adequate to handle the prescribed maximum amount of traffic.

2. The Customer selects the desired bundle.
3. The request is forwarded to the Orchestrator, which maps the request to the available infrastructure resources. VNFs are assigned either to the hub or the terminal NFVI resources.
4. The Orchestrator issues requests to the satellite VIM to i) instantiate the customer virtual network (in this case it can be just a QoS- enabled connectivity service and not a “virtual network” in the strict sense), ii) instantiate the VNFs at the hub NFVI resources and iii) interconnect (“chain”) the VNFs together.
5. The Orchestrator issues requests to the remote (terminal) VIM to instantiate remote edge VNFs, if any.

As a final step and as described in Scenario #2, customer-side VNF configuration may take place for setting VNF-specific parameters (e.g. setting the content filter VNF to filter out specific content etc.)

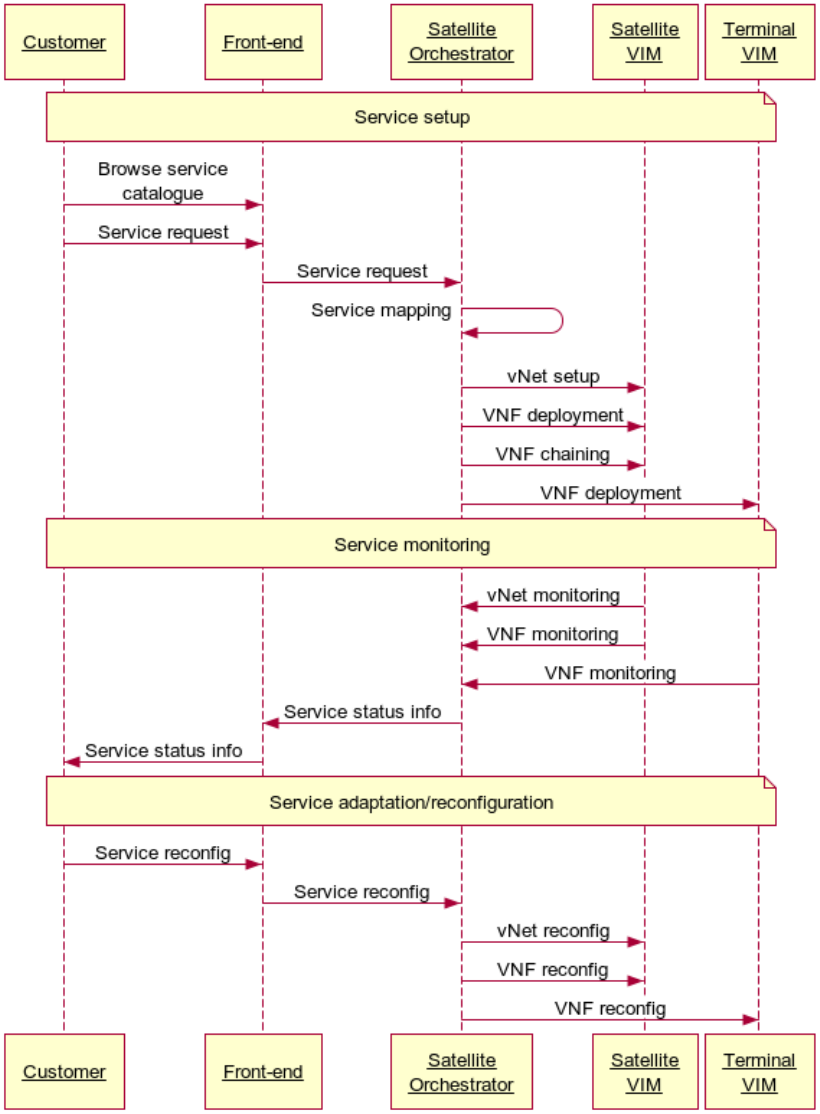

Figure 87. High-level sequence for Scenario #3 realisation

Service *monitoring* and *reconfiguration* are performed in a similar manner with the previous scenarios, yet without the involvement of the Federated Manager, which is absent in this scenario. Again, VNF monitoring involves collection and aggregation of both VM generic and VNF-specific metrics.

SLA and billing operations are not displayed in the diagram, but can be included in the same manner as Scenario #1.

6. Validation Requirements, Framework and Results

This chapter describes the CloudSat experimentation platform, as well as the validation and assessment of the use cases and reference architectures.

Following the analysis which took place in the previous chapters, there is enough input so as to conclude on specific recommendations for the CloudSat platform components in terms of cloud networking technologies and tools. It is recalled that the following technologies were identified as most promising for integration with satcom:

- IT virtualisation based on Openstack cloud management as NFV enabler
- Software Defined Networking, based on Openflow and OpenDaylight controller
- Management and Orchestration mechanisms, including Openstack and OpenDaylight for intra-domain management, as well as inter-domain federated management platforms, at higher layer, to enable multi-domain SDN/NFV network services

Thus, the experimentation platform needs to encompass all these technologies, also of course involving a satellite emulator component in order to emulate the satellite network.

6.1. CloudSat Experimentation Platform Overview

6.1.1. Requirements of the experimentation platform

The requirements which drive the design and implementation of the experimentation platform stem from the various integration scenarios described in TN2.2. Requirements address several functional capabilities, as follows:

- Virtualisation
- Elasticity
- Security
- Resiliency
- Programmability
- Management and Orchestration
- Infrastructure federation
- Terrestrial and satellite segment emulation

Every requirement has an implicit severity level, which is indicated by the verb used to express it, in accordance to IETF RFC 2119:

- SHALL corresponds to an absolute requirement, something that must be supported by the implementation.

- SHOULD corresponds to a recommended, but optional, requirement – paraphrasing RFC 2119, this means that "there may exist valid reasons in particular circumstances to ignore a particular item, but the full implications must be understood and carefully weighed before choosing a different course".

The following list summarises the main conclusions in terms of requirements:

- **Satellite segment emulation**. The CloudSat experimental platform SHALL be capable of providing emulation capabilities of the whole satellite service chain, i.e. including a satellite gateway, a satellite and a satellite terminal.
- **SDN compatibility.** The CloudSat experimental platform SHALL be compatible with the OpenFlow protocol in order to support the testing and validation of Openflow-based SDN control.
- **Virtualization of Infrastructure.** The CloudSat experimental platform SHALL include an appropriate cloud computing platform, which will offer resource pooling and infrastructure virtualization for supporting NFV services.
- **NFV service deployment.** The CloudSat platform SHALL include mechanisms for automated deployment of NFV services.
- **Resource isolation.** Resource isolation SHALL be provided for any network services provided on top of shared infrastructures.
- **Resource monitoring.** The CloudSat experimental platform SHALL periodically provide information about the resources consumed by the deployed network services. This information can be used to detect anomalies, resources failures, or severe performance degradation due to resource shortage
- **Service scaling.** Existing active network services SHOULD be able to scale up or down, upon a customer's demands for each service.
- **Federated management.** The CloudSat experimental platform SHALL include mechanisms to allow federated management of terrestrial and satellite resources.

6.1.2. Overall Architecture

The CloudSat prototype platform aims at testing and validating specific use case integrating scenarios of the cloud networking paradigm, considering a hybrid (i.e. terrestrial and satellite) service provision and demonstrating scientific and technological advancements. Figure 2 shows a very high-level view of the CloudSat prototype platform architecture.

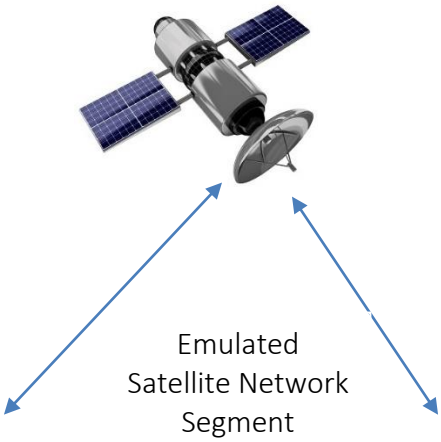

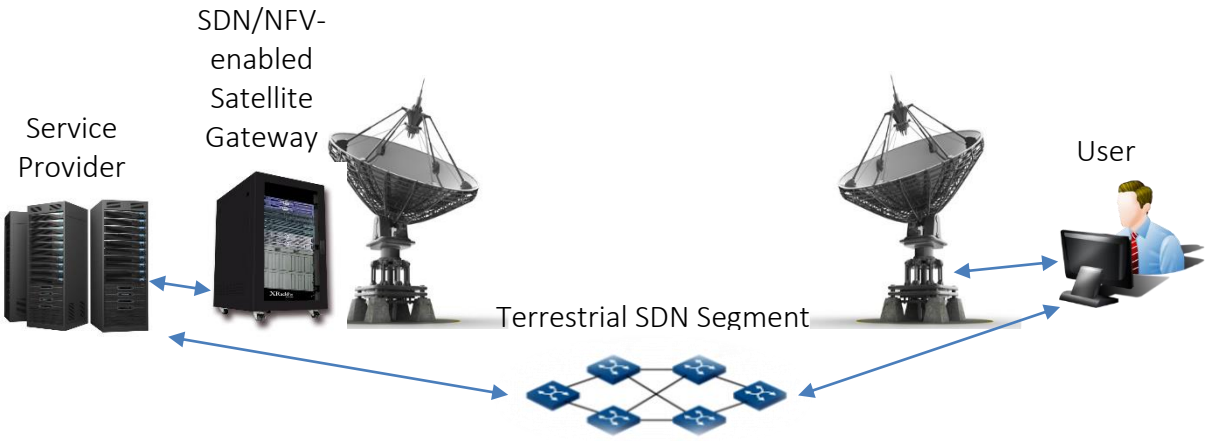

Figure 88. High-level CloudSat prototype platform architecture.

For the sake of clarity, a single Satellite Network Operator is considered, which provides a GEO satcom service via a transparent satellite. We will follow the SDN approach to achieve programmability in the provided network infrastructure and adopt the NFV concept to investigate the capability to insert virtual network services on-demand.

Having in mind the requirements expressed in the previous section, the CloudSat experimental testbed includes:

- Satcom emulator platform, based on OpenSAND (based on DVB-RCS and DVB-S2)
- SDN programmable network segment (Openflow-enabled), comprising both physical and virtual appliances
- Openstack infrastructure for VNF hosting

For the management and the orchestration of the infrastructure resources, the following two layers are provided:

- A management layer for managing the IT (Openstack) infrastructure as well as the SDN network
- An orchestration layer for orchestrating the IT and network resources, and for interconnecting virtual functions to achieve service chaining.

The physical network architecture is provided in the following figure:

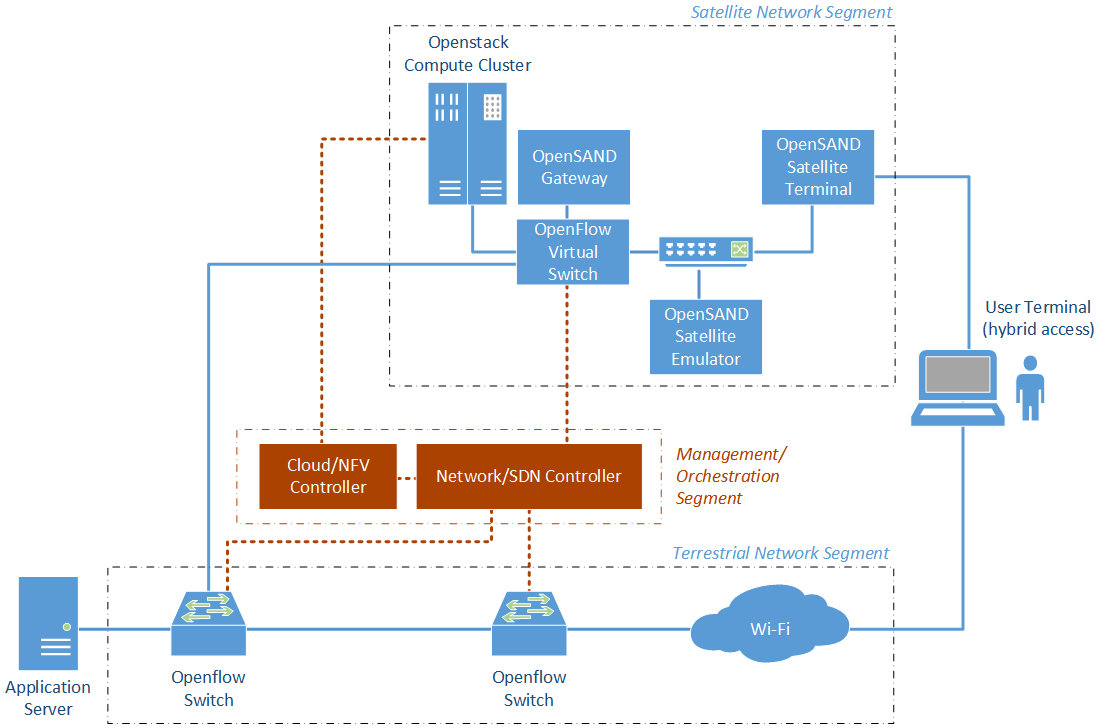

Figure 89. Physical Network architecture of the CloudSat Prototype Platform.

Figure 3 presents an architectural view of the physical testbed architecture for the CloudSat platform that satisfies the aforementioned requirements. The main elements composing the architecture are the cloud infrastructure, where the virtual network appliances are running, the SDN-compatible terrestrial network and the satellite network emulator.

The virtual resources provided in CloudSat project include virtual machines (VMs), which comprise virtualised network functions (VNFs, such as e.g. proxies, firewalls etc, see Sec. 6.6.3). These are combined with other virtualised network resources and/or physical resources in order to create the virtual networks. Resource virtualisation aims at better utilization of the underlying infrastructure in terms of (i) reusing a single physical or logical resource for multiple other network instances, and (ii) aggregating multiple resources in order to optimise resource usage.

In order to manage both virtual and physical resources effectively, an effective orchestration and network management system is needed.

The deployment of the software-based CloudSat testbed components has been completed in the premises of NCSR Demokritos. All components have been installed in a virtualisation-capable IT infrastructure based on high-capacity rack servers (Dell™ PowerEdge™ R510).

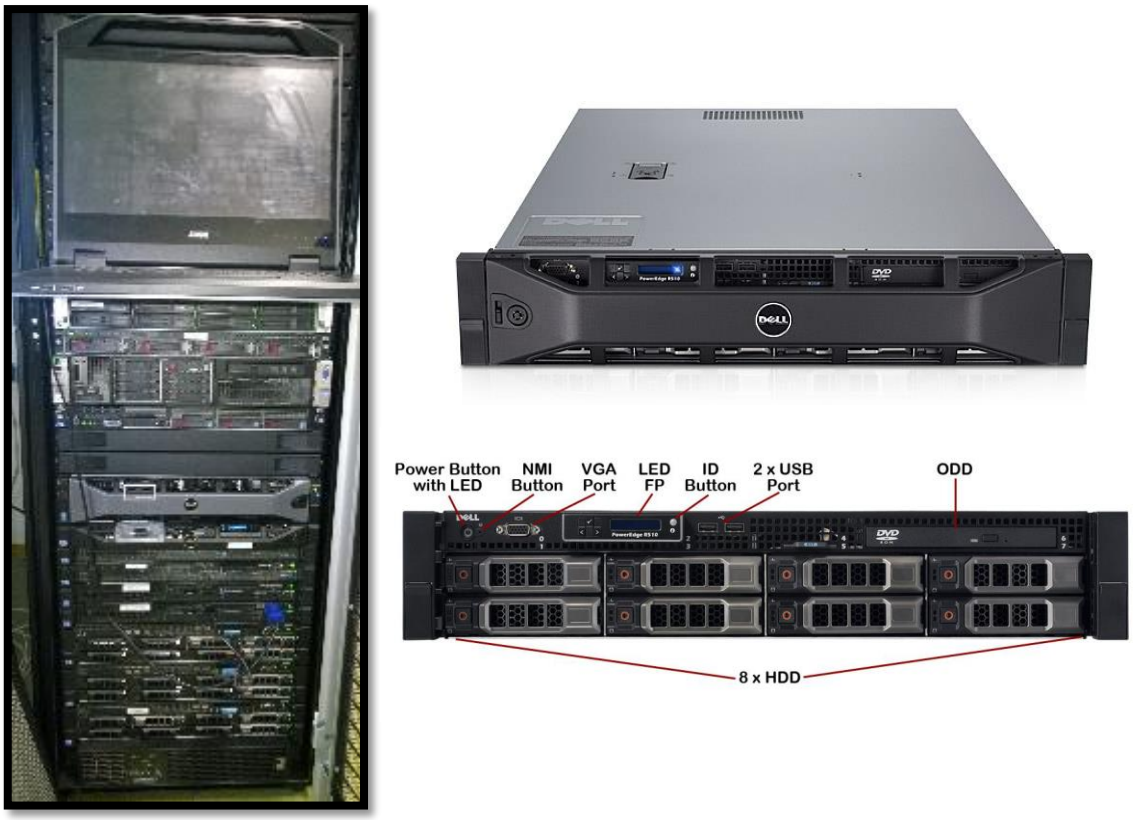

Figure 90. The rack at NCSR Demokritos premises that hosts the Dell PowerEdge R510 servers running the software components of the CloudSat testbed.

These servers have the following specifications:

- Intel 5500 chipset + Intel Xeon processor 5500 and 5600 series
- PCIe Generation 2
- DDR3 Memory Technology (Memory RAS feature—mirroring and sparing)
- LVDIMM memory support with Intel Xeon processors 5600 series
- iDRAC6 (Dell server remote management controller)
- Dell Management Console (provides a consolidated view of the IT environment)
- Virtualization (supports various virtualization applications)
- SSD advantage (support of SSD drives)

Chapters 4, 5 and 6 present in more details the software components of the CloudSat testbed.

6.1.3. Metrics and evaluation framework

The main purpose of the experimentation campaign in CloudSat is, to demonstrate the feasibility of the selected cloud networking use cases by showcasing the functional capabilities of the implemented architecture. These capabilities include:

- SDN-based QoS differentiation and enforcement
- SDN-based traffic redirection from satellite to terrestrial and vice versa

- SDN-based per-flow monitoring
- VNF instantiation at the terminal
- VNF traffic steering at the terminal
- VNF instantiation at the GTW
- VNF traffic steering and service chaining at the GTW
- Functional capabilities of the various VNFs (caching, transcoding, filtering etc. depending on the VNF)

Apart from the functional testing, several performance metrics are collected in order to assess the efficiency of the proposed virtualisation mechanisms. These metrics are shown in the table below.

Table 16. Metrics to be collected during experimentation

Metric	Units	Comments
SDN reconfiguration delay	msec	Applies to QoS enforcement and traffic redirection
ICMP round-trip time	msec	Used to assess the delay introduced by the satellite emulator, the SDN switches and the VNFs
Per-queue and per-interface bitrate	Kbps	Used to verify QoS enforcement and traffic redirection
SSIM (Structural Similarity) – video quality metric	- (absolute value)	Measures video quality degradation due to congestion
SDN Manager (federator) response time	msec	The delay from the congestion incident to the actual application of the new stream priority or delivery channel
VNF VM resource utilization (CPU, HDD, memory)	percent (%)	Applies to VNF workloads
VNF instantiation delay	sec	

6.2. Satcom Emulator

6.2.1. OpenSAND

This section briefly describes the subsystem of the testbed which emulates a satcom network (Figure 91)

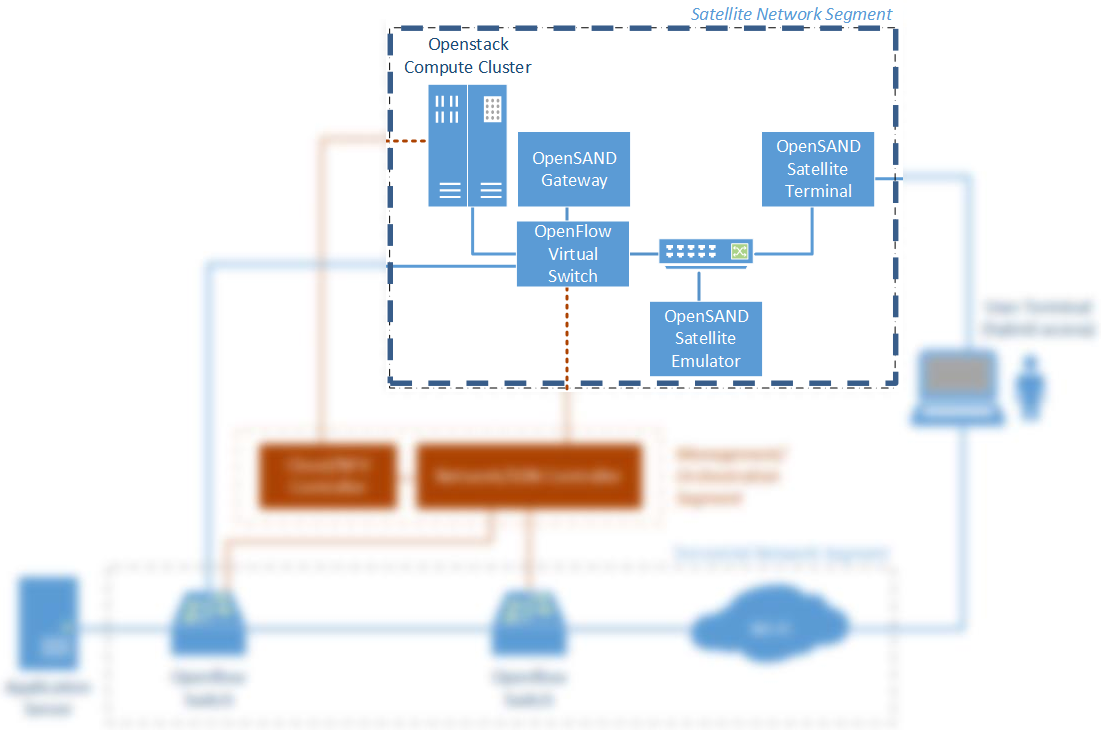

Figure 91. Satellite Network Segment.

For the emulation needs of the CloudSat platform, the OpenSAND emulator [OpenSAND], developed and supported by TAS-F, has been selected and deployed. OpenSAND provides an easy and flexible way to emulate satellite communication systems, based on a simple architecture, which is described on Figure 6, demonstrating the different components of OpenSAND software, namely:

- Satellite Terminal (ST),
- Satellite Emulator (SE),
- Gateway (GW).

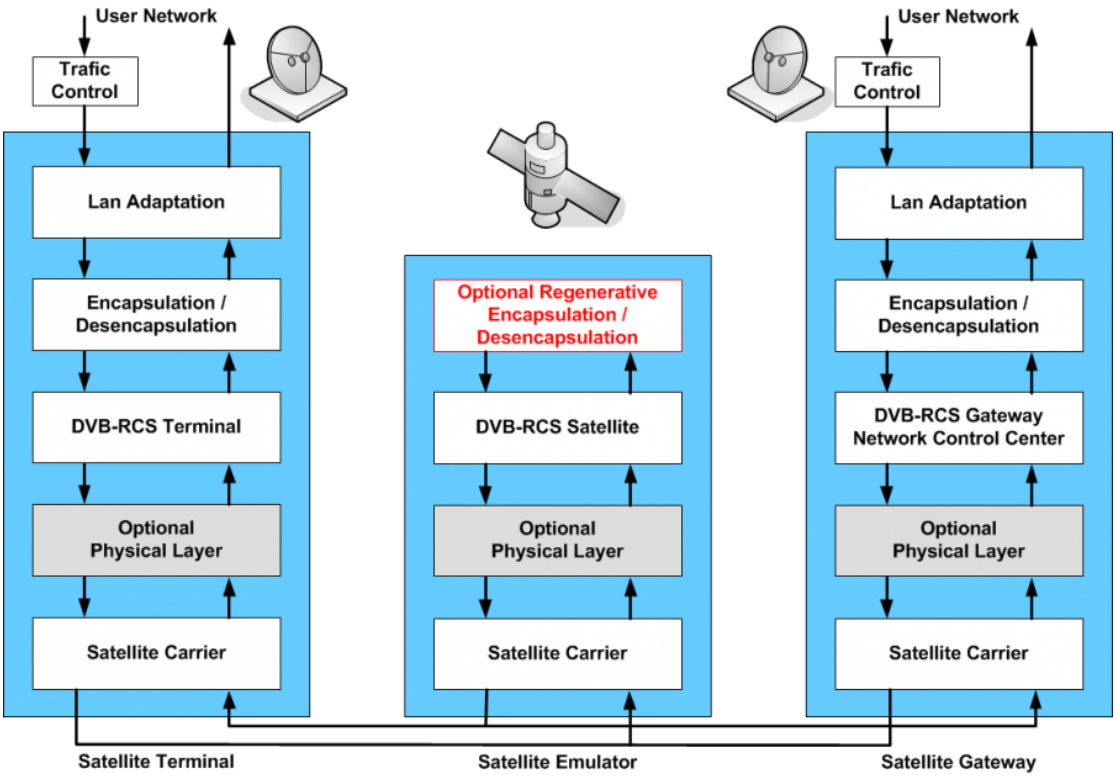

Figure 92. OpenSAND architecture

The SE is able to emulate either a transparent or regenerative satellite in combination with different encapsulation schemes depending on payload type, up/return link standard and installed plugins. Figure 7 provides two representative snapshots of the OpenSAND admin GUI, which allows to the user to set up the necessary configuration and monitor the data traffic.

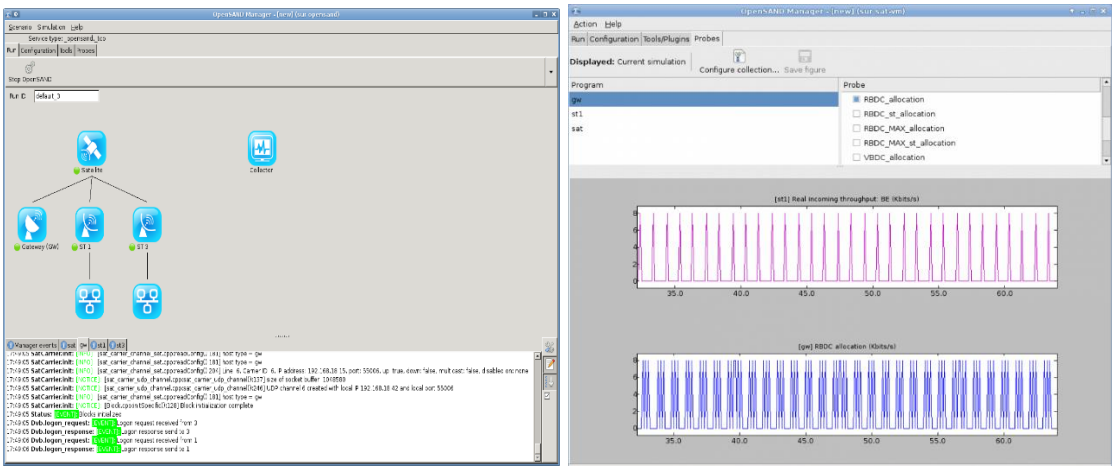

Figure 93. OpenSAND monitoring

6.3. IT virtualisation

This section refers to the IT virtualization component of the CloudSat architecture, which is highlighted in Figure 94 and required so as to support NFV services. The

Openstack Cloud platform is used [Openstack], whose configuration is briefly described in the following subsection.

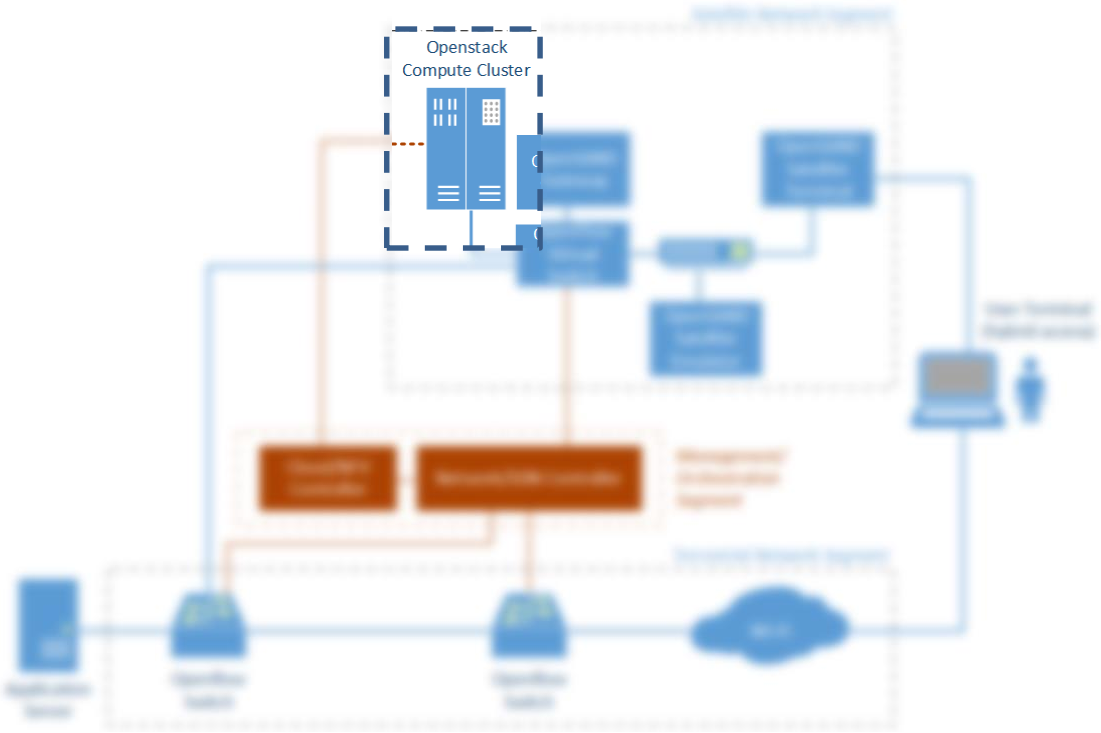

Figure 94. Openstack compute cluster for NFV services

6.3.1. OpenStack

OpenStack has a modular design that enables integration with legacy and third-party technologies. The Openstack architecture and components have been described in detail in Chapter 2. The following subsections briefly discuss the most important OpenStack components that are of high interest for the CloudSat platform, namely the Nova and Neutron components, which have been successfully deployed and tested, which are essential for the implementation of the integration scenarios selected.

6.3.1.1. OpenStack Nova

The primary component of the cloud operating environment, which is of high importance for the VNF deployment of CloudSat prototype platform is the Nova compute service. Nova orchestrates the creation and deletion of VM instances, which are used as carriers/enablers for the VNFs image deployment. In the CloudSat testbed, the following components of Nova are used:

- The nova-api accepts and responds to end-user compute API calls. It also initiates most of the orchestration activities (such as running an instance) as well as enforcing some policies.

- The nova-compute process is primarily a worker daemon that creates and terminates VM instances via hypervisor APIs. In CloudSat, the KVM hypervisor is used.
- The nova-scheduler process keeps a queue of VM instance requests and for each request it determines where the VM instance should run (specifically, which compute node it should run on).

6.3.1.2. OpenStack Neutron

OpenStack Neutron, is the OpenStack module focused on delivering Networking as a Service (NaaS), which for the needs of the CloudSat platform will be the networking interface between the Cloud platform and the SDN controller (i.e. OpenDayLight). Neutron makes easier to deliver networking as a service in the cloud and provides HTTP REST (Representational State Tranfer) APIs to manage network connections for the resources managed by other OpenStack services.

Neutron provides native multi-tenancy support (isolation, abstraction and full control over virtual networks), letting tenants create multiple private networks and control the IP addressing on them, and exposes vendor-specific network virtualisation and SDN technologies.

The core Neutron API to be used in CloudSat, includes support for Layer 2 networking and IP Address Management (IPAM), as well as an extension for a Layer 3 router construct that enables routing between Layer 2 networks and gateways to external networks. It is based on a simple model of virtual networks, subnet, and port abstractions to describe networking resources. A network is an isolated layer-2 segment, analogous to a VLAN in the physical networking world. More specifically, it is a broadcast domain reserved for the tenant that created it or explicitly configured as shared.

6.4. SDN Infrastructure

This section refers to the deployment of the SDN-enabled network segment of the CloudSat platform. For the needs of the validation purposes of the project, two SDN compatible switches have been tested and deployed: The Open vSwitch [OVS] (software virtual switch) and the Pica8 [Pica8] (physical switch).

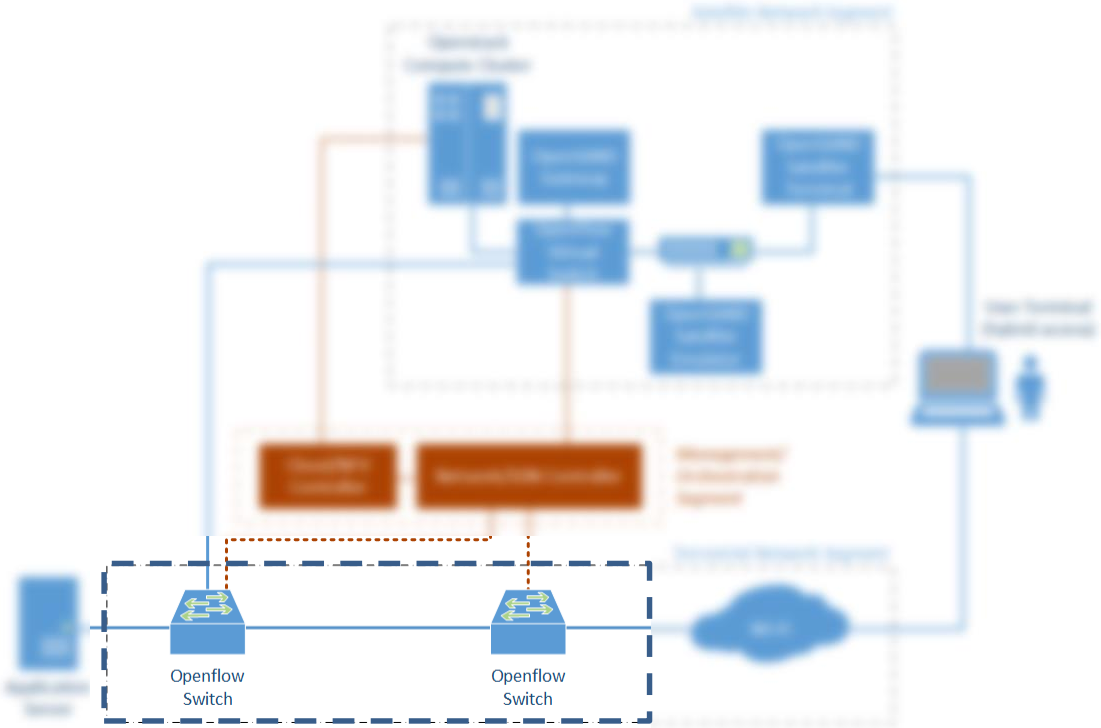

Figure 95. Terrestrial Network Segment.

6.4.1. Open vSwitch

Open vSwitch is a multilayer virtual switch licensed under the open source Apache 2.0 license. It is designed to enable massive network automation through programmatic extension, while still supporting standard management interfaces and protocols, among others OpenFlow for SDN programmability.

Open vSwitch is well suited to function as a virtual switch in VM environments. In addition to exposing standard control and visibility interfaces to the virtual networking layer, it has been designed to support distribution across multiple physical servers.

The current release of Open vSwitch supports the following features:

- Standard 802.1Q VLAN model with trunk and access ports
- NIC (Network Interface Card) bonding (i.e. port aggregation to increase bandwidth) with or without LACP (Link Aggregation Control Protocol) on upstream switch
- NetFlow, sFlow(R), and mirroring for increased visibility
- QoS (Quality of Service) configuration, plus policing
- Geneve (Generic Network Virtualisation Encapsulation), GRE, GRE over IPSEC, VXLAN, and LISP tunneling
- 802.1ag connectivity fault management
- OpenFlow 1.0 plus numerous extensions
- Transactional configuration database with C and Python bindings
- High-performance forwarding using a Linux kernel module

In CloudSat, Open vSwitch is used both in the terrestrial network emulation, as well as integrated with OpenSAND, in order to add SDN capabilities to the satellite segment.

6.4.2. Pica8 open switch

Figure 96. Pica8 open switch, installed at the CloudSat testbed

Pica8 is an SDN switching platform, which supports OpenFlow 1.4 through integration of Open vSwitch (OVS) v2.0; OVS runs as a process within Pica8's operating system, providing an OpenFlow interface for external programmability and setting the stage for SDN features, such as:

- Traffic engineering: OpenFlow 1.4 statistics analyze utilization to help determine the best path for application flows
- GRE tunnelling: Connect logical domains without disrupting the overall network fabric, and isolate sensitive traffic
- Network Taps: OpenFlow 1.4 can both dynamically program a network tap and adjust its characteristics, thereby greatly reducing CAPEX

6.5. Management and Orchestration

The management segment consists of the SDN network segment control, the cloud controller and the satellite network segment management.

A higher-layer federated management component is foreseen to couple and jointly manage all the controller domains in order to establish composite end-to-end services on the testbed infrastructure.

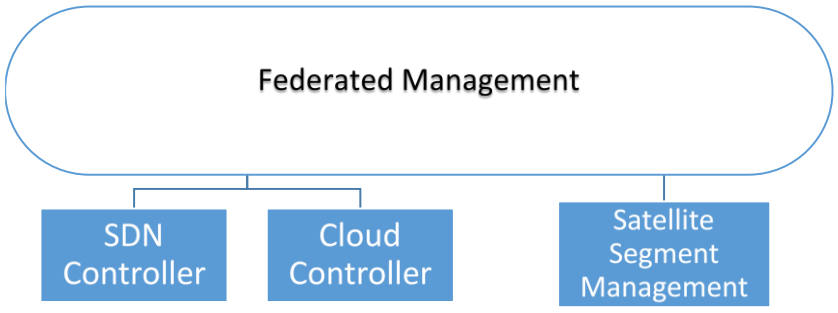

Figure 97. Hierarchical Structure of Federated Management and Orchestration.

The Network Controller of the SDN network, also known as SDN controller, manages the terrestrial and satellite SDN network elements. In the CloudSat testbed, the OpenDaylight controller is used (Sec. 6.5.2). OpenDaylight uses OpenFlow [OpenFlow], to communicate with the SDN infrastructure (Open vSwitch instances and Pica8 switch), receiving flow notifications and applying flow rules. OpenDaylight acts as the core component of the CloudSat SDN infrastructure, withholding all the intelligence; it is relaying information to the SDN infrastructure 'below' and the federated management 'above' (via northbound APIs).

The Cloud Controller is required for building and managing the distributed cloud computing resources within a datacenter. It is responsible for numerous tasks including (i) controlling processing, storage, and networking resources; (ii) performance monitoring (response times, latency, uptime, etc.); and (iii) security and compliance auditing and management. In the CloudSat testbed, the control APIs exposed by Openstack as used, as well as the Horizon Dashboard (Sec. 6.5.1) for interacting with the administrator.

For the management of the satellite emulator, the management capabilities of OpenSAND are exploited (Sec.6.5.3)

Finally, the federated management layer undertakes the coordination of the underlying management/control components for the automated provision and maintenance of cloud network services, including end-to-end system configuration and holistic resource management [NFVMAN]. It aims to address the two main issues: a) the establishment of the connectivity service with specific bandwidth and QoS constraints and b) the deployment of VNFs, as well as their chaining (Service Function Chaining, SFC)

Since, as explained in Chapter 3, no mature, open and well-established solution yet exists for federated management, in CloudSat we are using control scripts which interact with the underlying controllers via their APIs, and which are specific to the corresponding topology and use case.

6.5.1. Horizon OpenStack Dashboard (Cloud Controller)

The OpenStack dashboard [Horizon] used in the CloudSat testbed is an extensible web app that allows the system administrator to control their compute, storage and networking resources assigned to VNFs (Figure 98)

Via the Dashboard, it also becomes possible to monitor the VNFs running, start/stop them and re-assign resources.

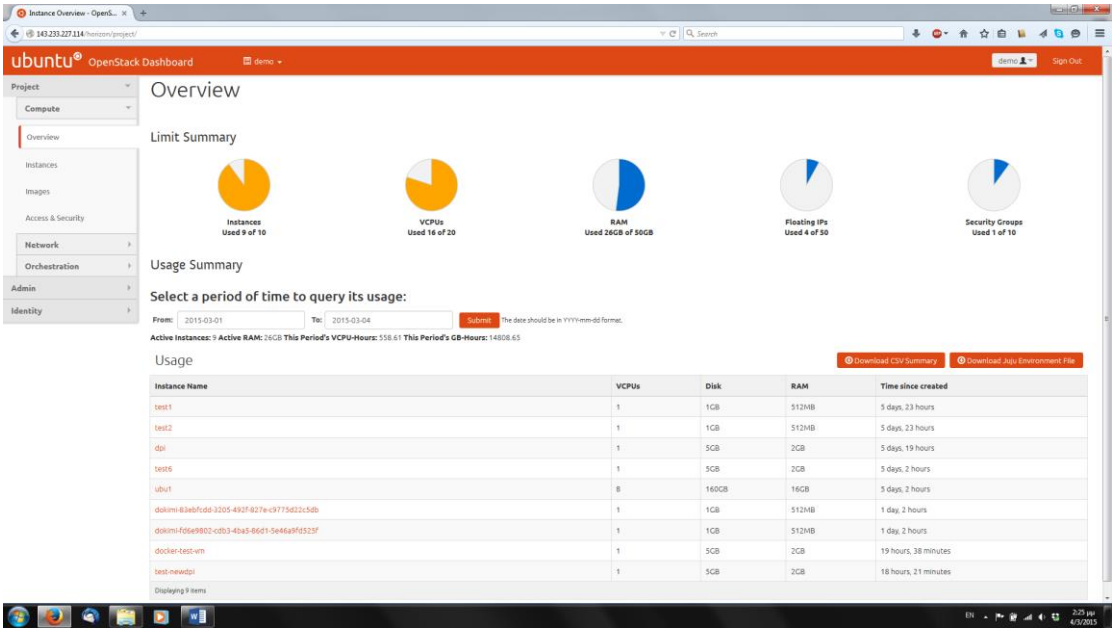

Figure 98. OpenStack Horizon Management Dashboard

In CloudSat, all three central dashboards are used, i.e. the “User Dashboard”, the “System Dashboard”, and the “Settings” dashboard.

The Horizon application also includes a set of API abstractions for the core OpenStack projects in order to provide a consistent, stable set of reusable methods for developers. Using these abstractions, developers working on Horizon don’t need to be intimately familiar with the APIs of each OpenStack project.

6.5.2. OpenDaylight (SDN controller)

OpenDayLight [ODL] is the SDN controller used in CloudSat, as proposed in Chapter 3. OpenDaylight is currently the newest and also largest SDN controller platform. Figure 99 depicts the management front-end for the OpenDaylight platform deployed in the CloudSat testbed.

In addition to the graphical front-end, OpenDaylight provides a flexible northbound interface using Representational State Transfer APIs (REST APIs). In the CloudSat testbed, these APIs are exploited by the federated management layer for automated network service control.

The southbound interface is capable of supporting multiple protocols (as separate plugins), e.g. OpenFlow 1.0, OpenFlow 1.3, BGP-LS, etc. These modules are dynamically linked into a Service Abstraction Layer (SAL). In CloudSat, the OpenFlow 1.3 plugin is used.

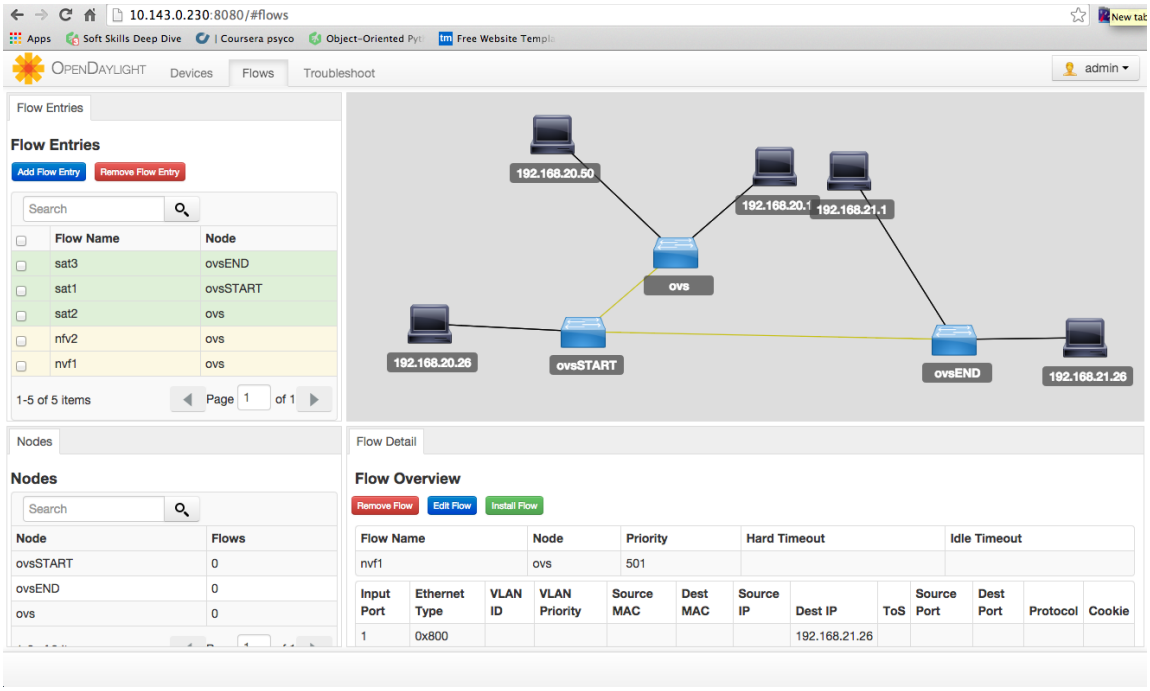

Figure 99. OpenDayLight Management GUI of CloudSat platform.

6.5.3. OpenSAND UI (Satellite Emulator Management)

OpenSAND provides an Admin GUI, through which appropriate configuration of the Satellite emulator can be performed. User through the interface has full control of the simulation process, while at the same time monitoring of the IP traffic over the satellite network is monitored. Figure 100 shows a snapshot of the GUI of the OpenSAND platform deployed for CloudSat.

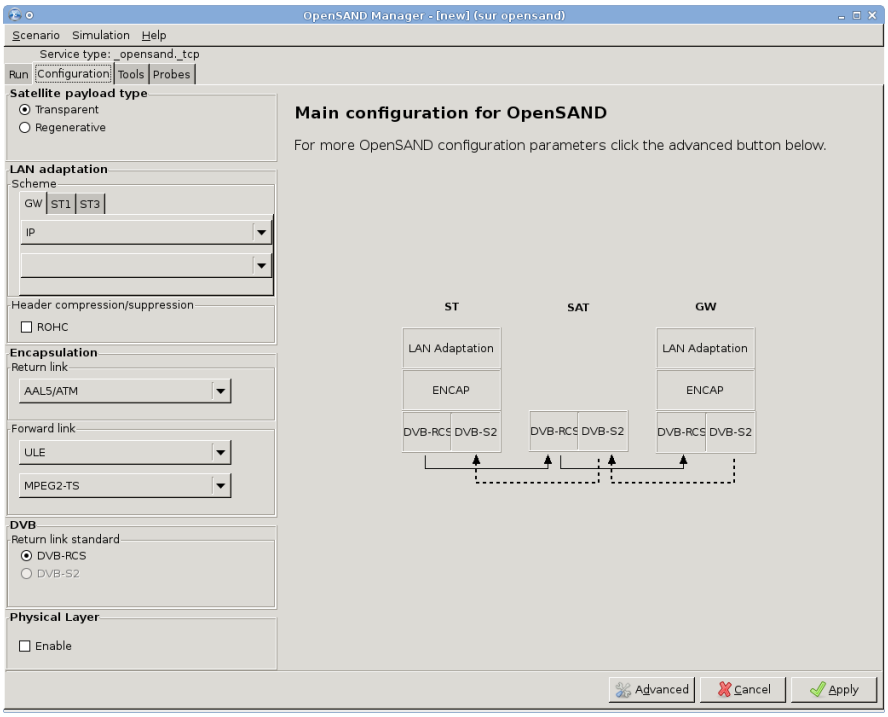

Figure 100. OpenSAND management GUI

6.6. Scenarios For Experimentation And Performance Analysis

In Chapter 4, an assessment of each of the described integration scenarios was performed in terms of technical added-value, technology readiness and market potentials. This analysis finally concluded to the selection of three Scenarios for further elaboration and experimentation. More specifically, the following scenarios were chosen with the aim to emphasize and demonstrate specific added-value characteristics of the SDN/NFV techniques:

- Scenario #1: Hybrid media distribution network as-a-Service

  This scenario aims at demonstrating the Federated management of the terrestrial and satellite segment together with the agility and network resource elasticity offered by the SDN and its management platform, the OpenDaylight. The scenario will show how the agility and reconfigurability introduced by SDN control greatly facilitates the efficient resource management and QoS enforcement within the network retaining seamlessly the delivered QoE level at the desired levels.

- Scenario #2: Dynamic backhauling with edge processing

  This scenario aims at demonstrating the SDN/NFV complementarity and flexibility in the case of the edge computing, which allows the continuation of a service, that under normal circumstances would have been seriously degraded due to the lack of resources. The demo of an SDN/NFV-enabled terminal is also presented.

- Scenario #3: Customer functions virtualization

  The last scenario demonstrates the automatic deployment of a network service over the OpenStack Cloud Computing platform, which is composed of three individual VNFs. The scenario demonstrates the SDN agility on the traffic steering of the selected flow through the VNFs that constitutes the Network Service, showing its positive impact on the satellite bandwidth saving, the flow optimization and the security coordination.

For the needs of the three scenarios an experimental topology has been deployed, which includes both the satellite and the terrestrial segment, where the satellite segment is emulated by OpenSand (three computing nodes) and the terrestrial segment is designated by two SDN-compatible Open Virtual Switches. It is important to point out, that the whole topology does not form a single network domain, but two discrete domains, where the satellite hub segment corresponds to the 192.168.20.0/24 network and the satellite terminal segment corresponds to the 192.168.21.0/24 network. Therefore the SDN-based traffic steering, the NFV orchestration and the federation is experimented over two different and discrete network domains, emulating the real challenges that are raised by the integration under one management entity of two different network domains.

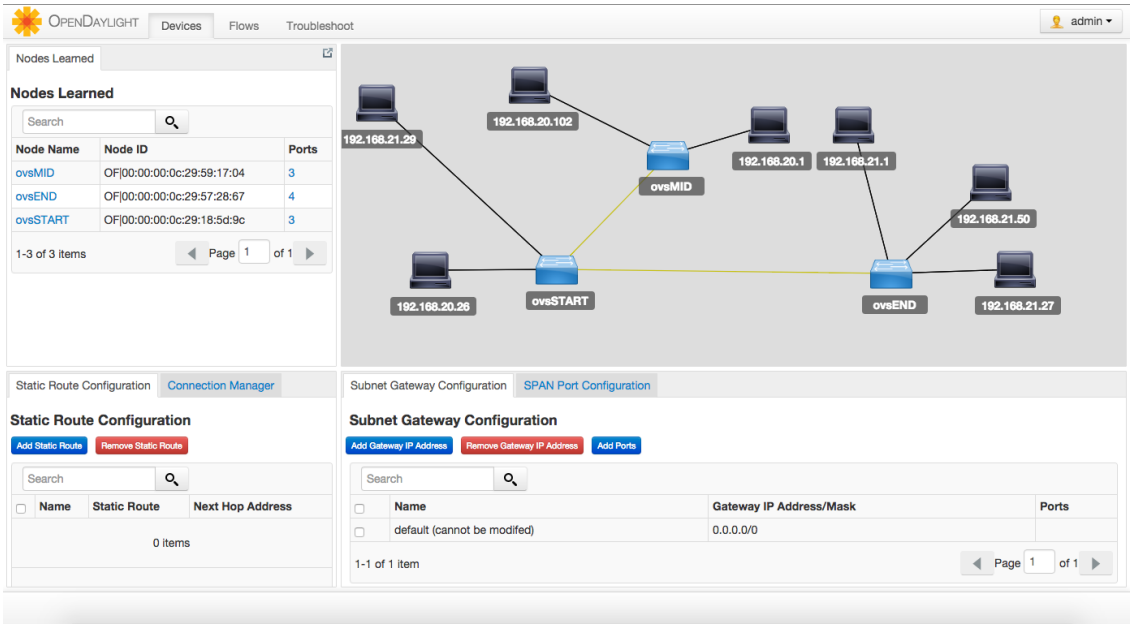

Figure 101. Experimental Topology for the Three Scenarios as visualised in OpenDayLight GUI

Figure 14 is a snapshot taken by the SDN controller and shows the unified network management interface that the OpenDaylight platform provides for the CloudSat experimental platform. As it can be observed by the figure, the SDN controller establishes a “federation umbrella” (super-controller) of both the terrestrial and the satellite segment on top of both infrastructures providing a single point of management and coordination.

6.6.1. Scenario #1: Hybrid media distribution network as-a-Service

6.6.1.1. Scenario Description

This scenario focuses on the federation of satellite and terrestrial domains and the provision of a hybrid satellite/terrestrial access network slice to a media service provider for content distribution. From a technological perspective, the scenario aims at demonstrating the agility and flexibility of SDN management over the federated infrastructure. Please refer to Chapter 4 for a detailed description of this scenario.

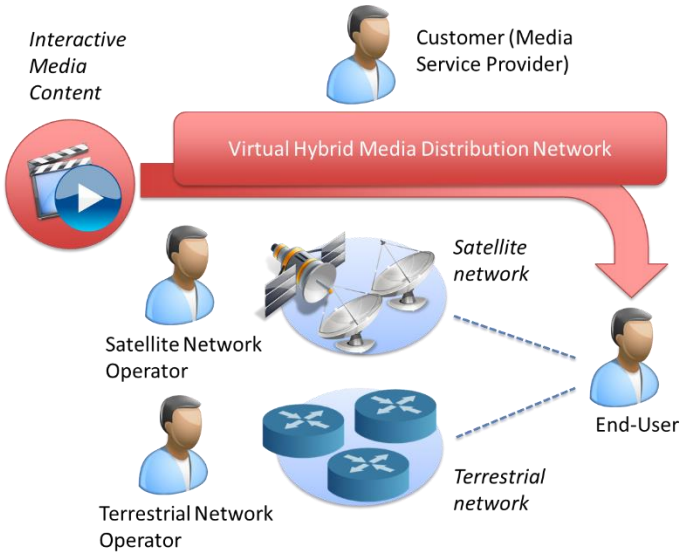

Figure 102. Hybrid media distribution network as-a-Service scenario

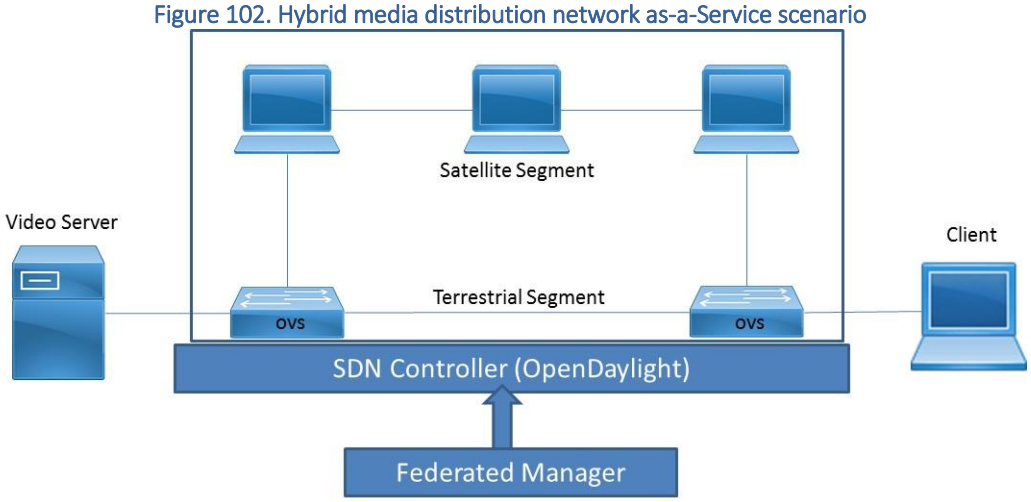

Figure 103. Hybrid media distribution network as-a-Service Experimental Topology

The experimental topology of this scenario is depicted in Figure 103, where at the ingress and egress points of the two segments (i.e. the Satellite and the Terrestrial) have been placed two SDN-compatible Open Virtual Switches (vSwitches), which are under the management and control of the OpenDaylight SDN controller. The use of SDN in this scenario permits the balancing of the load between the terrestrial and

satellite segment. Especially for the execution of the two demos of scenario 1, it should be pointed out that at the ingress OVS, the eth1 interface forwards to the satellite segment and the eth2 interface forwards to the terrestrial segment.

In the simplest approach, the MSP (Media Service Provider/content provider) just uses the hybrid virtual network as a “dumb pipe” (yet with specific SLA) to convey media streams. However, a significant added-value of the use of virtualization and programmability technologies would be to offer to the MSP elevated management and control capabilities on the hybrid virtual network. This means that the MSP may develop his/her own network control logic in order to dynamically configure the network at runtime, allocate resources and also influence routing/forwarding decisions as desired (i.e. divert streams from the terrestrial to the satellite channel and vice versa on-the-fly or adjust the load balancing between the two networks).

Furthermore, thanks to resource elasticity, the capacity and QoS offered to the MSP virtual network may fluctuate over time as coordinated by the Federator, enabling the MSP service to be up and down scaled on-demand or automatically, to react to the customers’ demand. This means that the MSP may dynamically request more capacity if needed (e.g. in case of highly popular content).

Therefore, for the Federation needs of this experimental scenario, an appropriate software module was developed to emulate the Federator/Orchestrator entity. This module exploits network programmability to achieve administrative federation of the satellite and network infrastructure segment and aggregates the monitoring data of the delivered service along with the utilization statistics of the network capacity, providing immediate decisions and actions on the load balancing of the delivered traffic between the terrestrial or the satellite segment.

This experimental scenario will execute two demos:

- Demo #1: SDN-based QoS policy enforcement
- Demo #2: SDN-based video stream steering

Both demos aim at presenting specific advantages of the SDN applicability on the terrestrial and satellite segment federation. Simultaneous hybrid service delivery with scalable media service is not considered in this demo, due to the synchronization difficulties that are introduced by the different delay of the two network segments (i.e. the satellite and the terrestrial), making the implementation of the scenario beyond the scope of this study, which aims at presenting the agility and the performance efficiency of SDN. Moreover, it should be pointed out that that there are not sufficiently mature implementations of scalable video suites (i.e. codecs, servers and clients) capable of providing a stable multipath scalable video service over such a highly heterogeneous environment based.

6.6.1.2. Demo #1: SDN-based QoS policy enforcement

This demo scenario considers that a unicast video service is initially delivered over the satellite network at the pre-defined QoE level. The storyline of this scenario is the following:

1. Video service initially delivered over the satellite network
2. Background network traffic degrades video quality
3. The Federator monitors and applies appropriate QoS policy
4. SDN-policy is applied (L2 queue prioritization) (L2 QoS)
5. Video quality is reinstated
6. Queue traffic is increased -> New Video Quality Degradation
7. Video service is shifted to an even higher priority queue (Network Resource Elasticity)
8. Video quality is reinstated

Figure 17 depicts the initiation of the unicast video service from the content server (192.168.20.26) over the satellite link towards the end-user 192.168.21.27 at port 33334 utilizing the MPEG-4 video codec using Simple Profile with spatial resolution 640x480, frame rate 24 fps and at ~1024 kbps.

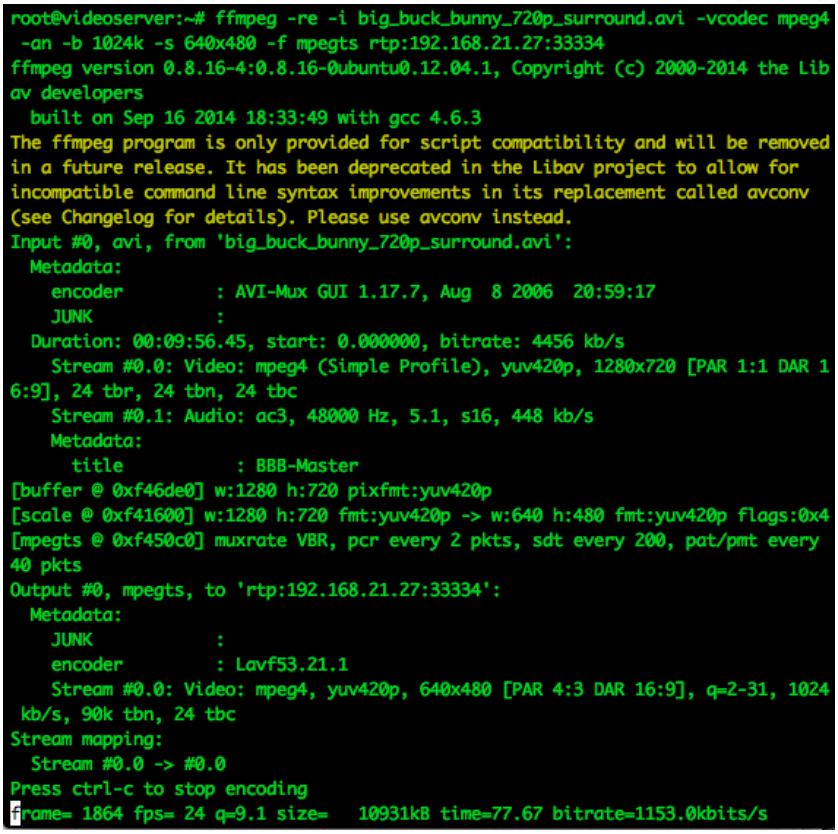

Figure 104. Initiation of the Unicast Video Streaming over the Satellite segment

The unicast video stream is forwarded over the satellite link. As shown below, the ICMP round-trip-time (RTT) to the video server, is approx. 530 msec.

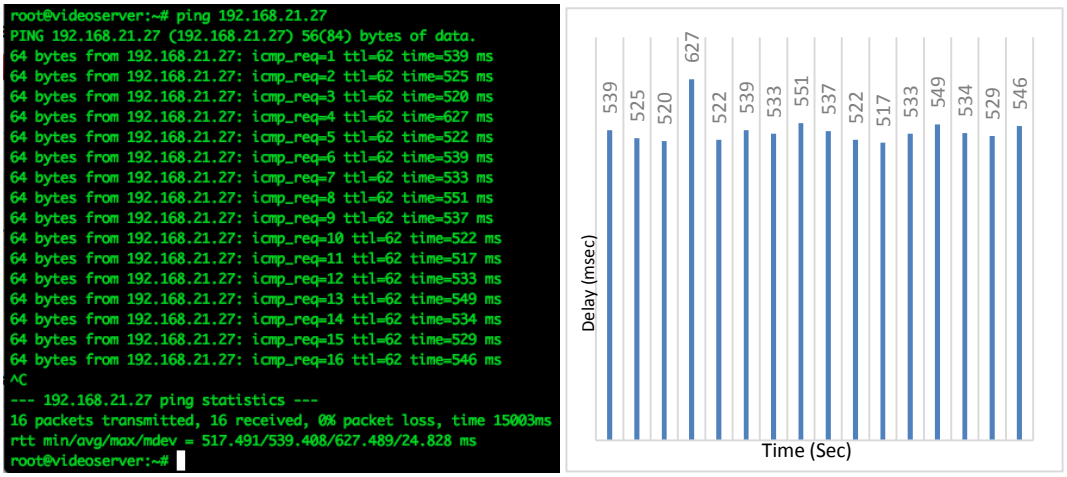

Figure 105. ICMP RTT to the video server over the satellite link

The unicast video traffic successfully passes through the first switch (labelled "ovsSTART"), where it is monitored by the Federator, as the following figure depicts, at a detected rate of 141.23 kB (i.e. approx. 1Mbps).

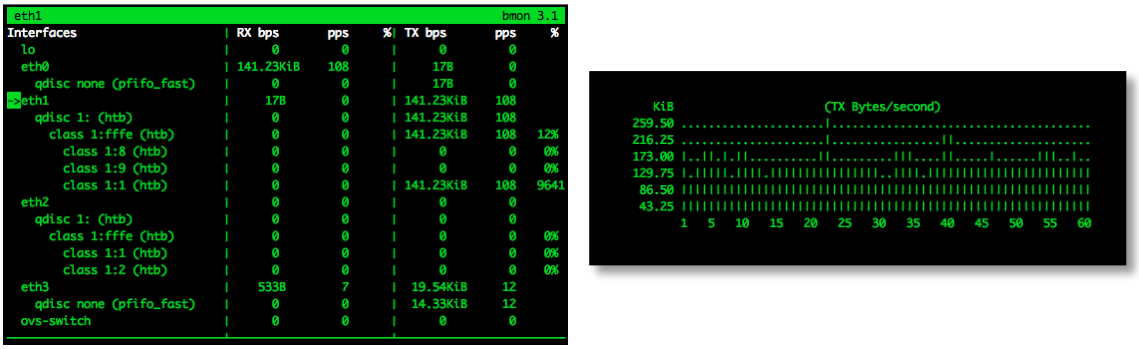

Figure 106. Video Traffic monitoring at Ingress OVS (Eth1, Class 1:1)

Figure 20 depicts the flawless playback of the unicast video service at the end-user premises (i.e. 192.168.21.27)

Figure 107. Normal media delivery over satellite

According to the story line, background traffic is added in the satellite channel in order to saturate it and thus force quality degradation (due to packet loss) of the delivered video service. For the experimental needs of the scenario, the maximum

available bandwidth of each interface has been reduced to 10 Mbit in order to be easily saturated with background traffic.

Towards flooding the satellite link with background traffic, synthetic UDP traffic is generated by a Linux virtual machine utilizing the iperf command. Iperf generates approximately traffic of 10Mbit, which is enough in order to flood the link and therefore create significant degradation to the delivered video service.

```
root@videoserver:~# iperf -c 192.168.21.27 -u -b 10M -t 120
------------------------------------------------------------
Client connecting to 192.168.21.27, UDP port 5001
Sending 1470 byte datagrams
UDP buffer size:  224 KByte (default)
------------------------------------------------------------
[  3] local 192.168.20.26 port 33547 connected with 192.168.21.27 port 5001
```

Figure 108. Background traffic is introduced over Satellite segment

Upon the introduction of the background traffic, significant degradation is observed in the perceived quality of the delivered service, which is made practically unviewable due to multiple error propagations. In terms of video quality assessment, the respective average QoE level (measured with the SSIM[15] metric) drops from the approx. 0.83 value down to 0.21.

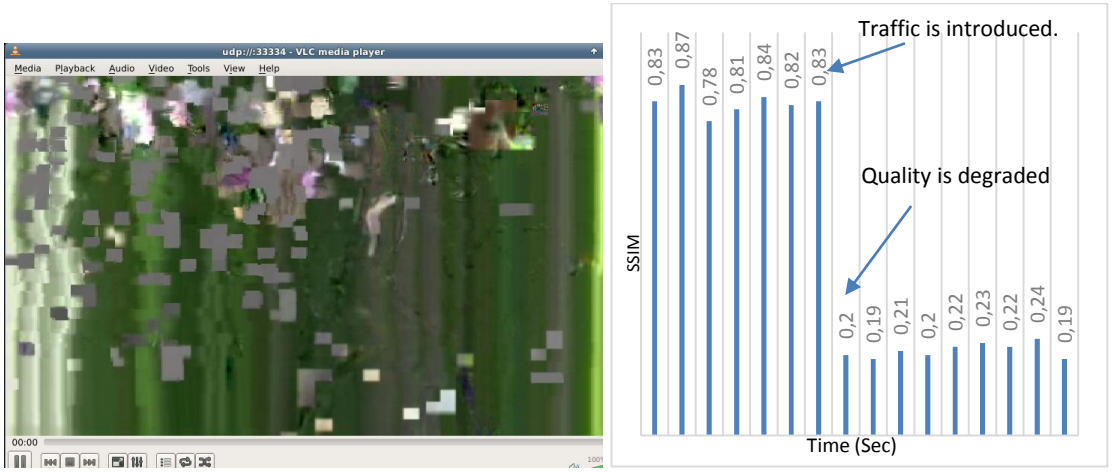

Figure 109. Video quality degradation due to packet loss

The background traffic together with the unicast video traffic is monitored by the Federator at the ovsSTART switch, as the following figure depicts, showing a total flow of 1222.68 kB (i.e. approx. 10Mbps).

[15] The Structural Similarity (SSIM) metric is a method for inferring the perceptual quality of a video stream. It is a full-reference metric, utilizing the average, variance and covariance of the luminosity values.

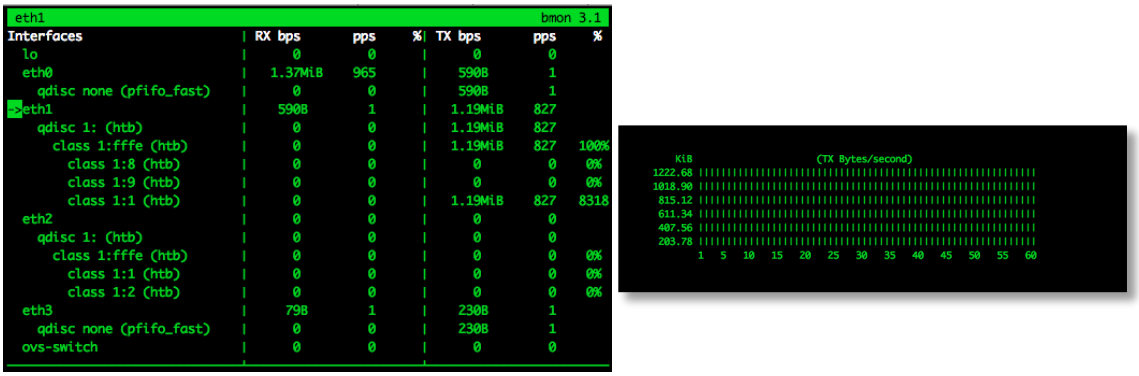

Figure 110. Network congestion at Ingress OVS (Eth1, Class 1:1) and Network Link Utilization Chart

The Federator automatically senses the network congestion by continuously monitoring the network condition, and immediately reacts by applying appropriate SDN-based QoS policy based on L2 queue prioritization of the unicast video service.

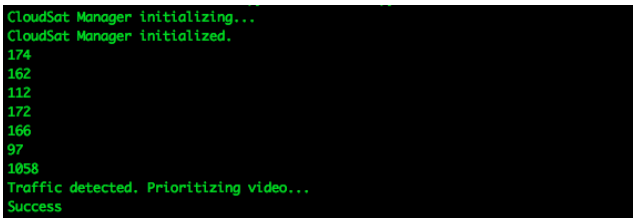

Figure 111. Federator detects traffic and applies QoS prioritization

This SDN-based network elasticity policy is also monitored by the federator and is depicted in the following figure, showing that in the Queue class 1:1 the background traffic is served, while in the Queue class 1:9 (which is of higher priority than class 1:1) the unicast video service is delivered flawlessly.

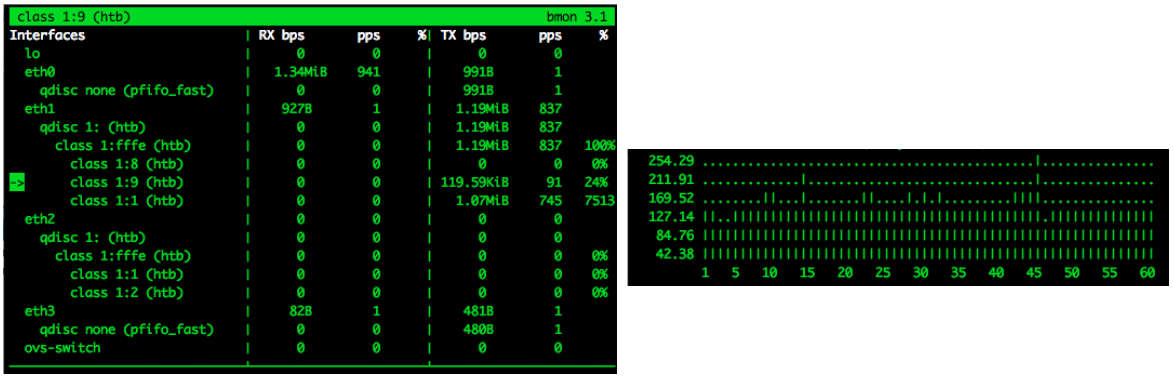

Figure 112. Traffic Classificaton (Eth1, Class 1:1 & Class 1:9) and Class 1:9 Network Link Utilization Chart

Respectively Figure 25 depicts the traffic chart of the unicast video service, without the deterioration of the background traffic, which has been differentiated and served by a different queue class.

A metric to be measured is the response time of the Federator, i.e. the delay from the congestion incident to the actual application of the new stream priority. Upon the execution of an experimental set of ten repetitions, the responsiveness of the federator to the traffic congestion, measured within the range of 1-2 secs. (An accurate measurement cannot be derived, since the response time is measured by the human administrator.)

Figure 26 depicts the improvement in the QoE level of the delivered unicast service in parallel with SSIM measurements during the SDN adaptation action. It should be noted in terms of analysing the achieved service stability, that by increasing or decreasing the background traffic, the prioritized video quality service is not deteriorated and the flawless video service delivery is retained.

Figure 113. Video Quality is reinstated

However, video quality degradation may also occur in cases that the prioritized queue does not have the appropriate bandwidth in order to handle the specific video service or other in-queue background activities may stress the available queue bandwidth utilization, resulting again to QoE degradation and quality distortion.

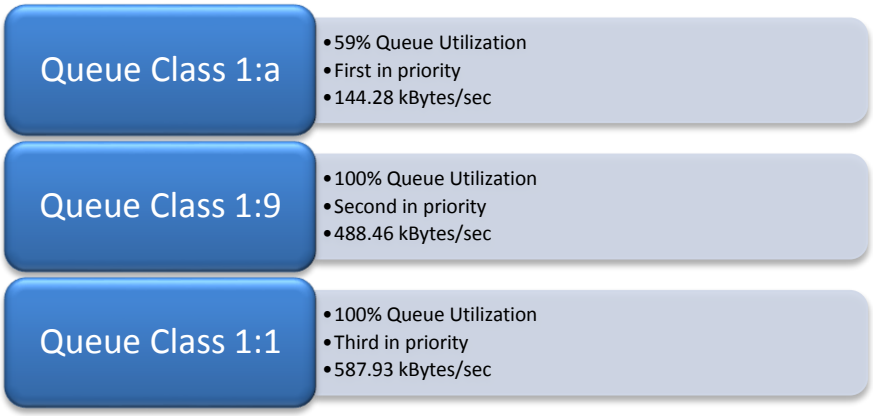

Figure 114. Service prioritisation with triple SDN-based eth1 QoS queues

In this case, the federator, which continues to monitor the delivery of the unicast service, will sense the service degradation and will again apply appropriate SDN-based QoS policies by upgrading the classification of the video service to a higher class (1:a), achieving again service prioritisation and quality improvement.

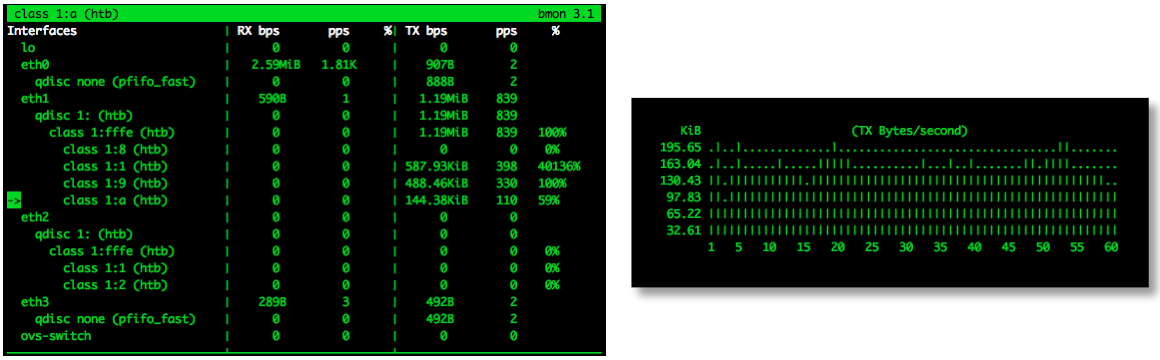

Figure 115. Traffic Classificaton (Class 1:1,Class 1:9,Class 1:a) and Class 1:a Network Link Utilization Chart

Figure 28, depicts this triple service prioritisation/isolation into three discrete SDN-based Layer 2 QoS queues of eth1, where both queue class 1:1 and class 1:9 are 100% utilized by background traffic, while class 1:a is utilized at 59% serving only the unicast video service.

6.6.1.3. Demo #2: SDN-based video stream steering

This demo scenario will present the SDN-based video stream steering between the satellite and the terrestrial segments, considering that a unicast video service is initially delivered over the terrestrial network at the pre-defined QoE level and then due to service degradation (or other triggering event) the service delivery is switched seamlessly via the satellite segment.

Thus, in this demo, the unicast media streams are load-balanced between the satellite and the terrestrial segment, according to the available network resources. We assume that the primary distribution channel should be the terrestrial one; the customer receives the media content over terrestrial and when insufficient terrestrial capacity is observed, the traffic is diverted by appropriate SDN multipath rules over the satellite segment.

The demo can be also executed in the reverse order, where the primary delivery channel is different and considers the switching of the video service from the satellite segment to the terrestrial segment. The storyline of this scenario is the following:

1. Video service delivered over terrestrial network
2. Background terrestrial network traffic degrades video quality
3. Federator monitors and applies appropriate traffic steering SDN rule
4. SDN-rule is applied (L2 forwarding over satellite)
5. Video quality is reinstated

Figure 29 depicts the initiation of the unicast video service from the content server (192.168.20.26) over the terrestrial link towards the end-user 192.168.21.27 at port 33334 utilizing the MPEG-4 video codec using Simple Profile with spatial resolution 640x480, frame rate 24 fps and bitrate ~1024 kbps.

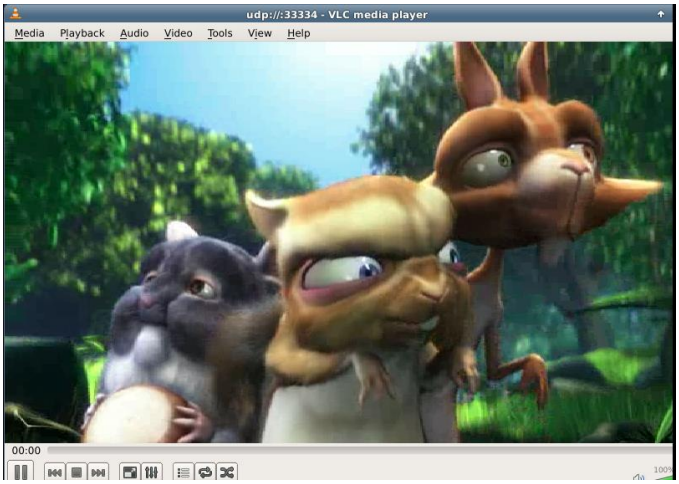

Figure 116. Normal media delivery over terrestrial

The unicast video traffic successfully passes through the entry switch (ovsSTART), where it is monitored by the Federator, as the following figure depicts, measuring a flow rate of 133.16 kB (i.e. approx. 1Mbps) at eth2, which is the port that leads to the terrestrial segment. During the delivery of the unicast media service the terrestrial network link, as depicted on Figure 30, utilizes approximately 11% of the overall available network bandwidth.

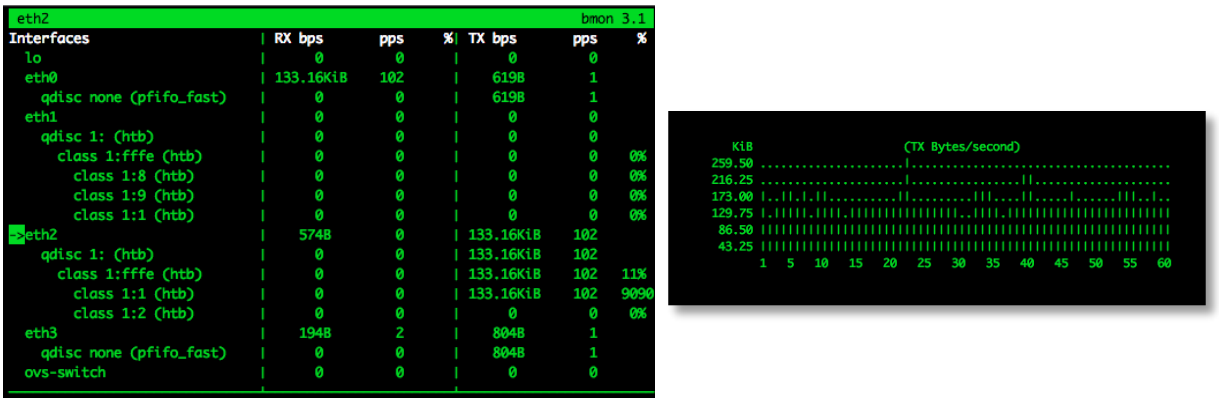

Figure 117. Media Delivery over terrestrial

The unicast media stream is delivered over the terrestrial link, it can be seen that the ICMP RTT to the video server is below 1 msec.

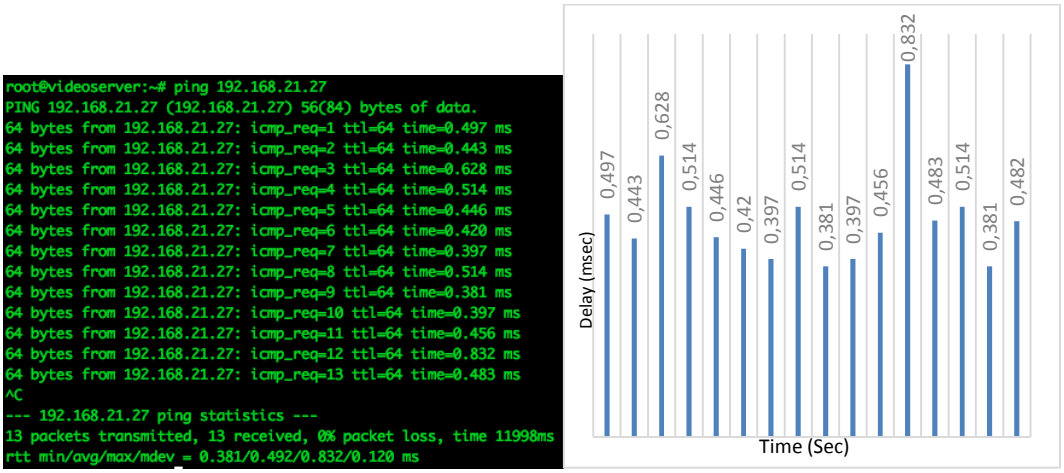

Figure 118. RTT to the video server over the terrestrial link

According to the story line, background traffic is added in the terrestrial channel in order to force quality degradation of the delivered video service. For the experimental needs of the scenario, the maximum available bandwidth of each interface has been reduced to 10 Mbit in order to be facilitated to traffic flooding of the channel with background traffic.

Towards flooding the terrestrial link with background traffic, synthetic UDP traffic is generated by a Linux virtual machine utilizing the iperf command similarly to the previous demo. The produced traffic creates approximately traffic of 10Mbit, which is enough in order to flood the link with 70% utilization and therefore creating quality degradation to the delivered video service due to network impairments, such as jitter, delay etc. This congested network link is depicted in Figure 32, where a traffic of 1.19 MB/sec is monitored at eth2 and the quality degradation at the SSIM metric is depicted on Figure 33.

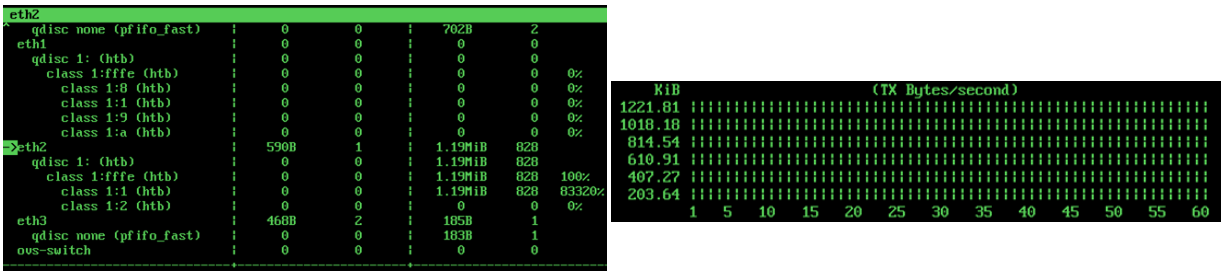

Figure 119. Background traffic introduced in terrestrial link

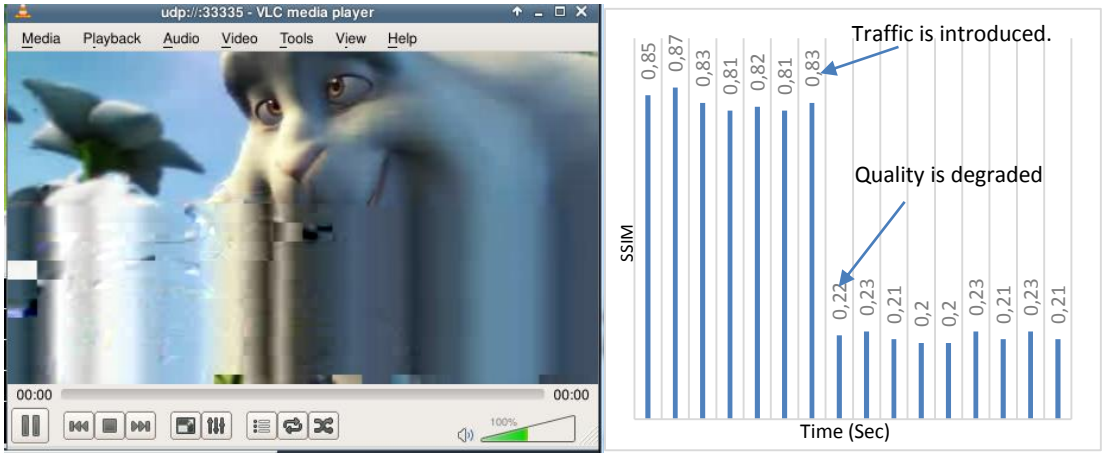

Figure 120. Severe video quality degradation

The Federator automatically senses the network congestion and immediately applies appropriate SDN-based traffic steering commands at the ingress OVS of the experimental topology (i.e. ovsSTART). The media service is seamlessly diverted to be delivered over the satellite link, while the rest background traffic is still conveyed over the terrestrial domain.

```
CloudSat Manager initializing...
CloudSat Manager initialized.
174
162
112
172
166
97
1058
Traffic detected. Prioritizing video...
Success
```

Figure 121. Federator detects traffic and applies SDN-based load balancing rule

The response time of the federator in this congestion is again in the order of 1-2sec.

Figure 36 visualises the activity of the Federator, which monitors both eth1 (i.e. the satellite domain port) and eth2 (i.e. the terrestrial domain port), where it is observed that the media service (approx. 144.93kB/sec) is delivered over the satellite link (through eth1 port) at queue class 1:9 (reassuring even prioritization among the other satellite flows) and the rest background traffic (approx. 732.34kB/sec) continues to be delivered over the terrestrial link (through eth2). The delivery of the media service over the satellite link is confirmed also by the measurement of the round-trip delay between the MSP and the client, which has been increased at approx. 535 msec.

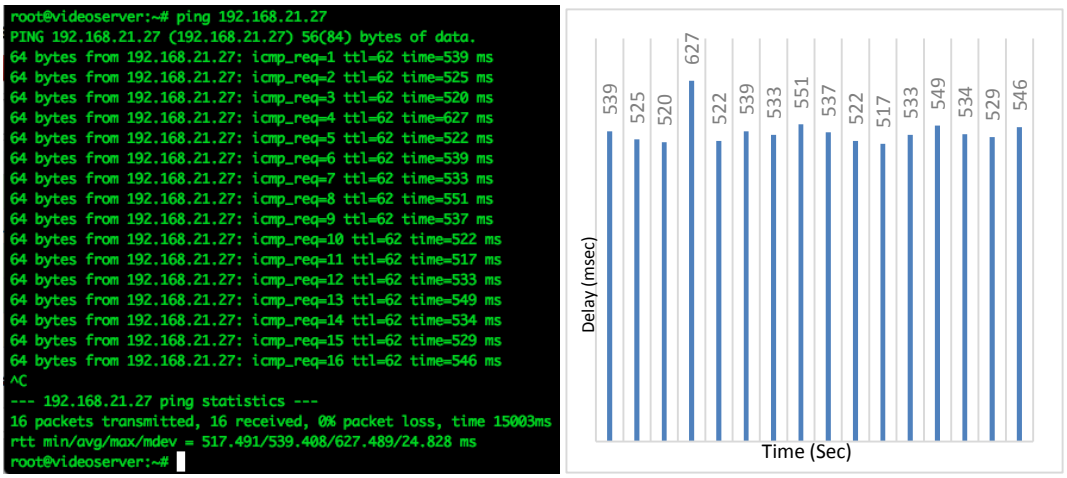

Figure 122. ICMP RTT to the video server over the satellite link

In this traffic steering demo, the network utilization of the satellite link is approx. 10%, while the network utilization of the terrestrial link is approx. 60%, as figure 34 depicts.

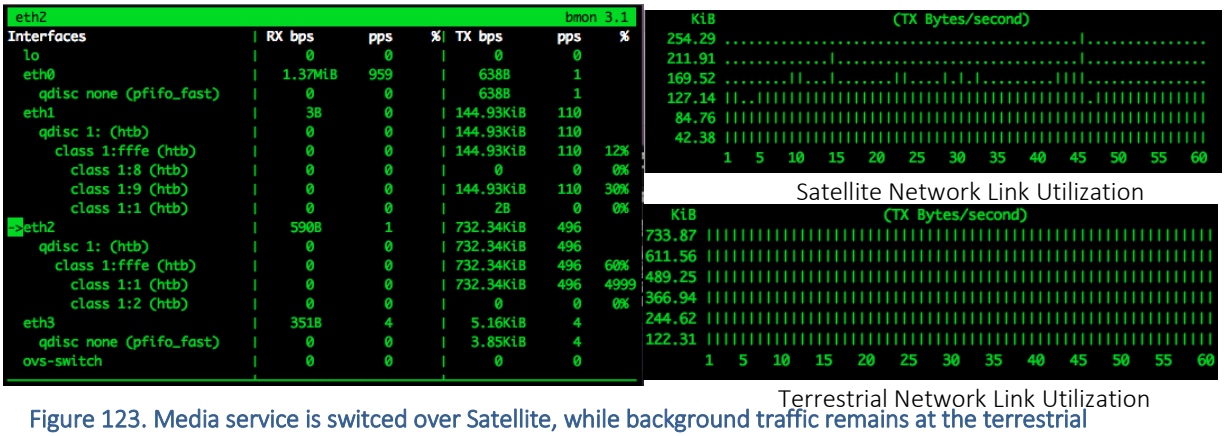

Figure 123. Media service is switced over Satellite, while background traffic remains at the terrestrial

Upon the traffic steering of the media service over the satellite link, the video quality of the media service (SSIM) is reinstated at approximately 0.85 from approx. 0.21 as it is depicted on Figure 37

Figure 124. Video Quality is reinstated

6.6.2. Scenario #2: Dynamic backhauling with edge processing

6.6.2.1. Scenario Description

The satellite edge-processing scenario assumes the extension of the Mobile Edge Computing (MEC) paradigm to the satellite domain; specifically, it foresees that the backhauling service is coupled with virtualization capabilities at the satellite terminal, able to host virtual traffic processors close to the end users (Figure 59). Such local traffic processing can achieve significant savings in satellite capacity. Please refer to Chapter 4 for a detailed description of this scenario.

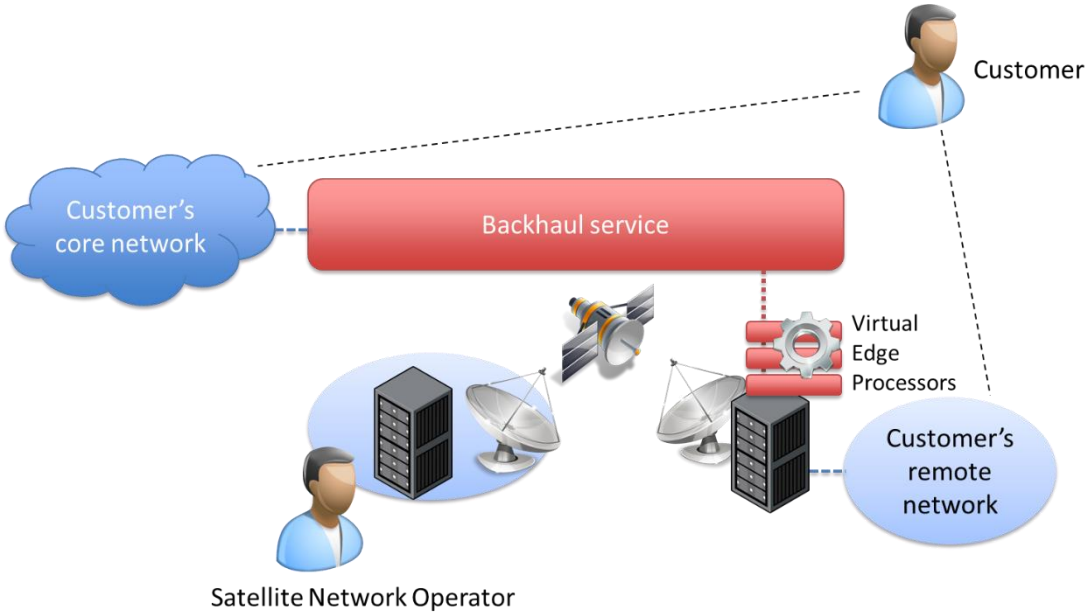

Figure 125. Dynamic backhauling with edge processing scenario

This demo focuses on experimenting the news aggregation case, where user generated content (e.g. videos) is transmitted over the satellite towards the news

aggregator server. Due to specific bandwidth availability, especially in case of multiple users, the generated video content may not be possible to be transmitted over the satellite link and therefore network congestion should result to quality degradation or even service interruption. By exploiting the Mobile Edge Computing capabilities of an appropriate VNF instantiated at the SDN/NFV-enabled Satellite Terminal, the generated streams can be dynamically transcoded and then transmitted back over satellite, minimizing the bandwidth utilization needed for the transmission of the video content and optimizing the video transmission given the total number of the video signals and the available bandwidth of the satellite link.

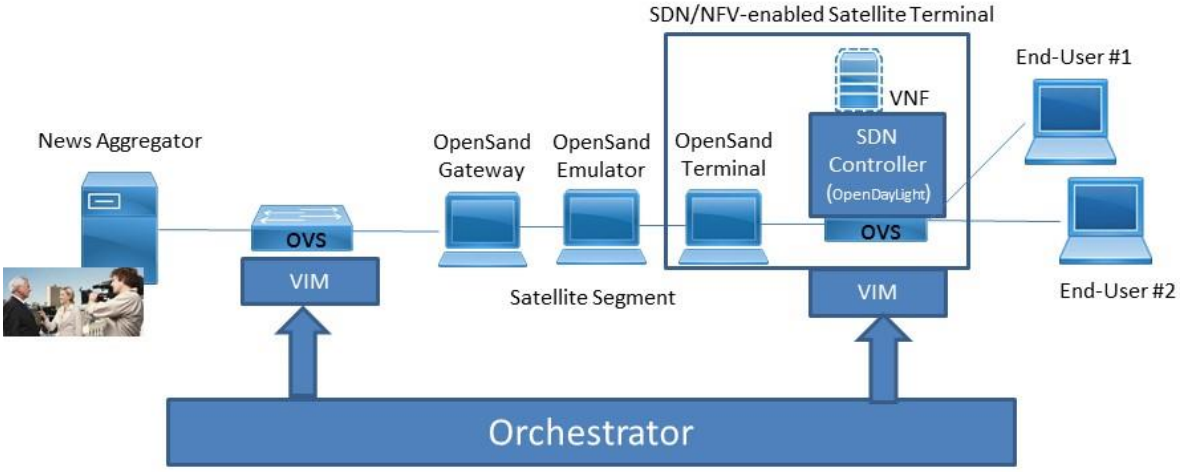

Figure 126. Dynamic backhauling with edge processing Experimental Topology

The experimental topology of this scenario is depicted in figure 39, where at the ingress and egress points of the satellite segments they have been placed two SDN-compatible Open Virtual Switches, which are under the management and control of the OpenDaylight SDN controller and the orchestrator.

With regard to edge processing, the NFV coupled with emerging Mobile Edge Computing (MEC) concepts for deployment of cloud resources at the network edge, are the key enabling technologies. The satellite terminal needs to encompass virtualized IT resources in order to host the traffic processors, as virtual network functions (VNFs). Thus, at the edge of the satellite segment, it is considered an SDN/NFV-enabled Satellite Terminal, which is capable of being controlled by the Orchestrator, instantiating appropriate VNFs and implementing upon the orchestrator mandates the required SDN rules for performing traffic steering as needed for the Service Function Chaining (SFC).

When it comes to management, since it is not advisable to deploy an entire cloud system (e.g. Openstack) on the terminal, in a more lightweight approach, the terminals can encompass plain IT virtualization (e.g. via a KVM hypervisor or even via Docker containers), without any cloud framework. This approach has the cost of reduced elasticity and management features. However, it saves IT resources and also relieves the satellite segment from excessive signaling, thus it would be more appropriate for edge VNFs (rather than for VNFs hosted at the Gateway, where Openstack-based management is still advisable).

Therefore, for the orchestration needs of this experimental scenario, an appropriate software module was developed, emulating an Orchestrator platform. This module achieves NFV instantiation at the SDN/NFV-enabled satellite terminal and applies the appropriate SDN rules at the OVSs for supporting the SFC through appropriate traffic steering actions.

The demo (Live news gathering with dynamic transcoding) aims at presenting specific advantages of the SDN/NFV applicability at the edge of the satellite segment, such as real-time service adaptation, minimization of the satellite link utilization, achieving scalability in case of multiple end-users by applying network resource elasticity per end-user.

6.6.2.2. Demo: Live news gathering with dynamic transcoding

This demo scenario considers that two end users wish to simultaneously transmit (professional?) user-generated video content back to the video news aggregator via the same (professional) satellite terminal. Therefore two discrete unicast flows are initiated by the end-users with final destination the remote news aggregator server. However, due to limited satellite bandwidth the two media streams exceed the available bandwidth in the satellite link resulting in network congestion and therefore degradation of the QoE of the transmitted media signals.

Towards facilitating the video transmission in a dynamic and transparent way for the end-users, the orchestrator monitors the quality degradation and (following the Mobile Edge Computing (MEC) paradigm) instantiates a transcoder as VNF at the SDN/NFV-enabled terminal and applies appropriate SDN-based traffic steering rules at the OVSs to transparently steer the two media flows through the VNF-based transcoder and then to be forwarded over the satellite in order to finally reach the news aggregator. The storyline of this scenario is the following:

1. Two end-users are sending (each one) a unicast media service to the news aggregator
2. Both media services are transmitted over the satellite network uplink
3. Total traffic of the two media services exceeds the available bandwidth of the satellite link resulting to network congestion and quality degradation.
4. Orchestrator monitors and instantiates at the SDN/NFV-enabled satellite terminal a VNF-based transcoder and appropriate SDN rules for the traffic steering in order to support the SFC.
5. Both media streams are transcoded (i.e. at lower bitrate/frame rate/resolution) in real time and transparently from the end-users
6. The two transcoded media services can fit in the available satellite bandwidth and are transmitted without any network congestion.
7. The QoE level of the two transmitted signals is re-instated at satisfactory levels and the two signals reach the remote news aggregator server without impairments.

Initially each user initiates a unicast video from its terminal back to the news aggregator over the satellite return link. Figure 40 depicts this initiation from the each one of the two terminals. The first terminal transmits its media service at the port 33334 and the second terminal at the port 33335.

```
root@enduser:~# ffmpeg -re -i big_buck_bunny_720p_surround.avi -vcodec mpeg4 -an
 -b 512k -s 320x280 -f mpegts rtp:192.168.20.26:33334
ffmpeg version 0.8.16-4:0.8.16-0ubuntu0.12.04.1, Copyright (c) 2000-2014 the Lib
av developers
  built on Sep 16 2014 18:33:49 with gcc 4.6.3
```

```
root@enduser:~# ffmpeg -re -i big_buck_bunny_720p_surround.avi -vcodec mpeg4 -an
 -b 512k -s 320x280 -f mpegts rtp:192.168.20.26:33335
ffmpeg version 0.8.16-4:0.8.16-0ubuntu0.12.04.1, Copyright (c) 2000-2014 the Lib
av developers
  built on Sep 16 2014 18:33:49 with gcc 4.6.3
```

Figure 127. Initiation of video flows

The bitrate of each video is approximately 256 kbps, so the two streams sub up at approximately 512 kbps, which exceeds the capacity of the emulated DVB-RCS return channel. The allocated return channel satellite terminal (RCST) capacity is constrained to approx. 500 kbps.

In our experimental case the total uploading of the two individual media services (i.e. 512 kbps each) exceeds the available uplink capacity resulting to severe quality degradation of both media services at the media aggregator side as it is depicted on Figure 41.

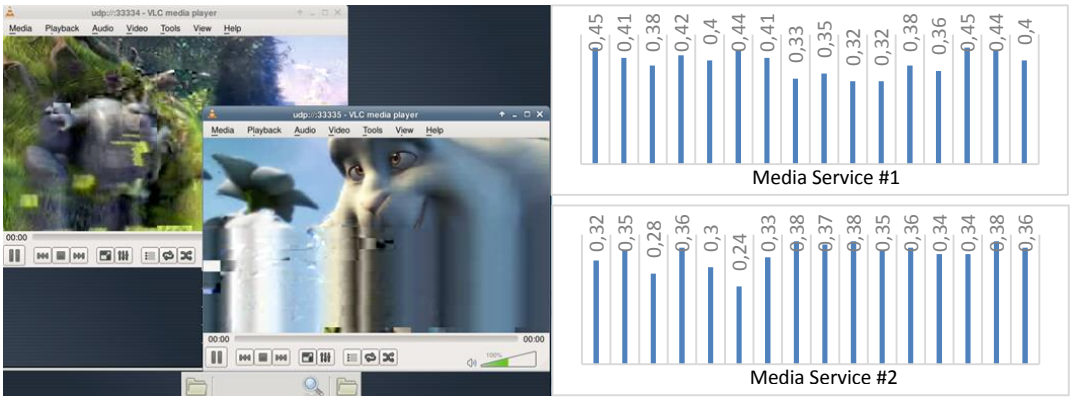

Figure 128. Degradation of the two streams caused by exceeding the capacity of satellite link

Towards dynamically and seamlessly improving the media service delivery, the orchestrator instantiates the transcoder VNF at the SDN/NFV-enabled satellite terminal, as it is depicted in Figure 42, and performs in real time and totally seamlessly to the end-user either spatial or temporal or transcoding (or combination of them) in both video signals. It should be pointed out that for each media service a different instance of the transcoder VNF is instantiated, which results on two different instances (one for the media service at port 33334 and one for the service at the port 33335) that are executed on the SDN/NFV-enabled satellite terminal.

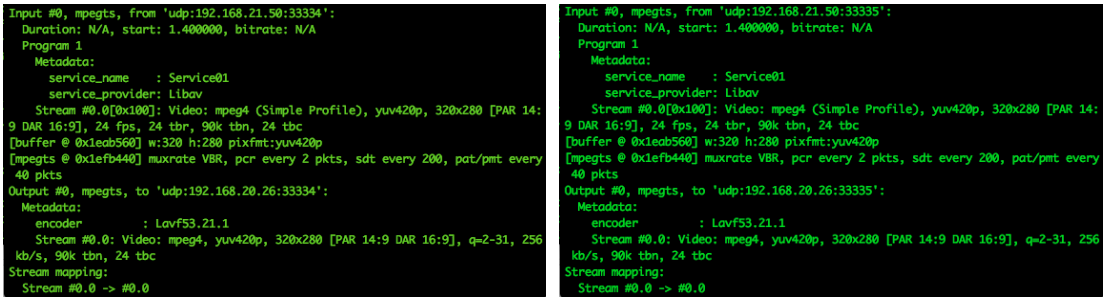

Figure 129. Initiation and operation of the two transcoder VNF instances

Figure 43 depicts the total traffic before and after the transcoding process, as monitored at the NIC of the VM, on which the two instances of the transcoder VNF have been instantiated.

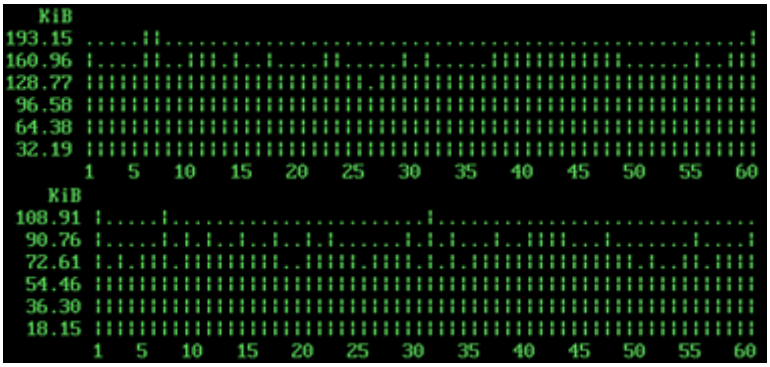

Figure 130. Total Media Service Traffic Before and After Transcoding

In terms of VNF performance evaluation, the VM that hosts the two instances of the transcoder VNF has been assigned 1GB of RAM, 15GB of HDD size and 1 CPU core (2GHz). As it is depicted on figure 44, during VNF operation, the VM has an overall system load of 5% in terms of CPU utilization, 9.4% HDD utilization and 36% memory usage, allowing the VM to operate on a stable and efficient status.

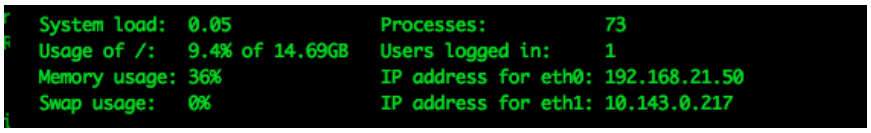

Figure 131. performance parameters (%CPU, %MEM, %HDD) of the VM hosting the VNFs

Considering the performance of each instance of the VNF running at the specific VM, Figure 45 depicts that each instance utilizes approximately 3% of the CPU (instance #1 3.3%, instance #2 2.7%), while each of them occupies only 0.8% of the available RAM. Thus, for the selected configuration, the two VNFs consume a relatively low amount of resources. Of course, this amount is only indicative and depends on the implementation of the transcoder and also on the bitrate and format of the streams being transcoded.

```
top - 17:13:30 up 8 days, 23:02,  3 users,  load average: 0.02, 0.02, 0.05
Tasks:  72 total,   2 running,  69 sleeping,   1 stopped,   0 zombie
Cpu(s):  7.5%us,  0.7%sy,  0.0%ni, 91.8%id,  0.0%wa,  0.0%hi,  0.0%si,  0.0%st
Mem:   1019512k total,   637924k used,   381588k free,   159824k buffers
Swap:  1044476k total,        0k used,  1044476k free,   366200k cached

  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
12674 root      20   0  108m 7992 4056 S  3.3  0.8   0:01.59 ffmpeg
12651 root      20   0  108m 8024 4036 S  2.7  0.8   0:19.87 ffmpeg
```

Figure 132. Performance parameters of the two VNF instances (%CPU, %MEM)

Upon the real time transcoding of the two media services from 512kbps down to 256 kbps each, the video quality is reinstated seamlessly (i.e. without requiring any interruption) for both signals as it is depicted in Figure 45.

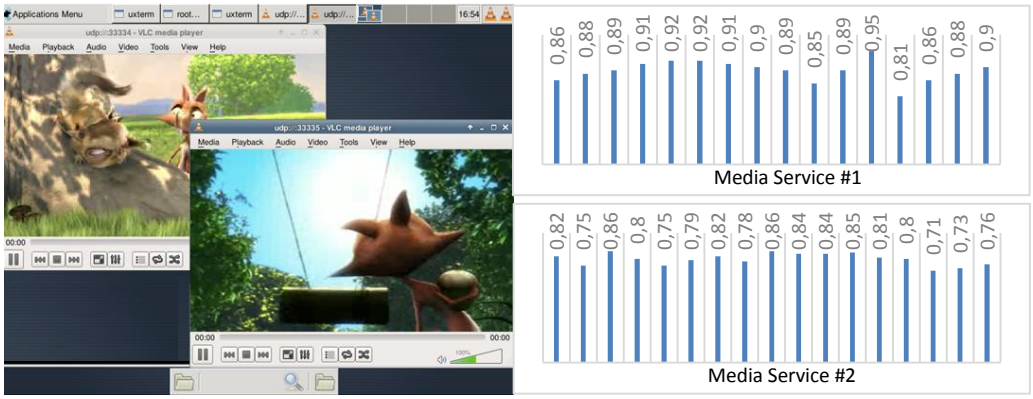

Figure 133. Recovery of the stream quality via transcoding

Finally it should be noted that the responsiveness of the system (measured from the congestion incident until the recovery of the video quality after the traffic steering via the VNFs) was measured approx. at 2 seconds, depending on the cache configuration of the client at the news aggregator server. This means that it took approx. 2 seconds for the system to detect the congestion and redirect the traffic via the edge VNF.

6.6.3. Scenario #3: Customer functions virtualization

This scenario is based on the VNF-as-a-Service (VNFaaS) paradigm and assumes the dynamic offering of virtual network appliances to satcom customers in the form of VNFs (e.g. firewalls, traffic filters, home gateway functionalities, media storage and processing etc.). According to their nature, these VNFs can be instantiated either at the satellite gateway or at VNF-enabled satellite terminals. Please refer to Chapter 4 for a detailed description of this scenario.

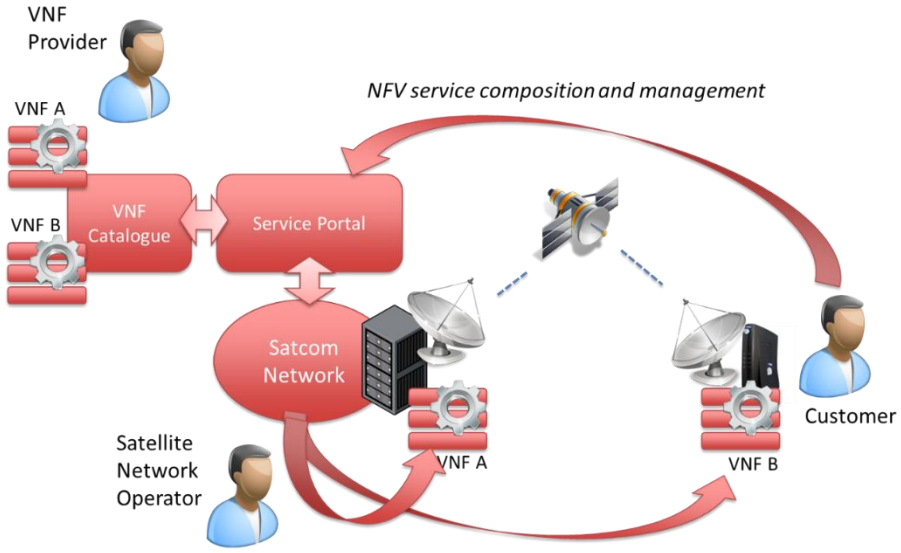

Figure 134. Customer functions virtualization scenario

This demo focuses on experimenting the automatic network service deployment, where the Satellite Network Operator deploys the VNFs and interconnects them, following a customer request. In a more interactive and dynamic approach, the customer composes the NFV service in a completely automated manner by accessing a service portal, browsing the VNF catalogue, selecting the VNFs which best match his/her needs and integrating them into a satcom service package.

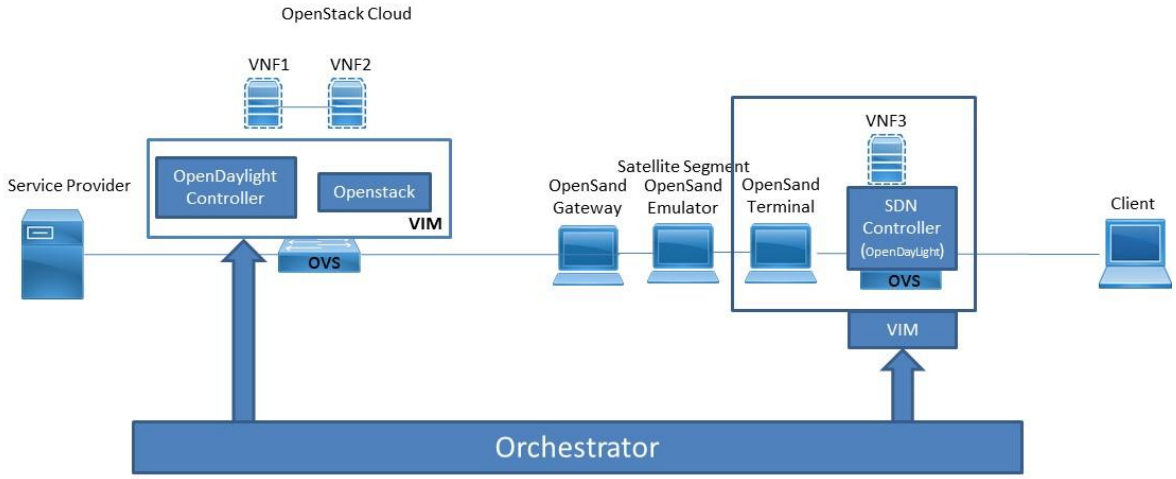

Figure 135. Customer functions virtualisation Experimental Topology

The experimental topology of this scenario is depicted in figure 48, where a NFVI PoP (NFV Infrastructure Point-of-Presence) is considered at the GTW side, together with

an SDN/NFV-enabled terminal in order to support the automatic deployment of the Network Service. The NFVI PoP is built on OpenDaylight SDN controller and Openstack Cloud computing platform. Commonly, the NFV management entities are deployed at the Gateway side, controlling NFV resources both local (at the Gateway) and remote (at the terminals). The NFV management through the orchestrator carries out procedures for:

- VNF instantiation, i.e. launching of the VNF images in the host machines
- Service function chaining (SFC), i.e. controlling the network to interconnect the various VNFs of the service and directing the customers' traffic through the VNFs
- Service monitoring, i.e. collecting and aggregating metrics from VNFs and virtual networks
- Service starting/stopping and teardown

The demo aims at presenting the benefits of the multiple VNF deployment in order to form a specific network service, such as firewalling and content filtering (GTW side), TCP acceleration (GTW side), and caching (Terminal side to cache traffic from external networks).

6.6.3.1. Demo: Chaining of three customer functions

This demo scenario considers that NFV capabilities are present at both ends of the satellite segment: on the GTW side a fully operational Openstack cloud platform integrated with the OpenDaylight controller, while at the terminal side a SDN/NFV-enabled terminal with VNF hosting capabilities, as also described in the previous scenario. The demo considers that the Customer composes a Network Service (NS) for optimizing and securing the satellite link utilization. To achieve this Network Service, comprising three VNFs, Service Function Chaining (SFC) needs to be employed as shown in the following figure:

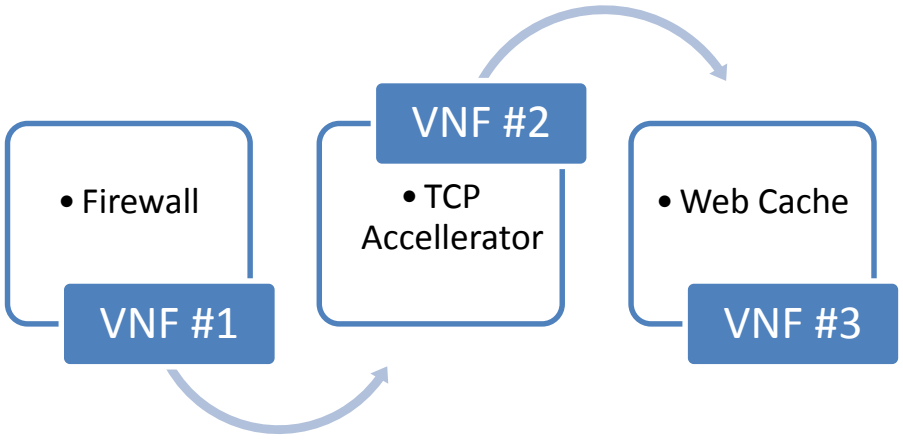

Figure 136. Service Function Chaining concept

Towards facilitating the SFC in a dynamic and transparent way for the end-users, the orchestrator instantiates the three VNFs, (i.e. VNF#1 and VNF#2 at the Openstack-based NFVI PoP and VNF#3 at the SDN/NFV-enabled terminal) and then applies appropriate SDN-based traffic steering rules at the OVSs in order to realise the

desired Service Function Chaining (SFC). In particular, the traffic is steered from the service provider through the VNF1 (virtualized firewall) VNF2, (virtualized TCP optimizer) and then is forwarded over the satellite in order to finally pass through VNF3 (virtualized web cache) at the terminal and then to reach the end-user.

It must be stated that all VNFs used in this demo are based on existing open-source software.

- The firewall VNF is based on the Linux iptables module with the Firewall builder GUI (http://www.fwbuilder.org/)
- The TCP optimiser is based on native Linux capabilities for TCP window adjustment controlled by user-defined scripts.
- The Web cache is based on the well-known squid platform (http://www.fwbuilder.org/)

Depending on the firewall rules, the requested service may be allowed to be delivered or not, then the TCP optimizer will seamlessly configure the optimized TCP settings for the satellite provision and then the Web Cache will cache the most frequently requested services, in order to minimize the satellite utilization. The storyline of this scenario is the following:

1. The orchestrator instantiates at the SDN/NFV-enabled satellite terminal and at the NFVI PoP the requested VNFs towards implementing the requested Network Service
2. The orchestrator applies appropriate SDN rules for the traffic steering so that user traffic traverses the VNFs as desired (SFC)
3. The end-user requests from the service provider a web page
4. The Service provider receives the requests and replies with the requested web page
5. The requested web page is filtered through the firewall and depending on the type of the service is the flow is allowed or rejected
6. The requested web page is passed through the traffic accelerator and is optimized for satellite transmission
7. At the reception side the web page is cached by the web cache and then is delivered at the user.
8. The user requests again the same web page content and this time is delivered directly from the cache, without requiring any satellite transmission with much shorter delay.

Initially the instantiation of the VNFs is performed.

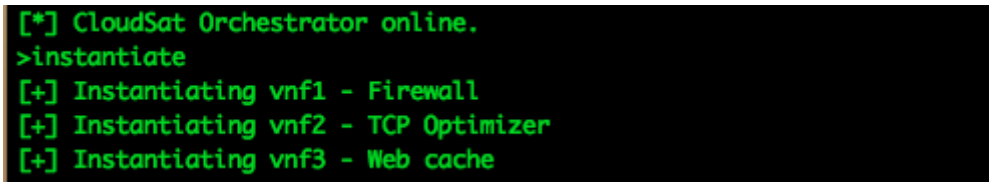

Figure 137. Instantiation of the three VNFs by the Orchestrator script

In terms of the response time of the instantiation process, two cases were benchmarked: starting the VNF VMs from either the PAUSED or the STOPPED state.

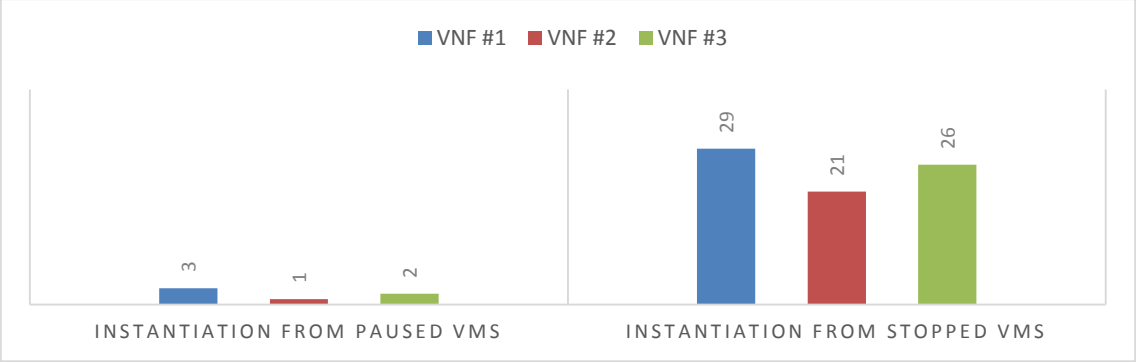

Figure 138. Instantiation Time needed for the three VNFs

The results are depicted in Figure 138, the following figure, showing that the instantiation of the VNFs from PAUSED status requires 2-5 seconds, while the instantiation of the VNFs from STOPPED status requires 23-29 seconds.

In the NFVI PoP, through the OpenStack Dashboard interface, the instantiation of the two GW-side VNFs is monitored and depicted on Figure 52. For each VNF and internal and external IP address is allocation, where the internal is used for in-cloud routing purposes, while the external is public for accessing the VNF outside from the cloud.

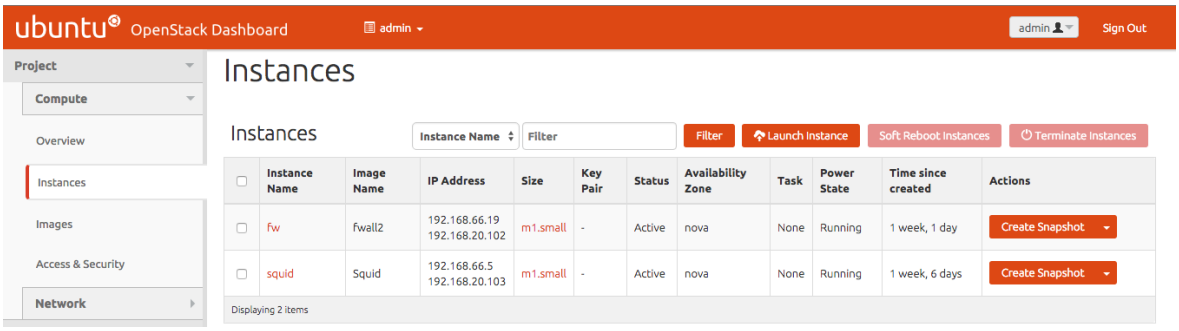

Figure 139. VNF information as seen via the Openstack Horizon dashboard

It should be noted that for both VNF VMs a standard m1.small flavour for instantiation was selected. Virtual hardware templates are called "flavours" in OpenStack, defining sizes for RAM, disk, number of cores, and so on. The default install provides five flavours:

```
$ nova flavor-list
+----+-----------+-----------+------+-----------+\+-------+-\+-------------+
| ID | Name      | Memory_MB | Disk | Ephemeral |/| VCPUs | /| extra_specs |
+----+-----------+-----------+------+-----------+\+-------+-\+-------------+
| 1  | m1.tiny   | 512       | 1    | 0         |/| 1     | /| {}          |
| 2  | m1.small  | 2048      | 10   | 20        |\| 1     | \| {}          |
| 3  | m1.medium | 4096      | 10   | 40        |/| 2     | /| {}          |
| 4  | m1.large  | 8192      | 10   | 80        |\| 4     | \| {}          |
| 5  | m1.xlarge | 16384     | 10   | 160       |/| 8     | /| {}          |
+----+-----------+-----------+------+-----------+\+-------+-\+-------------+
```

The following table lists the elements that can be set. Note in particular extra_specs, which can be used to define free-form characteristics, giving a lot of flexibility beyond just the size of RAM, CPU, and Disk.

ID	A unique numeric ID.
Name	A descriptive name, such as xx.size_name, is conventional but not required, though some third-party tools may rely on it.
Memory_MB	Virtual machine memory in megabytes.
Disk	Virtual root disk size in gigabytes. This is an ephemeral disk the base image is copied into. You don't use it when you boot from a persistent volume. The "0" size is a special case that uses the native base image size as the size of the ephemeral root volume.
Ephemeral	Specifies the size of a secondary ephemeral data disk. This is an empty, unformatted disk and exists only for the life of the instance.
Swap	Optional swap space allocation for the instance.
VCPUs	Number of virtual CPUs presented to the instance.
RXTX_Factor	Optional property that allows created servers to have a different bandwidth cap from that defined in the network they are attached to. This factor is multiplied by the rxtx_base property of the network. Default value is 1.0 (that is, the same as the attached network).
Is_Public	Boolean value that indicates whether the flavor is available to all users or private. Private flavors do not get the current tenant assigned to them. Defaults to True.
extra_specs	Additional optional restrictions on which compute nodes the flavor can run on. This is implemented as key-value pairs that must match against the corresponding key-value pairs on compute nodes. Can be used to implement things like special resources (such as flavors that can run only on compute nodes with GPU hardware).

The virtual networking configuration as it is depicted from the Openstack Dashboard, depicts the two instances of the two VNFs at the GTW-side NFVI PoP. The two VNFs are depicted on the same virtual local area network.

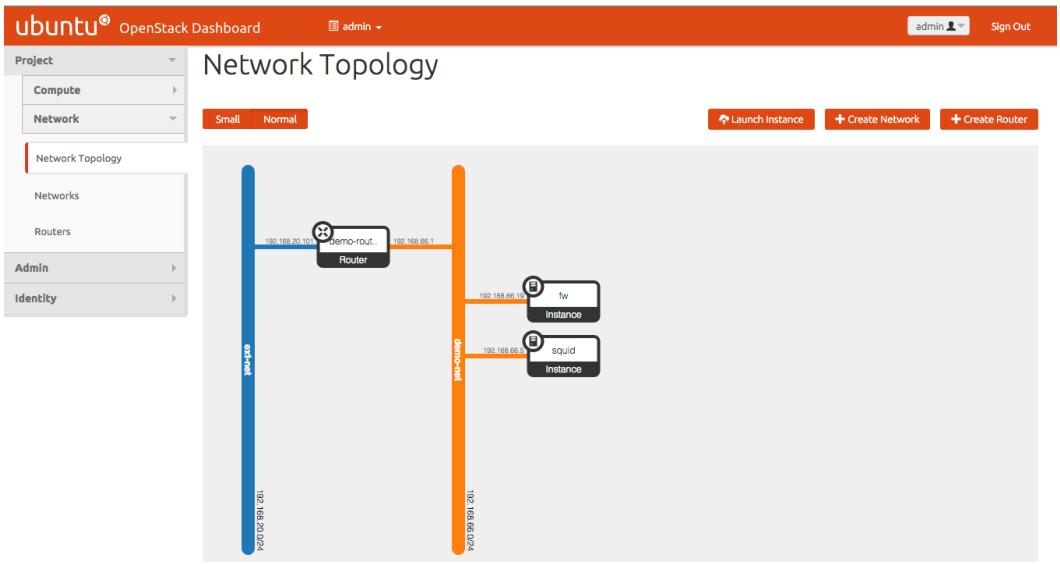

Figure 140. Network Topology of the two VNFs at the GTW-side NFVI-PoP

The orchestrator in order to support the SFC, applies appropriate SDN-based traffic steering commands at the NFVI PoP in order the traffic flow to be forwarded via the two VMs/VNFs as Figure 54 depicts.

```
cookie=0x0, duration=3052.825s, table=0, n_packets=1227, n_bytes=114417, in_port=18,dl_src=fa:16:3e:fa:5e:b3 actions=set_field:0x1->tun_id,load:0x1->NXM_NX_REG0[],goto_table:20
cookie=0x0, duration=624437.164s, table=0, n_packets=60561, n_bytes=5220389, in_port=5,dl_src=fa:16:3e:c0:a5:34 actions=set_field:0x1->tun_id,load:0x1->NXM_NX_REG0[],goto_table:20
cookie=0x0, duration=624437.172s, table=0, n_packets=107424, n_bytes=94947103, in_port=1,dl_src=fa:16:3e:fc:92:06 actions=set_field:0x1->tun_id,load:0x1->NXM_NX_REG0[],goto_table:20
cookie=0x0, duration=624437.164s, table=0, n_packets=118283, n_bytes=11230018, in_port=15,dl_src=fa:16:3e:35:f2:52 actions=set_field:0x1->tun_id,load:0x1->NXM_NX_REG0[],goto_table:20
cookie=0x0, duration=962.067s, table=90, n_packets=954, n_bytes=93492, priority=16399,ip,tun_id=0x1,dl_dst=fa:16:3e:35:f2:52,nw_src=192.168.20.26 actions=mod_nw_dst:192.168.66.5,outp
cookie=0x0, duration=35.49s, table=90, n_packets=36, n_bytes=3528, priority=16399,ip,tun_id=0x1,dl_dst=fa:16:3e:fa:5e:b3,nw_src=192.168.20.26 actions=mod_nw_dst:192.168.21.27,output:
cookie=0x0, duration=624437.164s, table=110, n_packets=730, n_bytes=54570, priority=16383,reg0=0x1,tun_id=0x1,dl_dst=01:00:00:00:00:00/01:00:00:00:00:00 actions=output:1,output:2,out
4,output:5,output:15,output:18
```

Figure 141. SDN-based Traffic Steering at VIM

One the VNFs have been appropriately instantiated, each VNF can be individually configured through its EMS (Element Management System). For example the Firewall VNF as Figure 55 depicts can be set up to allow or reject specific types of traffic.

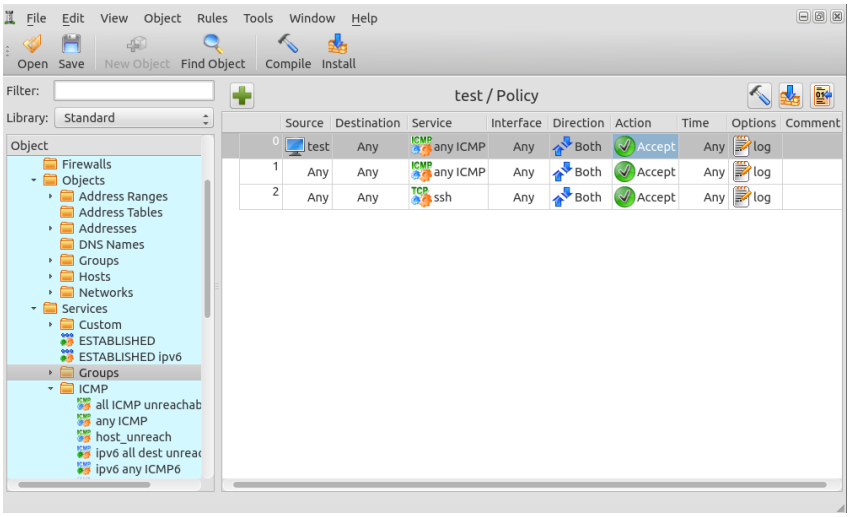

Figure 142. Configuration of Firewall Rules at VNF #1

Initially the user requests to receive a web page from the Service provider, where the user's request arrives successfully and the web page is to be delivered.

As the web page is delivered for first time, it passes through the firewall, through which is allowed and then its steering continues and the content flow comes through the VNF #2 (i.e. the TCP optimizer) and as Figure 56 depicts, TCP optimization can be applied.

```
root@squidsquid:/home/localadmin# sysctl -p
net.ipv4.ip_forward = 1
net.core.wmem_max = 12582912
net.core.rmem_max = 12582912
net.ipv4.tcp_rmem = 10240 87380 12582912
net.ipv4.tcp_wmem = 10240 87380 12582912
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
net.ipv4.tcp_no_metrics_save = 1
net.core.netdev_max_backlog = 5000
root@squidsquid:/home/localadmin#
```

Figure 143. TCP optimization of the requested content by VNF #2

Then the video traffic is forwarded over the satellite link and at the SDN/NFV-enabled terminal the flow is steered to be passed through the VNF #3 (i.e. the Web Cache). For new web pages, as the requested one, the web page is cached as figure 57 depicts.

```
root@transcodervnf:~# tail -f /var/log/squid3/access.log
1442582733.287   1129 192.168.21.27 TCP_REFRESH_UNMODIFIED/200 577 GET http://192.168.20.26/index.html - DIRECT/192.168.20.26 text/html
1442582735.693   1146 192.168.21.27 TCP_REFRESH_UNMODIFIED/200 576 GET http://192.168.20.26/test.html - DIRECT/192.168.20.26 text/html
```

Figure 144. Web Caching of the requested page at the VNF #3

Finally, the requested page, since it has been passed through the VNF #1, VNF #2 and VNF #3, is delivered at the end-user through the satellite link, experiencing 577 msec

overall latency. Figure 58 depicts the web page delivery together with the its timing showing the relevant delay.

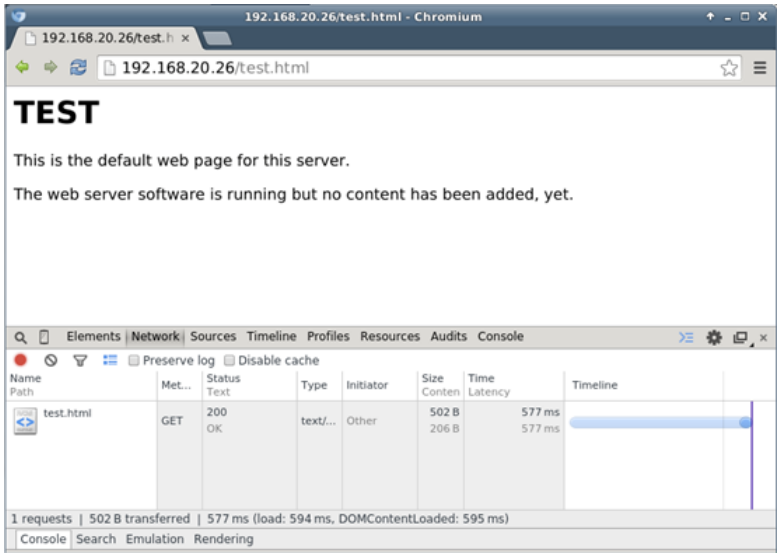

Figure 145. Web page delivery over the satellite with 577msec latency

Once the user decides to request for a second time the same web page, as the request is passed through the VNF #3, the web cache identifies that the specific VNF has been cached and it is not necessary to be requested from the original server. Figure 59 depicts the Web Cache matching at VNF #3 for the requested web page.

```
1442582738.085      0 192.168.21.27 TCP_MEM_HIT/200 584 GET http://192.168.20.26/test.html - NONE/- text/html
```

Figure 146. Web page delivery over the satellite with 577msec latency

Finally, the web cache servers the requested web page back to the user and the page is delivered directly without experiencing the satellite delay, as figure 60 depicts, but only ~4 msec latency.

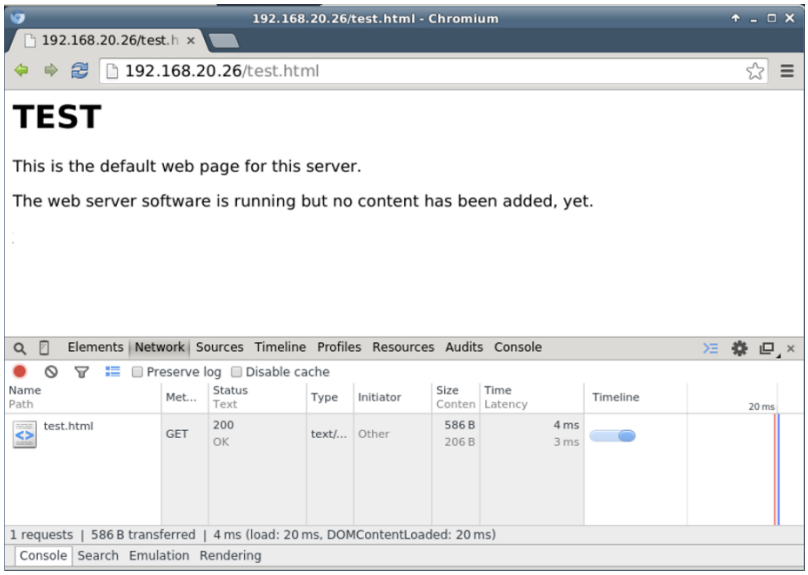

Figure 147. Web page delivery by the VNF #3 within ~4msec latency

This specific Network Service, as a set of three chained VNFs, achieves optimisation of the usage of the satellite resources by i) filtering content before the satellite segment,

ii) optimising TCP connections and iii) local caching at the terminal. These three VNFs are fully customisable and configurable by the customer and their parameters can be fine-tuned as desired.

The flexibility and agility of SDN/NFV allows arbitrary sets of VNFs to be composed and offered as Network Services to customers, serving a wide range of customer needs.

7. Analysis of Cost Effectiveness and Economic Gains vs Constraints

CloudSat studies solutions to alleviate certain limitations in the SatCom domain, utilizing the cloud networking model. In addition to the technical evaluation which was laid out in the previous chapters, this innovative potential must be further investigated and validated under a well-defined market, considering a detailed business and financial analysis framework, which will identify the pros and cons of the proposed business concept.

In this context, the objective of this chapter is to provide a high-level Financial and Cost-effectiveness analysis (CEA) of the CloudSat selected scenarios/architectures, while taking also into consideration the corresponding business and market dynamics.

Concerning the prevailing business and market environment of the proposed architectures, it is considered essential to first define and validate the existing business and market aspects which will highly influence the cost effectiveness of the proposed scenarios. A close and detailed look of the current market dynamics is undertaken by assessing the market (in terms of technological trends, key market players, recent actions, market barriers and drivers), aiming to model the most appropriate and suitable business model for CloudSat scenarios.

Concerning the financial aspects, the previously described business and market framework study will provide the basis of an efficient cost effectiveness analysis. Inevitably, the implementation and exploitation of the proposed architectures is limited by budgetary constraints. Businesses need to make choices among different alternatives in order to achieve optimized results given the limited financial resources available.

7.1. Market and Business Analysis

This section contains the market and business analysis of the CloudSat architecture considering the overall innovative characteristics of the proposed cloud networking model, without focusing exclusively on each one of the three selected scenarios presented in the previous Chapters. The input to the market and business analysis is mainly the functional capabilities that can benefit from CloudSat network model adoption, the importance of achieved added-value, as well as market and business data which define the framework of the analysis and the blueprints of the business modeling procedure. The outcome of this market and business analysis section is of relative importance to the outcomes of the rest sections, which consider technical details and gains, estimate of initial Investment cost, estimate of operating/running cost, the cost of deployment, and the market-business impact of CloudSat environment on the three selected case scenarios.

Considering the current market and business environment, nowadays, satellites have become part of a rich market environment characterized by the high degree of diversity of its market sectors of application such as communications, entertainment, business and finance, navigation, weather, climate and environmental monitoring, safety, space science and space exploration and many more [BOO].

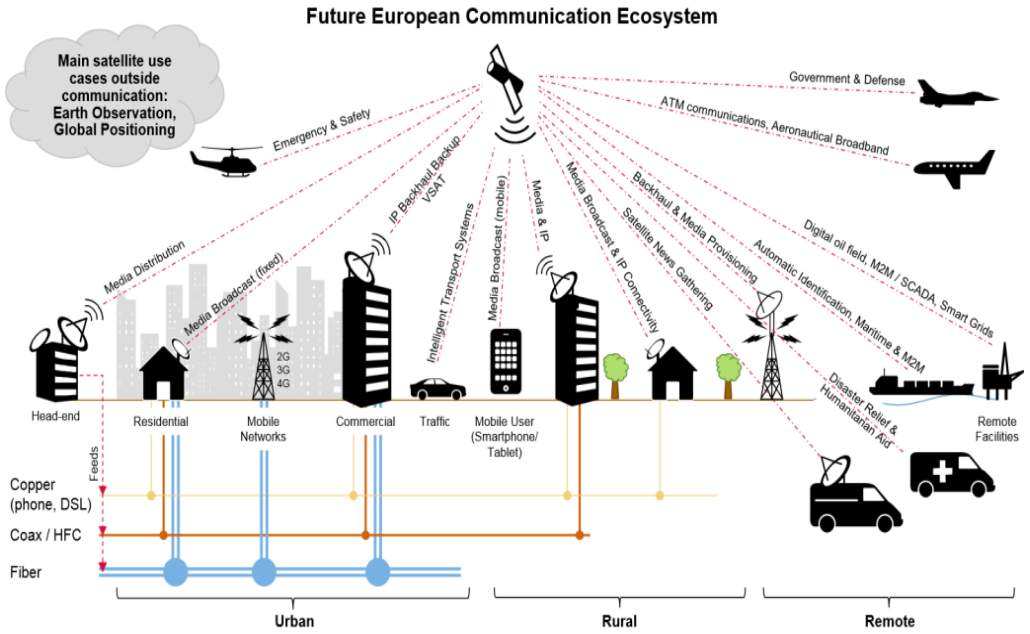

Figure 148. Future European Communication Ecosystem [BOO]

However, this diversity of the satellite communication applicability at a variety of sectors and fields, it is mainly limited by the low degree of Satellite network integration with other existing networks and communication platforms. Currently the satellite role is considered as autonomous and standalone, without the agility to integrate with other existing networks or communication technologies in the framework of the 5G ecosystem.

CloudSat market opportunity is inspired by this need, towards providing to the SatCom service providers the agility to deploy fast and on-demand the appropriate satellite links as part of an overall network ecosystem, providing novel business opportunities from the collaboration and the joint-ventures. The concepts of Cloud Networking and Virtualization technologies for application in satellite platforms are considered as the most promising techniques towards providing the missing characteristics to the SatCom providers in terms of functional aspects, such as networks federation and coupling, isolation and services, reconfigurability and programmability, mobility, resource elasticity, availability and resiliency, security and privacy, accounting and billing, performance and exploitation of satellite-specific capabilities. Integration aspects are also supported, such as integration and coexistence with terrestrial networks, improving the satcom market attractiveness and stakeholder interests.

The main feature of the CloudSat Network model is the virtualization and abstraction of network resources and their agile provision to the end-user as-a-Service, in a cloud-like manner, featuring dynamic resource pooling and elasticity. For this purpose, as also described in the previous Chapters, the CloudSat Network model relies on the combined application of specific component technologies and not on a single one.

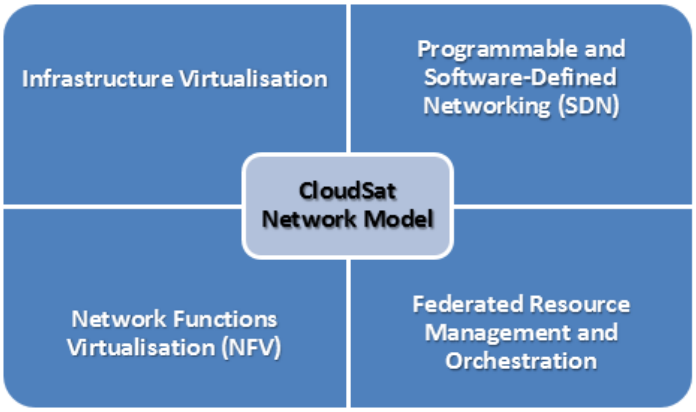

Figure 149. CloudSat Network model

The CloudSat Network model by taking advantage of this combined application of technologies, especially focusing on SDN and NFV, creates new market opportunities for the integration of satellite components with terrestrial future networks, in a continuously developing market.

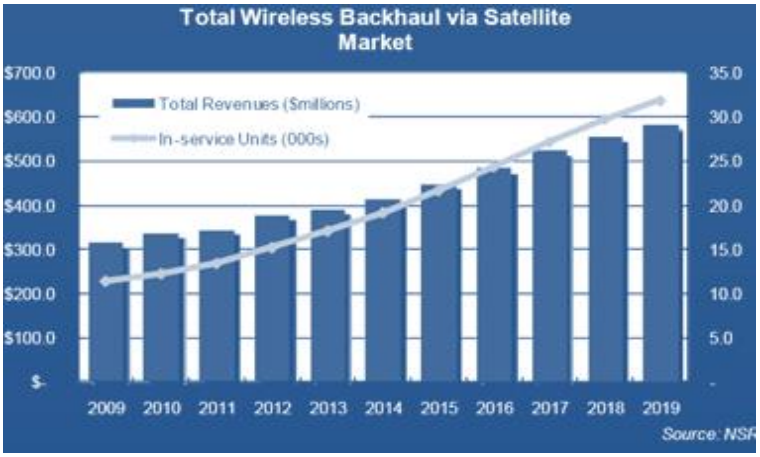

Figure 150. Future European Communication Ecosystem [NSR]

Over the last years, as per Figure 3, a steady increase in revenues and in service units is documented in the use of wireless backhaul via the satellite market, a trend that CloudSat takes into account in its business and market planning towards providing an agile collaboration and integration framework through federation and orchestration techniques with other network domains.

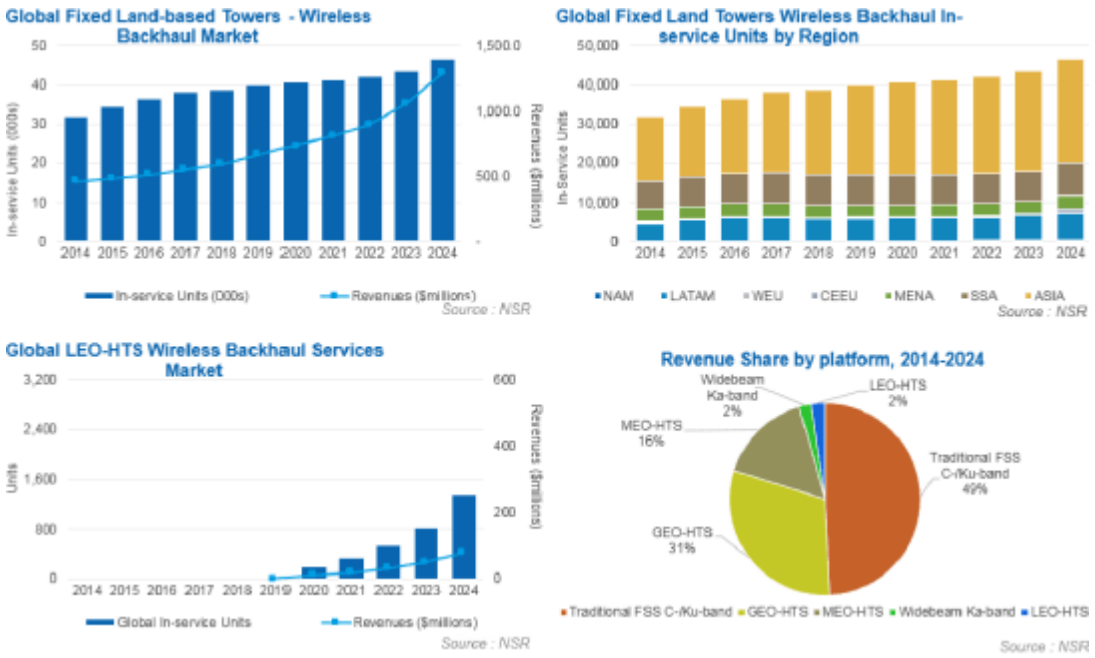

Figure 151. Future European Communication Ecosystem [NSR]

As Figure 4 depicts, Northern Sky Research (NSR) projects that revenues from wireless backhaul over satellite will climb from $1.7 billion in 2014 to $5.3 billion by 2024, showing a steady increase in the prominence of High Throughput Satellite (HTS) systems on this particular market. NSR notes that traditional Fixed Satellite Service (FSS) capacity in C band and Ku band has been the predominant solution for backhaul and trunking in land, with growing business in maritime and aviation. However, the firm notes “a clear migration” by fixed land towers backhaul and trunking markets toward Geostationary (GEO) HTS. Nascent Non-Geostationary (NGSO) HTS players, such as O3b Networks, are making inroads into backhaul, trunking and mobility markets as well [NSR].

Therefore the market trend towards offering blended services over satellite links is growing, providing the positive ecosystem for further agility by a virtualized-based satellite hub featuring SDN and NFV capabilities.

In the next section, the CloudSat market and business framework will be defined following the Business Model Canvas methodology of Osterwalder and Pigneur [BMC].

7.1.1. Methodology-Framework

The Business Model Canvas (BMC) methodology of analysis is used in order to define and present the Business and Market CloudSat framework. Detailed information concerning the BMC methodology is provided in Appendix II.

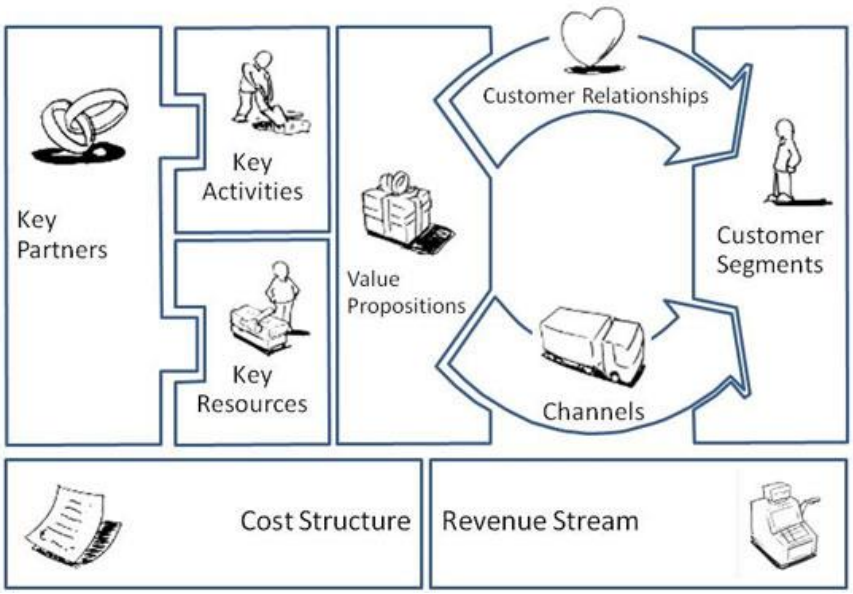

Figure 152. Business Model Canvas (PictorialView as per Osterwalder and Pigneur's) [BMC]

As per the above figure, the BMC methodology, offers a visual, one-page canvas providing a way of composing a business model with nine building blocks [BMC]:

- Customer Segments
- Value Propositions
- Channels
- Customer Relationships
- Revenue Streams
- Key Resources
- Key Activities
- Key Partnerships
- Cost Structure

This BMC canvas and its nine building blocks are used for the analysis of the CloudSat business model.

In parallel, the Business Model Canvas documents and analyzes the interactions of the nine building blocks with the rest internal and external forces of the business. Through this interaction process specific business modeling building blocks interact with the internal environment of the business such as employees, values, vision, culture, strategy, operations, internal stakeholders etc.

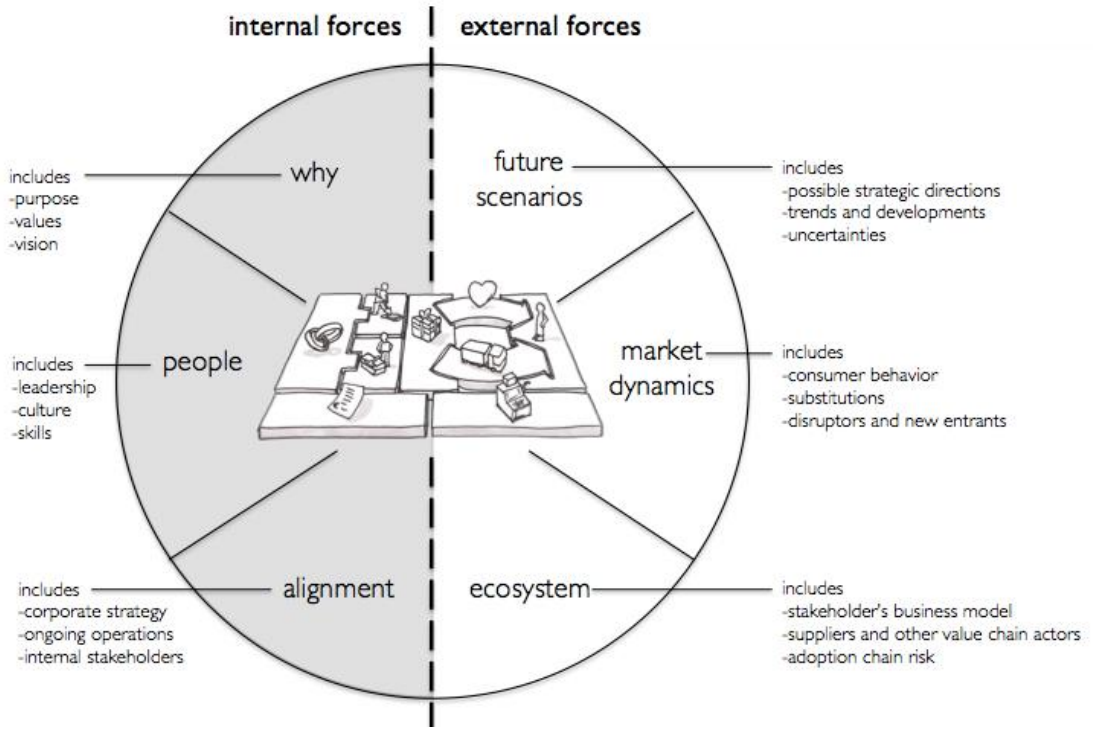

Figure 153. Business Model Canvas interaction with Internal and External forces [BMC]

In parallel, the rest business modeling building blocks interact with the external environment and its forces such as consumers, consumers' behavior, new entrants, suppliers, trends, uncertainties, risk, stakeholders, society, legislation, governments etc. The figure below depicts this type of interactions [BMC].

7.1.1.1. CloudSat Customer Segments

As per [BMC], the Customer Segments block is intended to capture the different groups of customers, which are the target for the organization described in the BMC. Segmentation is key, since it allows the identification and grouping of those potential customers that share similarities and can be served by the same product/service. It is very important that an organization identifies its potential segments and selects which one should be finally part of its value chain and able to be served at the end.

Figure 154. Customer Segments

In CloudSat the potential customer segments have been identified into the following entities that are also participating in CloudSat value chain: The Virtual Network Operators (VNOs) and the End-Users. These two roles were defined in Chapter 2 and briefly overviewed here for the sake of completeness.

The VNOs, also known as Tenants, are the “operators” of the network slice, therefore they are the front line customers of the CloudSat service. VNOs gain specific management, control and monitoring rights on the provisioned slices, having a unified view of the provisioned slice, regardless of the multiple domains on which it may be built. VNOs may exploit the network slice for own internal use or may also in turn act as Service Providers themselves and exploit the slice for offering a service to their customers. Several market entities which can act as Tenants/Customers (VNOs) are considered:

- Enterprises,
- Content Providers,
- Service Providers
- Etc.

The End-Users (EUs) are the customers of Tenants (VNOs) and receive the provisioned service over the slice. The existence of the slice is totally transparent to the EUs, who interact only with the offered service. The technical and business interactions between VNOs and EUs strictly depend on the service offered and are out of the scope of this study. Market entities than can act as End -Users (EUs) are considered

- Individuals,
- SMEs,
- Public sector,
- Governmental agents/services,
- Etc.

Particular attention should be drawn to multi-sided platforms, such as CloudSat, since target customers or participating actors are rather different but, at the same time, complementary. The business model, as well as the business value chain, cannot be in place if one of them is missing.

7.1.1.2. CloudSat Value Propositions

The Value Proposition building block is probably the most important building block since it captures which products and/or services are offered to a specific customer/business segment. As a fact it should be stressed that a key concept of marketing is that customers are also interested in the benefits and accompanying value of products/services and not solely in the product as a commodity item [MARK]. So, clearly identifying the proposed value proposition is essential to create differentiation and eventually capture customer interest and market share.

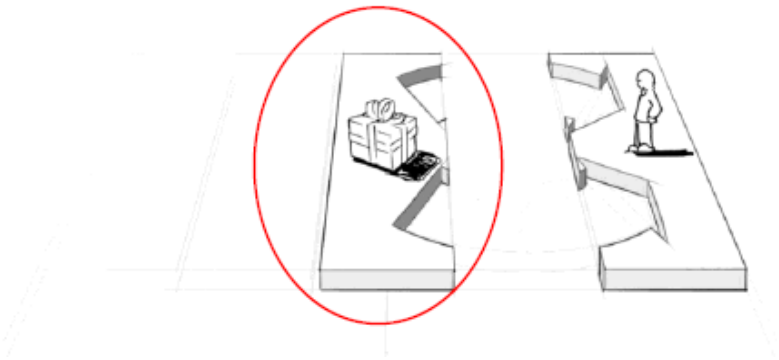

Figure 155. Value Proposition

The value of a product/service can be quantitative (e.g. price, speed of service) or qualitative (e.g. design, customer experience). In addition, it is important to be aware whether the service/product is addressing a new set of needs that customers previously didn't perceive, or whether it is addressing existing needs by offering better performances or a higher level of customization. Other factors to take into account when specifying the value proposition are, for examples, price, design, peace of mind (i.e. managed service), convenience, usability, and cost reduction [BMC].

Overall, it is important to understand whether the service/product is addressing a new set of needs that customers previously didn't perceive, or whether it is addressing existing needs by offering better performances or a higher level of customization. The newly formed CloudSat Network model creates a new Marketing Mix (4Ps) for the prevailing SAT business and market environment:

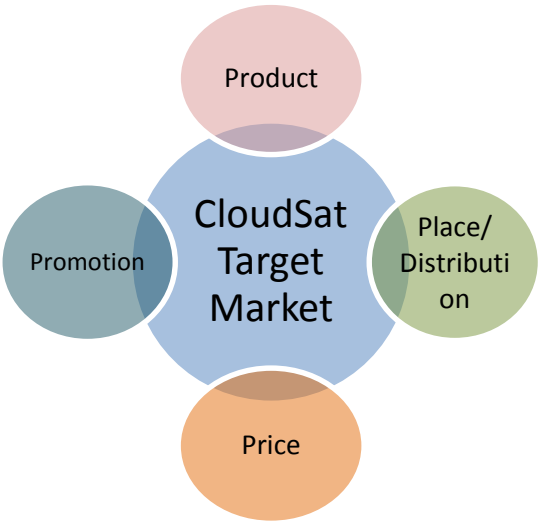

Figure 156. CloudSat Customer Segments

- ***Product:*** A new intelligent service of virtualization and abstraction of network resources and their provision to the end-user as-a-Service
- ***Place/Distribution:*** CloudSat Network model will follow the distribution and penetration pattern of the current network infrastructure models while enhancing its presence in even more market sectors (greater and more intense market penetration)

- ***Price:*** The Pricing models being adopted by the market exploitation of CloudSat Network model should contribute towards lowering prices, introducing new billing models and balancing prices and demand
- ***Promotion:*** All involved key players (InPs, VNSPs, VNFPs, Tenants/VNOs, Users) are providing and experiencing new services that need to be promoted and advertised for their characteristics, advantages and added-on value through new channels of communications, addressing a greater market share than the previously niche SAT market.

By the introduction and adaptation of the CloudSat new marketing mix (4Ps strategy), the SAT market is further enriched with new values bearing both quantitative and qualitative characteristics. In specific, added-on values by the new CloudSat marketing mix and CloudSat network model can be separated into two main categories: (a) the value propositions of Satcom providers and corresponding businesses and (b) the value propositions of Tenants / Customers. Following this split, the following lists briefly summarize the business benefits of the two categories.

- Value Propositions of Satcom Providers and Businesses
  - Enhanced exploitation of available networks resources leading to novel Value Added Services (VAS)
  - New market opportunities for Satcom businesses
  - Introducing a variety of new billing models
  - Easiness of deployment of new network services
  - Minimal delay and cost for upgrade and replacement of networking equipment due to softwarization
  - Cost efficient operations (via NFV)
- Value Propositions of Tenants / Customers
  - Improved QoS of scalable resources
  - New billing models (usage-based)
  - Enjoying NFV capabilities (shifting to operators' cloud infrastructure previously costly self-maintained (customized) hardware services)
  - Reducing cost of service (reducing CAPEX and OPEX while also shifting specific CAPEX costs to OPEX by using the cloud network services)
  - Fast responsiveness to changing demands, high adaptability and flexible customization.

Current market conditions and trends are in favor for the development of such value propositions as proposed by the CloudSat network model. For example, backhauling through satellite links is one of the most compelling scenarios for the integration of the satellite component in a terrestrial 4G/5G infrastructure. While mobile satellite backhauling solutions are a reality today, the high cost of satellite bandwidth and the technical complexities concerning the integration of satellite links into the mobile networks have largely restricted the use of satellite backhauling to remote or hard to

reach locations and low traffic settings (e.g., backhauling of a base station site mainly used for the delivery of voice communications and very low data rate services).

As another example, the real-time services market has experienced continuous growth over the last years due to its unique characteristics such as cost effectiveness, quality of service and rising demand for product and service differentiation by end consumers. This market penetration can be further enhanced by reinforcing the provision of real-time services through the use of satellite networks. On top of that, CloudSat solutions are expected to further reduce the adoption barriers of the satellite component by easing the integration and use of satellite backhauling services within 4G/5G mobile networks, facilitating also the provision of broadband service via satellite.

7.1.1.3. CloudSat Channels

Communication channels capture the company's interface with customers. This includes communication (e.g. advertisement), distribution and sales channels. All these components play different roles at different stages. BMC methodology calls these stages 'Channel Phases' and they identify five of them:

- *Awareness*. This stage involves raising customer awareness of a company's products/services. This might also include making customers aware of new needs. In the case of CloudSat, the awareness is related to the 5G new era and the agility that cloud networking offers in the collaboration and integration of SatCom industry with other terrestrial networking domains.
- *Evaluation*. This involves channel activities and resources that needs to help the potential customer to gain enough information to evaluate the company's offering and compare it against competing products. In the CloudSat case the advantages of the SDN/NFV techniques such as the resource elasticity and deployment agility, both in qualitative and quantitative terms can highlight the superior performance of the CloudSat networking model over competing products, which may be substitute products/services (e.g. SVNO without virtualization capabilities, Legacy Backhauling, etc).
- *Purchase*. This phase means dealing with the structure put in place to allow potential customers to purchase the specific products/services. In the case of CloudSat model, a user-friendly front-end GUI would enable the interaction between customers and the federated CloudSat architecture, visualizing service parameters, such as network service topology and service monitoring metrics, billing information and SLA status, facilitating operations like service advertisement (i.e. presentation of available NS and VNF catalogs), service composition, service deployment, monitoring, management and teardown. Alternatively, an API will be foreseen for expanding Orchestrator/Federator applicability to custom user applications, especially developed and tailored to automate service deployment and management processes, according to the customer's special needs.

- ***Delivery***. This phase is about delivering the value proposition to the customers via the considered products or services. In the case of CloudSat networking model the Orchestrator/Federator is deemed essential so as to enable coordinated management of inter-domain services (e.g. services spanning across the satellite and terrestrial segments), yet without violating the administrative independence of the involved domains, which may belong to different business entities/operators (i.e. multi-domain/multi-operator ecosystem), providing added value to the involved customers in the federation by this alliance.
- ***After sales***. This is also a very important phase especially since post sales experiences have a relevant influence in re-purchasing decisions. This phase is emphasized in CloudSat networking model, since network and satellite operators will no longer need to purchase dedicated hardware devices in order to build a service chain, but this action will be feasible in an agile manner through the provision of virtualized functions and resources in an elastic and dynamic way. Due to the resource elasticity, there will be no need for the operators to overprovision their data centers in order to achieve the wanted QoS level, reducing both CAPEX and OPEX costs and at the same time advancing the system performance and agility. If an application running on a VM required more bandwidth, for example, the administrator could move the VM to another physical server or provision another virtual machine on the original server to take part of the load. Having this flexibility will allow an IT department to respond in a more agile manner to changing business goals and network service demands, creating excellent after-sales conditions with minimum maintenance and operational costs.

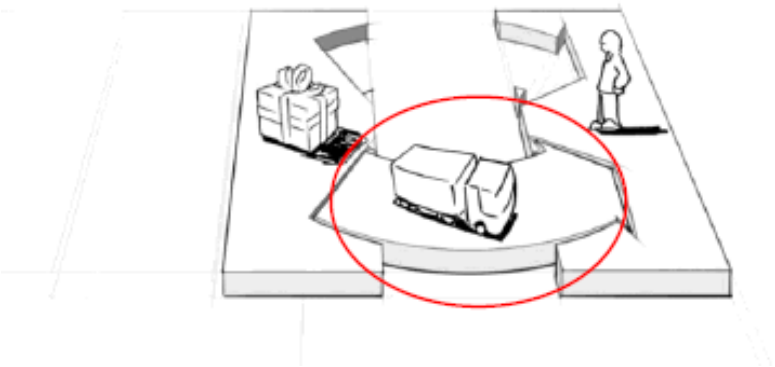

Figure 157. Channels

All these steps of CloudSat Channels are closely related with the key steps involved in customer's buying processes, because the CloudSat-enabled operator can utilise its channels or partners' channels or even a combination of both for promoting and communicating the CloudSat federated services. Thus, most of these activities described above can be also well combined and described from another perspective by the supply chain of CloudSat, showing how CloudSat channels are built over the respective value chain.

Provision of service should be well determined by a Supply chain strategy which mainly determines when product/service should be fabricated, delivered to distribution centers and made available in the retail channel. It also determines the way that the service will be delivered and through which channels, defining also its awareness and evaluation.

Towards this, CloudSat model replies to both pull and push strategies, considering JiT optimization and lean manufacturing approaches in terms of resources utilization. Morespecifically:

- Under ***pull strategies***, actual customer demand drives the process (in our case the customer is the tenant). Just-in-time (JiT) production relies on actual demand triggering the release of work into the system, and "pulling" work through the system to fill the demand order. This JiT approach is fully justified by the agility of the SDN/NFV schemes in cloudsat networking model, which allow the dynamic deployment of requested configuration variants when is needed without overprovision of resources.
- Under ***push strategies***, the process is driven by the amount of raw material available for production (i.e. Available Network Resources). Push strategy is not particularly responsive to changing customer demands, for it relies on forecasting future demand and scheduling the release of work into the system to meet expected demand. CloudSat networking model can satisfy this strategy by offering predefined product/services to the potential customers, e.g. a network slice of satellite operator with specific virtualization capabilities in terms of VNFs.

Summarizing both strategies, it is deduced that CloudSat channels follow an agile pull and push strategy which is further reinforced by the KANBAN pull strategy characteristics. KANBAN is a lean principle introduced in Japan in the 1950s [LPMS]. The KANBAN benefits and added on value can be summarized as following in the case of the CloudSat networking model:

- Just-in-time production is agile to changing customer demands (e.g. dynamic scale in/out of slice)
- Produces the right products at the right time and in the right amounts (e.g. resource elasticity, network slicing, VNF deployment)
- Allows rapid and low-cost changeovers to adapt production capacity (i.e. resource elasticity, network slicing) to the requested products (i.e. Network Services/VNFs).
- Production capacity efficiently adjusted through effective network slicing (resource elasticity and agility)
- Product/Service efficient offering through network slicing and Network Service/VNFs

CloudSat NFV/SDN characteristics provide the required automation towards transforming the CloudSat value chain five-actor model (InPs, VNFPs, VNSPs, Customers/Tenants, End-Users) to an agile PULL and PUSH model. In specific, CloudSat pull and push strategy objectives under the KANBAN framework are:

- Just in time delivery (high responsiveness)
- Controlled Work In Progress (WIP)
- Lean Manufacturing (optimal utilization)

As per Figure below, CloudSAT through the automation offered by the Federated Manager achieves an agile KANBAN pull strategy. The End-users receive services by the tenants, which in turn in order to provide these services request and receive a network slice by the Virtual Network Service Providers (VNSPs) in collaboration with the VNF providers, which enrich the allocated slice with VNFs and NSs. Up to this point the value chain of the CloudSat networking model is Pull based, therefore demand-driven. The VNSPs and the VNFPs are located on the edge of the value chain between the Pull and Push strategies, because both of them have a hybrid role in the value chain serving both Pull and Push requests. More specifically, both VNSPs and VNFPs can either provide their services (i.e. a network slice or VNFs) upon a request complementing the chain that was described previously (i.e. they will in turn request from the Infrastructure Provider the necessary resources needed for the network slicing) or alternatively in collaboration with the Infrastructure providers they can create a bundle of products ready to be sold based on market trends and forecasts.

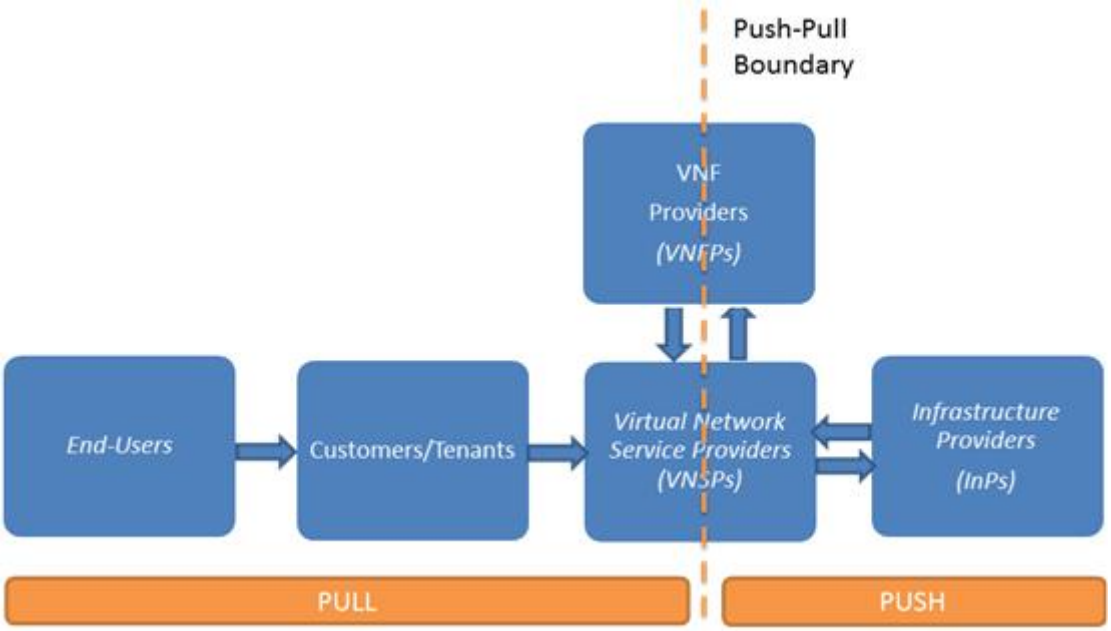

Figure 158. CloudSat Pull and Push Supply Chain

Therefore, the CloudSat value chain, creates value on both ways: Either by satisfying demands/requests on network slicing or by creating “off the shelf” products based on the market trends and customer demands. Either in both ways, the NFV/SDN techniques with the resources elasticity and the deployment agility contributes towards the lean manufacturing, which achieves minimization of the resource consumption for the provision of the requested/offered services at a specific level of performance. These properties constitute the KANBAN implementation, without the use of traditional cards, but with a modern and digital way [LPMS].

In specific:

- KANBAN uses the rate of demand to control the rate of production, passing demand from the end user through the supply chain
- KANBAN pull strategy tightly controls work-in progress (WIP) between each pair of actors
- The WIP is transferred to the next node of the chain, only when the rate of demand has been satisfied at the previous node.
- KANBAN model orchestrates the WIP at each node across the chain.

Similarly, as per CloudSat Pull and Push Supply chain model, the applicability of the KANBAN model towards achieving the lean manufacturing (i.e. the minimization of the use of resources per network slice in conjunction to the optimization of the allocation of the used ones within each network slice) is performed by the Federated Manager. The Federated Manager orchestrates work-in progress activities (WIP) by applying the KANBAN strategy characteristics:

- Customer-tailored products (Network Slicing + VNF Provision)
- Just in Time Production (Network Slicing + VNF Deployment)
- Customization and fine tuning (VNF Configuration per slice)
- Stabilize and rationalize the overall value chain process (Multi-tenancy)

Thus although from a technological perspective, the Federated Manager (or the Orchestrator depending on the use-case) perform the management of the VNFs and the traffic steering, from a business perspective these entities perform the value chain management towards achieving a lean manufacturing supply chain with optimized and agile resource allocation, implementing in a modern way the KANBAN model.

As per [BMC], these supply strategies, especially on the threshold line of Pull and Push strategies (i.e. on the VNFPs and VNSPs) create novel channels of communication, which are also called 'Channel Phases', where for the case of the CloudSat framework consists of the following five:

- Awareness. This stage involves raising customer awareness of the available VNFs that can be used by the VNSPs to compose the customer requests. This might also include making customers aware of new needs through the variety of the available VNFs.
- Evaluation. This involves channel activities and resources that needs to help the potential customer to gain enough information to evaluate the available VNF offering and compare it against competing hardware-based products, highlighting competing and substitute products/services.
- Purchase. This phase means dealing with the structure put in place to allow potential customers to purchase the specific VNFs and NSs products/services, considering licensing and pricing models.
- Delivery. This phase is about delivering the VNF and NS value proposition to the customers. In the case of VNF and NSs, the service agility and elasticity features deliver the value proposition to the customers.

- **Post-Sales Support**. This is also a very important phase especially since post sales experiences have a relevant influence in re-purchasing decisions. The VNSPs should provide customer support and tutorials of service function chaining (SFC).

7.1.1.4. CloudSat Customer Relationships

Linked to the Channels building block, the Customer Relationship block illustrates which type of relationship a company would like to establish with each customer segment. Different types of relationships can be envisaged from self-service, communities or even co-creation.

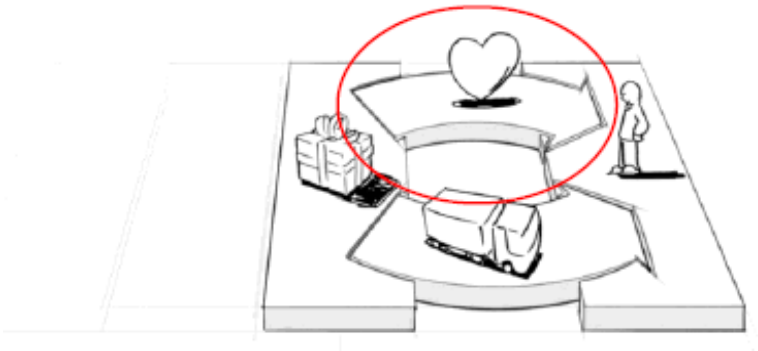

Figure 159. Customer Relationship (CR)

It is also important to stress those different types of customer relationships could be applied to different customer segments or whether the organization is focusing on customer acquisition, retention or service/product upselling.

As already stated, CloudSat further enhances the typical network model of four actors by following as part of its value chain a five-actor model (a double source of actors acting as providers i.e. the VNFPs and the VNSPs). All five actors actively participate through their operations and services in the formation of the CloudSat value chain in which the different types of interaction, interoperability and interdependence among them are depicted.

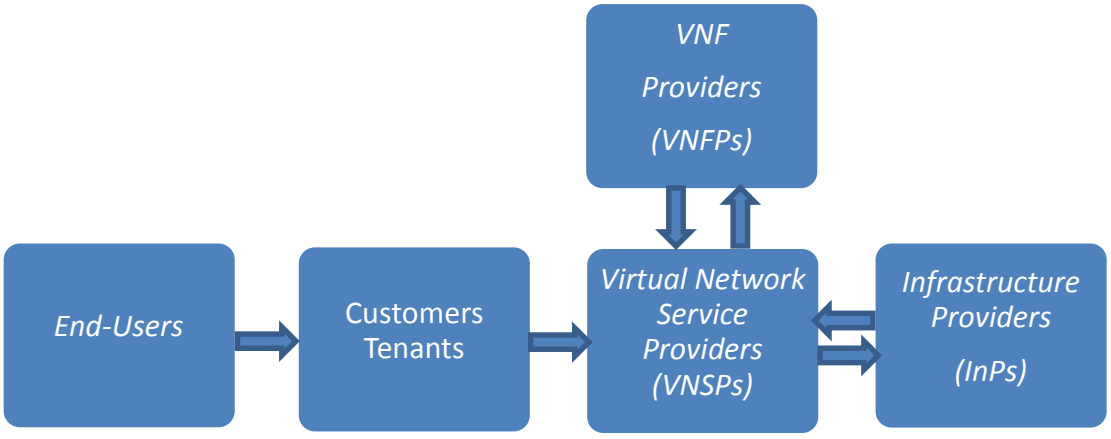

Figure 160. CloudSat value chain – Interaction, interoperability and interdependence of actors

In specific, the interactions within the CloudSat value chain are expected to be facilitated by a user-friendly brokerage interface, with which enterprise users interact for requesting cloud network services and for negotiating SLAs (provided by VNFPs and VNSPs utilizing InPs networks but also by Customers requesting specific service levels). After the establishment of the service, the customers/tenants are offered a management and monitoring front-end for the configuration, maintenance and fine-tuning of the assigned virtual slice.

In terms of limitations among actors interaction, it should be clarified that the customers cannot normally interact directly with the VNFPs, because the deployment and instantiation of the VNFs needs a virtualization capable network, which is offered by the VNSPs. Furthermore, the VNSPs will normally need to certify that the VNFs are suitable to be executed on the virtualized infrastructure. The end-users receive the service via the tenants, without being aware of the network slicing that has been performed and the virtualization of the network services. Respectively, the infrastructure provider is not limited to participate as a node to only one value chain. Simultaneously can be a node of multiple value chains, which are formed by competitive VNSPs and VNFPs.

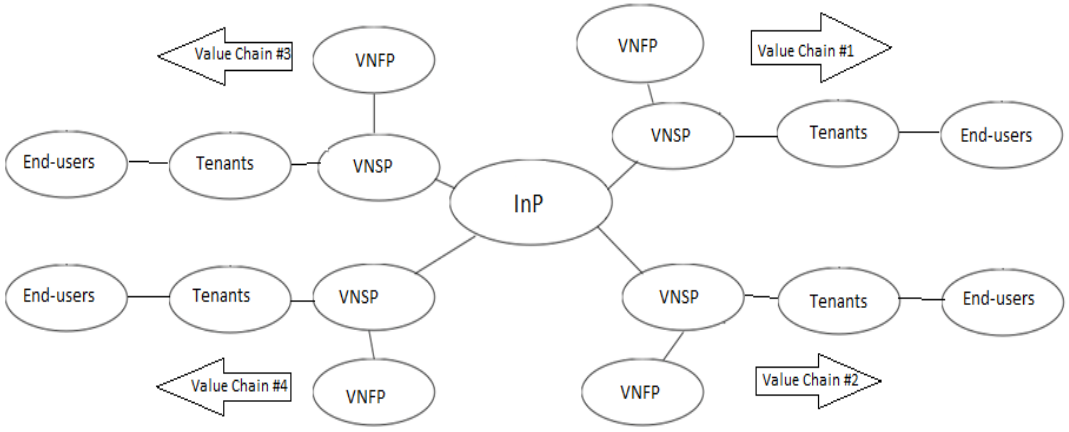

Figure 161. Multi-role of InP within multiple value chains

Therefore, in terms of limitations, it should be clarified that the Infrastructure provider is not limited to participate exclusively to one value chain, but on the contrary it may be the start of multiple chains.

It should be also noted that the VNFPs can -similarly to InPs- provide with VNFs multiple VNSPs, without being obliged to be limited to an exclusive collaboration within one value chain and one VNSP. However, this multiple role of the VNFP is not depicted in the figure above due to depiction restriction.

7.1.1.5. CloudSat Revenue Streams

This building block captures various sources of revenues that are envisaged by the organization with regard to the specific products/services that are offered to the customers.

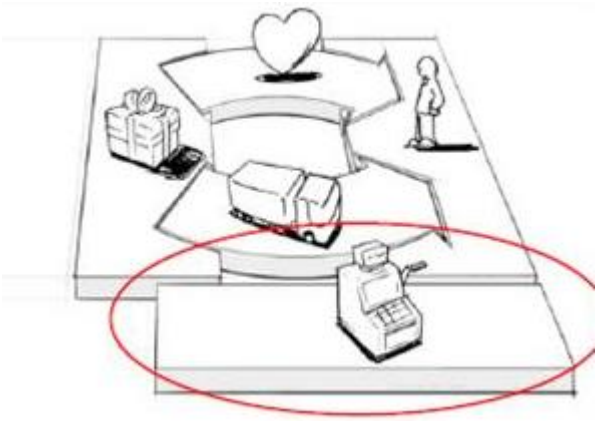

Figure 162. Revenue Streams (RS)

It is important to keep in mind that in multi-sided business model revenues streams are not generated from the end customers, but from one or more parties involved in the business model. Thus depending on the actor that runs the CloudSat model, the revenue streams are different.

More specifically, for the Infrastructure provider, which in this case could be the Satellite Operator, the main revenue stream is brought by the renting of the infrastructure to the various VNSPs. An important aspect to take into account here is whether a specific provision of infrastructure is suitable for a one-off payment approach or for recurring payments. Recurring payments can be created in different forms (e.g. subscription, usage base). Some of the possible pricing mechanisms are categorized in below figure [BMC][MARK].

PRICING MECHANISMS			
Fixed "Menu" Pricing **Predefined prices are based on static variables**		**Dynamic Pricing** **Prices change based on market conditions**	
List Price	Fixed process for individual products, services or other value propositions	**Negotiation (bargaining)**	Price negotiated between two or more partners, depending on negotiation power and/or negotiation skills
Product feature dependent	Price depends on the number or quality of value propositions features	**Yield management**	Price depends on inventory and time of purchase (normally used for limited resources)
Customer segment dependent	Price depends on the type and characteristic of the customer segment	**Real-time market**	Price is established dynamically based on supply and demand
Volume dependent	Price as a function of the quantity purchased	**Auctions**	Price determined by outcome of competitive bidding

Figure 163. Pricing Mechanisms

Detailed analysis concerning financial and pricing aspects of CloudSat is presented in section 3. However, it is important to point out that CloudSat follows a combination of Fixed “menu” and Dynamic pricing mechanisms (as the table above describes).

Actually, Cloudsat pricing relies on both:

- Close-to-static variables (such as bandwidth cost), which are not significantly affected by dynamic market conditions in a specific short period of time
- Product feature dependent pricing mechanism (such as VNF features per unit), which significantly fluctuates by VNF-specifications, VNF features per unit, market competition for the specific VNF etc.

Moreover, other schemes may be also applied from the Fixed “Menu” pricing models, such as List Price, where off-the-self products are sold on fixed prices as part of the Push supply chain strategy.

7.1.1.6. CloudSat Key Resources

The Key Resources building block captures all resources that are needed to make the entire business model possible. These resources are needed to create the offered products/services and the related value proposition, and to support the ‘Channels’, i.e. how to reach the customers, deliver the products/services and collect revenues.

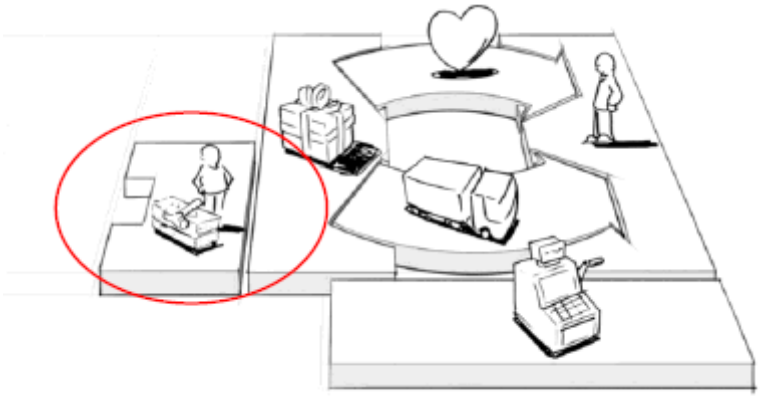

Figure 164. Key Resources

Key resources can be physical, financial, intellectual, or human. In addition, resources could be provided by external companies and/or partners, such in the case of CloudSat where the infrastructure is provided through a partnership. Thus, CloudSat resources and competences are technology oriented, focused on the establishment of the NS over the virtualization-capable infrastructure as well as the provision of QoS features across the network slicing.

For the implementation of the CloudSat networking model, the following architectural elements are considered, as essential resources. These components were extensively described in Chapter 5:

- NFVI-PoPs
- Network Elements
- Satellite Ground Segment Elements
- Virtualised Infrastructure Managers (VIMs).

- WAN Managers
- Satellite Hub Managers
- NFV Orchestrators
- Network Management Systems
- Federated Managers (Federators).

7.1.1.7. CloudSat Key Activities

This building block is close related to the previous one and it is meant to capture the key activities needed to offer the specific products/services.

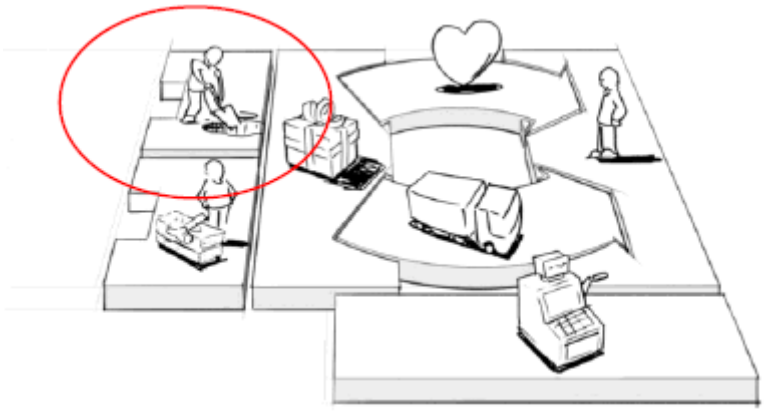

Figure 165. Key Activities

Osterwalder and Pigneur (2010) categorise these key activities into three main groups [BMC]:

- *Production*. Activities in this group relate to designing, making and delivering products/services. For the case of CloudSat, this activity can be summarized in the service composition, deployment and instantiation across the network slice.
- *Problem solving*. These activities are typically found in the service industry and are related to, for example, consulting or knowledge management. In the case of CloudSat, this phase is automated and performed by the Orchestrator (or the Federator depending on the use case), which gathers all the relevant information through monitoring actions and appropriate actions are taken according to the scalability strategy (e.g. scale up/out) for dealing with problems that may appear due to traffic congestion etc.
- *Platform/network*. This type of activities are related to platform management and promotion as well as service provisioning. They are very important in those business models in which a platform is a 'key resource', such as in CloudSat, where the system involves complex and heterogeneous infrastructure components whose management requires significant effort.

7.1.1.8. CloudSat Key Partnerships

This building block is utilised to capture suppliers and partners that are essential to create and deliver the specific products/services, which is very important to the business case of CloudSat which relies on partnerships for the provision of the cloud networking service. The BMC suggests four main categories of partnership:

- Strategic alliances between non-competitors
- Strategic partnerships between competitors
- Joint ventures to develop new businesses
- Buyer-supplier relationships to assure reliable supplies

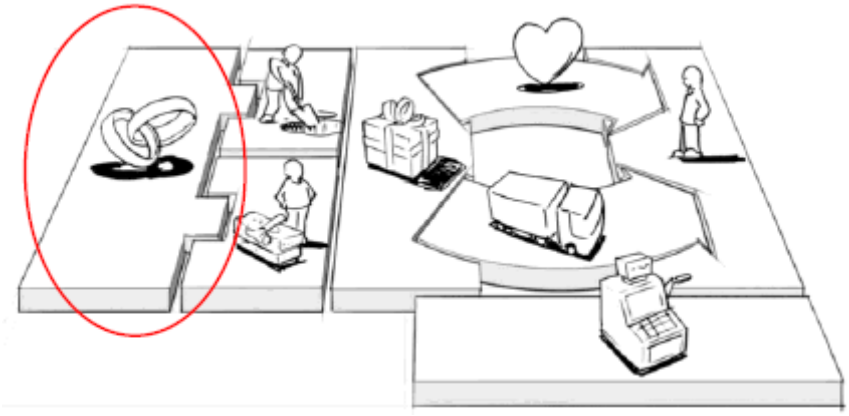

Figure 166. Key Partners

One important aspect to understand in the business model is the reason driving a specific partnership. Partnerships aimed at optimisation and economies of scales are driven mainly by cost savings objectives. Partnerships could be also driven by the objective of reducing risks in a new competitive environment by spreading them across the partners. Furthermore, partnerships are often created in order to acquire specific knowledge, technologies or to access new customer segments.

For the case of the CloudSat, from the above list, all types of partnerships may apply for a service provider (VNSP).

In specific, strategic alliances between competitors and non-competitors may help to shape a common standardized technological framework and introduce new technologies and innovations. Also, such partnerships may lead to multi-domain service scenarios, where services traverse multiple administrative domains.

In the same context, joint ventures may be created between e.g. terrestrial and satellite service providers in order to offer virtualized services with common federated management.

Finally, relationships with “suppliers” are also critical. VNF Providers can be seen as suppliers of virtualized network functions; enriching the network services with cutting-edge VNFs which are also reliable and certified are of key importance for the Service Provider. Relationships with equipment providers but also with InPs are also critical, since the actual infrastructure assets need to be technologically advanced,

reliable and always accompanied with adequate technical support, so that the Service Provider can actually be able to offer the desired service level to its customers.

The role of CloudSat key partners is considered of high importance since they could assist, through the applied automations, in expanding the reachability of the participating members, thus widening the targeted customer set. Also, a Trusted third party could act as an auction intermediary or as an Orchestrator of the service.

7.1.1.9. CloudSat Cost Structure

This building block is intended to capture the most relevant costs associated to implementing and operating a specific business model. The most important cost components can be identified by considering the information captured in the Key Resources, Key Activities, and Key Partnership building blocks of the BMC.

In broad terms, costs can be divided into fixed and variable. The former are costs that are incurred regardless of the volume of products manufactured or services that are delivered. The latter are costs that are proportionally linked to the volume of products manufactured and services delivered. The creation of large volume of products and/or services often benefits from 'economy of scale' as the average cost per unit falls as output rises. In addition, large organizations can benefit from 'economy of scale' as their operations can be shared and support multiple products and/or services.

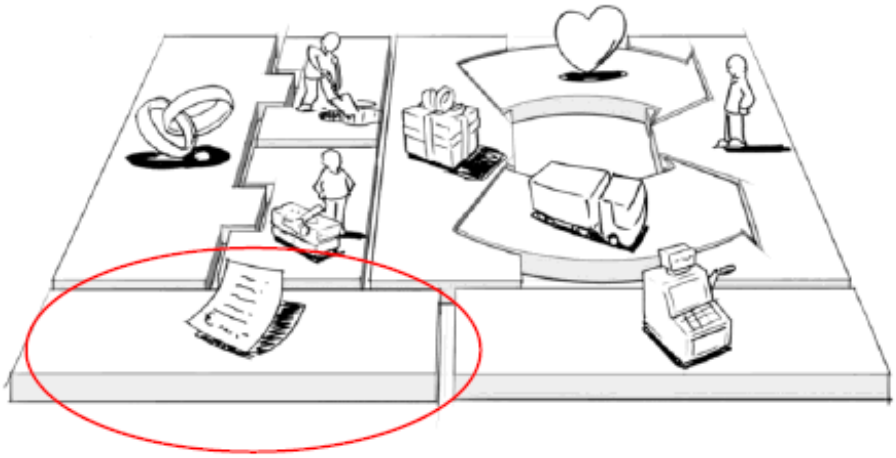

Figure 167. Cost Structure

Clarifying whether the organization's business model is 'cost-driven' or 'value-driven' is a very important aspect of an organization business model, and it has large implications on operations and costs.

Detailed analysis concerning financial and cost aspects of CloudSat is presented in the next sections.

7.1.2. CloudSat case scenarios

As per Chapter 5, the proposed CloudSat architecture is further refined to match the three integration scenarios selected for further study. These refinements are essentially subsets of the architecture, involving only specific entities and actors.

- Scenario #1 (Hybrid media distribution network as-a-Service) highlights the capabilities of the system to orchestrate multi-domain network services and also support customer-side network programmability. This scenario follows the GEO satellite system typical infrastructure.
- Scenario #2 (Dynamic backhauling with edge processing) stresses the feature of resource elasticity, as well as the capability to deploy satellite edge VNFs. This scenario follows the HTS/GEO satellite system typical infrastructure.
- Scenario #3 (Customer functions virtualization) focuses on the provision of VNFs as added-value services to enterprise but also non-professional satcom customers via a tailored customer portal. This scenario follows the LEO satellite system typical infrastructure.

7.2. Financial Analysis

Decision makers must make the most of scarce resources and at the same time respond to ever increasing demands for improved performance and new technology. The importance of investment management in information technology continues to increase. The failure rate of many IT investments and projects raises legitimate concerns about the value of those investments.

Investments in any sector are generally undertaken for one, or a combination, of four general purposes:

- Expansion or improvement in service or function of agency.
- Reduction of operating costs/increasing revenues.
- Research and development.
- Mandate

Benefits should clearly answer the question, “What does this investment provide the customer, public or agency?” Whether expressed in qualitative or quantitative terms, benefits should relate directly to the fulfilment of specific, expressed needs.

As a result, ICT investment proposals often require a rigorous business, market and financial analysis which will form a business case scenario in order to justify new IT investments. The business case, and associated feasibility studies, will include methods of assessing the costs and returns expected from the investment.

Generally, feasibility and financial analysis studies help to determine if potential solutions are viable and provide a basis of comparison and selection between alternatives. Technical feasibility studies focus on the technology of the solution and are used to determine a preferred IT solution from a technology perspective. An

economic feasibility study/analysis, such as a Cost-Benefit analysis (CBA), determines if a solution is economically sound and cost effective. Based upon these analyses, a new technology solution is proposed in the next step of the initiation process, and the results of the technical and economic feasibility studies are used to justify the proposed technology solution.

Cost-Benefit analysis (CBA), which is a form of economic evaluation, can be an aid to making such choices, together with other criteria, such as technological, market or business feasibility. A project proposal is said to be cost-effective if it produces relatively large benefits-gains for relatively low costs, compared to other ways of achieving the same goal. By evaluating both costs and effects of various options, their relative cost-effectiveness can be established. Efficiency is a related economic concept, but focuses more on the way in which inputs are transformed into outputs during the process of implementation. An intervention is said to be efficient if its implementation delivers the maximum amount of output (rather than benefits) given the amount of resources used in the intervention.

It is important to note that cost-effectiveness is a relative concept. A particular way of implementing a project is only cost-effective compared to other ways of providing the same outcomes. Or one particular proposal is considered more cost-effective than others aimed at similar outcomes. Moreover, whether the cost-effectiveness ratio is considered too high (e.g. high costs given results) depends on the overall budget (plus initial investment cost). If the budget is large, less cost-effective – but nevertheless effective – projects might still be considered as favourable.

Concerning CloudSat financial aspects, the previously described CloudSat business and market analysis provide the basis of an efficient cost effectiveness evaluation. Within this business, market and financial framework, the economic benefits/gains of CloudSat proposed architecture is financially analysed and assessed, providing a Cost-Benefit analysis, as a form of economic evaluation.

7.2.1. Methodology-Framework

The aim of this section is to familiarize the reader with the framework and terms used in the techno-economic analysis. Readers already familiar with financial frameworks of analysis are advised to proceed to the next section. Otherwise you are advised to proceed to Appendix III and get familiar with techno-economic analysis terms and processes.

7.2.2. CloudSat Financial Analysis

So, following the discounted Cash Flow Model (CFM) as presented in Appendix III, the structure of this section, presenting the financial framework of analysis, is summarized as per below:

- Identify the reasonable assumptions (general and scenario specific) of the financial analysis

- Define the benchmarking cases (as per GEO, MEO, LEO satellite systems) and estimate corresponding CAPEX costs and Cost reduction rates
- Summarize findings of benchmarking cases in a CAPEX Cost-Benefit Analysis (CBA)
- Define the three CloudSat case scenarios to be analyzed

and perform the following steps of financial analysis per CloudSat case/scenario:

- Estimate the initial Investment cost
    - Per component
    - FC/CAPEX
- Estimate the Operating/Running cost
    - Per component
    - VC/OPEX
- Perform high level financial analysis
    - Revenues
    - Cash Flow Model (CFM) and discounted Cash Flow
    - Financial Ratios
- In addition, for the three CloudSat cases, three different business environment scenarios of evaluations are used with their specific high level financial analysis variables/parameters
    - Optimistic (Blue Ocean)
    - Normal
    - Pessimistic (Red Ocean)
- Summarize findings of all three CloudSat cases in a Cost-Benefit Analysis (CBA)

7.2.2.1. Financial analysis assumptions

In order to form a sound framework for evaluating the CloudSat network model in a financial context, particular reasonable and realistic assumptions have to be made. More specifically, the assumptions, on which the proposed financial analysis is based, are the following.

General Assumptions:

- Two type of deployments (HW-based and SW-based) are examined in principle for CAPEX estimation based on the three satellite systems (GEO, MEO, LEO), resulting into three corresponding CAPEX financial analysis cases (GEO, MEO, LEO) that are used for benchmarking purposes (cost reduction rate per satellite system):
    1. Typical HW-based Satcom deployment without virtualization (which will be used as a benchmarking case) and

2. A full-stack deployment of SW-based virtualization-capable Satcom infrastructure made from scratch

- A 5-Year analysis period is applied.
- CloudSat related hardware and software investments as well as installation/configuration/license fees (for a 5-year period) are considered as part of calculated CAPEX.
- No extra investment costs involved (such as space leasing, personnel, salaries, extra training, compensations etc) are taken into account, since the study corresponds to a business case undertaken by an existing telco enterprise.
- As operating expenditures we consider licensing, installation, initial configuration, leasing of services/equipment, maintenance, repairs, unforeseen cases, marketing and promotion activities, supplies, property management, operations communications and bill/utilities expenses (throughout the specified period of analysis - 5 years)
- For the calculation of financial indicators we assume that CloudSat initial investment and CAPEX expenses take place at the Y0
- Cost, price, rates and charges suggestions for the CAPEX and OPEX calculations are result of market research and represent rates, cost and prices of actual market hardware, software products and/or services. Euribor, Eurozone rates and financial ratios have been used throughout the analysis unless otherwise stated.
- In NFV-enabled scenarios, the VNF costs represent the license fees associated for using the specific VNF in the system. In other words, flat-fee (and not usage-based) VNF pricing is assumed.
- Provided costs and prices are result of market research and corresponding resource references are provided per case.
- CAPEX costs are recorded as initial investment expense in Y0
- Depreciation (upon CAPEX) is applied as part of the income tax calculation only. It is not further analyzed or recorded as additional expense.
- For the income tax calculation CAPEX costs are not expensed in Y0 but capitalized in Y1-Y5 (therefore we have the corresponding CAPEX Depreciation Expense but only as part of the income tax expenses).
- The Modified Accelerated Cost Recovery System (MACRS) of depreciation is applied over the 5-year period for depreciation calculations.
- All amounts are in Euro currency. Current currency conversion rates have been applied whenever required.
- Analysis is made and presented by the perspective of well established business entities.
- Key actors have available the required infrastructure and the respective one that is needed for providing CloudSat services.

- Markets are characterized by economic stability (fixed rates and variables) unless otherwise stated
- Regulatory framework is not restrictive to CloudSat network model
- Data consistency and quality is ensured throughout the analysis

CloudSat Case Scenarios specific assumptions:

- For CloudSat Case Scenario #1:
    - CloudSat Case Scenario #1 is based on the GEO satellite systems
    - Leasing cost of 1 transponder (36MHz) is included in OPEX
    - Charge of services is done upon bandwidth utilization (MHz) of transponder and not per duration of usage
    - Revenues originate from two sources: (a) earnings as percentage of transponder utilization per year (basic satcom service, equal to the number of subscribers multiplied by the average price of subscription package) and (b) earnings from add-on services (SDN/NFV) calculated as fixed 15% percentage upon (a).
    - Initial transponder utilization rate is fixed at 30% for Y1 and dynamic for the years Y2-Y5 (as per the conditions of each scenario)
- For CloudSat Case Scenario #2:
    - CloudSat Case Scenario #2 is based on the HTS/GEO satellite systems
    - Satellite bandwidth leasing cost of 10MBps is included in OPEX
    - Revenues originate from two sources: (a) earnings as percentage of Satellite bandwidth resale usage (€/10Mbit/s) per year (including all add-on provisional services provided through this bandwidth) and (b) earnings from add-on services (SDN/NFV) calculated as fixed 15% percentage upon (a).
    - Initial Satellite bandwidth resale revenues percentage is fixed at an initial 30% for Y1 and dynamic for the years Y2-Y5 (as per the conditions of each scenario)
- For CloudSat Case Scenario #3:
    - CloudSat Case Scenario #3 is based on the LEO satellite systems and the financial analysis is performed on one gateway components and usage
    - Satellite bandwidth leasing cost is zero (owned satellites) and is included in OPEX
    - Revenues estimated figures are based on a typical LEO satellite system eg OneWeb annual revenues (average revenues over the period 2010-2014)

- The revenues of a LEO system such as OneWeb (originating from 50 gateways) are used for estimating the corresponding relative ratio of revenues for one gateway (OneWeb Revenues/50)
- The corresponding revenues of one OneWeb gateway include the provision of several services. It is estimated as assumption that the CloudSat Case Scenario #3 service represents the 1/10 of total provided services. Consequently, the revenues of CloudSat Case Scenario #3 service correspond to the 1/10 of the total gateway revenues (One gateway Revenues/10).
- Revenues originate from two sources: (a) earnings as percentage of the corresponding annual revenues originating from the utilization of the one gateway of the CloudSat Case Scenario #3 and (b) earnings from add-on services (SDN/NFV) calculated as fixed 15% percentage upon (a).
- Initial percentage of earnings upon the annual revenues originating from the utilization of the one gateway of the CloudSat Case Scenario #3 is set at 30% for Y1 and dynamic for the years Y2-Y5 (as per the conditions of each scenario)

It is fundamental for the accuracy and validity of the financial analysis herein, to have an accurate and trustworthy estimate of the cost reduction percentage resulting from the use of NFV technology and moving from a HW-based Satcom deployment without virtualization to a full-stack deployment of SW-based virtualization (NFV enabled).

Several things have to be considered when deploying SW-based virtualization (NFV), including accrued cost savings, business models, and architectural options. Cost must be considered over a full five-year lifecycle that would include both capital expenses (CAPEX) as well as operating expenses (OPEX). CAPEX reductions may become apparent immediately, while OPEX will improve over time. Further, NFV helps service providers avoid the traditional stair-step CAPEX needed to provision capacity in advance of expected demand. Instead, it follows a cloud utility-based business model where capacity is easily and cost-effectively added when you need it.

Within the above framework, three studies have been taken into account in order to quantify the cost benefits of an NFV based deployment, which CloudSat proposes, versus existing HW-based Satcom deployments without virtualization.

According to Hewlett-Packard white paper (2014) “The Reality of Cost Reduction”, a virtual customer premises equipment (vCPE) use case is analyzed and findings conclude that NFV deployment can reduce costs for about 18% to 24%, depending on the size of the deployment. Even with increased software costs, the reduction in hardware, installation, configuration, and power costs is more than enough to compensate [HWP].

	2	2a	2b	3	3a	3b
Branch Size	Small	Small	Small	Large	Large	Large
Router type	H/W	H/W	S/W	H/W	H/W	S/W
Branch services	WAN Accel, Caching, SBC	WAN Accel, Caching, SBC	WAN Accel, Caching, SBC	WAN Accel, Caching, SBC	WAN Accel, Caching, SBC	WAN Accel, Caching, SBC
HW Required	MSR930 3 Appliances	MSR930 Server	Server	MSR3012 3 Appliances	MSR3012 OAPv2	Server
SW Required	Included	Services vApps	VSR1001 + Services vApps	Included	Services vApps	VSR1008 + Services vApps
HW Cost	$11,100	$2,100	$1,500	$21,500	$6,000	$2,500
SW Cost	$0	$7,500	$8,000	$0	$15,000	$16,250
Installation/config time (hrs)	14	4.75	3	14	2.75	3
Installation/ config costs	$1,400	$475	$300	$1,400	$275	$300
Support Costs (3 yr)	$4,995	$4,320	$4,275	$9,675	$9,450	$8,438
Power costs (3 yr)	$1,440	$720	$360	$1,440	$360	$360
Total Costs	$18,935	$15,115	**$14,435**	$34,015	$31,085	**$27,848**

Typ small virtual appliance – $2,500
Typ small H/W appliance – $3,500
Typ large virtual appliance – $5,000
Typ large H/W appliance – $6,500
Small Server – $1,500
Larger Server – $2,500
Basic HW install hours – 2
Local SW config hours – 2
Remote SW config hours – 0.25
Installation labor/hr – $100
Support cost/year – 15%
Power consumption/box (kW) – 0.3
Hours of use/year – 4000
Power cost/kWh – $0.10

Figure 168. H/W vs S/W-NFV based Cost Model Comparison [HWP]

Similarly, the white paper study "Business Case for Moving DNS to the Cloud" (2014), performed by Alcatel-Lucent's Cloud Consulting team in cooperation with a Tier-1 service provider, proves how even a simple application like DNS can benefit enormously from running on an NFV platform. The cost savings analysis is related to three main cost drivers: CAPEX, OPEX infrastructure and OPEX processes. It shows the enormous benefits brought about by NFV even for a simple application such as DNS. In addition, apart from substantial monetary gains, the study determined that running the DNS application on Cloud simplifies complex processes such as healing, caling and software upgrading, which gives service providers greater agility and flexibility [ALC].

Figure 169 presents a summary of the 5-year total CAPEX and OPEX costs split in two scenarios. The first scenario (dedicated) shows the total costs that the service provider will need to incur to migrate from the PMO (Present Mode of Operations) to the FMO (Future Mode of Operations). The second scenario (shared) shows the new cost of operating DNS in a more efficient way. The difference between the two is the cost of idle capacity, which other applications could use in a shared scenario.

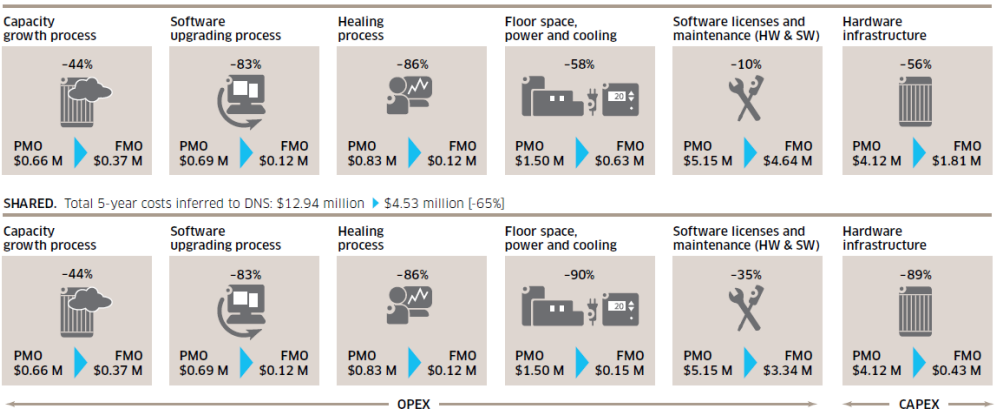

Figure 169. Summary of the 5-year total CAPEX and OPEX [ALC]

The analysis shows an impressive reduction in most of the process-related costs. In relative numbers, software upgrading and healing show an 83 percent and 86 percent reduction, respectively. In the shared scenario, the CAPEX and the OPEX infrastructure-related costs increase their contribution to the savings as the cost of idle capacity is not allocated to DNS. In total, concerning CAPEX and OPEX costs, we have cost savings of 41% and 65% respectively per case scenario [ALC].

Finally, as per ACG Research study (2013) "NFV Promises Cost Savings of Nearly 70%", commissioned by Affirmed Networks Inc. and VMware Inc, communication network operators could reduce capital expenditure by 68% and operating expenditure by 67% through the use of NFV. Explaining its findings in greater detail, the research says there are two main sources of cost savings: the use of x86-based hardware with the virtualized solution, compared with proprietary processing blades in the traditional one, and the service flexibility of the virtualized solution next to the traditional technology [ACG]. Another factor in the cost savings is a reduction in systems management thanks largely to the orchestration capabilities of the virtualized solution. NFV virtualization will cut service deployment time in half, in many cases, thanks to service orchestration, a reduction in the number of manual processes and greater automation. Figure 170 summarizes the finding of this study.

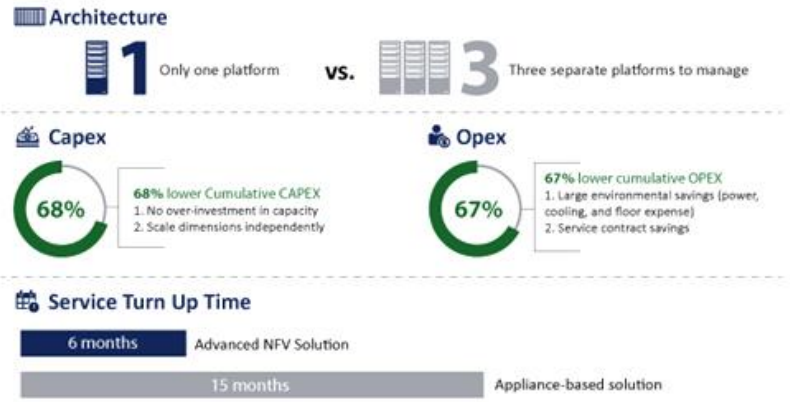

Figure 170. NFV CAPEX and OPEX savings [ACG]

Study	Cost Savings %
Hewlett-Packard (2014)	24%
Alcatel-Lucent (2014)	65%
ACG Research (2013)	67,5%
NFV Cost Reduction Rate	**52%**

Summarising the findings of the three studies, CloudSat adopts a NFV cost reduction rate of 52%, to be used as the average cost reduction of specific CAPEX components (with software capability) of a Satcom deployment with virtualization capabilities.

7.2.2.2. Benchmarking analysis case 1: GEO

This section provides the CAPEX financial analysis of a typical GEO satellite system HW-based satcom deployment without use of virtualization versus GEO Full-stack Satcom deployment of SW-based virtualization. This benchmarking process is a good indicator of proving the actual cost reduction that occurs when a Satcom infrastructure deployment is GEO SW-based and virtualization-capable as CloudSat proposes.

Below figure summarizes the major differences among LEO, MEO and GEO satellite systems. It is important to point out the difference in the terrestrial gateway cost which highly affects the CAPEX estimates of this analysis.

Parameter	LEO	MEO	GEO
Satellite Height	700 to 1400 Km	10,000 to 15,000 Km	36,000 Km
Orbital Period	10-40 minutes	2-8 hours	24 hours
Number of Satellites	40 +	10 to 15	3 to 4
Satellite Life	3 to 7 yrs	10 to 15 yrs	10 to 15 yrs
Space Segment Cost	High	Low	Medium
Terrestrial Gateway Cost	***High***	***Medium***	***Low***
Propagation Loss	Low	High	Highest

Figure 171. Satellite systems differences [STC], [LMG]

Typical GEO HW-based Satcom deployment without virtualization

In order to establish a typical GEO HW-based Satcom deployment, a typical satellite hub infrastructure should be developed. This infrastructure is part of the initial investment cost which is composed by the following components (Figure 172). The cost of all components results from market research and represents cost and prices of actual market hardware, software products and/or services. For reference and validation purposes, next to most of the components a reference link is provided.

Typical satellite hub infrastructure (GEO HW based)	Cost	Market Research References
Outdoor unit (eg Antenna, RF front-end etc)	70,000.00 €	http://www.satsig.net/ivsatcos.htm
Modulator	250,000.00 €	
Encapsulator/Multiplexer		
Return Link sub-system		
GW Management & Access Control		

Terrestrial interface subsystem (firewall) 5 year license	25,000.00 €	http://www.mcafeeworks.com/Firewall-Enterprise-S1104.asp
Terrestrial interface subsystem (PEP)	4,500.00 €	http://www.satcomresources.com/iDirect-SkyCelerator-Network-Accelerator-Model-1150
Terrestrial interface subsystem (router)	10,000.00 €	http://www.router-switch.com/Price-cisco-routers-cisco-router-3800-series_c40
TOTAL	359,500.00 €	

Figure 172. GEO H/W based deployment CAPEX

In specific, for the cost computation of terrestrial interface subsystem (firewall) with 5 year license, the lowest cost options have been selected (especially for the support and maintenance 5-year contract) and the breakdown of its cost is as follows [MCA]:

McAfee Firewall Enterprise S1104 Hardware	2,500.00 €
McAfee Firewall Enterprise - Standard, 5 Years	5,000.00 €
McAfee Firewall Enterprise Support - Standard 5 years	15,000.00 €
Intrusion Prevention 5 years	2,500 €
TOTAL	**25,000 €**

Figure 173. GEO Terrestrial interface subsystem (firewall) with 5 year license - H/W based [MCA]

Since this case is used for benchmarking purposes we will not proceed with detailed analysis/interpretation of findings at this point. All results will be used for comparison purposes against the corresponding findings of the GEO SW-based Satcom deployment case, following in the next section.

GEO SW-based deployment

This section provides the financial analysis of a GEO full stack SW-based Satcom deployment with use of virtualization through NFV. The results of this analysis will be compared against the corresponding results of GEO HW-based Satcom deployment as presented in the previous section. This benchmarking process will prove the actual cost reduction that occurs when a Satcom infrastructure deployment is SW-based and virtualization-capable, as CloudSat defines, compared to the typical HW-based infrastructures without the use of virtualization.

In order to establish a GEO Satcom deployment of SW-based virtualization, a typical satellite hub infrastructure should be developed but with software capability. The satellite hub infrastructure presented in the previous section, is used again but with several HW-based components/services being replaced by SW-based. In addition, extra components are introduced in order the additional software needs to be served efficiently.

A major prerequisite for achieving an accurate CAPEX calculation for such a deployment is the precise cost estimate of the involved software components. Based on the three business studies presented and analyzed, the NFV Cost Reduction Rate of 52% will be applied for the cost reduction calculation. In specific the cost of all software enabled hardware components of the typical GEO HW-based satellite hub infrastructure will be discounted by 52% in order to estimate the corresponding investment cost of a GEO Satcom deployment of SW-based virtualization. However, further down a sensitivity analysis is performed in order to assess the impact of the reduction rate to the total estimated cost.

Below figure presents the cost reduction process for those hardware components with enabled software capability plus the additional ones needed for such a deployment.

Typical satellite hub infrastructure (GEO HW-based)	HW-based Cost	SW Capability	NFV Cost Reduction Rate	GEO SW-based virtualization Cost
Outdoor unit (eg Antenna, RF front-end etc)	70,000.00 €	NO	-	70,000.00 €
Modulator	250,000.00 €	YES	52%	120,000.00 €
Encapsulator/Multiplexer				
Return Link sub-system				
GW Management & Access Control				
Terrestrial interface subsystem (firewall) with 5 year license	25,000.00 €	YES	-	17,500.00 €
Terrestrial interface subsystem (PEP)	4,500.00 €	YES	-	2,500.00 €
Terrestrial interface subsystem (router)	10,000.00 €	NO	-	10,000.00 €
Plus additional components/services required for a SW-based virtualization deployment				
Satellite Hub Generic Server HW (2 servers each)	-	-	-	10,000.00 €
Satellite Hub Generic Server Software Configuration	-	-	-	3,500.00 €
NFVI PoP SW + License + Configuration	-	-	-	3,500.00 €
NFVI PoP HW (5 servers each)	-	-	-	25,000.00 €
Orchestrator (HW/SW)				30,000.00 €
TOTAL	**359,500.00 €**			**292,000.00 €**
			GEO CAPEX cost reduction rate	**18.78%**

Figure 174. GEO Cost reduction Process

The initial cost of all components results from market research and represents cost and prices of actual market hardware, software products and/or services which are discounted (wherever applicable) by the NFV Cost Reduction Rate of 52%.

In specific, for the cost computation of terrestrial interface subsystem (firewall) with 5 year license, the hardware McAfee Firewall Enterprise S1104 has been excluded, the support cost has been reduced and the breakdown of its cost is now as follows (reduced by 7,500€)[MCA]:

McAfee Firewall Enterprise - Standard, 5 Years	5,000.00 €
McAfee Firewall Enterprise Support - Standard 5 years	10,000.00 €
Intrusion Prevention 5 years	2,500 €
TOTAL	**17,500 €**

Figure 175. GEO Terrestrial interface subsystem (firewall) with 5 year license - S/W based [MCA]

The initial investment cost (CAPEX) of this SW-based virtualization deployment is summarized in below figure, pointing out the cost per component and its corresponding market research references.

GEO SW-based virtualization satellite hub infrastructure	Cost	Market Research References
Outdoor unit (eg Antenna, RF front-end etc)	70,000.00 €	http://www.satsig.net/ivsatcos.htm
Modulator	120,000.00 €	
Encapsulator/Multiplexer		
Return Link sub-system		
GW Management & Access Control		
Satellite Hub Generic Server HW (2 servers each)	10,000.00 €	Hewlett-Packard white paper (2014) "The Reality of Cost Reduction" http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0 IBM report "A Deep Dive Into The Cost Benefits of KVM and Open Virtualization"
Satellite Hub Generic Server Software Configuration	3,500.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0
Terrestrial interface subsystem (firewall) 5 year license	17,500.00 €	http://www.mcafeeworks.com/Firewall-Enterprise-S1104.asp
Terrestrial interface subsystem (PEP)	2,500.00 €	http://www.satcomresources.com/iDirect-SkyCelerator-Network-Accelerator-

		Model-1150
Terrestrial interface subsystem (router)	10,000.00 €	http://www.router-switch.com/Price-cisco-routers-cisco-router-3800-series_c40
NFVI PoP SW + License + Configuration	3,500.00 €	Hewlett-Packard white paper (2014) "The Reality of Cost Reduction"
NFVI PoP HW (5 servers each)	25,000.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0 IBM report "A Deep Dive Into The Cost Benefits of KVM and Open Virtualization"
Orchestrator (HW/SW)	30,000.00 €	
TOTAL	292,000.00 €	

Figure 176. GEO S/W based deployment - cost justification

Although detailed evaluation of financial results of the two deployments (HW-based without virtualization vs SW-based virtualization) will be presented (CAPEX based Cost Benefit Analysis of Benchmarking cases), it is worth mentioning that initial investment cost (CAPEX) has been **reduced from 359,500€ to 292,000€, a cost reduction of 18.78%,** as result of the SW-based virtualization capability of the GEO Satcom deployment.

7.2.2.3. Benchmarking analysis case 2: MEO

This section provides the CAPEX financial analysis of a typical MEO satellite system HW-based satcom deployment without use of virtualization versus MEO Full-stack Satcom deployment of SW-based virtualization. This benchmarking process is a good indicator of proving the actual cost reduction that occurs when a Satcom infrastructure deployment is MEO SW-based and virtualization-capable.

For this MEO benchmarking case analysis [O3B]:

- The characteristics of a typical MEO operator have been used (e.g. O3B)
- Both the HW and SW based configuration is based on the use of 10 terrestrial gateways and their corresponding components.
- One Gateway cost is set at 340,000 €

Typical MEO HW-based Satcom deployment without virtualization

In order to establish a typical MEO HW-based Satcom deployment, a typical MEO satellite hub infrastructure should be developed (we assume 10 terrestrial gateways). This infrastructure is part of the initial investment cost which is composed by the following components (See below Figure). The cost of all components results from market research and represents cost and prices of actual market hardware, software

products and/or services. For reference and validation purposes, next to most of the components a reference link is provided.

Typical satellite hub infrastructure for 10 gateways (MEO HW based)	Cost per Unit	Units for 10 gateways	Cost	Market Research References
Outdoor units	100,000.00 €	10	1,000,000.00€	http://www.o3bnetworks.com/
Modulator	340,000.00 €	10	3,400,000.00€	http://www.satsig.net/ivsatcos.htm
Encapsulator/Multiplexer				
Return Link sub-system				
GW Management & Access Control				
Fibre Leased Lines (1 GBps annual cost + initial installation)	30,000.00 €	10	300,000.00 €	http://www.hso.co.uk/leased-lines/ http://business.bt.com/broadband-and-internet/leased-lines/
Terrestrial interface subsystem (firewall) 5 year license	25,000.00 €	10	250,000.00 €	http://www.mcafeeworks.com/Firewall-Enterprise-S1104.asp
Terrestrial interface subsystem (PEP)	4,500.00 €	10	45,000.00 €	http://www.satcomresources.com/iDirect-SkyCelerator-Network-Accelerator-Model-1150
Terrestrial interface subsystem (router)	10,000.00 €	10	100,000.00 €	http://www.router-switch.com/Price-cisco-routers-cisco-router-3800-series_c40
TOTAL			5,095,000.00€	

Figure 177. MEO H/W based deployment CAPEX

Since this case is used for benchmarking purposes we will not proceed with detailed analysis/interpretation of findings at this point. All results will be used for comparison purposes against the corresponding findings of the MEO SW-based Satcom deployment case, following in the next section.

MEO SW-based deployment

This section provides the financial analysis of a MEO full stack SW-based Satcom deployment with use of virtualization through NFV. This benchmarking process will prove the actual cost reduction that occurs when a Satcom infrastructure deployment is SW-based and virtualization-capable, as CloudSat defines, compared to the typical HW-based infrastructures without the use of virtualization.

A major prerequisite for achieving an accurate CAPEX calculation for MEO deployment is the precise cost estimate of the involved GW software components. Based on the three NFV cost reduction studies presented and analyzed, the average NFV Cost Reduction Rate of 52% (as calculated for terrestrial network components) will be applied for the NFV cost reduction calculation of the MEO HW-based GWs.

However, for the MEO case, beyond the cost reduction of the a single MEO GW, it must be mentioned that in order to establish a MEO Satcom deployment worldwide, a group of an average 10 MEO gateways should be considered for virtualization at distributed NFVI-PoPs at the geographical areas of interest. However, instead of considering an 1-to-1 cardinality between each GW and NFVI PoP, one of the added values of the NFV approach is the ability to collocate more than one NFV-based MEO GWs within the same cloud data centre (i.e. NFVI PoP), achieving by this way multiple cost reduction. In this direction, we consider in our analysis that the initially 10 deployed GWs will be virtually instantiated at 7 NFVI PoPs. The relative ratio of using 7 NFVI-PoP units (instead of 10) results to an additional 30% cost reduction to the initial cost needed for 10 units in terms of network equipment and network elements.

However, this concentration of the SW-based GWs to less NFVI-PoPs has as a result the cost increase of the Fibre Leased Lines needed for connectivity in comparison to the cost estimated for 10 HW-based GWs. The reason for the increased connectivity cost is that in the virtualization case, it is created the need to connect the ODUs to the NFVI-PoPs and then each NFVI-PoP to the network. Therefore, we consider an increase by 50% in the initial connectivity cost (estimating 15 connections instead of 10 connections).

Typical satellite hub infrastructure (MEO HW-based)	HW-based Cost (for 10 gateways)	SW Capability	Cost Reduction/Increase Rate	MEO SW-based virtualization Cost
Outdoor units	1,000,000.00€	NO	-	1,000,000.00 €
Modulator	3,400,000.00€	YES	NFV Cost Reduction - 52%	1,632,000.00 €
Encapsulator/Multiplexer				
Return Link sub-system				
GW Management & Access Control				
Fibre Leased Lines (1 GBps annual cost + initial	300,000.00 €	NO	(for 15 units)	450,000.00 €

installation)			+50%	
Terrestrial interface subsystem (firewall) with 5 year license	250,000.00 €	YES	(17,500*7) -30%	122,500.00 €
Terrestrial interface subsystem (PEP)	45,000.00 €	YES	(2,500*7) -30%	17,500.00 €
Terrestrial interface subsystem (router)	100,000.00 €	YES	(7 units) -30%	70,000.00 €
Plus additional components/services required for a SW-based virtualization deployment				
Satellite Hub Generic Server HW (2 servers at 10,000€)	-	-	(7 units)	70,000.00 €
Satellite Hub Generic Server Software Configuration (3,000€)	-	-	(7 units)	24,500.00 €
NFVI PoP SW + License + Configuration(3,000€)	-	-	(7 units)	24,500.00 €
NFVI PoP HW (5 servers at 25,000€)	-	-	(7 units)	175,000.00 €
Orchestrator (HW/SW)	-	-	-	30,000.00 €
TOTAL	**5,095,00.00 €**			**3,616,000.00 €**
			MEO CAPEX cost reduction rate	**29.03%**

Figure 178. MEO Cost reduction Process

The initial cost values of all the components results from market research and represents cost and prices of actual market hardware, software products and/or services.

The market research references of the SW-based initial investment cost (CAPEX) used in the benchmarking table above are provided in the following table.

MEO SW-based virtualization satellite hub infrastructure	Cost	Market Research References
Outdoor units	1,000,000.00 €	
Modulator	1,632,000.00 €	http://www.o3bnetworks.com/ http://www.satsig.net/ivsatcos.htm
Encapsulator/Multiplexer		
Return Link sub-system		
GW Management & Access Control		
Fibre Leased Lines (1 GBps annual cost + initial installation)	450,00000 €	http://www.hso.co.uk/leased-lines/ http://business.bt.com/broadband-and-internet/leased-lines/
Satellite Hub Generic Server HW (2 servers)	70,000.00 €	Hewlett-Packard white paper (2014) "The Reality of Cost Reduction" http://www8.hp.com/h20195/V2/getp

		df.aspx/4AA5-2160ENW.pdf?ver=1.0 IBM report "A Deep Dive Into The Cost Benefits of KVM and Open Virtualization"
Satellite Hub Generic Server Software Configuration	24,500.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0
Terrestrial interface subsystem (firewall) 5 year license	122,500.00 €	http://www.mcafeeworks.com/Firewall-Enterprise-S1104.asp
Terrestrial interface subsystem (PEP)	17,500.00 €	http://www.satcomresources.com/iDirect-SkyCelerator-Network-Accelerator-Model-1150
Terrestrial interface subsystem (router)	70,000.00 €	http://www.router-switch.com/Price-cisco-routers-cisco-router-3800-series_c40
NFVI PoP SW + License + Configuration	24,500.00 €	Hewlett-Packard white paper (2014) "The Reality of Cost Reduction"
NFVI PoP HW (5 servers each)	175,000.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0 IBM report "A Deep Dive Into The Cost Benefits of KVM and Open Virtualization"
Orchestrator (HW/SW)	30,000.00 €	
TOTAL	3,616,000.00 €	

Figure 179. MEO S/W based deployment - cost justification

Although detailed evaluation of financial results of the two MEO deployments (HW-based without virtualization vs SW-based virtualization) will be presented (CAPEX based Cost Benefit Analysis of Benchmarking cases), it is worth mentioning that initial investment cost (CAPEX) has been **reduced from 5,095,00.00 € to 3,616,000.00 €, a cost reduction of 29.03%,** as result of the SW-based virtualization capability of the MEO Satcom deployment and the centralization of the GW function to less in number NFVI-PoPs.

7.2.2.4. Benchmarking analysis case 3: LEO

This section provides the CAPEX financial analysis of a typical LEO satellite system HW-based satcom deployment versus SW-based deployment with virtualization capabilities.

For this LEO benchmarking case analysis [OWB]:

- The characteristics of a typical LEO operator are considered (e.g. OneWeb).

- Both the HW and SW based configuration is based on the use of 50 terrestrial gateways and their corresponding components.
- One terrestrial gateway cost is set at 750,000 € (LEO satellite system)

Typical LEO HW-based Satcom deployment without virtualization

In order to establish a typical LEO HW-based Satcom deployment, we assume the deployment of 50 terrestrial gateways worldwide. This infrastructure is part of the initial investment cost which is composed by the following components (See below Figure 172). The cost of all components results from market research and represents cost and prices of actual market hardware, software products and/or services. For reference and validation purposes, next to most of the components a reference link is provided.

Typical satellite hub infrastructure for 50 gateways (LEO HW based)	Cost per Unit	Units for 50 gateways	Cost	Market Research References
Outdoor units	100,000.00 €	50	5,000,000.00€	http://www.oneweb.net
Modulator	750,000.00 €	50	37,500,000.00€	https://www.iridium.com/
Encapsulator/Multiplexer				
Return Link sub-system				
GW Management & Access Control				http://www.satsig.net/ivsatcos.htm
Fibre Leased Lines (1 GBps annual cost + initial installation)	30,000.00 €	50	1,500,000.00 €	http://www.hso.co.uk/leased-lines/ http://business.bt.com/broadband-and-internet/leased-lines/
Terrestrial interface subsystem (firewall) 5 year license	25,000.00 €	50	1,250,000.00 €	http://www.mcafeeworks.com/Firewall-Enterprise-S1104.asp
Terrestrial interface subsystem (PEP)	4,500.00 €	50	225,000.00 €	http://www.satcomresources.com/iDirect-SkyCelerator-Network-Accelerator-

				Model-1150
Terrestrial interface subsystem (router)	10,000.00 €	50	500,000.00 €	http://www.router-switch.com/Price-cisco-routers-cisco-router-3800-series_c40
TOTAL			45,975,000.00€	

Figure 180. LEO H/W based deployment CAPEX

Since this case is used for benchmarking purposes we will not proceed with detailed analysis/interpretation of findings at this point. All results will be used for comparison purposes against the corresponding findings of the LEO SW-based Satcom deployment case, following in the next section.

LEO SW-based deployment

This section provides the financial analysis of a LEO full stack SW-based Satcom deployment with use of virtualization through NFV. The results of this analysis will be compared against the corresponding results of LEO HW-based Satcom deployment as presented in the previous section. This benchmarking process will prove the actual cost reduction that occurs when a Satcom infrastructure deployment is SW-based and virtualization-capable, as CloudSat defines, compared to the typical HW-based infrastructures without the use of virtualization.

In order to establish a LEO Satcom deployment of SW-based virtualization, a typical LEO satellite hub infrastructure (with 50 terrestrial gateways) should be developed but with virtualization capability. The satellite hub infrastructure previously presented, is used again but with several HW-based components/services being replaced by SW-based ones. In addition, extra components are introduced in order the additional software needs to be served efficiently.

A major prerequisite for achieving an accurate CAPEX calculation for such a deployment is the precise cost estimate of the involved software components. Based on the three business studies presented and analyzed, the NFV Cost Reduction Rate of 52% will be applied for the cost reduction calculation. In specific the cost of all software enabled hardware components of the typical LEO HW-based satellite hub infrastructure will be discounted by 52% in order to estimate the corresponding investment cost of a LEO Satcom deployment of SW-based virtualization.

Below figure presents the cost reduction process for those hardware components with enabled software capability plus the additional ones needed for such a LEO deployment. Please note that specific components (eg leased lines, terrestrial interface, servers, NFVI PoP SW/HW etc) have a cost value reduced by 50%. This is due to the available virtualization services (NFV) that allow these components to

serve more gateways and consequently to need less units in our MEO SW-based virtualization installation (25 instead of 50 units). The relative ratio of using 25 units instead of 50 is equivalent to a 50% cost reduction to the initial cost of 50 units.

In addition the Fibre Leased Lines cost, as estimated for the 50 gateways, is increased by 50% (using 75 units instead of 50) since due to the virtualization services several clouds should be used and more communication lines are required (not just for the 50 gateways).

Typical satellite hub infrastructure (LEO HW-based)	HW-based Cost (for 50 gateways)	SW Capability	Cost Reduction/Increase Rate	LEO SW-based virtualization Cost
Outdoor units	5,000,000.00€	NO	-	5,000,000.00 €
Modulator Encapsulator/Multiplexer Return Link sub-system GW Management & Access Control	37,500,000.00€	YES	NFV Cost Reduction - 52%	18,000,000.00 €
Fibre Leased Lines (1 GBps annual cost + initial installation)	1,500,000.00 €	NO	(75 units) +50%	2,250,000.00 €
Terrestrial interface subsystem (firewall) with 5 year license	1,250,000.00 €	YES	(17,500*25) -50%	437,500.00 €
Terrestrial interface subsystem (PEP)	225,000.00 €	YES	(2,500*25) -50%	62,500.00 €
Terrestrial interface subsystem (router)	500,000.00 €	YES	-50%	250,000.00 €
Plus additional components/services required for a SW-based virtualization deployment				
Satellite Hub Generic Server HW (2 servers at 10,000€)	-	-	(25 units) -50%	250,000.00 €
Satellite Hub Generic Server Software Configuration (3,500€)	-	-	(25 units) -50%	87,500.00 €
NFVI PoP SW + License + Configuration(3,500€)	-	-	(25 units) -50%	87,500.00 €
NFVI PoP HW (5 servers at 25,000€)	-	-	(25 units) -50%	625,000.00 €
Orchestrator (HW/SW)	-	-	-	30,000.00 €
TOTAL	**45,975,00.00€**			**27,080,000.00€**
			LEO CAPEX cost reduction rate	**41.10%**

Figure 181. LEO Cost reduction Process

The initial cost of all components results from market research and represents cost and prices of actual market hardware, software products and/or services which are discounted (wherever applicable) by the NFV Cost Reduction Rate of 52%.

The initial investment cost (CAPEX) of this SW-based virtualization deployment is summarized in below figure, pointing out the cost per component and its corresponding market research references.

LEO SW-based virtualization satellite hub infrastructure (50 terrestrial gateways)	Cost	Market Research References
Outdoor units	5,000,000.00 €	http://www.oneweb.net https://www.iridium.com/ http://www.satsig.net/ivsatcos.htm
Modulator	18,000,000.00 €	
Encapsulator/Multiplexer		
Return Link sub-system		
GW Management & Access Control		
Fibre Leased Lines (1 GBps annual cost + initial installation)	2,250,00000 €	http://www.hso.co.uk/leased-lines/ http://business.bt.com/broadband-and-internet/leased-lines/
Satellite Hub Generic Server HW (2 servers)	250,000.00 €	Hewlett-Packard white paper (2014) "The Reality of Cost Reduction" http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0 IBM report "A Deep Dive Into The Cost Benefits of KVM and Open Virtualization"
Satellite Hub Generic Server Software Configuration	87,500.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0
Terrestrial interface subsystem (firewall) 5 year license	437,500.00 €	http://www.mcafeeworks.com/Firewall-Enterprise-S1104.asp
Terrestrial interface subsystem (PEP)	62,500.00 €	http://www.satcomresources.com/iDirect-SkyCelerator-Network-Accelerator-Model-1150
Terrestrial interface subsystem (router)	250,000.00 €	http://www.router-switch.com/Price-cisco-routers-cisco-router-3800-series_c40
NFVI PoP SW + License + Configuration	87,500.00 €	Hewlett-Packard white paper (2014) "The Reality of Cost Reduction" http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0 IBM report "A Deep Dive Into The Cost Benefits of KVM and Open Virtualization"
NFVI PoP HW (5 servers each)	625,000.00 €	

Orchestrator (HW/SW)	30,000.00 €	
TOTAL	27,080,000.00€	

Although detailed evaluation of financial results of the two LEO deployments (HW-based without virtualization vs SW-based virtualization) will be presented (CAPEX based Cost Benefit Analysis of Benchmarking cases), it is worth mentioning that initial investment cost (CAPEX) has been **reduced from 45,975,00.00€ to 27,080,000.00€, a cost reduction of 41.10%,** as result of the SW-based virtualization capability of the LEO Satcom deployment.

7.2.2.5. CAPEX-based Cost Benefit Analysis of GEO/MEO/LEO Cases

In this section, a benchmarking analysis of CAPEX figures is presented by comparing the CAPEX financial analysis results of a HW-based vs a SW-based with virtualization deployment for the three satellite systems (GEO, MEO, LEO).

Below figure summarizes and compares the CAPEX cost the benchmarking cases (HW and SW) for the three satellite systems (GEO, MEO, LEO). Amounts refer to the total CAPEX amount per case.

Satellite System	CAPEX (HW-based case)	CAPEX (SW-Based with virtualization case)	Cost Reduction Percentage (%)
GEO	359,500.00 €	292,000.00 €	18.78%
MEO (10 gateways)	5,095,000.00 €	3,616,000.00 €	29.03%
LEO (50 gateways)	45,975,000.00 €	27,080,000.00 €	41.10%

Figure 182. Benchmarking cases – Comparison of CAPEX financial figures

In specific, below figures show graphically the CAPEX amount cost reduction per benchmarking case, revealing that significantly higher amounts are saved when the deployment type requires many GWs, such as MEO, but even more in LEO case.

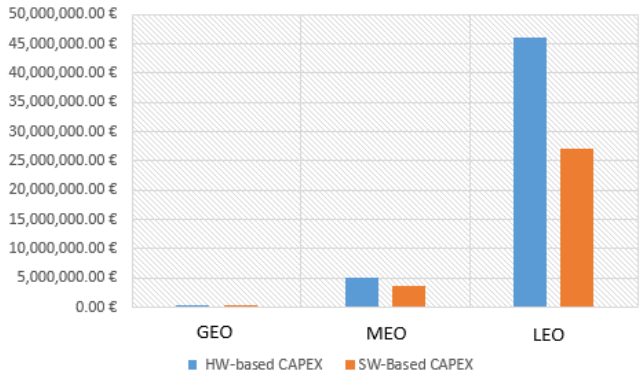

Figure 183. Benchmarking case GEO/MEO/LEO – CAPEX cost reduction

Overall, we conclude that introducing a SW-based Satcom deployment with virtualization capabilities reduces the initial investment cost (CAPEX) compared to a HW one, no matter the selected satellite system (GEO, MEO, LEO). By using a NFV Cost Reduction factor of 52% we achieve cost reduction of 18.78%, 29.03% and 41.10% for GEO, MEO and LEO respectively.

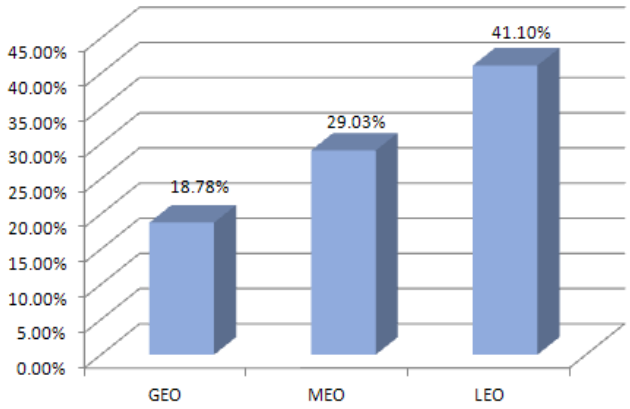

Figure 184. Benchmarking cases GEO, MEO, LEO – Cost Reduction (%)

In order to verify the sensitivity of the CAPEX financial results and the corresponding cost reduction percentages, a sensitivity analysis will be performed per GEO/MEO/LEO case, using different NFV Cost Reduction factors.

Below figures show how different NFV Cost Reduction percentages affect the cost reduction results for benchmarking analysis case 1 - GEO.

GEO	NFV Cost Reduction percentages								
	30%	35%	40%	45%	52%	55%	60%	65%	70%
CAPEX Cost Reduction (in %)	3.48%	6.95%	10.43%	13.91%	18.78%	20.86%	24.34%	27.82%	31.29%

Figure 185. Sensitivity Analysis: Benchmarking case GEO – CAPEX cost reduction

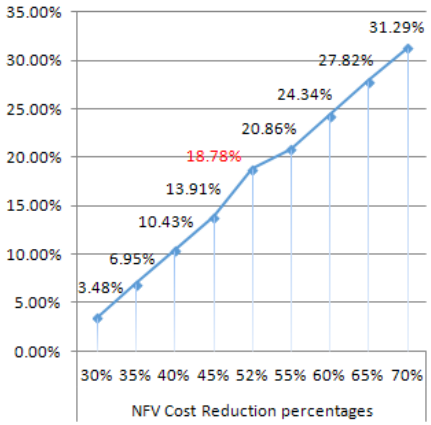

Figure 186. Sensitivity Analysis: Benchmarking case GEO – CAPEX cost reduction graph

Similarly, below figures show how different NFV Cost Reduction percentages affect the cost reduction results for benchmarking analysis case 2 - MEO.

MEO	NFV Cost Reduction percentages								
	30%	35%	40%	45%	52%	55%	60%	65%	70%
CAPEX Cost Reduction (in %)	14.35%	17.68%	21.02%	24.36%	29.03%	31.03%	34.37%	37.70%	41.04%

Figure 187. Sensitivity Analysis: Benchmarking case MEO – CAPEX cost reduction

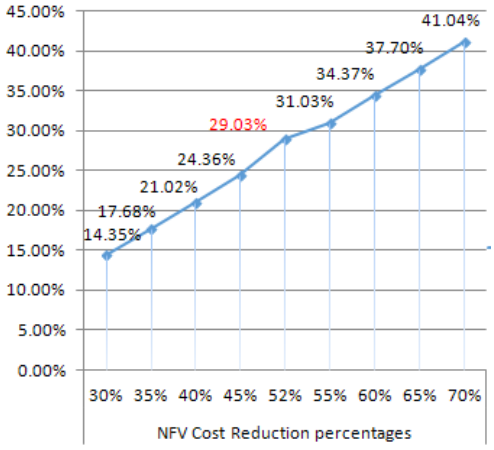

Figure 188. Sensitivity Analysis: Benchmarking case GEO – CAPEX cost reduction graph

Finally, below figures show how different NFV Cost Reduction percentages affect the cost reduction results for benchmarking analysis case 3 - LEO.

LEO	NFV Cost Reduction percentages								
	30%	35%	40%	45%	52%	55%	60%	65%	70%
CAPEX Cost Reduction (in %)	23.15%	27.23%	31.31%	35.39%	41.10%	43.55%	47.62%	51.70%	55.78%

Figure 189. Sensitivity Analysis: Benchmarking case LEO – CAPEX cost reduction

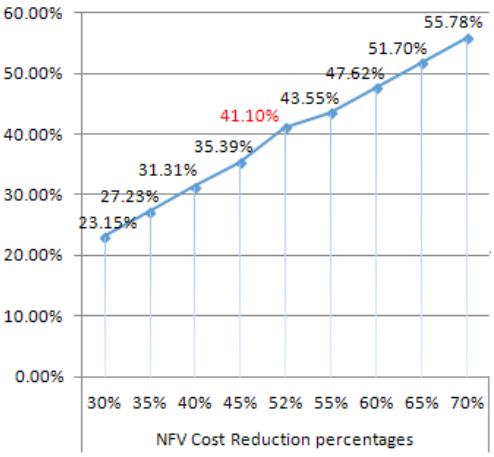

Figure 190. Sensitivity Analysis: Benchmarking case LEO – CAPEX cost reduction graph

Summarising the above findings of the sensitivity analysis for the three benchmarking cases we conclude that introducing a SW-based Satcom deployment with virtualization capabilities reduces the initial investment cost (CAPEX) compared to a HW one, for any satellite system (GEO, MEO, LEO) when we use a NFV Cost Reduction factor ranging from 30% to 70% and above.

The following table and figure summarize the findings of the sensitivity analysis both quantitatively and graphically for better understanding and conception of the potential cost reduction trends.

CAPEX Cost Reduction (in %)	NFV Cost Reduction percentages								
	30%	35%	40%	45%	52%	55%	60%	65%	70%
GEO	3.48%	6.95%	10.43%	13.91%	18.78%	20.86%	24.34%	27.82%	31.29%
MEO	14.35%	17.68%	21.02%	24.36%	29.03%	31.03%	34.37%	37.70%	41.04%
LEO	23.15%	27.23%	31.31%	35.39%	41.10%	43.55%	47.62%	51.70%	55.78%

Figure 191. Sensitivity Analysis: Benchmarking cases GEO, MEO, LEO

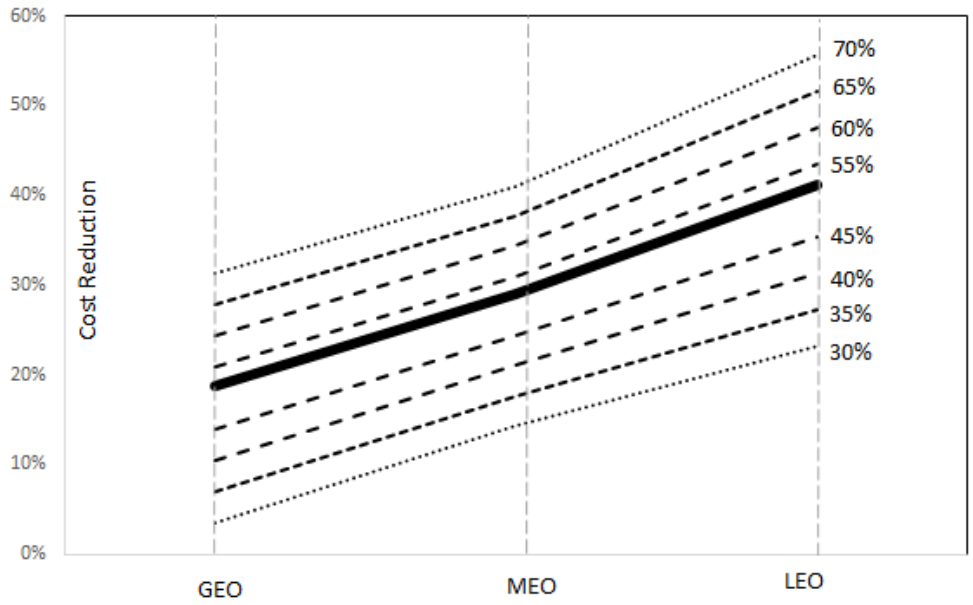

Figure 192. Sensitivity Analysis Graph: Benchmarking cases GEO, MEO, LEO

7.2.3. Financial Analysis of the three CloudSat case scenarios

In this section, the three CloudSat case scenarios will be financially analyzed and evaluated under the conditions of three business and market environment scenarios: Normal, Optimistic and Pessimistic.

A main addition compared to the benchmarking scenarios analysis of previous sections (HW-based without virtualization vs SW-based with virtualization) is that part of CAPEX will be financed through a 200,000 Euro bank loan with constant/fixed 2.2% interest rate (Euribor interest rate of 0.2% plus 2%) paid in 5 years with annual constant capital installments.

This extra source of financing creates more realistic business, finance and market conditions for the evaluation of the proposed CloudSat case scenarios/deployments since accounting-wise CAPEX figure is reduced by the loan amount which is then equally expensed (with interest) as OPEX over the 5-year period of analysis.

Details about this source of financing are provided in the table below.

	Y1	Y2	Y3	Y4	Y5	TOTAL
CAPITAL PAID	40,000.00 €	40,000.00 €	40,000.00 €	40,000.00 €	40,000.00 €	**200,000.00 €**
INTEREST PAID	4,400.00 €	3,557.88 €	2,697.23€	1,817.65 €	918.71 €	**13,391.46**
ANNUAL INSTALLMENT	**44,400.00 €**	**43,557.88 €**	**42,697.23€**	**41,817.65 €**	**40,918.71 €**	**213,391.46€**
200,000€ bank loan with constant/fixed 2.2% interest rate (Euribor interest rate of 0.2% plus 2%) paid in 5 years with annual constant capital installments						

Figure 193. Bank Loan Analysis – fixed interest rate for 5-year period

As have been stated in the reasonable assumptions section of this financial analysis no extra investment costs such as space leasing, personnel, salaries, extra training, compensations etc are involved in OPEX calculation. The variables that are taken into account as operating/running costs throughout the specified period of analysis (5 years) other than the ones included in CAPEX (the cost of initial investment, licensing, installation, initial configuration) are the maintenance, repairs and unforeseen expenses, marketing and promotion activities, supplies, property management, operations and bill/utilities expenses as well as then main Satellite bandwidth leasing cost of the 1 Transponder (for CloudSat Case Scenario 1/GEO) or 10Mbps leasing cost (for CloudSat Case Scenario 2/GEO HTS) or zero leasing cost (for CloudSat Case Scenario 3/LEO).

Most of these OPEX variables have been calculated as a percentage upon CAPEX and are either static or dynamic over the 5-year period. Below figure summarizes the OPEX variables and their characteristics for our H/W based deployment without virtualization scenario.

OPEX variables	TYPE	Characteristics over the 5-year period
CloudSat Case Scenario 1: Annual Satellite bandwidth leasing cost (1 Transponder)	Static	Fixed 1,100,000€ per year
CloudSat Case Scenario 2: Annual Satellite 10MBps leasing cost (2000Euro per Mbps/month)	Static	Fixed 240,000€ per year
CloudSat Case Scenario 3: Annual Satellite leasing cost	Static	Zero (due to owned satellites)
Maintenance and Repair/Replacement/Unforeseen expenses	Dynamic	1.5% of CAPEX increased by 10% annually
Marketing/Advertising/Promotion expenses	Dynamic	4% of CAPEX reduced by 50% annually
Property management, operations, communication and bill/utilities expenses	Static	0.5% of CAPEX
Supply expenses	Static	Fixed 4,000€ per year

Escalation of Costs	Dynamic	0.60% as result of Eurozone inflation rate (0.10%) plus 0.50%

Figure 194. OPEX variables for H/W based deployment without virtualization

In project management, escalation of costs is considered as a financial estimate risk of costs calculation, and that is why it should be included in project estimates and budgets. Cost escalation variable is defined as changes in the cost of specific goods or services in a given economy over a period. This is similar to the concepts of inflation and deflation except that escalation is specific to a category of products or services. While escalation includes general inflation related to the money supply, it is also driven by changes in technology, practices, and particularly supply-demand imbalances that are specific to a good or service in a given economy.

In technological related projects and economies characterized by financial stability, the escalation of costs rate used in financial analysis is usually set as per the given inflation rate or slightly above (so as to include slight inflation rate fluctuations over the period of financial analysis) and is applied to the total OPEX figure. In our financial analysis, escalation of cost rate is set at 0.60% (Eurozone inflation rate 0.10% plus 0.50%).

Similar to the escalation of cost factor, as already mentioned in the OPEX estimation paragraph, the Escalation of Benefits factor should be used in the Revenues calculation. The escalation of benefits factor takes into account any fluctuation in the price levels of the goods/services that positively affect the attained revenues. In our financial analysis, escalation of benefits rate is set at 1.6% (plus 1% above the escalation of cost factor).

In addition, please note that below additional financial assumptions also apply to the financial analysis of the three CloudSat case scenarios. Due to the SW-based configuration and virtualization attributes, costs of

(a) Maintenance/Repair/Replacement and

(b) Property Management/Operations, Communications and Bills,

are reduced by 50% compared to a HW-based configuration. The reason, as has been already mentioned in the previous sections, is that software compared to hardware is less costly for maintenance, rarely requires replacement and overall is less costly, 'bulky' and power consuming. This 50% reduction (rounded) has been based in the use of the NFV Cost Reduction Rate which was previously analyzed.

Cost Escalation Factor	0.60%	(0,50 above Eurozone inflation rate)
Benefit Escalation Factor	1.60%	(1 unit above cost escalation factor)
Corporate Tax Rate (Europe Average 2015)	25.90%	Europe Average

Discount Rate	5.00%	Normal Case Scenario
Interest Rate for Loans	2.20%	(Euribor plus 2%)
Maintenance Repair/Replacement Expenses	1.50% of CAPEX	Reduced by 50% compared to a HW-based configuration where it is 3% of CAPEX
Annual Increase	10.00%	
Marketing/Advertising/ Promotion Expenses	4.00% of CAPEX	
Annual Decrease	50.00%	
Property Management, Operations, Communication and Bill Expenses	0.50% of CAPEX	Reduced by 50% compared to a HW-based configuration where it is 1% of CAPEX
Revenues from add-on services (SDN/NFV)	15%	Upon scenario's revenues

Figure 195. CFM financial analysis case assumptions

Overall, high level analysis will be based on the estimated CAPEX and OPEX of each CloudSat case scenario as well as on the estimated Revenues over the 5-year period. CAPEX, OPEX and Revenues figures are adjusted to the CloudSat case scenarios specifics (eg different components) but overall they remain unaffected by the applied business and market environment scenarios/conditions (normal, optimistic and pessimistic) and only specific variables will be altered per case so as the corresponding business and market conditions to be simulated per case. Similarly, depreciation and rest variable and rate figures remain unchanged for all cases, unless otherwise stated. In specific, for each of the evaluation CloudSat case scenarios/conditions (normal, optimistic and pessimistic) the following variables will be only re-adjusted:

- Annual Revenues (by modifying the % of annual Revenues increase depending on the CloudSat case scenario)
- Discount rate (for modifying the ratios affected by Discounted Cash Flow and NPV)

All the rest variables and financial analysis case assumptions will be considered as *ceteris paribus* case unless otherwise stated.

In specific, per scenario the following rates apply (where N: Normal, O: Optimistic and P: Pessimistic, and for Annual revenues the percentage refers to Annual revenues change).

Variables	Scenarios #1, #2, #3		
	N	O	P
Annual Revenues (annual percentage change)	12%	20%	4%
Discount Rate	5%	2%	10%

Figure 196. Rates per CloudSat scenario and per three environment cases

Summarizing all the estimated figures of CAPEX, OPEX and Revenues, for each CloudSat case scenario, and applying the CFM model of analysis, we perform the cash flow high level financial analysis and calculate the corresponding financial ratios per case.

7.2.3.1. Scenario #1 - GEO - Hybrid media distribution network as-a-Service

Initial cost investment (CAPEX) components remain mainly similar to the ones described (GEO SW-based virtualization deployment) apart from the use of two additional configuration variant components (see below figure).

CloudSat Scenario #1 satellite hub infrastructure	Cost	Market Research References
Outdoor unit (eg Antenna, RF front-end etc)	70,000.00 €	http://www.satsig.net/ivsatcos.htm
Modulator	120,000.00 €	
Encapsulator/Multiplexer		
Return Link sub-system		
GW Management & Access Control		
Satellite Hub Generic Server HW (2 servers)	10,000.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0 IBM report “A Deep Dive Into The Cost Benefits of KVM and Open Virtualization”
Satellite Hub Generic Server Software Configuration	3,500.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0
Terrestrial interface subsystem (firewall) 5 year license	17,500.00 €	http://www.mcafeeworks.com/Firewall-Enterprise-S1104.asp
Terrestrial interface subsystem (PEP)	2,500.00 €	http://www.satcomresources.com/iDirect-SkyCelerator-Network-Accelerator-Model-1150

Terrestrial interface subsystem (router)	10,000.00 €	http://www.router-switch.com/Price-cisco-routers-cisco-router-3800-series_c40
NFVI PoP SW + License + Configuration	3,500.00 €	Hewlett-Packard white paper (2014) "The Reality of Cost Reduction"
NFVI PoP HW (5 servers)	25,000.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0 IBM report "A Deep Dive Into The Cost Benefits of KVM and Open Virtualization"
CloudSat Scenario #1 Configuration Variant Components		
SDN Switches (Two units)	20,000.00 €	http://store.netgate.com/Pica8--C188.aspx
Federator	60,000.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0
TOTAL (actual)	**342,000.00 €**	
Bank Loan paid in 5 years with annual constant installments and constant/fixed interest rate	**-200,000.00 €**	
TOTAL (accounting)	**142,000.00 €**	

Figure 197. CloudSat case scenario #1: CAPEX Estimation

Below Figure summarizes the OPEX expenses of CloudSat case scenario 1.

OPEX	Y0	Y1	Y2	Y3	Y4	Y5
Satellite bandwidth leasing cost (1 Transponder)	0 €	-1,100,000.00 €	-1,100,000.00 €	-1,100,000.00 €	-1,100,000.00 €	-1,100,000.00 €
Maintenance and Repair/Replacement/Unforeseen expenses (1.5% of CAPEX increased by 10% annually)	0 €	**-5,130.00 €**	**-5,643.00 €**	**-6,207.30 €**	**-6,828.03 €**	**-7,510.83 €**
Marketing/Advertising/ Promotion expenses (4% of CAPEX reduced by 50% annually)	0 €	**-13,680.00 €**	**-6,840.00 €**	**-3,420.00 €**	**-1,710.00 €**	**-855.00 €**
Property management, operations, communication and bill/utilities expenses (0.5% of CAPEX)	0 €	**-1,710.00 €**	**-1,710.00 €**	**-1,710.00 €**	**-1,710.00 €**	**-1,710.00 €**
Supply expenses	0 €	**-4,000.00 €**	**-4,000.00 €**	**-4,000.00 €**	**-4,000.00 €**	**-4,000.00 €**
Loan Payments		**-40,000.00 €**	**-40,000.00 €**	**-40,000.00 €**	**-40,000.00 €**	**-40,000.00 €**
Interest Rate Expense		**-4,400.00 €**	**-3,557.88 €**	**-2,697.23 €**	**-1,817.65 €**	**-918.71 €**

Escalation of Costs	0 €	-7,013.52 €	-13,982.83 €	-20,969.94 €	-27,996.29 €	-35,068.14 €
TOTAL	0 €	-1,175,933.52 €	-1,175,733.71 €	-1,179,004.47 €	-1,184,061.96 €	-1,190,062.68 €

Figure 198. CloudSat case scenario #1: OPEX Estimation

Both CAPEX and OPEX figures remain unchanged during the evaluation of this scenario under the Normal, Optimistic and Pessimistic cases. In addition, as already stated in the assumptions section, as source of Revenues we consider the earnings from the annual utilization of the transponder, that is the revenues from the subscription fees for the satcom service (unless otherwise stated in the scenario).

At this point additional assumptions are made concerning the transponder's characteristics and utilization (references are provided for the validity of information) being used as the main source of revenues (along with revenues from add-on services SDN/NFV and the escalation of benefits factor).

ASSUMPTIONS	FIGURES	DETAILS	REFERENCES
Transponder Capacity	36 MHz		http://spacenews.com/rising-transponder-prices-mask-regional-disparity/ http://aias.iit.demokritos.gr/imosan/deliverables/IMOSAN_D19I_Report_on_Business_Plan.pdf http://www.satsig.net/ivsatcos.htm http://www.itso.int/images/stories/Capacity-Building/Dakar-2015/Satellite-Economics-beyond-the-cost-per-MHz.pdf http://www.globalcomsatphone.com/hughesnet/satellite/costs.html
Annual revenues of 1 MHz Utilization	90,000€		
Initial (Y1) Transponder utilization rate	30%	Utilization during the first year (Y1) of operation	
Annual Transponder utilization rate of change	12%	This is an average rate taking into account both the potential loss of customers percentage plus the newcomers/return percentage over the specific year (eg 18% newcomers/return and 6% loss of customers). This rate change applies to the utilization figure of the previous year.	

Figure 199. Transponder's characteristics and utilization assumptions

a) Scenario #1 – "Normal" case analysis

Under Normal case conditions, the two conditional variables are assumed to have the below values:

Variables	Scenario #1
	NORMAL
Annual Revenues (annual percentage change)	12%
Discount Rate	5%

Figure 200. CloudSat scenario #1: Normal case variables

Consequently the transponder utilization assumptions are:

ASSUMPTIONS	FIGURES
Transponder Capacity	36 MHz
Annual revenues of 1 MHz Utilization	90,000€
Initial (Y1) Transponder utilization rate	30%
Annual Transponder utilization rate of change	12%

Figure 201. CloudSat scenario #1 Normal case - Transponder utilization

The transponder utilization revenues are:

	Y1	Y2	Y3	Y4	Y5
Revenues per year	972,000.00 €	1,089,000.00 €	1,219,500.00 €	1,366,200.00 €	1,530,000.00 €
Utilization of Transponder (MHz)	10.80	12.10	13.55	15.18	17.00
% Utilization of Transponder (MHz)	30%	34%	38%	42%	47%

Figure 202. CloudSat scenario #1 Normal case- Transponder utilization Revenues

Therefore, the total revenues are calculated as per below figure.

	Y1	Y2	Y3	Y4	Y5
Revenues from transponder usage (basic satcom service fees)	972,000.00 €	1,089,000.00 €	1,219,500.00 €	1,366,200.00 €	1,530,000.00 €
Revenues from add-on services (SDN/NFV) as 15% of Revenues from transponder	145,800.00 €	163,350.00 €	182,925.00 €	204,930.00 €	229,500.00 €
Escalation of Benefits	17,884.80 €	40,395.80 €	68,399.21 €	102,991.42 €	145,336.97 €
TOTAL	**1,135,684.80 €**	**1,292,745.80 €**	**1,470,824.21 €**	**1,674,121.42 €**	**1,904,836.97 €**

Figure 203. CloudSat scenario #1 Normal case - Revenues

Summarizing all the estimated figures of CAPEX, OPEX and Revenues for CloudSat scenario #1 Normal case and applying the CFM model of analysis, we perform the cash flow financial analysis and calculate the financial ratios as per the figure on the next page. Please note that Discount rate 5% has been applied.

CloudSat Scenario #1 - Normal conditions of Evaluation									
	Y0	1	2	3	4	5	Totals		
CAPEX									
Outdoor unit (eg Antenna, RF front-end etc)	-70,000.00 €	0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	-70,000.00 €		
Modulator									
Encapsulator/Multiplexer	-120,000.00 €						-120,000.00 €		
Return Link sub-system									
GW Management & Access control									
Satellite Hub Generic Server HW (2 servers)	-10,000.00 €						-10,000.00 €		
Satellite Hub Generic Server Software Configuration	-3,500.00 €						-3,500.00 €		
Terrestrial interface subsystem (firewall) 5 year licence	-17,500.00 €						-17,500.00 €		
Terrestrial interface subsystem (PEP)	-2,500.00 €						-2,500.00 €		
Terrestrial interface subsystem (router)	-10,000.00 €						-10,000.00 €		
NFVI PoP SW+Licence+Configuration	-3,500.00 €						-3,500.00 €		
NFVI PoP HW (5 servers)	-25,000.00 €						-25,000.00 €		
SDN Switches (Two units)	-20,000.00 €						-20,000.00 €		
Federator	-60,000.00 €								
Bank Loan paid in 5 years with annual constant installments and constant/fixed interest rate	200,000.00 €						200,000.00 €		
Total CAPEX (for accounting)	**-142,000.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**-142,000.00 €**		
Actual CAPEX (excluding bank loan)	**-342,000.00 €**								
OPEX (VC Operating and Maintenance Costs)									
Satellite bandwidth leasing cost (1 Transponder)		-1,100,000.00 €	-1,100,000.00 €	-1,100,000.00 €	-1,100,000.00 €	-1,100,000.00 €	-5,500,000.00 €		
Maintenance and Repair/Replacement/Unforeseen expenses (1.5% of CAPEX increased by 10% annually)		-5,130.00 €	-5,643.00 €	-6,207.30 €	-6,828.03 €	-7,510.83 €	-31,319.16 €		
Marketing/Advertising/Promotion expenses (4% of CAPEX reduced by 50% annually)		-13,680.00 €	-6,840.00 €	-3,420.00 €	-1,710.00 €	-855.00 €	-26,505.00 €		
Property management, operations, communications and bill/utilities expenses (0.5% of CAPEX)		-1,710.00 €	-1,710.00 €	-1,710.00 €	-1,710.00 €	-1,710.00 €	-8,550.00 €		
Supply expenses		-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-20,000.00 €		
Loan Payments		-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-200,000.00 €		
Interest Rate Expense		-4,400.00 €	-3,557.88 €	-2,697.23 €	-1,817.65 €	-918.71 €	-13,391.46 €		
Escalation of Costs		-7,013.52 €	-13,982.83 €	-20,969.94 €	-27,996.29 €	-35,068.14 €	-105,030.72 €		
Total Costs	**0.00 €**	**-1,175,933.52 €**	**-1,175,733.71 €**	**-1,179,004.47 €**	**-1,184,061.96 €**	**-1,190,062.68 €**	**-5,904,796.34 €**		
Revenue and Operating Benefits									
Revenues from transponder usage		972,000.00 €	1,089,000.00 €	1,219,500.00 €	1,366,200.00 €	1,530,000.00 €	6,176,700.00 €		
Revenues from add-on services (SDN/NFV) as 15% of Revenues from transponder		145,800.00 €	163,350.00 €	182,925.00 €	204,930.00 €	229,500.00 €	926,505.00 €		
Escalation of Benefits		17,884.80 €	40,395.80 €	68,399.21 €	102,991.42 €	145,336.97 €	375,008.20 €		
Total Benefits and Revenue	**0.00 €**	**1,135,684.80 €**	**1,292,745.80 €**	**1,470,824.21 €**	**1,674,121.42 €**	**1,904,836.97 €**	**7,478,213.20 €**		
Contribution Margin % (Sales - VC/Sales)		**-3.54%**	**9.05%**	**19.84%**	**29.27%**	**37.52%**	**21.04%**		
Contribution Margin Amt (Sales - VC)		**-40,248.72 €**	**117,012.09 €**	**291,819.74 €**	**490,059.46 €**	**714,774.29 €**	**1,573,416.85 €**		
Cash Flow Before Taxes	**-142,000.00 €**	**-40,248.72 €**	**117,012.09 €**	**291,819.74 €**	**490,059.46 €**	**714,774.29 €**	**1,431,416.85 €**	INCLUDING CAPEX (Y0)	
EBITDA		**-44,648.72 €**	**113,454.21 €**	**289,122.51 €**	**488,241.81 €**	**713,855.58 €**	**1,560,025.39 €**	EXCLUDING CAPEX (Y0)	
Breakeven Year			*						
Income Tax Calculation									
Depreciation Expense		-28,400.00 €	-45,440.00 €	-27,264.00 €	-16,358.40 €	-16,358.40 €	-133,820.80 €		
Operating and Maintenance Cost (OPEX)		-1,175,933.52 €	-1,175,733.71 €	-1,179,004.47 €	-1,184,061.96 €	-1,190,062.68 €	-5,904,796.34 €		
Revenue and Operating Benefits		1,135,684.80 €	1,292,745.80 €	1,470,824.21 €	1,674,121.42 €	1,904,836.97 €	7,478,213.20 €		
Net Income Taxes	**0.00 €**	**17,780.02 €**	**-18,537.17 €**	**-68,519.94 €**	**-122,688.57 €**	**-180,889.71 €**	**-372,855.38 €**		
Cash Flow After Taxes	**-142,000.00 €**	**-22,468.70 €**	**98,474.92 €**	**223,299.80 €**	**367,370.88 €**	**533,884.57 €**	**1,058,561.48 €**	INCLUDING CAPEX (Y0)	
Discounted Cash Flow (After Tax)	**-142,000.00 €**	**-21,398.76 €**	**89,319.65 €**	**192,894.77 €**	**302,236.94 €**	**418,312.53 €**	**839,365.12 €**	INCLUDING CAPEX (Y0)	

Business Case Results:	
NPV of Cash Flow	**839,365.12 €**
IRR	**71.9%**
Profitability Index	**6.91**
Simple Payback	**2 Years 4 Months**
Discounted Payback	**2 Years 5 Months**
ROI	**26%**

Case Assumptions:		
Cost Escalation Factor	**0.60%**	(0,50 above Eurozone inflation rate)
Benefit Escalation Factor	**1.60%**	(1 above cost escalation factor)
Corporate Tax Rate (Europe Average 2015)	**25.90%**	
Discount Rate	**5.00%**	
Interest Rate for Loans	**2.20%**	(Euribor plus 2%)
Maintenance Repair/Replacement Exp	**1.50%**	of CAPEX
Annual Increase	**10.00%**	
Marketing/Advertising/Promotion Exp	**4.00%**	of CAPEX
Annual Decrease	**50.00%**	
Property Mng, operations and bill Exp	**0.50%**	of CAPEX
Revenues from add-on services (SDN/NFV)	**15.00%**	

CAPEX Costs are NOT expensed in Y0 but CAPITALIZED in Y1-Y5 (Depreciation Expense)

Figure 204. CloudSat scenario #1 Normal case: CFM Financial analysis and financial ratios

Under normal financial, business and market conditions, CloudSat scenario #1 is a viable solution with appropriate financial results and ratios within a 5-years period. Findings are summarized in below figure:

Financial variables and ratios	Values
NPV of Cash Flow	839,365.12 €
IRR	71.90%
Profitability Index	6.91
Simple Payback	2 Years 4 Months
Discounted Payback	2 Years 5 Months
ROI	26%
Breakeven Year (excluding CAPEX)	Middle Year 2

Figure 205. CloudSat scenario #1 Normal case: Financial variables and ratios results

This investment becomes positive in cash flows within year 2 (Breakeven Year) and after 2 years and 5 months reaches its Payback point (having make up for CAPEX and OPEX so far). The discounted cash flow after taxes at the end of 5th year has a NPV of 839,365.12€ (total profit of the 5 years at present value). This means that the investment is profitable within the 5 years period with an IRR 71.90%, Profitability index 6.91 and a ROI of 26% upon initial investment and variable costs.

b) Scenario #1 – "Optimistic" case analysis

Under Optimistic case conditions, the two conditional variables are assumed to have the below values:

Variables	Scenario #1
	OPTIMISTIC
Annual Revenues (annual percentage change)	20%
Discount Rate	2%

Figure 206. CloudSat scenario #1: Optimistic case variables

Consequently the transponder utilization assumptions are:

ASSUMPTIONS	FIGURES
Transponder Capacity	36 MHz
Annual revenues of 1 MHz Utilization	90,000€
Initial (Y1) Transponder utilization rate	30%
Annual Transponder utilization rate of change	20%

Figure 207. CloudSat scenario #1 Optimistic case - Transponder utilization

The transponder utilization revenues are:

	Y1	Y2	Y3	Y4	Y5
Revenues per year	972,000.00 €	1,166,400.00 €	1,399,500.00 €	1,679,400.00 €	2,015,100.00 €
Utilization of Transponder (MHz)	10.80	12.96	15.55	18.66	22.39
% Utilization of Transponder (MHz)	30%	36%	43%	52%	62%

Figure 208. CloudSat scenario #1 Optimistic case- Transponder utilization Revenues

Therefore, the total revenues are calculated as per below figure.

	Y1	Y2	Y3	Y4	Y5
Revenues from transponder usage (basic satcom service fees)	972,000.00 €	1,166,400.00 €	1,399,500.00 €	1,679,400.00 €	2,015,100.00 €
Revenues from add-on services (SDN/NFV) as 15% of Revenues from transponder	145,800.00 €	174,960.00 €	209,925.00 €	251,910.00 €	302,265.00 €
Escalation of Benefits	17,884.80 €	43,266.91 €	78,495.03 €	126,602.10 €	191,417.34 €
TOTAL	**1,135,684.80 €**	**1,384,626.91 €**	**1,687,920.03 €**	**2,057,912.10 €**	**2,508,782.34 €**

Figure 209. CloudSat scenario #1 Optimistic case - Revenues

Summarizing all the estimated figures of CAPEX, OPEX and Revenues for CloudSat scenario #1 Optimistic case and applying the CFM model of analysis, we perform the cash flow financial analysis and calculate the financial ratios as per the figure on the next page. Please note that Discount rate 2% has been applied.

CloudSat Scenario #1 - Optimistic conditions of Evaluation								
	Y0	1	2	3	4	5	Totals	
CAPEX								
Outdoor unit (eg Antenna, RF front-end etc)	-70,000.00 €	0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	-70,000.00 €	
Modulator								
Encapsulator/Multiplexer								
Return Link sub-system	-120,000.00 €						-120,000.00 €	
GW Management & Access control								
Satellite Hub Generic Server HW (2 servers)	-10,000.00 €						-10,000.00 €	
Satellite Hub Generic Server Software Configuration	-3,500.00 €						-3,500.00 €	
Terrestrial interface subsystem (firewall) 5 year licence	-17,500.00 €						-17,500.00 €	
Terrestrial interface subsystem (PEP)	-2,500.00 €						-2,500.00 €	
Terrestrial interface subsystem (router)	-10,000.00 €						-10,000.00 €	
NFVI PoP SW+Licence+Configuration	-3,500.00 €						-3,500.00 €	
NFVI PoP HW (5 servers)	-25,000.00 €						-25,000.00 €	
SDN Switches (Two units)	-20,000.00 €						-20,000.00 €	
Federator	-60,000.00 €							
Bank Loan paid in 5 years with annual constant installments and constant/fixed interest rate	200,000.00 €						200,000.00 €	
Total CAPEX (for accounting)	**-142,000.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**-142,000.00 €**	
Actual CAPEX (excluding bank loan)	**-342,000.00 €**							
OPEX (VC Operating and Maintenance Costs)								
Satellite bandwidth leasing cost (1 Transponder)		-1,100,000.00 €	-1,100,000.00 €	-1,100,000.00 €	-1,100,000.00 €	-1,100,000.00 €	-5,500,000.00 €	
Maintenance and Repair/Replacement/Unforeseen expenses (1.5% of CAPEX increased by 10% annually)		-5,130.00 €	-5,643.00 €	-6,207.30 €	-6,828.03 €	-7,510.83 €	-31,319.16 €	
Marketing/Advertising/Promotion expenses (4% of CAPEX reduced by 50% annually)		-13,680.00 €	-6,840.00 €	-3,420.00 €	-1,710.00 €	-855.00 €	-26,505.00 €	
Property management, operations, communication and bill/utilities expenses (0.5% of CAPEX)		-1,710.00 €	-1,710.00 €	-1,710.00 €	-1,710.00 €	-1,710.00 €	-8,550.00 €	
Supply expenses		-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-20,000.00 €	
Loan Payments		-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-200,000.00 €	
Interest Rate Expense		-4,400.00 €	-3,557.88 €	-2,697.23 €	-1,817.65 €	-918.71 €	-13,391.46 €	
Escalation of Costs		-7,013.52 €	-13,982.83 €	-20,969.94 €	-27,996.29 €	-35,068.14 €	-105,030.72 €	
Total Costs	**0.00 €**	**-1,175,933.52 €**	**-1,175,733.71 €**	**-1,179,004.47 €**	**-1,184,061.96 €**	**-1,190,062.68 €**	**-5,904,796.34 €**	
Revenue and Operating Benefits								
Revenues from transponder usage		972,000.00 €	1,166,400.00 €	1,399,500.00 €	1,679,400.00 €	2,015,100.00 €	7,232,400.00 €	
Revenues from add-on services (SDN/NFV) as 15% of Revenues from transponder		145,800.00 €	174,960.00 €	209,925.00 €	251,910.00 €	302,265.00 €	1,084,860.00 €	
Escalation of Benefits		17,884.80 €	43,266.91 €	78,495.03 €	126,602.10 €	191,417.34 €	457,666.18 €	
Total Benefits and Revenue	**0.00 €**	**1,135,684.80 €**	**1,384,626.91 €**	**1,687,920.03 €**	**2,057,912.10 €**	**2,508,782.34 €**	**8,774,926.18 €**	
Contribution Margin % (Sales - VC/Sales)		**-3.54%**	**15.09%**	**30.15%**	**42.46%**	**52.56%**	**32.71%**	
Contribution Margin Amt (Sales - VC)		**-40,248.72 €**	**208,893.20 €**	**508,915.56 €**	**873,850.14 €**	**1,318,719.66 €**	**2,870,129.83 €**	
Cash Flow Before Taxes	**-142,000.00 €**	**-40,248.72 €**	**208,893.20 €**	**508,915.56 €**	**873,850.14 €**	**1,318,719.66 €**	**2,728,129.83 €**	INCLUDING CAPEX (Y0)
EBITDA		**-44,648.72 €**	**205,335.32 €**	**506,218.33 €**	**872,032.49 €**	**1,317,800.94 €**	**2,856,738.37 €**	EXCLUDING CAPEX (Y0)
Breakeven Year			*					
Income Tax Calculation								
Depreciation Expense		-28,400.00 €	-45,440.00 €	-27,264.00 €	-16,358.40 €	-16,358.40 €	-133,820.80 €	
Operating and Maintenance Cost (OPEX)		-1,175,933.52 €	-1,175,733.71 €	-1,179,004.47 €	-1,184,061.96 €	-1,190,062.68 €	-5,904,796.34 €	
Revenue and Operating Benefits		1,135,684.80 €	1,384,626.91 €	1,687,920.03 €	2,057,912.10 €	2,508,782.34 €	8,774,926.18 €	
Net Income Taxes	**0.00 €**	**17,780.02 €**	**-42,334.38 €**	**-124,747.75 €**	**-222,090.36 €**	**-337,311.57 €**	**-708,704.04 €**	
Cash Flow After Taxes	**-142,000.00 €**	**-22,468.70 €**	**166,558.82 €**	**384,167.81 €**	**651,759.78 €**	**981,408.09 €**	**2,019,425.79 €**	INCLUDING CAPEX (Y0)
Discounted Cash Flow (After Tax)	**-142,000.00 €**	**-22,028.14 €**	**160,091.14 €**	**362,009.91 €**	**602,125.29 €**	**888,891.54 €**	**1,849,089.74 €**	INCLUDING CAPEX (Y0)

Business Case Results:	
NPV of Cash Flow	**1,849,089.74 €**
IRR	**103.0%**
Profitability Index	**14.02**
Simple Payback	**1 Years 12 Months**
Discounted Payback	**2 Years 0 Months**
ROI	**47%**

Case Assumptions:		
Cost Escalation Factor	**0.60%**	(0,50 above Eurozone inflation rate)
Benefit Escalation Factor	**1.60%**	(1 above cost escalation factor)
Corporate Tax Rate (Europe Average 2015)	**25.90%**	
Discount Rate	**2.00%**	
Interest Rate for Loans	**2.20%**	(Euribor plus 2%)
Maintenance Repair/Replacement Exp	**1.50%**	of CAPEX
Annual Increase	**10.00%**	
Marketing/Advertising/Promotion Exp	**4.00%**	of CAPEX
Annual Decrease	**50.00%**	
Property Mng, operations and bill Exp	**0.50%**	of CAPEX
Revenues from add-on services (SDN/NFV)	**15.00%**	

CAPEX Costs are NOT expensed in Y0 but CAPITALIZED in Y1-Y5 (Depreciation Expense)

Figure 210. CloudSat scenario #1 Optimistic case: CFM Financial analysis and financial ratios

Under optimistic financial, business and market conditions, CloudSat scenario #1 is an extremely viable and profitable investment with very attractive financial figures within a 5-years period. Findings are summarized in below figure:

Financial variables and ratios	Values
NPV of Cash Flow	1,849,089.74 €
IRR	103.0%
Profitability Index	14.02
Simple Payback	1 Year 11-12 Months
Discounted Payback	2 Years 0 Months
ROI	47%
Breakeven Year (excluding CAPEX)	Early Year 2

Figure 211. CloudSat scenario #1 Optimistic case: Financial variables and ratios results

This investment is very promising since it becomes positive in cash flows within early year 2 (Breakeven Year) and after 2 years reaches its Payback point (having make up for CAPEX and OPEX so far). The discounted cash flow after taxes at the end of 5th year has a NPV of 1,849,089.33€ (total profit of the 5 years at present value). This means that the investment is profitable within the 5 years period with an IRR 103%, Profitability index 14.02 and a ROI of 47% upon initial investment, variable costs and operations.

c) Scenario #1 – "Pessimistic" case analysis

Under Pessimistic case conditions (more unstable economy and market/business environment), the two conditional variables are assumed to have the below values:

Variables	Scenario #1
	PESSIMISTIC
Annual Revenues (annual percentage change)	4%
Discount Rate	10%

Figure 212. CloudSat scenario #1: Pessimistic case variables

Consequently the transponder utilization assumptions are:

ASSUMPTIONS	FIGURES
Transponder Capacity	36 MHz
Annual revenues of 1 MHz Utilization	90,000€
Initial (Y1) Transponder utilization rate	30%
Annual Transponder utilization rate of change	4%

Figure 213. CloudSat scenario #1 Pessimistic case - Transponder utilization

The transponder utilization revenues are:

	Y1	Y2	Y3	Y4	Y5
Revenues per year	972,000.00 €	1,010,700.00 €	1,051,200.00 €	1,093,500.00 €	1,137,600.00 €
Utilization of Transponder (MHz)	10.80	11.23	11.68	12.15	12.64
% Utilization of Transponder (MHz)	30%	31%	32%	34%	35%

Figure 214. CloudSat scenario #1 Pessimistic case- Transponder utilization Revenues

Therefore, the total revenues are calculated as per below figure.

	Y1	Y2	Y3	Y4	Y5
Revenues from transponder usage (basic satcom service fees)	972,000.00 €	1,010,700.00 €	1,051,200.00 €	1,093,500.00 €	1,137,600.00 €
Revenues from add-on services (SDN/NFV) as 15% of Revenues from transponder	145,800.00 €	151,605.00 €	157,680.00 €	164,025.00 €	170,640.00 €
Escalation of Benefits	17,884.80 €	37,491.31 €	58,959.61 €	82,433.84 €	108,062.31 €
TOTAL	**1,135,684.80 €**	**1,199,796.31 €**	**1,267,839.61 €**	**1,339,958.84 €**	**1,416,302.31 €**

Figure 215. CloudSat scenario #1 Pessimistic case - Revenues

Summarizing all the estimated figures of CAPEX, OPEX and Revenues for CloudSat scenario #1 Pessimistic case and applying the CFM model of analysis, we perform the cash flow financial analysis and calculate the financial ratios as per the figure on the next page. Please note that Discount rate 10% has been applied.

CloudSat Scenario #1 - Pessimistic conditions of Evaluation									
	Y0	1	2	3	4	5	Totals		
CAPEX									
Outdoor unit (eg Antenna, RF front-end etc)	-70,000.00 €	0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	-70,000.00 €		
Modulator									
Encapsulator/Multiplexer	-120,000.00 €						-120,000.00 €		
Return Link sub-system									
GW Management & Access control									
Satellite Hub Generic Server HW (2 servers)	-10,000.00 €						-10,000.00 €		
Satellite Hub Generic Server Software Configuration	-3,500.00 €						-3,500.00 €		
Terrestrial interface subsystem (firewall) 5 year licence	-17,500.00 €						-17,500.00 €		
Terrestrial interface subsystem (PEP)	-2,500.00 €						-2,500.00 €		
Terrestrial interface subsystem (router)	-10,000.00 €						-10,000.00 €		
NFVI PoP SW+Licence+Configuration	-3,500.00 €						-3,500.00 €		
NFVI PoP HW (5 servers)	-25,000.00 €						-25,000.00 €		
SDN Switches (Two units)	-20,000.00 €						-20,000.00 €		
Federator	-60,000.00 €								
Bank Loan paid in 5 years with annual constant installments and constant/fixed interest rate	200,000.00 €						200,000.00 €		
Total CAPEX (for accounting)	**-142,000.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**-142,000.00 €**		
Actual CAPEX (excluding bank loan)	**-342,000.00 €**								
OPEX (VC Operating and Maintenance Costs)									
Satellite bandwidth leasing cost (1 Transponder)		-1,100,000.00 €	-1,100,000.00 €	-1,100,000.00 €	-1,100,000.00 €	-1,100,000.00 €	-5,500,000.00 €		
Maintenance and Repair/Replacement/Unforeseen expenses (1.5% of CAPEX increased by 10% annually)		-5,130.00 €	-5,643.00 €	-6,207.30 €	-6,828.03 €	-7,510.83 €	-31,319.16 €		
Marketing/Advertising/Promotion expenses (4% of CAPEX reduced by 50% annually)		-13,680.00 €	-6,840.00 €	-3,420.00 €	-1,710.00 €	-855.00 €	-26,505.00 €		
Property management, operations, communication and bill/utilities expenses (0.5% of CAPEX)		-1,710.00 €	-1,710.00 €	-1,710.00 €	-1,710.00 €	-1,710.00 €	-8,550.00 €		
Supply expenses		-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-20,000.00 €		
Loan Payments		-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-200,000.00 €		
Interest Rate Expense		-4,400.00 €	-3,557.88 €	-2,697.23 €	-1,817.65 €	-918.71 €	-13,391.46 €		
Escalation of Costs		-7,013.52 €	-13,982.83 €	-20,969.94 €	-27,996.29 €	-35,068.14 €	-105,030.72 €		
Total Costs	**0.00 €**	**-1,175,933.52 €**	**-1,175,733.71 €**	**-1,179,004.47 €**	**-1,184,061.96 €**	**-1,190,062.68 €**	**-5,904,796.34 €**		
Revenue and Operating Benefits									
Revenues from transponder usage		972,000.00 €	1,010,700.00 €	1,051,200.00 €	1,093,500.00 €	1,137,600.00 €	5,265,000.00 €		
Revenues from add-on services (SDN/NFV) as 15% of Revenues from transponder		145,800.00 €	151,605.00 €	157,680.00 €	164,025.00 €	170,640.00 €	789,750.00 €		
Escalation of Benefits		17,884.80 €	37,491.31 €	58,959.61 €	82,433.84 €	108,062.31 €	304,831.88 €		
Total Benefits and Revenue	**0.00 €**	**1,135,684.80 €**	**1,199,796.31 €**	**1,267,839.61 €**	**1,339,958.84 €**	**1,416,302.31 €**	**6,359,581.88 €**		
Contribution Margin % (Sales - VC/Sales)		**-3.54%**	**2.01%**	**7.01%**	**11.63%**	**15.97%**	**7.15%**		
Contribution Margin Amt (Sales - VC)		**-40,248.72 €**	**24,062.60 €**	**88,835.14 €**	**155,896.88 €**	**226,239.63 €**	**454,785.53 €**		
Cash Flow Before Taxes	**-142,000.00 €**	**-40,248.72 €**	**24,062.60 €**	**88,835.14 €**	**155,896.88 €**	**226,239.63 €**	**312,785.53 €**	INCLUDING CAPEX (Y0)	
EBITDA		**-44,648.72 €**	**20,504.72 €**	**86,137.92 €**	**154,079.24 €**	**225,320.92 €**	**441,394.07 €**	EXCLUDING CAPEX (Y0)	
Breakeven Year				*					
Income Tax Calculation									
Depreciation Expense		-28,400.00 €	-45,440.00 €	-27,264.00 €	-16,358.40 €	-16,358.40 €	-133,820.80 €		
Operating and Maintenance Cost (OPEX)		-1,175,933.52 €	-1,175,733.71 €	-1,179,004.47 €	-1,184,061.96 €	-1,190,062.68 €	-5,904,796.34 €		
Revenue and Operating Benefits		1,135,684.80 €	1,199,796.31 €	1,267,839.61 €	1,339,958.84 €	1,416,302.31 €	6,359,581.88 €		
Net Income Taxes	**0.00 €**	**17,780.02 €**	**5,536.75 €**	**-15,946.93 €**	**-36,140.47 €**	**-54,359.24 €**	**-83,129.87 €**		
Cash Flow After Taxes	**-142,000.00 €**	**-22,468.70 €**	**29,599.35 €**	**72,888.22 €**	**119,756.42 €**	**171,880.39 €**	**229,655.67 €**	INCLUDING CAPEX (Y0)	
Discounted Cash Flow (After Tax)	**-142,000.00 €**	**-20,426.09 €**	**24,462.27 €**	**54,762.00 €**	**81,795.24 €**	**106,724.20 €**	**105,317.62 €**	INCLUDING CAPEX (Y0)	
Business Case Results:				**Case Assumptions:**					
NPV of Cash Flow	**105,317.62 €**			**Cost Escalation Factor**		**0.60%**	(0,50 above Eurozone inflation rate)		
IRR	**25.3%**			**Benefit Escalation Factor**		**1.60%**	(1 above cost escalation factor)		
Profitability Index	**1.74**			**Corporate Tax Rate (Europe Average 2015)**		**25.90%**			
Simple Payback	**3 Years 6 Months**			**Discount Rate**		**10.00%**			
Discounted Payback	**4 Years 0 Months**			**Interest Rate for Loans**		**2.20%**	(Euribor plus 2%)		
ROI	**7%**			**Maintenance Repair/Replacement Exp**		**1.50%**	of CAPEX		
				Annual Increase		**10.00%**			
				Marketing/Advertising/Promotion Exp		**4.00%**	of CAPEX		
				Annual Decrease		**50.00%**			
				Property Mng, operations and bill Exp		**0.50%**	of CAPEX		
				Revenues from add-on services (SDN/NFV)		**15.00%**			
				CAPEX Costs are NOT expensed in Y0 but CAPITALIZED in Y1-Y5 (Depreciation Expense)					

Figure 216. CloudSat scenario #1 Pessimistic case: CFM Financial analysis and financial ratios

Under pessimistic financial, business and market conditions, CloudSat scenario #1 remains still viable but later in time (close to 4th year). Findings are summarized in below figure:

Financial variables and ratios	Values
NPV of Cash Flow	105,317.62 €
IRR	25.3%
Profitability Index	1.74
Simple Payback	3 Years 6 Months
Discounted Payback	4 Years 0 Months
ROI	7%
Breakeven Year (excluding CAPEX)	Late Year 3

Figure 217. CloudSat scenario #1 Pessimistic case: Financial variables and ratios results

This investment has reached a marginal profitability within the end of year 3. In specific, it acquires positive cash flows within late year 3 (Breakeven Year) and after 4 years it reaches its Payback point (having make up for CAPEX and OPEX so far).

This investment is still viable but by the end of the 5-year period of analysis. The discounted cash flow after taxes at the end of 5th year has a NPV of 105,317.62€ (total profit of the 5 years at present value). This means that the investment is profitable within the 5 years period with an IRR 25.3% , Profitability index 1.74 and a ROI of 7% upon initial investment, variable costs and operations. To sum up, even under pessimistic conditions, we consider this case viable too, with slow pace of profitability/turnover, provided that its life is extended to year 4 and further.

7.2.3.2. Scenario #2 - HTS/GEO - Dynamic backhauling with edge processing

Initial cost investment (CAPEX) components remain mainly similar to the ones described (GEO SW-based virtualization deployment) apart from the use of four additional configuration variant components (Figure 218).

CloudSat Scenario #2 satellite hub infrastructure	Cost	Market Research References
Outdoor units	70,000.00 €	http://www.satsig.net/ivsatcos.htm
Modulator	120,000.00 €	
Encapsulator/Multiplexer		
Return Link sub-system		
GW Management & Access		

Control		
Satellite Hub Generic Server HW (2 servers)	10,000.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0 IBM report "A Deep Dive Into The Cost Benefits of KVM and Open Virtualization"
Satellite Hub Generic Server Software Configuration	3,500.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0
Terrestrial interface subsystem (firewall) 5 year license	17,500.00 €	http://www.mcafeeworks.com/Firewall-Enterprise-S1104.asp
Terrestrial interface subsystem (PEP)	2,500.00 €	http://www.satcomresources.com/iDirect-SkyCelerator-Network-Accelerator-Model-1150
Terrestrial interface subsystem (router)	10,000.00 €	http://www.router-switch.com/Price-cisco-routers-cisco-router-3800-series_c40
NFVI PoP SW + License + Configuration	3,500.00 €	Hewlett-Packard white paper (2014) "The Reality of Cost Reduction"
NFVI PoP HW (5 servers)	25,000.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0 IBM report "A Deep Dive Into The Cost Benefits of KVM and Open Virtualization"
CloudSat Scenario #2 Configuration Variant Components		
SDN Switches (Two units)	20,000.00 €	http://store.netgate.com/Pica8--C188.aspx
NFV compatible terminal (upgrade)	5,000.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0
VNF1 Transcoder	20,000.00 €	http://www.streamingmedia.com/Articles/Editorial/Featured-Articles/High-End-Video-Transcoder-Shootout-Elemental-Vs.-Telestream-100978.aspx
NFV Orchestrator (HW/SW)	30,000.00 €	
TOTAL (actual)	337,000.00 €	
Bank Loan paid in 5 years with annual constant installments and constant/fixed interest rate	-200,000.00 €	
TOTAL (accounting)	137,000.00 €	

Figure 218. CloudSat case scenario #2: CAPEX Estimation

Please note additional assumptions being valid for CloudSat Case Scenario #2 OPEX and Revenues calculation:

- CloudSat Case Scenario #2 is based on the HTS/GEO satellite systems

- Satellite bandwidth leasing cost of 10MBps is included in OPEX (OPEX Satellite bandwidth leasing cost €/10Mbit/s/year equals to 2,000€x10Mbit/s x12months = 240,000€)
- Revenues originate from two sources: (a) earnings as percentage of Satellite bandwidth resale usage (€/10Mbit/s) per year (including all add-on provisional services) and (b) earnings from add-on services (SDN/NFV) calculated as fixed 15% percentage upon (a).
- Initial Satellite bandwidth resale revenues percentage is fixed at an initial 30% for Y1 and dynamic for the years Y2-Y5 (as per the conditions of each scenario)

Below Figure summarizes the OPEX expenses of CloudSat case scenario #2.

OPEX	Y0	Y1	Y2	Y3	Y4	Y5
Satellite bandwidth leasing cost (€/10Mbit/s/year)	0 €	-240,000.00 €	-240,000.00 €	-240,000.00 €	-240,000.00 €	-240,000.00 €
Maintenance and Repair/Replacement/Unforeseen expenses (1.5% of CAPEX increased by 10% annually)	0 €	-5,055.00 €	-5,560.50 €	-6,116.55 €	-6,728.21 €	-7,401.03 €
Marketing/Advertising/ Promotion expenses (4% of CAPEX reduced by 50% annually)	0 €	-13,480.00 €	-6,740.00 €	-3,370.00 €	-1,685.00 €	-842.50 €
Property management, operations and bill/utilities expenses (0.5% of CAPEX)	0 €	-1,685.00 €	-1,685.00 €	-1,685.00 €	-1,685.00 €	-1,685.00 €
Supply expenses	0 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €
Loan Payments		-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €
Interest Rate Expense		-4,400.00 €	-3,557.88 €	-2,697.23 €	-1,817.65 €	-918.71 €
Escalation of Costs	0 €	-1,851.72 €	-3,629.38 €	-5,393.87 €	-7,166.15 €	-8,952.20 €
TOTAL	**0 €**	**-310,471.72 €**	**-305,172.75 €**	**-303,262.65 €**	**-303,082.00 €**	**-303,799.44 €**

Figure 219. CloudSat case scenario #2: OPEX Estimation

a) Scenario #2 – "Normal" case analysis

Under Normal case conditions, the two conditional variables are assumed to have the below values:

Variables	Scenario #2
	NORMAL
Annual Revenues (annual percentage change)	12%
Discount Rate	5%

Figure 220. CloudSat scenario #2: Normal case variables

Revenues originate from two sources:

(a) earnings as percentage of Satellite bandwidth resale usage (€/10Mbit/s) per year (including ALL add-on provisional services being offered through the leased bandwidth) reselling at 850,000€ annually and (b) earnings from add-on services (SDN/NFV) calculated as fixed 15% percentage upon (a). In specific:

ASSUMPTIONS	FIGURES
Revenues originating from leased Satellite bandwidth /10Mbit/s/year	
Estimated Annual Revenues for CloudSat Scenario 2 from reselling leased Satellite bandwidth of 10Mbit/s/year (including all add-on provisional services being offered)	€ 850,000.00
Initial (Y1) Revenues (% upon Estimated Annual Revenues)	30.00%
Annual Revenue increase rate change	12.00%

Figure 221. CloudSat scenario #2 Normal case – Satellite bandwidth utilization Revenues

The satellite bandwidth utilization revenues are:

	Y1	Y2	Y3	Y4	Y5
Revenues per year	255,000.00 €	285,600.00 €	319,872.00 €	358,256.64 €	401,247.44 €
% on Total Estimated Annual Revenues of CloudSat Scenario 2	30.00%	33.60%	37.63%	42.15%	47.21%

Figure 222. CloudSat scenario #2 Normal case- Satellite bandwidth utilization Revenues

Therefore, the total revenues are calculated as per below figure.

	Y1	Y2	Y3	Y4	Y5
Revenues	255,000.00 €	285,600.00 €	319,872.00 €	358,256.64 €	401,247.44 €
Revenues from add-on services (SDN/NFV) as 15% of Revenues	38,250.00 €	42,840.00 €	47,980.80 €	53,738.50 €	60,187.12 €
Escalation of Benefits	4,692.00 €	10,594.16 €	17,940.95 €	27,007.29 €	38,115.09 €
TOTAL	**297,942.00 €**	**339,034.16 €**	**385,793.75 €**	**439,002.43 €**	**499,549.64 €**

Figure 223. CloudSat scenario #2 Normal case - Revenues

Summarizing all the estimated figures of CAPEX, OPEX and Revenues for CloudSat scenario #2 Normal case and applying the CFM model of analysis, we perform the cash flow financial analysis and calculate the financial ratios as per the figure on the next page. Please note that Discount rate 5% has been applied.

CloudSat Scenario #2 - Normal conditions of Evaluation								
	Y0	1	2	3	4	5	Totals	
CAPEX								
Outdoor units	-70,000.00 €	0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	-70,000.00 €	
Modulator								
Encapsulator/Multiplexer	-120,000.00 €						-120,000.00 €	
Return Link sub-system								
GW Management & Access control								
Satellite Hub Generic Server HW (2 servers)	-10,000.00 €						-10,000.00 €	
Satellite Hub Generic Server Software Configuration	-3,500.00 €						-3,500.00 €	
Terrestrial interface subsystem (firewall) 5 year licence	-17,500.00 €						-17,500.00 €	
Terrestrial interface subsystem (PEP)	-2,500.00 €						-2,500.00 €	
Terrestrial interface subsystem (router)	-10,000.00 €						-10,000.00 €	
NFVI PoP SW+Licence+Configuration	-3,500.00 €						-3,500.00 €	
NFVI PoP HW (5 servers)	-25,000.00 €						-25,000.00 €	
SDN Switches (Two units)	-20,000.00 €						-20,000.00 €	
NFV compatible terminal (upgrade)	-5,000.00 €						-5,000.00 €	
VNF1 Transcoder	-20,000.00 €						-20,000.00 €	
NFV Orchestrator (HW/SW)	-30,000.00 €						-30,000.00 €	
Bank Loan paid in 5 years with annual constant installments and constant/fixed interest rate	200,000.00 €						200,000.00 €	
Total CAPEX (for accounting)	**-137,000.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**-137,000.00 €**	
Actual CAPEX (excluding bank loan)	**-337,000.00 €**							
OPEX (VC Operating and Maintenance Costs)								
Satellite bandwidth leasing cost (€/10Mbit/s/year)		-240,000.00 €	-240,000.00 €	-240,000.00 €	-240,000.00 €	-240,000.00 €	-1,200,000.00 €	
Maintenance and Repair/Replacement/Unforeseen expenses (1.5% of CAPEX increased by 10% annually)		-5,055.00 €	-5,560.50 €	-6,116.55 €	-6,728.21 €	-7,401.03 €	-30,861.28 €	
Marketing/Advertising/Promotion expenses (4% of CAPEX reduced by 50% annually)		-13,480.00 €	-6,740.00 €	-3,370.00 €	-1,685.00 €	-842.50 €	-26,117.50 €	
Property management, operations, communications and bill/utilities expenses (0.5% of CAPEX)		-1,685.00 €	-1,685.00 €	-1,685.00 €	-1,685.00 €	-1,685.00 €	-8,425.00 €	
Supply expenses		-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-20,000.00 €	
Loan Payments		-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-200,000.00 €	
Interest Rate Expense		-4,400.00 €	-3,557.88 €	-2,697.23 €	-1,817.65 €	-918.71 €	-13,391.46 €	
Escalation of Costs		-1,851.72 €	-3,629.38 €	-5,393.87 €	-7,166.15 €	-8,952.20 €	-26,993.32 €	
Total Costs	**0.00 €**	**-310,471.72 €**	**-305,172.75 €**	**-303,262.65 €**	**-303,082.00 €**	**-303,799.44 €**	**-1,525,788.57 €**	
Revenue and Operating Benefits								
Revenues		255,000.00 €	285,600.00 €	319,872.00 €	358,256.64 €	401,247.44 €	1,619,976.08 €	
Revenues from add-on services (SDN/NFV) as 15% of Revenues		38,250.00 €	42,840.00 €	47,980.80 €	53,738.50 €	60,187.12 €	242,996.41 €	
Escalation of Benefits		4,692.00 €	10,594.16 €	17,940.95 €	27,007.29 €	38,115.09 €	98,349.49 €	
Total Benefits and Revenue	**0.00 €**	**297,942.00 €**	**339,034.16 €**	**385,793.75 €**	**439,002.43 €**	**499,549.64 €**	**1,961,321.98 €**	
Contribution Margin % (Sales - VC/Sales)		**-4.21%**	**9.99%**	**21.39%**	**30.96%**	**39.19%**	**22.21%**	
Contribution Margin Amt (Sales - VC)		**-12,529.72 €**	**33,861.41 €**	**82,531.10 €**	**135,920.42 €**	**195,750.20 €**	**435,533.41 €**	
Cash Flow Before Taxes	**-137,000.00 €**	**-12,529.72 €**	**33,861.41 €**	**82,531.10 €**	**135,920.42 €**	**195,750.20 €**	**298,533.41 €**	INCLUDING CAPEX (Y0)
EBITDA		**-16,929.72 €**	**30,303.53 €**	**79,833.87 €**	**134,102.78 €**	**194,831.49 €**	**422,141.95 €**	EXCLUDING CAPEX (Y0)
Breakeven Year			*					
Income Tax Calculation								
Depreciation Expense		-27,400.00 €	-43,840.00 €	-26,304.00 €	-15,782.40 €	-15,782.40 €	-129,108.80 €	
Operating and Maintenance Cost (OPEX)		-310,471.72 €	-305,172.75 €	-303,262.65 €	-303,082.00 €	-303,799.44 €	-1,525,788.57 €	
Revenue and Operating Benefits		297,942.00 €	339,034.16 €	385,793.75 €	439,002.43 €	499,549.64 €	1,961,321.98 €	
Net Income Taxes	**0.00 €**	**10,341.80 €**	**2,584.46 €**	**-14,562.82 €**	**-31,115.75 €**	**-46,611.66 €**	**-79,363.98 €**	
Cash Flow After Taxes	**-137,000.00 €**	**-2,187.92 €**	**36,445.86 €**	**67,968.28 €**	**104,804.67 €**	**149,138.54 €**	**219,169.44 €**	INCLUDING CAPEX (Y0)
Discounted Cash Flow (After Tax)	**-137,000.00 €**	**-2,083.74 €**	**33,057.47 €**	**58,713.56 €**	**86,223.06 €**	**116,853.95 €**	**155,764.31 €**	INCLUDING CAPEX (Y0)

Business Case Results:	
NPV of Cash Flow	**155,764.31 €**
IRR	**27.6%**
Profitability Index	**2.14**
Simple Payback	**3 Years 4 Months**
Discounted Payback	**3 Years 7 Months**
ROI	**25%**

Case Assumptions:		
Cost Escalation Factor	**0.60%**	(0,50 above Eurozone inflation rate)
Benefit Escalation Factor	**1.60%**	(1 above cost escalation factor)
Corporate Tax Rate (Europe Average 2015)	**25.90%**	
Discount Rate	**5.00%**	
Interest Rate for Loans	**2.20%**	(Euribor plus 2%)
Maintenance Repair/Replacement Exp	**1.50%**	of CAPEX
Annual Increase	**10.00%**	
Marketing/Advertising/Promotion Exp	**4.00%**	of CAPEX
Annual Decrease	**50.00%**	
Property Mng, operations and bill Exp	**0.50%**	of CAPEX
Revenues from add-on services (SDN/NFV)	**15.00%**	

CAPEX Costs are NOT expensed in Y0 but CAPITALIZED in Y1-Y5 (Depreciation Expense)

Figure 224. CloudSat scenario #2 Normal case: CFM Financial analysis and financial ratios

Under Normal financial, business and market conditions, CloudSat scenario #2 is a viable solution with appropriate financial results and ratios within a 5-years period. Findings are summarized in below figure:

Financial variables and ratios	Values
NPV of Cash Flow	155,764.31 €
IRR	27.6%
Profitability Index	2.14
Simple Payback	3 Years 4 Months
Discounted Payback	3 Years 7 Months
ROI	25%
Breakeven Year (excluding CAPEX)	Late Year 2

Figure 225. CloudSat scenario #2 Normal case: Financial variables and ratios results

This investment becomes positive in cash flows within year 2 (Breakeven Year) and after 3 years and 7 months reaches its Payback point (having make up for CAPEX and OPEX so far). The discounted cash flow after taxes at the end of 5th year has a NPV of 155,764.31€ (total profit of the 5 years at present value). This means that the investment is profitable within the 5 years period with an IRR 27.6% , Profitability index 2.14 and a ROI of 25% upon initial investment and variable costs.

b) Scenario #2 – "Optimistic" case analysis

Under Optimistic case conditions, the two conditional variables are assumed to have the below values:

Variables	Scenario #2
	OPTIMISTIC
Annual Revenues (annual percentage change)	20%
Discount Rate	2%

Figure 226. CloudSat scenario #2: Optimistic case variables

Consequently the estimated revenues are:

ASSUMPTIONS	FIGURES
Revenues originating from leased Satellite bandwidth /10Mbit/s/year	
Estimated Annual Revenues for CloudSat Scenario 2 from reselling leased Satellite bandwidth of 10Mbit/s/year (including all add-on provisional services being offered)	€ 850,000.00

Initial (Y1) Revenues (% upon Estimated Annual Revenues)	30.00%
Annual Revenue increase rate change	20.00%

Figure 227. CloudSat scenario #2 Optimistic case – Satellite bandwidth utilization Revenues

The satellite bandwidth utilization revenues are:

	Y1	Y2	Y3	Y4	Y5
Revenues per year	255,000.00 €	306,000.00 €	367,200.00 €	440,640.00 €	528,768.00 €
% on Total Estimated Annual Revenues of CloudSat Scenario 2	30.00%	36.00%	43.20%	51.84%	62.21%

Figure 228. CloudSat scenario #2 Optimistic case- Satellite bandwidth utilization Revenues

Therefore, the total revenues are calculated as per below figure.

	Y1	Y2	Y3	Y4	Y5
Revenues	255,000.00 €	306,000.00 €	367,200.00 €	440,640.00 €	528,768.00 €
Revenues from add-on services (SDN/NFV) as 15% of Revenues	38,250.00 €	45,900.00 €	55,080.00 €	66,096.00 €	79,315.20 €
Escalation of Benefits	4,692.00 €	11,350.89 €	20,595.48 €	33,217.79 €	50,228.46 €
TOTAL	**297,942.00 €**	**363,250.89 €**	**442,875.48 €**	**539,953.79 €**	**658,311.66 €**

Figure 229. CloudSat scenario #2 Optimistic case – Revenues

Summarizing all the estimated figures of CAPEX, OPEX and Revenues for CloudSat scenario #2 Optimistic case and applying the CFM model of analysis, we perform the cash flow financial analysis and calculate the financial ratios as per the figure on the next page. Please note that Discount rate 2% has been applied.

CloudSat Scenario #2 - Optmistic conditions of Evaluation									
	Y0	1	2	3	4	5	Totals		
CAPEX									
Outdoor units	-70,000.00 €	0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	-70,000.00 €		
Modulator									
Encapsulator/Multiplexer	-120,000.00 €						-120,000.00 €		
Return Link sub-system									
GW Management & Access control									
Satellite Hub Generic Server HW (2 servers)	-10,000.00 €						-10,000.00 €		
Satellite Hub Generic Server Software Configuration	-3,500.00 €						-3,500.00 €		
Terrestrial interface subsystem (firewall) 5 year licence	-17,500.00 €						-17,500.00 €		
Terrestrial interface subsystem (PEP)	-2,500.00 €						-2,500.00 €		
Terrestrial interface subsystem (router)	-10,000.00 €						-10,000.00 €		
NFVI PoP SW+Licence+Configuration	-3,500.00 €						-3,500.00 €		
NFVI PoP HW (5 servers)	-25,000.00 €						-25,000.00 €		
SDN Switches (Two units)	-20,000.00 €						-20,000.00 €		
NFV compatible terminal (upgrade)	-5,000.00 €						-5,000.00 €		
VNF1 Transcoder	-20,000.00 €						-20,000.00 €		
NFV Orchestrator (HW/SW)	-30,000.00 €						-30,000.00 €		
Bank Loan paid in 5 years with annual constant installments and constant/fixed interest rate	200,000.00 €						200,000.00 €		
Total CAPEX (for accounting)	**-137,000.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**-137,000.00 €**		
Actual CAPEX (excluding bank loan)	**-337,000.00 €**								
OPEX (VC Operating and Maintenance Costs)									
Satellite bandwidth leasing cost (€/10Mbit/s/year)		-240,000.00 €	-240,000.00 €	-240,000.00 €	-240,000.00 €	-240,000.00 €	-1,200,000.00 €		
Maintenance and Repair/Replacement/Unforeseen expenses (1.5% of CAPEX increased by 10% annually)		-5,055.00 €	-5,560.50 €	-6,116.55 €	-6,728.21 €	-7,401.03 €	-30,861.28 €		
Marketing/Advertising/Promotion expenses (4% of CAPEX reduced by 50% annually)		-13,480.00 €	-6,740.00 €	-3,370.00 €	-1,685.00 €	-842.50 €	-26,117.50 €		
Property management, operations, communication and bill/utilities expenses (0.5% of CAPEX)		-1,685.00 €	-1,685.00 €	-1,685.00 €	-1,685.00 €	-1,685.00 €	-8,425.00 €		
Supply expenses		-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-20,000.00 €		
Loan Payments		-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-200,000.00 €		
Interest Rate Expense		-4,400.00 €	-3,557.88 €	-2,697.23 €	-1,817.65 €	-918.71 €	-13,391.46 €		
Escalation of Costs		-1,851.72 €	-3,629.38 €	-5,393.87 €	-7,166.15 €	-8,952.20 €	-26,993.32 €		
Total Costs	**0.00 €**	**-310,471.72 €**	**-305,172.75 €**	**-303,262.65 €**	**-303,082.00 €**	**-303,799.44 €**	**-1,525,788.57 €**		
Revenue and Operating Benefits									
Revenues		255,000.00 €	306,000.00 €	367,200.00 €	440,640.00 €	528,768.00 €	1,897,608.00 €		
Revenues from add-on services (SDN/NFV) as 15% of Revenues		38,250.00 €	45,900.00 €	55,080.00 €	66,096.00 €	79,315.20 €	284,641.20 €		
Escalation of Benefits		4,692.00 €	11,350.89 €	20,595.48 €	33,217.79 €	50,228.46 €	120,084.61 €		
Total Benefits and Revenue	**0.00 €**	**297,942.00 €**	**363,250.89 €**	**442,875.48 €**	**539,953.79 €**	**658,311.66 €**	**2,302,333.81 €**		
Contribution Margin % (Sales - VC/Sales)		**-4.21%**	**15.99%**	**31.52%**	**43.87%**	**53.85%**	**33.73%**		
Contribution Margin Amt (Sales - VC)		**-12,529.72 €**	**58,078.13 €**	**139,612.83 €**	**236,871.78 €**	**354,512.22 €**	**776,545.24 €**		
Cash Flow Before Taxes	**-137,000.00 €**	**-12,529.72 €**	**58,078.13 €**	**139,612.83 €**	**236,871.78 €**	**354,512.22 €**	**639,545.24 €**	INCLUDING CAPEX (Y0)	
EBITDA		**-16,929.72 €**	**54,520.26 €**	**136,915.60 €**	**235,054.14 €**	**353,593.51 €**	**763,153.78 €**	EXCLUDING CAPEX (Y0)	
Breakeven Year			*						
Income Tax Calculation									
Depreciation Expense		-27,400.00 €	-43,840.00 €	-26,304.00 €	-15,782.40 €	-15,782.40 €	-129,108.80 €		
Operating and Maintenance Cost (OPEX)		-310,471.72 €	-305,172.75 €	-303,262.65 €	-303,082.00 €	-303,799.44 €	-1,525,788.57 €		
Revenue and Operating Benefits		297,942.00 €	363,250.89 €	442,875.48 €	539,953.79 €	658,311.66 €	2,302,333.81 €		
Net Income Taxes	**0.00 €**	**10,341.80 €**	**-3,687.68 €**	**-29,346.99 €**	**-57,262.15 €**	**-87,731.02 €**	**-167,686.04 €**		
Cash Flow After Taxes	**-137,000.00 €**	**-2,187.92 €**	**54,390.46 €**	**110,265.84 €**	**179,609.63 €**	**266,781.20 €**	**471,859.20 €**	INCLUDING CAPEX (Y0)	
Discounted Cash Flow (After Tax)	**-137,000.00 €**	**-2,145.02 €**	**52,278.41 €**	**103,905.97 €**	**165,931.54 €**	**241,631.95 €**	**424,602.84 €**	INCLUDING CAPEX (Y0)	

Business Case Results:	
NPV of Cash Flow	**424,602.84 €**
IRR	**46.6%**
Profitability Index	**4.10**
Simple Payback	**2 Years 9 Months**
Discounted Payback	**2 Years 10 Months**
ROI	**46%**

Case Assumptions:		
Cost Escalation Factor	**0.60%**	(0,50 above Eurozone inflation rate)
Benefit Escalation Factor	**1.60%**	(1 above cost escalation factor)
Corporate Tax Rate (Europe Average 2015)	**25.90%**	
Discount Rate	**2.00%**	
Interest Rate for Loans	**2.20%**	(Euribor plus 2%)
Maintenance Repair/Replacement Exp	**1.50%**	of CAPEX
Annual Increase	**10.00%**	
Marketing/Advertising/Promotion Exp	**4.00%**	of CAPEX
Annual Decrease	**50.00%**	
Property Mng, operations and bill Exp	**0.50%**	of CAPEX
Revenues from add-on services (SDN/NFV)	**15.00%**	

CAPEX Costs are NOT expensed in Y0 but CAPITALIZED in Y1-Y5 (Depreciation Expense)

Figure 230. CloudSat scenario #2 Optimistic case: CFM Financial analysis and financial ratios

Under optimistic financial, business and market conditions, CloudSat scenario #2 is much more viable and profitable investment (in comparison to normal conditions scenario) with very attractive financial figures within a 5-years period. Findings are summarized in below figure:

Financial variables and ratios	Values
NPV of Cash Flow	424,602.84 €
IRR	46.6%
Profitability Index	4.10
Simple Payback	2 Years 9 Months
Discounted Payback	2 Years 10 Months
ROI	46%
Breakeven Year (excluding CAPEX)	Early Year 2

Figure 231. CloudSat scenario #2 Optimistic case: Financial variables and ratios results

Favourable financial conditions of the optimistic case make this investment more promising since it becomes positive in cash flows within first months of year 2 (Breakeven Year) and after 2 years and 10 months reaches its Payback point (having make up for CAPEX and OPEX so far). The discounted cash flow after taxes at the end of 5th year has a NPV of 424,602.84€ (total profit of the 5 years at present value). This means that the investment is well profitable within the 5 years period with an IRR 46.6%, Profitability index 4.10 and a ROI of 46% upon initial investment, variable costs and operations.

c) Scenario #2 – "Pessimistic" case analysis

Under Pessimistic case conditions (more unstable economy and market/business environment), the two conditional variables are assumed to have the below values:

Variables	Scenario #2
	PESSIMISTIC
Annual Revenues (annual percentage change)	4%
Discount Rate	10%

Figure 232. CloudSat scenario #2: Pessimistic case variables

Consequently the estimated revenues are:

ASSUMPTIONS	FIGURES
Revenues originating from leased Satellite bandwidth /10Mbit/s/year	
Estimated Annual Revenues for CloudSat Scenario 2 from reselling leased Satellite bandwidth of 10Mbit/s/year (including all add-on provisional services being offered)	€ 850,000.00
Initial (Y1) Revenues (% upon Estimated Annual Revenues)	30.00%
Annual Revenue increase rate change	4.00%

Figure 233. CloudSat scenario #2 Pessimistic case – Satellite bandwidth utilization Revenues

The satellite bandwidth utilization revenues are:

	Y1	Y2	Y3	Y4	Y5
Revenues per year	255,000.00 €	265,200.00 €	275,808.00 €	286,840.32 €	298,313.93 €
% on Total Estimated Annual Revenues of CloudSat Scenario 2	30.00%	31.20%	32.45%	33.75%	35.10%

Figure 234. CloudSat scenario #2 Pessimistic case- Satellite bandwidth utilization Revenues

Therefore, the total revenues are calculated as per below figure.

	Y1	Y2	Y3	Y4	Y5
Revenues	255,000.00 €	265,200.00 €	275,808.00 €	286,840.32 €	298,313.93 €
Revenues from add-on services (SDN/NFV) as 15% of Revenues	38,250.00 €	39,780.00 €	41,371.20 €	43,026.05 €	44,747.09 €
Escalation of Benefits	4,692.00 €	9,837.43 €	15,469.49 €	21,623.55 €	28,337.28 €
TOTAL	**297,942.00 €**	**314,817.43 €**	**332,648.69 €**	**351,489.92 €**	**371,398.31 €**

Figure 235. CloudSat scenario #2 Pessimistic case – Revenues

Summarizing all the estimated figures of CAPEX, OPEX and Revenues for CloudSat scenario #2 Pessimistic case and applying the CFM model of analysis, we perform the cash flow financial analysis and calculate the financial ratios as per the figure on the next page. Please note that Discount rate 10% has been applied.

CloudSat Scenario #2 -Pessimistic conditions of Evaluation								
	Y0	1	2	3	4	5	Totals	
CAPEX								
Antenna & RF front-end	-70,000.00 €	0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	-70,000.00 €	
Modulator								
Encapsulator/Multiplexer	-120,000.00 €						-120,000.00 €	
Return Link sub-system								
GW Management & Access control								
Satellite Hub Generic Server HW (2 servers)	-10,000.00 €						-10,000.00 €	
Satellite Hub Generic Server Software Configuration	-3,500.00 €						-3,500.00 €	
Terrestrial interface subsystem (firewall) 5 year licence	-17,500.00 €						-17,500.00 €	
Terrestrial interface subsystem (PEP)	-2,500.00 €						-2,500.00 €	
Terrestrial interface subsystem (router)	-10,000.00 €						-10,000.00 €	
NFVI PoP SW+Licence+Configuration	-3,500.00 €						-3,500.00 €	
NFVI PoP HW (5 servers)	-25,000.00 €						-25,000.00 €	
SDN Switches (Two units)	-20,000.00 €						-20,000.00 €	
NFV compatible terminal (upgrade)	-5,000.00 €						-5,000.00 €	
VNF1 Transcoder	-20,000.00 €						-20,000.00 €	
NFV Orchestrator (HW/SW)	-30,000.00 €						-30,000.00 €	
Bank Loan paid in 5 years with annual constant installments and constant/fixed interest rate	200,000.00 €						200,000.00 €	
Total CAPEX (for accounting)	**-137,000.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**-137,000.00 €**	
Actual CAPEX (excluding bank loan)	**-337,000.00 €**							
OPEX (VC Operating and Maintenance Costs)								
Satellite bandwidth leasing cost (€/10Mbit/s/year)		-240,000.00 €	-240,000.00 €	-240,000.00 €	-240,000.00 €	-240,000.00 €	-1,200,000.00 €	
Maintenance and Repair/Replacement/Unforeseen expenses (1.5% of CAPEX increased by 10% annually)		-5,055.00 €	-5,560.50 €	-6,116.55 €	-6,728.21 €	-7,401.03 €	-30,861.28 €	
Marketing/Advertising/Promotion expenses (4% of CAPEX reduced by 50% annually)		-13,480.00 €	-6,740.00 €	-3,370.00 €	-1,685.00 €	-842.50 €	-26,117.50 €	
Property management, operations, communication and bill/utilities expenses (0.5% of CAPEX)		-1,685.00 €	-1,685.00 €	-1,685.00 €	-1,685.00 €	-1,685.00 €	-8,425.00 €	
Supply expenses		-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-20,000.00 €	
Loan Payments		-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-200,000.00 €	
Interest Rate Expense		-4,400.00 €	-3,557.88 €	-2,697.23 €	-1,817.65 €	-918.71 €	-13,391.46 €	
Escalation of Costs		-1,851.72 €	-3,629.38 €	-5,393.87 €	-7,166.15 €	-8,952.20 €	-26,993.32 €	
Total Costs	**0.00 €**	**-310,471.72 €**	**-305,172.75 €**	**-303,262.65 €**	**-303,082.00 €**	**-303,799.44 €**	**-1,525,788.57 €**	
Revenue and Operating Benefits								
Revenues		255,000.00 €	265,200.00 €	275,808.00 €	286,840.32 €	298,313.93 €	1,381,162.25 €	
Revenues from add-on services (SDN/NFV) as 15% of Revenues		38,250.00 €	39,780.00 €	41,371.20 €	43,026.05 €	44,747.09 €	207,174.34 €	
Escalation of Benefits		4,692.00 €	9,837.43 €	15,469.49 €	21,623.55 €	28,337.28 €	79,959.76 €	
Total Benefits and Revenue	**0.00 €**	**297,942.00 €**	**314,817.43 €**	**332,648.69 €**	**351,489.92 €**	**371,398.31 €**	**1,668,296.35 €**	
Contribution Margin % (Sales - VC/Sales)		**-4.21%**	**3.06%**	**8.83%**	**13.77%**	**18.20%**	**8.54%**	
Contribution Margin Amt (Sales - VC)		**-12,529.72 €**	**9,644.68 €**	**29,386.04 €**	**48,407.91 €**	**67,598.87 €**	**142,507.79 €**	
Cash Flow Before Taxes	**-137,000.00 €**	**-12,529.72 €**	**9,644.68 €**	**29,386.04 €**	**48,407.91 €**	**67,598.87 €**	**5,507.79 €**	INCLUDING CAPEX (Y0)
EBITDA		**-16,929.72 €**	**6,086.80 €**	**26,688.82 €**	**46,590.27 €**	**66,680.16 €**	**129,116.32 €**	EXCLUDING CAPEX (Y0)
Breakeven Year				*				
Income Tax Calculation								
Depreciation Expense		-27,400.00 €	-43,840.00 €	-26,304.00 €	-15,782.40 €	-15,782.40 €	-129,108.80 €	
Operating and Maintenance Cost (OPEX)		-310,471.72 €	-305,172.75 €	-303,262.65 €	-303,082.00 €	-303,799.44 €	-1,525,788.57 €	
Revenue and Operating Benefits		297,942.00 €	314,817.43 €	332,648.69 €	351,489.92 €	371,398.31 €	1,668,296.35 €	
Net Income Taxes	**0.00 €**	**10,341.80 €**	**8,856.59 €**	**-798.25 €**	**-8,450.01 €**	**-13,420.47 €**	**-3,470.34 €**	
Cash Flow After Taxes	**-137,000.00 €**	**-2,187.92 €**	**18,501.27 €**	**28,587.79 €**	**39,957.90 €**	**54,178.40 €**	**2,037.45 €**	INCLUDING CAPEX (Y0)
Discounted Cash Flow (After Tax)	**-137,000.00 €**	**-1,989.02 €**	**15,290.30 €**	**21,478.43 €**	**27,291.79 €**	**33,640.53 €**	**-41,287.97 €**	INCLUDING CAPEX (Y0)
Business Case Results:				**Case Assumptions:**				
NPV of Cash Flow	**-41,287.97 €**			**Cost Escalation Factor**		**0.60%**	(0,50 above Eurozone inflation rate)	
IRR	**0.4%**			**Benefit Escalation Factor**		**1.60%**	(1 above cost escalation factor)	
Profitability Index	**0.70**			**Corporate Tax Rate (Europe Average 2015)**		**25.90%**		
Simple Payback	**4 Years 12 Months**			**Discount Rate**		**10.00%**		
Discounted Payback	**Above 5 years**			**Interest Rate for Loans**		**2.20%**	(Euribor plus 2%)	
ROI	**8%**			**Maintenance Repair/Replacement Exp**		**1.50%**	of CAPEX	
				Annual Increase		**10.00%**		
				Marketing/Advertising/Promotion Exp		**4.00%**	of CAPEX	
				Annual Decrease		**50.00%**		
				Property Mng, operations and bill Exp		**0.50%**	of CAPEX	
				Revenues from add-on services (SDN/NFV)		**15.00%**		
				CAPEX Costs are NOT expensed in Y0 but CAPITALIZED in Y1-Y5 (Depreciation Expense)				

Figure 236. CloudSat scenario #2 Pessimistic case: CFM Financial analysis and financial ratios

Under pessimistic financial, business and market conditions, CloudSat scenario #2 remains still viable but it is considered a very limited case. Since the discount rate used is high (10%), the fact that the investment becomes limited profitable by the end of year 5, the net present value of that profit in year 0 is negative. Findings are summarized in below figure:

Financial variables and ratios	Values
NPV of Cash Flow	-41,287.97 €
IRR	0.4%
Profitability Index	0.70
Simple Payback	4 Years 11-12 Months
Discounted Payback	Above 5 years
ROI	8%
Breakeven Year (excluding CAPEX)	Year 3

Figure 237. CloudSat scenario #2 Pessimistic case: Financial variables and ratios results

This pessimistic case makes the CloudSat scenario #2 to be a marginal case of profitable investment for the 5-year period analysis. This scenario becomes positive in cash flows within year 3 (Breakeven Year) and after 4 years and 11 months reaches its Simple Payback point (having make up for CAPEX and OPEX so far) but the net cash flows of 54,178.40 € at the end of year 5 are not enough. When the net cash flows are discounted with 10% then their NPV is -41,287.97. This negative figure proves that this investment will become more attractive only after at least 6 years of operation, a case scenario that at the end of the 6th year provide profits with positive net present value. This is shown also by the fact that the Discounted Payback point is greater than 5 years. However the fact that we have positive net cash flows of 54,178.40 € at the end of year 5 means that the investment has a good potential but requires more time to run. In addition it has an IRR 0.4%, Profitability index 0.70 and a ROI of 8% upon initial investment, variable costs and operations.

In summary, if all financial evaluations take place at the end of year 5 then we have a profitable investment with all ratios being positive. If we consider the time value of money and discount all amounts to present then we have a marginal case analysis with negative NPV of cash flow which requires additional time (during year 6) in order to balance and provide positive present value figures. Overall, even under the current pessimistic conditions of this case, we consider this scenario financially viable for becoming profitable with positive present value criteria provided that it is runs beyond the limit of 5 years.

7.2.3.3. Scenario #3 - LEO - Customer functions virtualization

Initial cost investment (CAPEX) components for scenario 3 are modified in relation to the ones described (following LEO SW-based virtualization deployment).

Please note additional assumptions being valid for CloudSat Case Scenario #3 CAPEX, OPEX and Revenues calculation [OWB]:

- CloudSat Case Scenario #3 is based on the LEO satellite systems characteristics
- The CloudSat Case Scenario #3 is based on the characteristics and figures of a typical LEO satellite system with 50 terrestrial gateways (such as OneWeb) but its financial analysis is performed using the figures corresponding to one gateway (installation/operation costs and revenues) for more detailed analysis
- As per the LEO benchmarking case, one terrestrial gateway cost is estimated at 750,000 € which is discounted to 360,000€ (using the NFV Cost Reduction of 52%)
- Satellite bandwidth leasing cost is zero (owned satellites) and is included in OPEX
- Revenues estimated figures are based on a typical LEO system's revenues (e.g. OneWeb) with annual revenues (average annual revenues over the period 2010-2014): 350,000,000€
- The revenues of a typical LEO system originating from 50 gateways are used for estimating the corresponding relative ratio of revenues for one gateway (Revenues/50): 350,000,000€/50=7,000,000€ per gateway
- The corresponding revenues of one gateway include the provision of several additional services being offered (other than the CloudSat specific ones). It is estimated as assumption that the CloudSat Case Scenario #3 service (as a novel service) represents the 1/10 in the revenues of the total provided services. Consequently, the revenues of CloudSat Case Scenario #3 service correspond to the 1/10 of the total gateway revenues (One gateway Revenues/10): 7,000,000€/10 = 700,000€ annual CloudSat revenues (of one gateway)
- Revenues originate from two sources: (a) earnings as percentage of the corresponding annual revenues originating from the utilization of the one gateway of the CloudSat Case Scenario #3 as described above and (b) earnings from add-on services (SDN/NFV) calculated as fixed 15% percentage upon (a).
- Initial percentage of earnings upon the annual revenues originating from the utilization of the one gateway of the CloudSat Case Scenario #3 is set at 30% for Y1 and dynamic for the years Y2-Y5 (as per the conditions of each scenario)

Figure below presents the CAPEX components of CloudSat case scenario 3 (referring to one gateway of LEO satellite system along with the market research references).

CloudSat Scenario #3 satellite hub infrastructure	Cost (1 gateway)	Market Research References
Outdoor units	100,000.00 €	http://www.satsig.net/ivsatcos.htm http://www.oneweb.net
Modulator	360,000.00 € (1 gateway)	
Encapsulator/Multiplexer		
Return Link sub-system		
GW Management & Access Control		
Satellite Hub Generic Server HW (2 servers)	10,000.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0 IBM report "A Deep Dive Into The Cost Benefits of KVM and Open Virtualization"
Fibre Leased Lines (1 GBps annual cost + initial installation)	30,000.00 €	http://www.hso.co.uk/leased-lines/ http://business.bt.com/broadband-and-internet/leased-lines/
Satellite Hub Generic Server Software Configuration	3,500.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0
CloudSat Scenario #3 Configuration Variant Components		
SDN Switches (Two units)	20,000.00 €	http://store.netgate.com/Pica8--C188.aspx
NFVI PoP (2 Units) HW (5 servers each): 25000€ each	50,000.00 €	http://www8.hp.com/h20195/V2/getpdf.aspx/4AA5-2160ENW.pdf?ver=1.0
NFVI PoP (2 Units) SW+ Licence + Configuration 3500€ each	7,000.00 €	http://www.streamingmedia.com/Articles/Editorial/Featured-Articles/High-End-Video-Transcoder-Shootout-Elemental-Vs.-Telestream-100978.aspx
VNF 1 (firewall) 5000€ license fee per year	25,000.00 €	http://www.techrepublic.com/article/evaluating-the-real-cost-of-an-enterprise-firewall/
VNF 2 (web cache) 1500€ fee per year	7,500.00 €	https://www.untangle.com/shop/Web-Cache
VNF 3 (TCP Optimizer) 500€ fee per year	2,500.00 €	TCP Optimizer – free/open source
NFV Orchestrator (HW/SW)	30,000.00 €	
TOTAL (actual)	**645,500.00 €**	
Bank Loan paid in 5 years with annual constant instalments and constant/fixed interest rate	**-200,000.00 €**	
TOTAL (accounting)	**445,500.00 €**	

Figure 238. CloudSat case scenario #3: CAPEX Estimation

Below Figure summarizes the OPEX expenses of CloudSat case scenario #3. Please note that Satellite bandwidth leasing cost is zero since owned satellites are used.

OPEX	Y0	Y1	Y2	Y3	Y4	Y5
Satellite bandwidth leasing cost	0 €	0 €	0 €	0 €	0 €	0 €
Maintenance and Repair/Replacement/Unforeseen expenses (1.5% of CAPEX increased by 10% annually)	0 €	-9,682.50 €	-10,650.75 €	-11,715.83 €	-12,887.41 €	-14,176.15 €
Marketing/Advertising/ Promotion expenses (4% of CAPEX reduced by 50% annually)	0 €	-25,820.00 €	-12,910.00 €	-6,455.00 €	-3,227.50 €	-1,613.75 €
Property management, operations and bill/utilities expenses (0.5% of CAPEX)	0 €	-3,227.50 €	-3,227.50 €	-3,227.50 €	-3,227.50 €	-3,227.50 €
Supply expenses	0 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €
Loan Payments		-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €
Interest Rate Expense		-4,400.00 €	-3,557.88 €	-2,697.23 €	-1,817.65 €	-918.71 €
Escalation of Costs	0 €	-522.78 €	-894.83 €	-1,233.09 €	-1,577.97 €	-1,941.24 €
TOTAL	0 €	-87,652.78 €	-75,240.96 €	-69,328.64 €	-66,738.02 €	-65,877.35 €

Figure 239. CloudSat case scenario #3: OPEX Estimation

a) Scenario #3 – "Normal" case analysis

Under Normal case conditions, the two conditional variables are assumed to have the below values:

Variables	Scenario #3
	NORMAL
Annual Revenues (annual percentage change)	12%
Discount Rate	5%

Figure 240. CloudSat scenario #3: Normal case variables

As already clarified in the assumptions section, the LEO revenues (originating from 50 gateways) are used for estimating the corresponding relative ratio of revenues for one gateway (Revenues/50) which is 350,000,000€/50=7,000,000€ revenues per gateway. However, since the corresponding revenues of one gateway include the provision of several additional services being offered (other than the CloudSat

specific), it is set as assumption that the CloudSat Case Scenario #3 service represents a small portion of total provided services of the gateway. Consequently, the revenues of CloudSat Case Scenario #3 service correspond to the 1/10 of the total gateway revenues (One gateway Revenues/10) meaning 7,000,000€/10 = 700,000€ annual CloudSat revenues (of one gateway). However, the initial revenues of 350,000,000€ correspond to a well-established operator, while in our scenario we analyze the case of a new entry operator (eg CloudSat operator) offering novel services. Thus, it is expected for CloudSat that the total annual revenues for the first years of operation (Y1-Y5) will be lower than the ones referred above.

So in our analysis for this CloudSat scenario the estimated revenues of the new entry CloudSat operator will be calculated as follows: (a) percentage of the revenues (30% for Y1 and dynamic for Y2-Y5) of the above well-established operator (eg percentage upon 700,000€) and (b) additional earnings from add-on services (SDN/NFV) calculated as fixed 15% percentage upon (a).

Consequently the revenues assumptions are:

ASSUMPTIONS	FIGURES	REMARKS
LEO system Annual Revenues (50 Terrestrial Gateways)	350,000,000.00 €	Used as base figure of Annual Revenues Estimates resulting in 7,000,000 revenues per gateway
1 gateway annual revenues	7,000,000.00 €	350,000,000.00 € / 50
CloudSat Scenario 3 - Number of Terrestrial Gateways as part of this financial analysis	1	
Estimated Annual Revenues for CloudSat Scenario 3 (1 gateway)	700,000.00 €	Based on the assumption that 1/10 of total gateway revenues are originating from the new CloudSat Scenario 3 services (7,000,000/10)
Initial (Y1) Revenues (% upon Estimated Annual Revenues)	30.00%	
Annual Revenues increase rate change	12.00%	An average percentage taking into account changes in revenues due to the potential loss of customers and/or the newcomers percentage over the year (eg 18% newcomers and 6% loss of customers)

Figure 241. CloudSat scenario #3 Normal case – Revenues Assumptions

The CloudSat case scenario #3 revenues are:

	Y1	Y2	Y3	Y4	Y5
Revenues per year	210,000.00 €	235,200.00 €	263,424.00 €	295,034.88 €	330,439.07 €
% on Annual Revenues of CloudSat Scenario 3	30.00%	33.60%	37.63%	42.15%	47.21%

Figure 242. CloudSat scenario #3 Normal case- Revenues

Therefore, the total revenues are calculated as per below figure.

	Y1	Y2	Y3	Y4	Y5
Revenues	210,000.00 €	235,200.00 €	263,424.00 €	295,034.88 €	330,439.07 €
Revenues from add-on services (SDN/NFV) as 15% of Revenues	31,500.00 €	35,280.00 €	39,513.60 €	44,255.23 €	49,565.86 €
Escalation of Benefits	3,864.00 €	8,724.60 €	14,774.90 €	22,241.30 €	31,388.90 €
TOTAL	245,364.00 €	279,204.60 €	317,712.50 €	361,531.41 €	411,393.82 €

Figure 243. CloudSat scenario #2 Normal case - Revenues

Summarizing all the estimated figures of CAPEX, OPEX and Revenues for CloudSat scenario #3 Normal case and applying the CFM model of analysis, we perform the cash flow financial analysis and calculate the financial ratios as per the figure on the next page. Please note that Discount rate 5% has been applied.

CloudSat Scenario #3 - Normal conditions of Evaluation								
	Y0	1	2	3	4	5	Totals	
CAPEX								
Outdoor units	-100,000.00 €	0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	-100,000.00 €	
Modulator								
Encapsulator/Multiplexer	-360,000.00 €						-360,000.00 €	
Return Link sub-system								
GW Management & Access control								
Satellite Hub Generic Server HW (2 servers)	-10,000.00 €						-10,000.00 €	
Satellite Hub Generic Server Software Configuration	-3,500.00 €						-3,500.00 €	
SDN Switches (Two units)	-20,000.00 €						-20,000.00 €	
NFVI PoP (2 Units) HW (5 servers): 25000 each	-50,000.00 €						-50,000.00 €	
NFVI PoP (2 Units) SW+Licence+Configuration 3500 each	-7,000.00 €						-7,000.00 €	
VNF 1 (firewall) 5000 licence fee per year (x 5 servers)	-25,000.00 €						-25,000.00 €	
VNF 2 (web cache) 1500 fee per year (x 5 servers)	-7,500.00 €						-7,500.00 €	
VNF 3 (TCP Optimizer) 500 fee per year (x 5 servers)	-2,500.00 €						-2,500.00 €	
NFV Orchestrator (HW/SW)	-30,000.00 €						-30,000.00 €	
Fibre Leased Lines (1 GBps annual cost + initial installation)	-30,000.00 €						-30,000.00 €	
	0.00 €						0.00 €	
Bank Loan paid in 5 years with annual constant installments and constant/fixed interest rate	200,000.00 €						200,000.00 €	
Total CAPEX (for accounting)	**-445,500.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**-445,500.00 €**	
Actual CAPEX (excluding bank loan)	**-645,500.00 €**							
OPEX (VC Operating and Maintenance Costs)								
Satellite bandwidth leasing cost		0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	
Maintenance and Repair/Replacement/Unforeseen expenses (1.5% of CAPEX increased by 10% annually)		-9,682.50 €	-10,650.75 €	-11,715.83 €	-12,887.41 €	-14,176.15 €	-59,112.63 €	
Marketing/Advertising/Promotion expenses (4% of CAPEX reduced by 50% annually)		-25,820.00 €	-12,910.00 €	-6,455.00 €	-3,227.50 €	-1,613.75 €	-50,026.25 €	
Property management, operations, communications and bill/utilities expenses (0.5% of CAPEX)		-3,227.50 €	-3,227.50 €	-3,227.50 €	-3,227.50 €	-3,227.50 €	-16,137.50 €	
Supply expenses		-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-20,000.00 €	
Loan Payments		-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-200,000.00 €	
Interest Rate Expense		-4,400.00 €	-3,557.88 €	-2,697.23 €	-1,817.65 €	-918.71 €	-13,391.46 €	
Escalation of Costs		-522.78 €	-894.83 €	-1,233.09 €	-1,577.97 €	-1,941.24 €	-6,169.91 €	
Total Costs	**0.00 €**	**-87,652.78 €**	**-75,240.96 €**	**-69,328.64 €**	**-66,738.02 €**	**-65,877.35 €**	**-364,837.75 €**	
Revenue and Operating Benefits								
Revenues		210,000.00 €	235,200.00 €	263,424.00 €	295,034.88 €	330,439.07 €	1,334,097.95 €	
Revenues from add-on services (SDN/NFV) as 15% of Revenues		31,500.00 €	35,280.00 €	39,513.60 €	44,255.23 €	49,565.86 €	200,114.69 €	
Escalation of Benefits		3,864.00 €	8,724.60 €	14,774.90 €	22,241.30 €	31,388.90 €	80,993.70 €	
Total Benefits and Revenue	**0.00 €**	**245,364.00 €**	**279,204.60 €**	**317,712.50 €**	**361,531.41 €**	**411,393.82 €**	**1,615,206.34 €**	
Contribution Margin % (Sales - VC/Sales)		**64.28%**	**73.05%**	**78.18%**	**81.54%**	**83.99%**	**77.41%**	
Contribution Margin Amt (Sales - VC)		**157,711.22 €**	**203,963.65 €**	**248,383.86 €**	**294,793.39 €**	**345,516.47 €**	**1,250,368.58 €**	
Cash Flow Before Taxes	**-445,500.00 €**	**157,711.22 €**	**203,963.65 €**	**248,383.86 €**	**294,793.39 €**	**345,516.47 €**	**804,868.58 €**	INCLUDING CAPEX (Y0)
EBITDA		**153,311.22 €**	**200,405.77 €**	**245,686.63 €**	**292,975.74 €**	**344,597.76 €**	**1,236,977.12 €**	EXCLUDING CAPEX (Y0)
Breakeven Year			*					
Income Tax Calculation								
Depreciation Expense		-89,100.00 €	-142,560.00 €	-85,536.00 €	-51,321.60 €	-51,321.60 €	-419,839.20 €	
Operating and Maintenance Cost (OPEX)		-87,652.78 €	-75,240.96 €	-69,328.64 €	-66,738.02 €	-65,877.35 €	-364,837.75 €	
Revenue and Operating Benefits		245,364.00 €	279,204.60 €	317,712.50 €	361,531.41 €	411,393.82 €	1,615,206.34 €	
Net Income Taxes	**0.00 €**	**-17,770.31 €**	**-15,903.54 €**	**-42,177.60 €**	**-63,059.19 €**	**-76,196.47 €**	**-215,107.11 €**	
Cash Flow After Taxes	**-445,500.00 €**	**139,940.91 €**	**188,060.10 €**	**206,206.26 €**	**231,734.19 €**	**269,320.00 €**	**589,761.47 €**	INCLUDING CAPEX (Y0)
Discounted Cash Flow (After Tax)	**-445,500.00 €**	**133,277.06 €**	**170,576.06 €**	**178,128.72 €**	**190,648.29 €**	**211,019.27 €**	**438,149.40 €**	INCLUDING CAPEX (Y0)

Business Case Results:	
NPV of Cash Flow	**438,149.40 €**
IRR	**32.2%**
Profitability Index	**1.98**
Simple Payback	**2 Years 7 Months**
Discounted Payback	**2 Years 10 Months**
ROI	**153%**

Case Assumptions:		
Cost Escalation Factor	**0.60%**	(0,50 above Eurozone inflation rate)
Benefit Escalation Factor	**1.60%**	(1 above cost escalation factor)
Corporate Tax Rate (Europe Average 2015)	**25.90%**	
Discount Rate	**5.00%**	
Interest Rate for Loans	**2.20%**	(Euribor plus 2%)
Maintenance Repair/Replacement Exp	**1.50%**	of CAPEX
Annual Increase	**10.00%**	
Marketing/Advertising/Promotion Exp	**4.00%**	of CAPEX
Annual Decrease	**50.00%**	
Property Mng, operations and bill Exp	**0.50%**	of CAPEX
Revenues from add-on services (SDN/NFV)	**15.00%**	

CAPEX Costs are NOT expensed in Y0 but CAPITALIZED in Y1-Y5 (Depreciation Expense)

Figure 244. CloudSat scenario #3 Normal case: CFM Financial analysis and financial ratios

Under normal financial, business and market conditions, CloudSat scenario #3 is a viable solution with appropriate financial results and ratios within a 5-years period. Findings are summarized in below figure:

Financial variables and ratios	Values
NPV of Cash Flow	438,149.40 €
IRR	32.2%
Profitability Index	1.98
Simple Payback	2 Years 7 Months
Discounted Payback	2 Years 10 Months
ROI	153%
Breakeven Year (excluding CAPEX)	Year 1

Figure 245. CloudSat scenario #3 Normal case: Financial variables and ratios results

This investment becomes positive in cash flows within year 1 (Breakeven Year) since we have very low OPEX figures, and after 2 years and 10 months reaches its Payback point (having make up for CAPEX and OPEX so far). The discounted cash flow after taxes at the end of 5th year has a NPV of 438,149.40€ (total profit of the 5 years at present value). This means that the investment is profitable within the 5 years period with an IRR 32.2%, Profitability index 1.98 and a ROI of 153% upon initial investment and variable costs. Please note that the high value of ROI points out the fact that from Year 1 we have a very high Cash Flow since Revenue figures are much more greater than corresponding OPEX figures (due to the zero leasing cost of satellite bandwidth).

b) Scenario #3 – "Optimistic" case analysis

Under Optimistic case conditions, the two conditional variables are assumed to have the below values:

Variables	Scenario #3
	Optimistic
Annual Revenues (annual percentage change)	20%
Discount Rate	2%

Figure 246. CloudSat scenario #3: Optimistic case variables

Consequently the revenues assumptions are:

ASSUMPTIONS	FIGURES	REMARKS
LEO system Annual Revenues (50 Terrestrial Gateways)	350,000,000.00 €	Used as base figure of Annual Revenues Estimates resulting in 7,000,000 revenues per gateway
1 gateway annual revenues	7,000,000.00 €	350,000,000.00 € / 50
CloudSat Scenario 3 - Number of Terrestrial Gateways as part of this financial analysis	1	
Estimated Annual Revenues for CloudSat Scenario 3 (1 gateway)	700,000.00 €	Based on the assumption that 1/10 of total gateway revenues are originating from the new CloudSat Scenario 3 services (7,000,000/10)
Initial (Y1) Revenues (% upon Estimated Annual Revenues)	30.00%	
Annual Revenues increase rate change	20.00%	An average percentage taking into account changes in revenues due to the potential loss of customers and/or the newcomers percentage over the year (eg 28% newcomers and 8% loss of customers)

Figure 247. CloudSat scenario #3 Optimistic case – Revenues Assumptions

The CloudSat case scenario #3 revenues are:

	Y1	Y2	Y3	Y4	Y5
Revenues per year	210,000.00 €	252,000.00 €	302,400.00 €	362,880.00 €	435,456.00 €
% on Annual Revenues of CloudSat Scenario 3	30.00%	36.00%	43.20%	51.84%	62.21%

Figure 248. CloudSat scenario #3 Optimistic case- Revenues

Therefore, the total revenues are calculated as per below figure.

	Y1	Y2	Y3	Y4	Y5
Revenues	210,000.00 €	252,000.00 €	302,400.00 €	362,880.00 €	435,456.00 €
Revenues from add-on services (SDN/NFV) as 15% of Revenues	31,500.00 €	37,800.00 €	45,360.00 €	54,432.00 €	65,318.40 €
Escalation of Benefits	3,864.00 €	9,347.79 €	16,960.98 €	27,355.82 €	41,364.61 €
TOTAL	245,364.00 €	299,147.79 €	364,720.98 €	444,667.82 €	542,139.01 €

Figure 249. CloudSat scenario #3 Optimistic case - Revenues

Summarizing all the estimated figures of CAPEX, OPEX and Revenues for CloudSat scenario #3 Optimistic case and applying the CFM model of analysis, we perform the

cash flow financial analysis and calculate the financial ratios as per the figure on the next page. Please note that Discount rate 2% has been applied.

CloudSat Scenario #3 - Optimistic conditions of Evaluation								
	Y0	1	2	3	4	5	Totals	
CAPEX								
Outdoor units	-100,000.00 €	0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	-100,000.00 €	
Modulator								
Encapsulator/Multiplexer	-360,000.00 €						-360,000.00 €	
Return Link sub-system								
GW Management & Access control								
Satellite Hub Generic Server HW (2 servers)	-10,000.00 €						-10,000.00 €	
Satellite Hub Generic Server Software Configuration	-3,500.00 €						-3,500.00 €	
SDN Switches (Two units)	-20,000.00 €						-20,000.00 €	
NFVI PoP (2 Units) HW (5 servers): 25000 each	-50,000.00 €						-50,000.00 €	
NFVI PoP (2 Units) SW+Licence+Configuration 3500 each	-7,000.00 €						-7,000.00 €	
VNF 1 (firewall) 5000 licence fee per year (x 5 servers)	-25,000.00 €						-25,000.00 €	
VNF 2 (web cache) 1500 fee per year (x 5 servers)	-7,500.00 €						-7,500.00 €	
VNF 3 (TCP Optimizer) 500 fee per year (x 5 servers)	-2,500.00 €						-2,500.00 €	
NFV Orchestrator (HW/SW)	-30,000.00 €						-30,000.00 €	
Fibre Leased Lines (1 GBps annual cost + initial installation)	-30,000.00 €						-30,000.00 €	
	0.00 €						0.00 €	
Bank Loan paid in 5 years with annual constant installments and constant/fixed interest rate	200,000.00 €						200,000.00 €	
Total CAPEX (for accounting)	**-445,500.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**-445,500.00 €**	
Actual CAPEX (excluding bank loan)	**-645,500.00 €**							
OPEX (VC Operating and Maintenance Costs)								
Satellite bandwidth leasing cost		0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	
Maintenance and Repair/Replacement/Unforeseen expenses (1.5% of CAPEX increased by 10% annually)		-9,682.50 €	-10,650.75 €	-11,715.83 €	-12,887.41 €	-14,176.15 €	-59,112.63 €	
Marketing/Advertising/Promotion expenses (4% of CAPEX reduced by 50% annually)		-25,820.00 €	-12,910.00 €	-6,455.00 €	-3,227.50 €	-1,613.75 €	-50,026.25 €	
Property management, operations, communications and bill/utilities expenses (0.5% of CAPEX)		-3,227.50 €	-3,227.50 €	-3,227.50 €	-3,227.50 €	-3,227.50 €	-16,137.50 €	
Supply expenses		-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-20,000.00 €	
Loan Payments		-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-200,000.00 €	
Interest Rate Expense		-4,400.00 €	-3,557.88 €	-2,697.23 €	-1,817.65 €	-918.71 €	-13,391.46 €	
Escalation of Costs		-522.78 €	-894.83 €	-1,233.09 €	-1,577.97 €	-1,941.24 €	-6,169.91 €	
Total Costs	**0.00 €**	**-87,652.78 €**	**-75,240.96 €**	**-69,328.64 €**	**-66,738.02 €**	**-65,877.35 €**	**-364,837.75 €**	
Revenue and Operating Benefits								
Revenues		210,000.00 €	252,000.00 €	302,400.00 €	362,880.00 €	435,456.00 €	1,562,736.00 €	
Revenues from add-on services (SDN/NFV) as 15% of Revenues		31,500.00 €	37,800.00 €	45,360.00 €	54,432.00 €	65,318.40 €	234,410.40 €	
Escalation of Benefits		3,864.00 €	9,347.79 €	16,960.98 €	27,355.82 €	41,364.61 €	98,893.21 €	
Total Benefits and Revenue	**0.00 €**	**245,364.00 €**	**299,147.79 €**	**364,720.98 €**	**444,667.82 €**	**542,139.01 €**	**1,896,039.61 €**	
Contribution Margin % (Sales - VC/Sales)		**64.28%**	**74.85%**	**80.99%**	**84.99%**	**87.85%**	**80.76%**	
Contribution Margin Amt (Sales - VC)		**157,711.22 €**	**223,906.83 €**	**295,392.34 €**	**377,929.80 €**	**476,261.66 €**	**1,531,201.85 €**	
Cash Flow Before Taxes	**-445,500.00 €**	**157,711.22 €**	**223,906.83 €**	**295,392.34 €**	**377,929.80 €**	**476,261.66 €**	**1,085,701.85 €**	INCLUDING CAPEX (Y0)
EBITDA		**153,311.22 €**	**220,348.95 €**	**292,695.11 €**	**376,112.15 €**	**475,342.95 €**	**1,517,810.39 €**	EXCLUDING CAPEX (Y0)
Breakeven Year			*					
Income Tax Calculation								
Depreciation Expense		-89,100.00 €	-142,560.00 €	-85,536.00 €	-51,321.60 €	-51,321.60 €	-419,839.20 €	
Operating and Maintenance Cost (OPEX)		-87,652.78 €	-75,240.96 €	-69,328.64 €	-66,738.02 €	-65,877.35 €	-364,837.75 €	
Revenue and Operating Benefits		245,364.00 €	299,147.79 €	364,720.98 €	444,667.82 €	542,139.01 €	1,896,039.61 €	
Net Income Taxes	**0.00 €**	**-17,770.31 €**	**-21,068.83 €**	**-54,352.79 €**	**-84,591.52 €**	**-110,059.48 €**	**-287,842.93 €**	
Cash Flow After Taxes	**-445,500.00 €**	**139,940.91 €**	**202,838.00 €**	**241,039.55 €**	**293,338.28 €**	**366,202.19 €**	**797,858.93 €**	INCLUDING CAPEX (Y0)
Discounted Cash Flow (After Tax)	**-445,500.00 €**	**137,196.97 €**	**194,961.56 €**	**227,136.95 €**	**270,999.22 €**	**331,680.60 €**	**716,475.31 €**	INCLUDING CAPEX (Y0)

Business Case Results:	
NPV of Cash Flow	**716,475.31 €**
IRR	**38.9%**
Profitability Index	**2.61**
Simple Payback	**2 Years 5 Months**
Discounted Payback	**2 Years 6 Months**
ROI	**187%**

Case Assumptions:		
Cost Escalation Factor	**0.60%**	(0,50 above Eurozone inflation rate)
Benefit Escalation Factor	**1.60%**	(1 above cost escalation factor)
Corporate Tax Rate (Europe Average 2015)	**25.90%**	
Discount Rate	**2.00%**	
Interest Rate for Loans	**2.20%**	(Euribor plus 2%)
Maintenance Repair/Replacement Exp	**1.50%**	of CAPEX
Annual Increase	**10.00%**	
Marketing/Advertising/Promotion Exp	**4.00%**	of CAPEX
Annual Decrease	**50.00%**	
Property Mng, operations and bill Exp	**0.50%**	of CAPEX
Revenues from add-on services (SDN/NFV)	**15.00%**	

CAPEX Costs are NOT expensed in Y0 but CAPITALIZED in Y1-Y5 (Depreciation Expense)

Figure 250. CloudSat scenario #3 Optimistic case: CFM Financial analysis and financial ratios

Under optimistic financial, business and market conditions, CloudSat scenario #3 is an extremely viable and profitable investment with very attractive financial figures within the 5-year period. Findings are summarized in below figure:

Financial variables and ratios	Values
NPV of Cash Flow	716,475.31 €
IRR	38.9%
Profitability Index	2.61
Simple Payback	2 Years 5 Months
Discounted Payback	2 Years 6 Months
ROI	187%
Breakeven Year (excluding CAPEX)	Early Year 1

Figure 251. CloudSat scenario #3 Optimistic case: Financial variables and ratios results

The favourable financial conditions of the optimistic case make this investment very promising early in time since it becomes positive in cash flows within first months of year 1 (Breakeven Year) and after 2 years and 6 months reaches its Payback point (having make up for CAPEX and OPEX so far). The discounted cash flow after taxes at the end of 5th year has a NPV of 716,475.31€ (total profit of the 5 years at present value). This means that the investment is profitable within the 5 years period with an IRR 38.9%, Profitability index 2.61 and a ROI of 187% upon initial investment and variable costs. Please note that the high value of ROI points out the fact that from Year 1 we have a very high Cash Flow since Revenue figures are much more greater than corresponding OPEX figures (due to the zero leasing cost of satellite).

c) Scenario #3 – "Pessimistic" case analysis

Under Pessimistic case conditions (more unstable economy and market/business environment), the two conditional variables are assumed to have the below values:

Variables	Scenario #3
	PESSIMISTIC
Annual Revenues (annual percentage change)	4%
Discount Rate	10%

Figure 252. CloudSat scenario #3: Pessimistic case variables

Consequently the revenues assumptions are:

ASSUMPTIONS	FIGURES	REMARKS

LEO system Annual Revenues (50 Terrestrial Gateways)	350,000,000.00 €	Used as base figure of Annual Revenues Estimates resulting in 7,000,000 revenues per gateway
1 gateway annual revenues	7,000,000.00 €	350,000,000.00 € / 50
CloudSat Scenario 3 - Number of Terrestrial Gateways as part of this financial analysis	1	
Estimated Annual Revenues for CloudSat Scenario 3 (1 gateway)	700,000.00 €	Based on the assumption that 1/10 of total gateway revenues are originating from the new CloudSat Scenario 3 services (7,000,000/10)
Initial (Y1) Revenues (% upon Estimated Annual Revenues)	30.00%	
Annual Revenues increase rate change	4.00%	An average percentage taking into account changes in revenues due to the potential loss of customers and/or the newcomers percentage over the year (eg 10% newcomers and 6% loss of customers)

Figure 253. CloudSat scenario #3 Pessimistic case – Revenues Assumptions

The CloudSat case scenario #3 revenues are:

	Y1	Y2	Y3	Y4	Y5
Revenues per year	210,000.00 €	218,400.00 €	227,136.00 €	236,221.44 €	245,670.30 €
% on Annual Revenues of CloudSat Scenario 3	30.00%	31.20%	32.45%	33.75%	35.10%

Figure 254. CloudSat scenario #3 Pessimistic case- Revenues

Therefore, the total revenues are calculated as per below figure.

	Y1	Y2	Y3	Y4	Y5
Revenues	210,000.00 €	218,400.00 €	227,136.00 €	236,221.44 €	245,670.30 €
Revenues from add-on services (SDN/NFV) as 15% of Revenues	31,500.00 €	32,760.00 €	34,070.40 €	35,433.22 €	36,850.54 €
Escalation of Benefits	3,864.00 €	8,101.42 €	12,739.58 €	17,807.63 €	23,336.59 €
TOTAL	245,364.00 €	259,261.42 €	273,945.98 €	289,462.28 €	305,857.43 €

Figure 255. CloudSat scenario #3 Pessimistic case - Revenues

Summarizing all the estimated figures of CAPEX, OPEX and Revenues for CloudSat scenario #3 Pessimistic case and applying the CFM model of analysis, we perform the cash flow financial analysis and calculate the financial ratios as per the figure on the next page. Please note that Discount rate 10% has been applied.

CloudSat Scenario #3 - Pessimistic conditions of Evaluation								
	Y0	1	2	3	4	5	Totals	
CAPEX								
Outdoor units	-100,000.00 €	0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	-100,000.00 €	
Modulator								
Encapsulator/Multiplexer	-360,000.00 €						-360,000.00 €	
Return Link sub-system								
GW Management & Access control								
Satellite Hub Generic Server HW (2 servers)	-10,000.00 €						-10,000.00 €	
Satellite Hub Generic Server Software Configuration	-3,500.00 €						-3,500.00 €	
SDN Switches (Two units)	-20,000.00 €						-20,000.00 €	
NFVI PoP (2 Units) HW (5 servers): 25000 each	-50,000.00 €						-50,000.00 €	
NFVI PoP (2 Units) SW+Licence+Configuration 3500 each	-7,000.00 €						-7,000.00 €	
VNF 1 (firewall) 5000 licence fee per year (x 5 servers)	-25,000.00 €						-25,000.00 €	
VNF 2 (web cache) 1500 fee per year (x 5 servers)	-7,500.00 €						-7,500.00 €	
VNF 3 (TCP Optimizer) 500 fee per year (x 5 servers)	-2,500.00 €						-2,500.00 €	
NFV Orchestrator (HW/SW)	-30,000.00 €						-30,000.00 €	
Fibre Leased Lines (1 GBps annual cost + initial installation)	-30,000.00 €						-30,000.00 €	
	0.00 €						0.00 €	
Bank Loan paid in 5 years with annual constant installments and constant/fixed interest rate	200,000.00 €						200,000.00 €	
Total CAPEX (for accounting)	**-445,500.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**0.00 €**	**-445,500.00 €**	
Actual CAPEX (excluding bank loan)	**-645,500.00 €**							
OPEX (VC Operating and Maintenance Costs)								
Satellite bandwidth leasing cost		0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	0.00 €	
Maintenance and Repair/Replacement/Unforeseen expenses (1.5% of CAPEX increased by 10% annually)		-9,682.50 €	-10,650.75 €	-11,715.83 €	-12,887.41 €	-14,176.15 €	-59,112.63 €	
Marketing/Advertising/Promotion expenses (4% of CAPEX reduced by 50% annually)		-25,820.00 €	-12,910.00 €	-6,455.00 €	-3,227.50 €	-1,613.75 €	-50,026.25 €	
Property management, operations, communications and bill/utilities expenses (0.5% of CAPEX)		-3,227.50 €	-3,227.50 €	-3,227.50 €	-3,227.50 €	-3,227.50 €	-16,137.50 €	
Supply expenses		-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-4,000.00 €	-20,000.00 €	
Loan Payments		-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-40,000.00 €	-200,000.00 €	
Interest Rate Expense		-4,400.00 €	-3,557.88 €	-2,697.23 €	-1,817.65 €	-918.71 €	-13,391.46 €	
Escalation of Costs		-522.78 €	-894.83 €	-1,233.09 €	-1,577.97 €	-1,941.24 €	-6,169.91 €	
Total Costs	**0.00 €**	**-87,652.78 €**	**-75,240.96 €**	**-69,328.64 €**	**-66,738.02 €**	**-65,877.35 €**	**-364,837.75 €**	
Revenue and Operating Benefits								
Revenues		210,000.00 €	218,400.00 €	227,136.00 €	236,221.44 €	245,670.30 €	1,137,427.74 €	
Revenues from add-on services (SDN/NFV) as 15% of Revenues		31,500.00 €	32,760.00 €	34,070.40 €	35,433.22 €	36,850.54 €	170,614.16 €	
Escalation of Benefits		3,864.00 €	8,101.42 €	12,739.58 €	17,807.63 €	23,336.59 €	65,849.21 €	
Total Benefits and Revenue	**0.00 €**	**245,364.00 €**	**259,261.42 €**	**273,945.98 €**	**289,462.28 €**	**305,857.43 €**	**1,373,891.11 €**	
Contribution Margin % (Sales - VC/Sales)		**64.28%**	**70.98%**	**74.69%**	**76.94%**	**78.46%**	**73.44%**	
Contribution Margin Amt (Sales - VC)		**157,711.22 €**	**184,020.46 €**	**204,617.34 €**	**222,724.26 €**	**239,980.08 €**	**1,009,053.36 €**	
Cash Flow Before Taxes	**-445,500.00 €**	**157,711.22 €**	**184,020.46 €**	**204,617.34 €**	**222,724.26 €**	**239,980.08 €**	**563,553.36 €**	INCLUDING CAPEX (Y0)
EBITDA		**153,311.22 €**	**180,462.58 €**	**201,920.11 €**	**220,906.61 €**	**239,061.37 €**	**995,661.90 €**	EXCLUDING CAPEX (Y0)
Breakeven Year				*				
Income Tax Calculation								
Depreciation Expense		-89,100.00 €	-142,560.00 €	-85,536.00 €	-51,321.60 €	-51,321.60 €	-419,839.20 €	
Operating and Maintenance Cost (OPEX)		-87,652.78 €	-75,240.96 €	-69,328.64 €	-66,738.02 €	-65,877.35 €	-364,837.75 €	
Revenue and Operating Benefits		245,364.00 €	259,261.42 €	273,945.98 €	289,462.28 €	305,857.43 €	1,373,891.11 €	
Net Income Taxes	**0.00 €**	**-17,770.31 €**	**-10,738.26 €**	**-30,842.07 €**	**-44,393.29 €**	**-48,862.55 €**	**-152,606.47 €**	
Cash Flow After Taxes	**-445,500.00 €**	**139,940.91 €**	**173,282.20 €**	**173,775.27 €**	**178,330.97 €**	**191,117.53 €**	**410,946.89 €**	INCLUDING CAPEX (Y0)
Discounted Cash Flow (After Tax)	**-445,500.00 €**	**127,219.01 €**	**143,208.43 €**	**130,559.94 €**	**121,802.45 €**	**118,668.95 €**	**195,958.78 €**	INCLUDING CAPEX (Y0)

Business Case Results:		**Case Assumptions:**	
NPV of Cash Flow	**195,958.78 €**	**Cost Escalation Factor**	**0.60%** (0,50 above Eurozone inflation rate)
IRR	**25.2%**	**Benefit Escalation Factor**	**1.60%** (1 above cost escalation factor)
Profitability Index	**1.44**	**Corporate Tax Rate (Europe Average 2015)**	**25.90%**
Simple Payback	**2 Years 9 Months**	**Discount Rate**	**10.00%**
Discounted Payback	**3 Years 4 Months**	**Interest Rate for Loans**	**2.20%** (Euribor plus 2%)
ROI	**123%**	**Maintenance Repair/Replacement Exp**	**1.50%** of CAPEX
		Annual Increase	**10.00%**
		Marketing/Advertising/Promotion Exp	**4.00%** of CAPEX
		Annual Decrease	**50.00%**
		Property Mng, operations and bill Exp	**0.50%** of CAPEX
		Revenues from add-on services (SDN/NFV)	**15.00%**
		CAPEX Costs are NOT expensed in Y0 but CAPITALIZED in Y1-Y5 (Depreciation Expense)	

Figure 256. CloudSat scenario #3 Pessimistic case: CFM Financial analysis and financial ratios

Under pessimistic financial, business and market conditions, CloudSat scenario #1 remains still viable but later in time (close to 3rd year). Findings are summarized in below figure:

Financial variables and ratios	Values
NPV of Cash Flow	195,958.78 €
IRR	25.2%
Profitability Index	1.44
Simple Payback	2 Years 9 Months
Discounted Payback	3 Years 4 Months
ROI	123%
Breakeven Year (excluding CAPEX)	Year 1

Figure 257. CloudSat scenario #3 Pessimistic case: Financial variables and ratios results

This investment reaches profitability within the 5-year period. In specific, it acquires positive cash flows from the first year (Breakeven Year), due to the very low OPEX, and after 3 years and 4 months it reaches its Payback point (having make up for CAPEX and OPEX so far).

This investment is still viable by the end of the 5-year period of analysis. The discounted cash flow after taxes at the end of 5th year has a NPV of 195,958.78€ (total profit of the 5 years at present value). This means that the investment is profitable within the 5 years period with an IRR 25.2%, Profitability index 1.44 and a ROI of 123% (due to the very low OPEX) upon initial investment, variable costs and operations. To sum up, even under pessimistic conditions, we consider this case viable too, with great profitability before the completion of the 5th year.

Overall, it is worth noticing that in CloudSat case scenario #3, all CFM financial and ratios results confirm that no matter the conditions (Normal, Optimistic, Pessimistic) the investment is viable early in time with great profitability figures. Most financial ratios have high values even under the pessimistic conditions fact which confirms that when the OPEX of an investment is low, this investment can attain high rates of return and profitability after the first couple of years in which it will have make up for the entire CAPEX and running OPEX figures.

7.2.4. CloudSat Case Scenarios Cost-Benefit Analysis (CBA)

In the previous sections we financially analyzed and evaluated the three CloudSat case scenarios under three different business environment scenarios: Normal, Optimistic and Pessimistic. In this section the results will be compared and evaluated in order to derive to meaningful and financially valuable conclusions concerning the viability and benefits of CloudSat concept.

The figure below summarizes and compares the basic financial analysis figures as result of the CFM methodology for the three CloudSat case scenarios under the "Normal" condition of evaluation. Amounts refer to the total amount of each financial component at the end of the 5-year period (or at the Y0 for CAPEX).

Financial Component	Scenario #1	Scenario #2	Scenario #3
CAPEX (including bank loan)	-342,000.00 €	-337,000.00 €	-645,500.00 €
OPEX	-5,904,796.34 €	-1,525,788.57 €	-364,837.75 €
Revenues	6,176,700.00 €	1,619,976.08 €	1,334,097.95 €
Revenues from add-on services	926,505.00 €	242,996.41 €	200,114.69 €
Escalation of Benefits Amount	375,008.20 €	98,349.49 €	80,993.70 €
Contribution Margin	21.04%	22.21%	77.41%
Contribution Amount	1,573,416.85 €	435,533.41 €	1,250,368.58 €
Cash Flow Before Taxes (including CAPEX)	1,431,416.85 €	298,533.41 €	804,868.58 €
EBITDA (excluding CAPEX)	1,560,025.39 €	422,141.95 €	1,236,977.12 €
Net Income Taxation	-372,855.38 €	-79,363.98 €	-215,107.11 €
Cash Flow After Taxes	1,058,561.48 €	219,169.44 €	589,761.47 €
Discounted Cash Flow	839,365.12 €	155,764.31 €	438,149.40 €

Figure 258. CloudSat Case Scenarios – Comparison of financial figures at the end of Y5

By analyzing the above figures we notice that all three scenarios are financially viable despite the different configuration/characteristics they have. All three CloudSat case scenarios are financially viable under a 5-year analysis period even if the worst-case scenarios (pessimistic conditions) apply.

Below figure compares the financial ratios of the three CloudSat case scenarios under the normal conditions of evaluation.

Financial variables and ratios	Scenario 1	Scenario 2	Scenario 3
NPV of Cash Flow	839,365.12 €	155,764.31 €	438,149.40 €
IRR	71.90%	27.6%	32.2%
Profitability Index	6.91	2.14	1.98
Simple Payback	2 Years 4 Months	3 Years 4 Months	2 Years 7 Months
Discounted Payback	2 Years 5 Months	3 Years 7 Months	2 Years 10 Months
ROI	26%	25%	153%
Breakeven Year (excluding CAPEX)	Middle Year 2	Late Year 2	Year 1

Figure 259. CloudSat Case Scenarios – Financial ratios comparison under Normal conditions

The financial ratios also confirm that all three cases are financially viable over the 5-year period with very close financial indices. Scenario #1 is a bit more promising, having greater profitability at present value (839,365.12 €), higher profitability index (6.91) and sooner Payback period (2 years and 5 months).

As analyzed, even under the pessimistic scenario conditions, all three cases remain viable even though the conditions are marginal (when cash flow discounted at present value at a high discount rate of 10%).

Financial variables and ratios	Scenario 1	Scenario 2	Scenario 3
Cash Flow	229,655.67 €	2,037.45 €	410,946.89 €
NPV of Cash Flow	105,317.62 €	-41,287.97 €	195,958.78 €
IRR	25.3%	0.4%	25.2%
Profitability Index	1.74	0.70	1.44
Simple Payback	3 Years 6 Months	4 Years 11-12 Months	2 Years 9 Months
Discounted Payback	4 Years 0 Months	Above 5 years	3 Years 4 Months
ROI	7%	8%	123%
Breakeven Year (excluding CAPEX)	Late Year 3	Year 3	Year 1

Figure 260. CloudSat Case Scenarios – Financial ratios comparison under Pessimistic conditions

For a 5-year period of analysis and under pessimistic market and financial conditions, the results of the above financial analysis are satisfactory for all the three CloudSat case scenarios. If we exclude the negative figures of the NPV of Cash flow and the discounted Payback period for Scenario 2, all the rest ratios are positive, confirming the financial viability of any Satcom deployment following the CloudSat proposed

configuration. This view is strengthened even more by the fact that Cash Flows at the end of year 5 are positive for all cases.

Scenario 3 is proven to be the most resistant of all even under pessimistic financial, market and business conditions. Since its OPEX cost is minimized because of the zero satellite leasing cost it can better cope with the financial difficulties originating from the high discount rate and inadequate annual increase of revenues. Similarly Scenario 2, although being a marginal case under pessimistic conditions for a 5-year period, it has positive cash flow and ratios proving that it can turn into a profitable investment with a minor extension of time (during year 6).

Obviously, the evaluation of benefits is even more promising under optimistic conditions where discount rates are minimized (2%) and revenues increase much more dynamically. Below figure summarizes the financial ratios of the three CloudSat case scenarios under optimistic market and financial conditions.

Financial variables and ratios	Scenario 1	Scenario 2	Scenario 3
NPV of Cash Flow	1,849,089.74 €	424,602.84 €	716,475.31 €
IRR	103.0%	46.6%	38.9%
Profitability Index	14.02	4.10	2.61
Simple Payback	1 Year 11-12 Months	2 Years 9 Months	2 Years 5 Months
Discounted Payback	2 Years 0 Months	2 Years 10 Months	2 Years 6 Months
ROI	47%	46%	187%
Breakeven Year (excluding CAPEX)	Early Year 2	Early Year 2	Early Year 1

Figure 261. CloudSat Case Scenarios – Financial ratios comparison under Optimistic conditions

Similar conclusions concerning the financial viability of the three CloudSat case scenarios are reached if we examine in parallel the financial ratios per case conditions for each scenario.

Below figure compares the financial values of Scenario #1 under the three different conditions.

Scenario #1 Financial ratios	Normal	Optimistic	Pessimistic
Cash Flow	1,058,561.48 €	2,019,425.79 €	229,655.67 €
NPV of Cash Flow	839,365.12 €	1,849,089.74 €	105,317.62 €
IRR	71.90%	103.0%	25.3%
Profitability Index	6.91	14.02	1.74
Simple Payback	2 Years 4 Months	1 Year 11-12 Months	3 Years 6 Months
Discounted Payback	2 Years 5 Months	2 Years 0 Months	4 Years 0 Months

Scenario #1 Financial ratios	Normal	Optimistic	Pessimistic
ROI	26%	47%	7%
Breakeven Year (excluding CAPEX)	Middle Year 2	Early Year 2	Late Year 3

Figure 262. CloudSat Case Scenario #1 – Financial ratios comparison under three conditions

All ratios, under any of the three conditions, show promising viability and high expectations even in the pessimistic case that has been already justified in previous paragraphs. The same beneficial conclusion is reached by the graphical representation of Cash flows and Discounted Cash flows over the 5-year period (per year representation) and in total in the figures below.

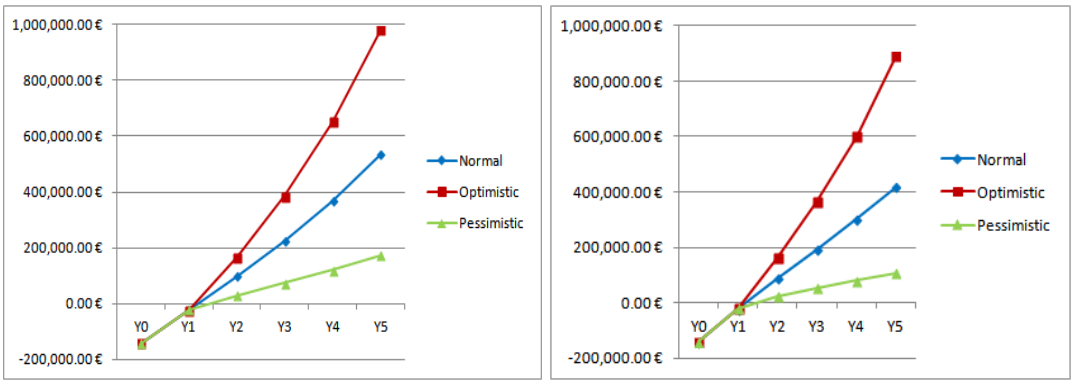

Figure 263. CloudSat Case Scenario #1 – Cash flows and Discounted Cash flow per year under three conditions

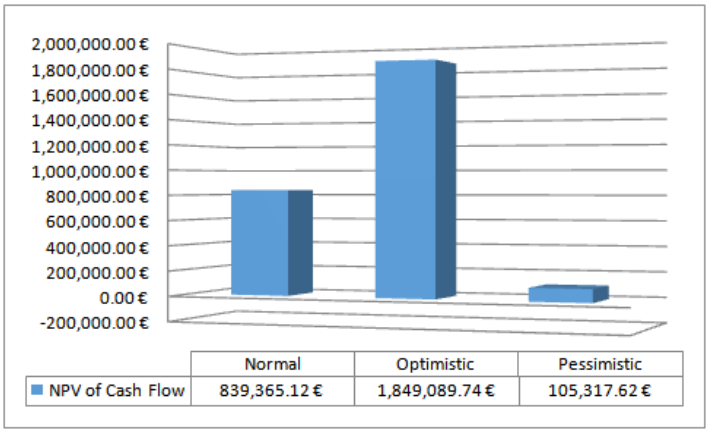

Figure 264. CloudSat Case Scenario #1 – NPV of Cash flow under three conditions

Similarly below are compared the financial values of Scenario #2 under the three different conditions.

Scenario #2 Financial ratios	Normal	Optimistic	Pessimistic
Cash Flow	219,169.44 €	471,859.20 €	2,037.45 €
NPV of Cash Flow	155,764.31 €	424,602.84 €	-41,287.97 €

IRR	27.6%	46.6%	0.4%
Profitability Index	2.14	4.10	0.70
Simple Payback	3 Years 4 Months	2 Years 9 Months	4 Years 11-12 Months
Discounted Payback	3 Years 7 Months	2 Years 10 Months	Above 5 years
ROI	25%	46%	8%
Breakeven Year (excluding CAPEX)	Late Year 2	Early Year 2	Year 3

Figure 265. CloudSat Case Scenario #2 – Financial ratios comparison under three conditions

All financial ratios, under any of the three conditions, show viability with high expectations with the exception of the NPV of Cash flow in the pessimistic case that has been already justified in previous paragraphs. The same is further elaborated by the graphical representation of Cash flows and Discounted Cash flows over the 5-year period (per year representation) and in total in the figures below.

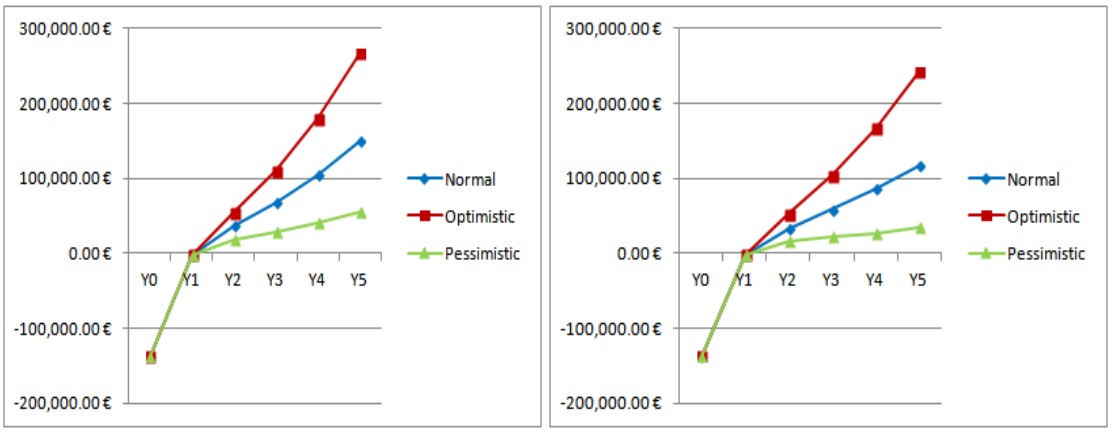

Figure 266. CloudSat Case Scenario #2 – Cash flows and Discounted Cash flow per year under three conditions

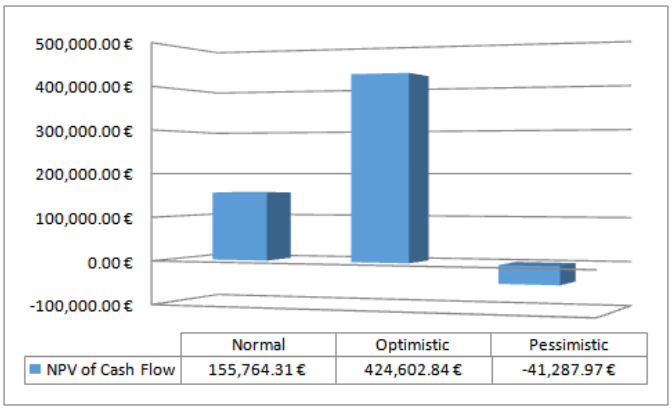

Figure 267. CloudSat Case Scenario #2 – Total Discounted Cash flow under three conditions

Finally below are compared the financial values of Scenario #3 under the three different conditions.

Scenario #3 Financial ratios	Normal	Optimistic	Pessimistic

Cash Flow	589,761.47 €	797,858.93 €	410,946.89 €
NPV of Cash Flow	438,149.40 €	716,475.31 €	195,958.78 €
IRR	32.2%	38.9%	25.2%
Profitability Index	1.98	2.61	1.44
Simple Payback	2 Years 7 Months	2 Years 5 Months	2 Years 9 Months
Discounted Payback	2 Years 10 Months	2 Years 6 Months	3 Years 4 Months
ROI	153%	187%	123%
Breakeven Year (excluding CAPEX)	Year 1	Early Year 1	Year 1

Figure 268. CloudSat Case Scenario #3 – Financial ratios comparison under three conditions

Above financial ratios, under any of the three conditions, show viability with great expectations with the exception of the NPV of Cash flow in the pessimistic case that has been already justified in previous paragraphs. The same is further elaborated by the graphical representation of Cash flows and Discounted Cash flows over the 5-year period (per year representation) and in total in the figures below.

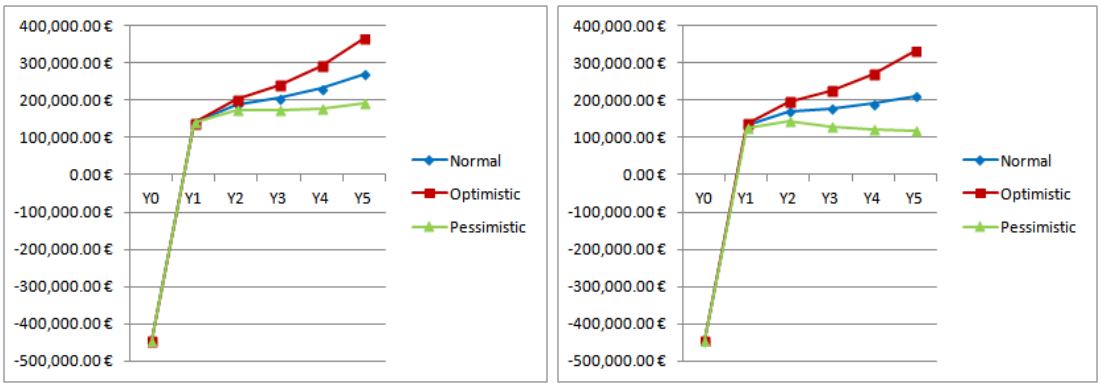

Figure 269. CloudSat Scenario #3 – Cash flows & Discounted Cash flow per year under three conditions

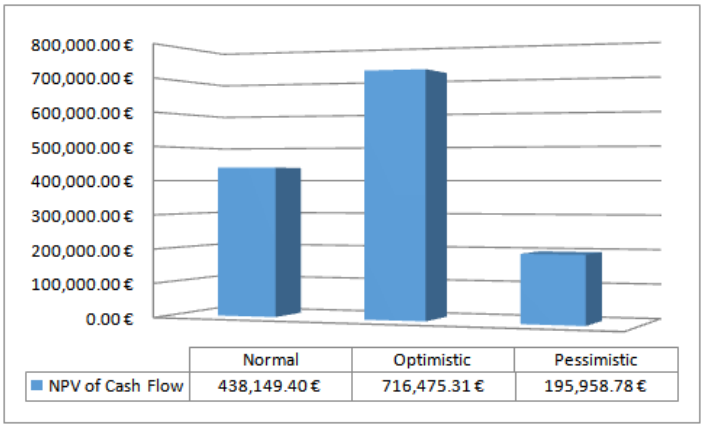

Figure 270. CloudSat Case Scenario #3 – Total Discounted Cash flow under three conditions

In summary, all comparison results and graphs lead us to the conclusion that any SW-based Satcom deployment with virtualization services (SDN/NFV) are financially viable provided that they operate under the same business and market environment. All three CloudSat case scenarios are financially viable under the 5-year analysis period even if the worse-case scenarios (pessimistic conditions) apply.

A summary of all the parameters analysed and considered in the CBA is provided in the table below, showing the impact of each aspect considering each part of Cost benefit analysis:

- Direct and indirect benefits
- Positive and negative impacts
- Economical gains and limitations/constraints
- Evaluation of social affect
- Conclusions on the financial viability and suitability of CloudSat

CloudSat Outcomes	Benefit	Gain	Impact (Business side)	Social Impact
Initial investment (CAPEX) cost reduction	Direct Benefit	Economical Gain	Positive impact on OPEX, Revenues, Taxation, Profit	
Reduction in OPEX	Direct Benefit	Economical Gain	Positive impact on Revenues, Taxation, Profit	
Increase in Revenues	Indirect Benefit	Economical Gain	Positive impact on the business, Negative impact on Taxation	Positive
Short Payback period	Indirect Benefit	Economical Gain	Positive impact on Finance	
Positive Financial Ratios	Indirect Benefit	Economical Gain	Positive impact on Finance	Positive
Reduced System Maintenance Cost	Indirect Benefit	Economical Gain	Positive impact on Finance	
Less resources utilization	Indirect Benefit	Economical Gain	Positive impact on Finance	
Accommodates increases in workload or demand without additional costs	Indirect Benefit	Economical Gain	Positive impact on Finance and Market	
Less prone to Hardware Failures	Indirect Benefit	Economical Gain	Positive impact on Finance and Market	Positive
Provision of new services	Indirect Benefit	Social Gain	Positive impact on Market	Positive
Improved QoS	Indirect Benefit	Social Gain	Positive impact on Market	Positive
Improved Effectiveness	Indirect Benefit	Social Gain	Positive impact on Market	Positive
Improved Ability to Deliver	Indirect Benefit	Social Gain	Positive impact on	Positive

CloudSat Outcomes	Benefit	Gain	Impact (Business side)	Social Impact
			Market	
Increased Security Risks	Limitation/ Constraint		Negative Impact	
Increased Software dependence	Limitation/ Constraint		Negative Impact	
Prone to Software Bugs and Software dependencies	Limitation/ Constraint		Negative Impact	
Dependencies to other technologies (SDN/NFV)	Limitation/ Constraint		Negative Impact	

Figure 271. CloudSat Cost-Benefit Outcomes

8. Future Work Recommendations, Technology Development Roadmap and Standards Evolution

Based on the findings and conclusions from all the phases of the study, as well as the lessons learnt throughout the project and the evolution of the software network landscape, this concluding chapter outlines key recommendations for future work, also including a roadmap for technology development and interaction with standardisation and relevant R&D initiative.

8.1. Lessons learnt from CloudSat scenarios with current Virtualisation technologies

8.1.1. SWOT analysis for integration of cloud networking enablers in satcom

Chapter 3 investigated the dimensions of suitability –regarding integration with satcom- for the most pertinent enabler technologies in the cloud networking domain. Software-Defined Networking (SDN) and Network Functions Virtualisation (NFV) were eventually selected as the most promising technologies which would help to integrate the cloud networking model into satellite networks and optimally facilitate the seamless integration between satellite and terrestrial.

Based on the analysis done on the study deliverables, the discussion on the integrated architectures, the findings of the experimentation campaign as well as the conclusions of the techno-economic analysis, we can summarize in a SWOT matrix the main strengths, weaknesses, opportunities and threats associated with the integration of SDN (Table 17) and NFV (Table 18) in satcom networks.

These tables summarise and consolidate only the main technical and business arguments which were derived from the project work; individual/minor aspects discussed in the previous Chapters are not repeated here.

Table 17.Main S/W/O/Ts associated with SDN integration in satcom

STRENGTHS	WEAKNESSES

• SDN greatly simplifies satcom network management and also enables vendor-agnostic control. • SDN significantly promotes integration with terrestrial (via unified management) and smoother inclusion of satcom in the 5G landscape. • Mobility and QoS management is simplified. • SDN can be adopted with minimal interventions at certain network points (as shown during experimentation)	• SDN (in its current form) cannot control multiple access procedures as well as PHY parameters. • SDN switching agility not well suited for single-beam configurations. • SDN capabilities onboard require additional resources/ power/ physical space.
OPPORTUNITIES	**THREATS**
• SDN controllers such as OpenDaylight are rapidly evolving and becoming more and more mature. • SDN is considered a key ingredient of 5G networks. • SDN market is rapidly growing (CAGR > 50% expected by 2023)	• SDN has not yet found its place in production ISP networks. • The SDN controller landscape is still evolving. • The SDN protocol landscape is also not very stable; Openflow is a de facto standard but NETCONF/YANG and OVSDB are also gaining ground (although they can be used complementarily with OF) • Openflow is rapidly evolving as a protocol and significant changes exist from version to version, which raises issues with regard to its integration in long-term infrastructures (such as satellite payloads).

Table 18.Main S/W/O/Ts associated with NFV integration in satcom

STRENGTHS	WEAKNESSES
• Satcom sector gains access to most	• VNFs are generally less reliable than

benefits of NFV (market openness, reduced TTM, service agility, reconfigurability, frequent upgrades etc.) • New service offerings (VNFaaS) for satcom • CAPEX/OPEX reduction at an estimate of 25% for a satcom SP • SVNO scenario significantly enhanced with virtualisation of core functions • Cloud-RAN scenario especially attractive for multi-GW LEO configurations (in the longer term) or HTS • NFV can be adopted in the short term (before SDN) for certain components (as also demonstrated during experimentation) • NFV services can be instantiated within seconds (as demonstrated during experimentation)	hardware functions, which is quite critical for core satcom functions. • VNF security and performance issues are also critical for core satcom functions. • NFV capabilities onboard require additional resources/ power/ physical space. • NFV at the terminal is also subject to resource constraints
OPPORTUNITIES	**THREATS**
• IT virtualisation technologies (required for NFV) are stable and proven. • The NFV domain currently has a very strong momentum and is backed by a very wide industrial and academic communities • The progress of NFV relies on several open-source community projects, eliminating the danger of vendor lock-in; OpenStack and OPNFV are rapidly evolving • NFV market is rapidly growing (CAGR > 60% expected by 2023)	• NFV technology development landscape not stable yet (OPNFV still at early stage and with continuously evolving roadmap, several non-interoperable Orchestrators at alpha versions etc.) • ETSI NFV “standards” currently are quite high-level and do not ensure interoperability between different platforms • Well-established satcom vendors are quite sceptic with regard to the market openness

Going a bit deeper, the next section investigates the techno-economic feasibility of such an integration from the perspective of the different actors of the value chain.

8.1.2. Techno-Economic feasibility for the different actors perspective

For a new technology, bridging the gap from low TRL (1-4) to high TRL (5-9), generally requires additional levels of financial investment in order to develop or to secure the maturity for operational and commercial environments. The bridging step – or rather the successive steps – are critical, not only to validate the technology itself, but also to get tangible hints and evidence that the technology fits the targeted ecosystem(s), and at the appropriate timeframes.

This section discusses the techno-economic feasibility for SDN and NFV adoption from the different actors perspective, based on the CloudSat value chain identified in Chapter 7.

8.1.2.1. Practical derivation of the CloudSat value chain

Chapter 7 presented the generic value chain for a softwarised satcom ecosystem, as shown in Figure 272 below.

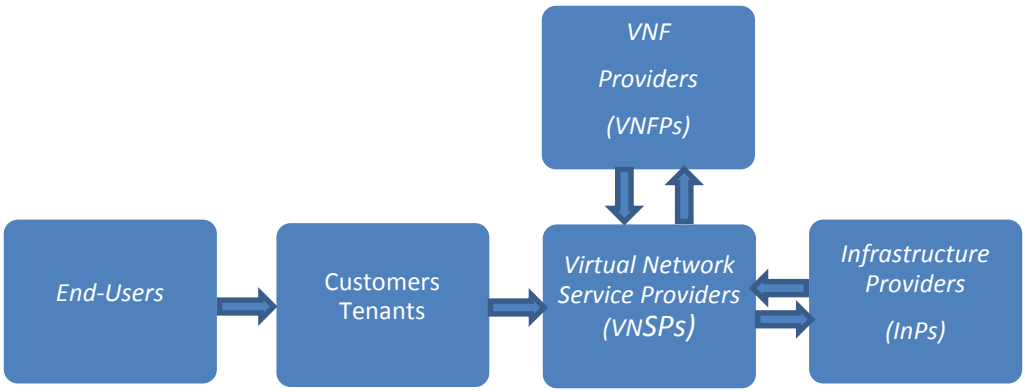

Figure 272. Generic CloudSat value chain – Interaction, interoperability, and interdependance of actors

From this model, at least two concrete chains could be derived to fit the approaches we know today.

The first variation considers a direct transposition of the current model to be applicable at a short term.

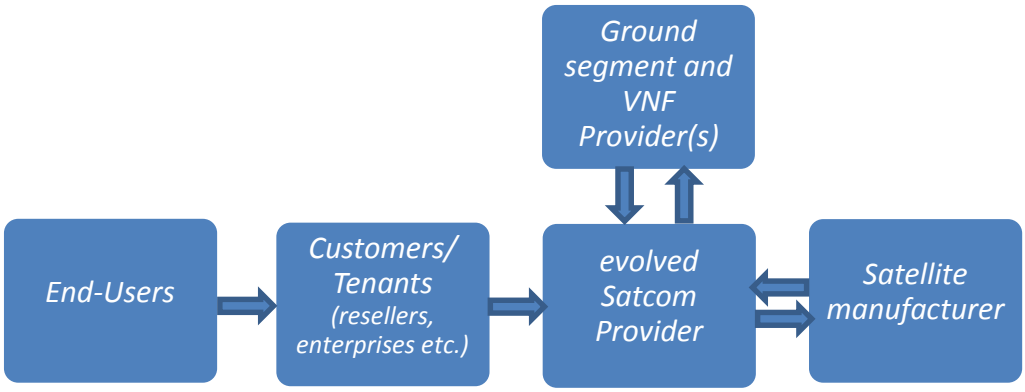

Figure 273. Variation #1 of the CloudSat value chain

In this model, the satcom provider owns and manages NCC and GW. It sells satellite bandwidth, at service level (e.g. IP data rates) to its customers who remain unaffected by the satellite access capabilities. The Customer role can be undertaken by e.g. satcom resellers, enterprise customers etc. Further, the satcom provider offers new features (SDN and/or NFV).

For the second variation shown in Figure 274, the proposed value chain positions a SVNO in conformance with the DVB-RCS2 model [DVBRCS2]. While the SVNO now manages low-level resources, the hardware centralized infrastructure (GW/NCC) management, operations and supervision normally remains at the charge of the SNO.

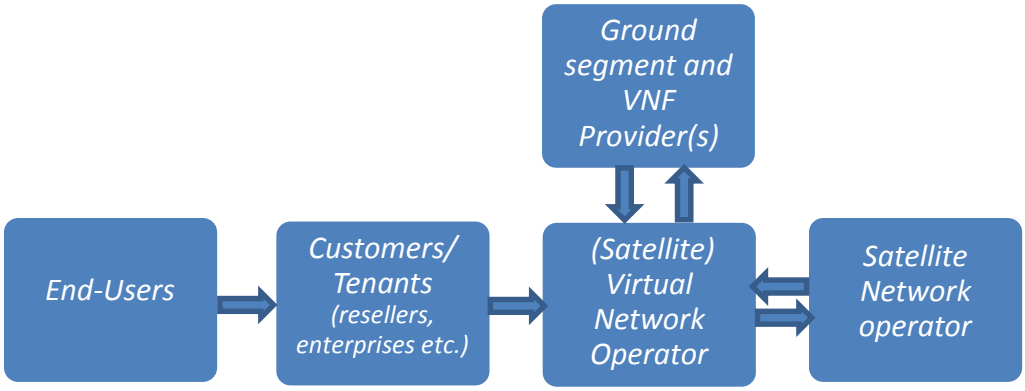

Figure 274. Variation #2 of the CloudSat value chain

Finally, note that the concept of **Virtual Satellite operator (VSO)** could also be considered for the long term, i.e. the support of sharing satellite resources with Virtualization technologies for satellite payloads. This case is more described in section 8.2.1.2. .

8.1.2.2. Satcom network vendors and integrators

From a pure satcom network vendor perspectives, we see different/complimentary levels of interest for SDN and NFV.

NFV

Benefits and Opportunities

NFV introduces different types of benefits. We propose to review the main benefits applicable for satcom, with *direct and/or indirect* tagging. A *direct* benefit is applicable for the own operations of the vendors. In addition, since the profitability of Service Providers must remain a strong concern for the network providers, *indirect benefits* shall also be mentioned. Those benefits shall largely participate to the increase of the competitiveness of the network vendors.

- Improved ease of product evolution *(direct benefits)*
    - However this may be partially balanced by the necessary support of multiple and/or incompatible versions of NFV infrastructure and management components
- Acceleration of Assembly, Integrations and Tests (AIT) phases thanks to Hardware abstraction *(direct benefits)*
- Enable concurrent/customized versions of features to be deployable in multi-Service Providers environments on shared infrastructure *(direct benefits).* We can cite the following examples:
    - Dedicated Access Control Layer rules per SP
    - Customized Traffic optimizations per SP/per customer
    - Customized billing model per SP/ per services
- Ease the lifecycle support of complete solution portfolios to address the various needs and markets, with common platform architecture basis. For example, this can be applied to different categories of satellite Hubs and terminals, from low-cost access solutions to high-end/professional/VSAT trunking) *(direct benefits).*For example, using a common architecture with customized and virtualized instances customized for each specific products, presenting the same interfaces and logics, increase the level of reusability and bring opportunities to offer more customization – to some extent.
- Access to new services by SPs *(indirect benefits).* Note that these concepts can often be associated to PaaS/SaaS approaches
    - core network services in satcom systems (i.e. basic traffic management, forwarding)
    - value-added network services (firewall, PEP, but also CDN, EPC, IMS functions, etc. that could be virtualized as components of service-oriented network architectures)
    - Also high-level services (example Applicative-level acceleration, specific applications servers, etc.)
- High Availability, Reliability and Manageability *(direct and indirect benefits)*

    - Replacement of failing “unit” can be operated and managed much easily, via simple software update (as long as Hardware and Virtualization Layer service is operational)

- Support on-demand execution of dedicated processing at network and higher layer (L4-L7), to mitigate the impact of congestions *(direct and indirect benefits)*. The system can react in short time (few seconds) with various strategies, depending on for example on
    - total traffic load
    - traffic repartition between services
- Integration of customer NFV functions at key satcom locations (at centralized nodes at/being satellite Gateway, and at satellite terminals) than must run on qualified appliances *(indirect benefits).*

Uncertainties, weaknesses, risks and threats

On the other hand, since basically NFV is not integrated in today’s satcom technologies, detailed business case studies need to be performed to further analyse the techno-economic feasibility, given also the exact requirements/constraints of each vendor as well as the technical specification of the products in market (e.g. satcom hubs). These studies will essentially be a vendor-specific specialisation of the generic study presented in Chapter 7.

In specific, some important variations is expected for the different vendors. This is due to the fact that the above functions may present for the vendors very heterogeneous different business scales, with assumptions of CAPEX/OPEX revenues, and projections, that may vary by one or several orders of magnitude. Also, companies involved in low production volumes would need, in proportion, higher amounts of investment, compared to the big competitors. Equivalently, the break-even point will be achieved at different times. For the smaller companies, virtualization is not necessary indicated as products and associated business can sometimes be more specialized and less adapted to a broad range of services.

If the support of on-demand (“on-off”) services based on NFV proves to be a good opportunity, this must be accompanied by suitable and stable management and monetization tools to support it. Since the business models related to this flexibility is far from being clear today, there might be additional delays before NFV can be fully ready for integration in commercial satcom systems.

One could wonder whether NFV could present any degree of threat for satcom vendors, regarding competitiveness, in the way terrestrial vendors could access (part of) the future satellite market. This question certainly exceeds the sole case of NFV when thinking about technology reuse, integration and interoperability of satellite and terrestrial systems. The 5G system development and deployment where NFV technologies are expected to play an important role, could illustrate well this situation, taking into account the risk of satcom constraining its role to a minimal data piping service. Since the development of terrestrial access radio technologies and

services are extremely rapid against satcom, this might be the only way satcom systems remain attractive in the long term.

From a more technical perspective, NFV may also present some drawbacks due to higher technical complexity requiring at the infrastructure.

- Higher resources (CPU/RAM) are necessary to support all the NFV middleware, in comparison to non-software/virtualized version. While this matter is of lesser importance for the candidate NFV functions located in the core of the network, there may be strong issues or even impossibility to deploy NFV for low-cost and very low-cost terminals. Residential broadband access (with individual terminals), and large-scale Machine-Type Communications (IoT/M2M) are two important use cases that might be affected by this limitation.
- In addition, the sensitivity to massive I/O operations and the possibilities for VNFs to bypass the virtualization layer for a direct access to network interfaces also appear as strong requirements. Although latest virtualisation technologies can support this, we need to make sure that no bottlenecks are introduced due to virtualisation. Indeed, it is known today that current OS virtualization induce overheads but without real penalty in most production environment with high resource (datacenters, etc.). The spatial environment has important difference here. However, on-going studies already demonstrated direct access and very limited overhead were achievable if specifically addressed by the hardware[16] ()

We also notice that Security brings additional constraints for satcom access network integrators and the satcom equipment providers with respect to the implementation of firewalling and access protection at key location(s) of the ground segment, at least at border/interconnection nodes. This means that any VNF controlled by e.g. customers will need to comply with the overall security policy of the SP. Additional access rules might need to be applied to prevent risk of misconfiguration by the VNF user or even software bugs.

Finally, NFV must still be considered today as a young technology. Despite of sounding concepts and existing implementations (such as OPNFV), we can expect that the massive adoption – or not– from the main terrestrial network vendors and operators will send strong signals to other communities like the satcom industry.

Related to this, some different flavours and/or evolutions of NFV might also appear and compete, with the risk of introducing important delay for adoption in satcoms. The risk lays in the definition of the interfaces (and the underlying model) between the local implementation of the Virtualization layer (client) and the VNFs on one hand, and on the definition of the Virtualization layer itself between Virtualisation clients and the NFV management entities (namely the NFV orchestrator, the VNF Manager and the Virtualized Infrastructure Manager) on the other hand.

[16] See studies on Flight Processor Virtualization, such as:
https://istcolloq.gsfc.nasa.gov/Fall2013/presentations/cudmore.pdf

SDN

Benefits and Opportunities

SDN shall be considered as relatively mature concept, with the key enabling technology (Openflow) already supported by several commercial products.

Following the same classification as previously, we can identify the following key benefits for satcom network vendors/integrators.

- Interoperability and support by most SDN management platforms *(direct benefit)*
- Enable concurrent/customized versions of isolated control planes *(indirect benefits)*. This is mainly related to deployments where the access system serves multiple service providers and/or multiple customers for each service providers. Also, native support of (separated) control planes may imply transport solutions that avoid tunnelling approaches that create systematic overhead header, for each single frame sent over the air
- Ease integration of satcom networks in converged satcom/terrestrial architectures such as 5G, CDN, IMS, etc. *(indirect benefits)*
- Potential interest for hybrid/multi-access satcom systems for supporting the unification of their control plane with common devices *(indirect benefits)*
    - Heterogeneous user planes and routing/switching technologies can possibly be used for each access

Uncertainties, weaknesses, risks and threats

An important uncertainty is associated with the impact of SDN on the Network Management functions to be delivered to satcom providers. Any SDN-enabled network will have to make its NMS evolve (as also stated in ETSI NFV architecture). Given that new capabilities will be offered by the network the NMS complexity will grow as well. This includes network measurement monitoring, attribution of static configurations parameters, etc.

Last, depending on the depth of integration of the SDN routers devices in the system, we observe that different levels of costs and investments would be needed (according to the fact that new developments activities are needed or not), whether the SDN devices remain external or not to the “core” hub components (mux, mod/demod, RL controller etc.).

Conclusions

The techno-economical equation identifies direct benefits, associated with risks and large uncertainties, but also indirect benefits which are much harder to quantify but could be high. We do believe NFV and SDN could be, with an increased gain of momentum in the network terrestrial context, future “Must-have” requirements expressed by the service providers.

So far, NFV support for the products of satcom network vendors and integrators will provide them higher degree of competitiveness and key differentiators. We also believe NFV could find applicability at short/mid-term.

From another hand, SDN provides to our view mostly focused and indirect benefits that may not justify integration at the core hub components. However, integration with SDN devices at e.g. entry points of the satcom access network would certainly be relevant approaches, at least at a first stage.

A final remark is about the strong expectations from satcom network vendors/integrators in terms of requirements which need to be expressed by the satcom service providers themselves. Sharing the views between solution providers and network operators shall be essential to anticipate, as much as possible, the development and integration efforts for supporting these new technologies. 2020 appears a good target for this.

8.1.2.3. (Evolved) Satcom service providers

NFV

Benefits and opportunities

Most of the satcom service providers' benefits associated with NFV adoption can be derived from the benefits referred to. These are:

- Widening of the services portfolio *(direct benefit)*. Essentially, this implies moving to a novel service paradigm from plain connectivity to added-value virtualised in-network services.
- CAPEX and OPEX reduction *(direct benefit)*, as discussed in Chapter 7. The exact amount of reduction depends mainly on the number of components being virtualised (see roadmap below) as well as the actual cost of the virtualised components. It also depends on the ability to mix-and-match virtualised appliances from multiple vendors.
- On-demand resource allocation and elasticity for virtualised services *(direct benefit)*, which leads to better utilisation of hardware resources.
- Acceleration of technology evolution *(indirect benefit)* thanks to softwarisation. This means that the SP can always benefit from the latest technology and access frequent updates.

Uncertainties, weaknesses, risks and threats

The two most important factors, which introduce uncertainty and risks for the SP regarding NFV adoption for the time being are the lack of maturity of the NFV landscape (standards and products) and the resiliency/availability issues of virtualised network appliances.

Although cloud technologies have become quite mature during the last years and OpenStack is already an industry standard, NFV Management and Orchestration

(MANO) technologies are still at early stage. This means that fully automated NFV service deployment and chaining in a standardised and vendor-neutral manner, especially for customer VNFs (VNFaaS), should not be foreseen in the very short term. However, static virtualisation of certain core SP functions can already be applied i.e. replacement of hardware routers or firewalls with virtual ones.

The second issue, regarding resiliency and availability, is common with all softwarised services – in the sense that software appliances are generally less reliable than hardware ones. This issue is being considered and is gradually mitigated as cloud and NFV technologies evolve – but will probably never cease to exist. An efficient solution would be the application of redundancy policies, i.e. instantiation of failover VNFs, to which traffic can be rapidly redirected (e.g. using SDN) within milliseconds if the primary VNF fails. This can be achieved at a fraction of the cost needed e.g. to purchase and maintain hardware failover units. Generally, however, availability should be considered a critical issue mainly for core SP functions whose failure could affect potentially thousands of customers. It is less critical for customer functions, who –depending on their SLA- could probably tolerate some unavailability.

Migration roadmap

Although a clean-slate "hard" migration from a hardware-based to a fully virtualisation-capable satcom infrastructure is always an option for SPs, it would probably be safer to consider an evolutionary approach. Such an approach is visualised in Figure 275 below. The approach is to break down the evolution/migration into three phases (short/medium/long term) and to identify the components of the satcom network which could be virtualised in each phase.

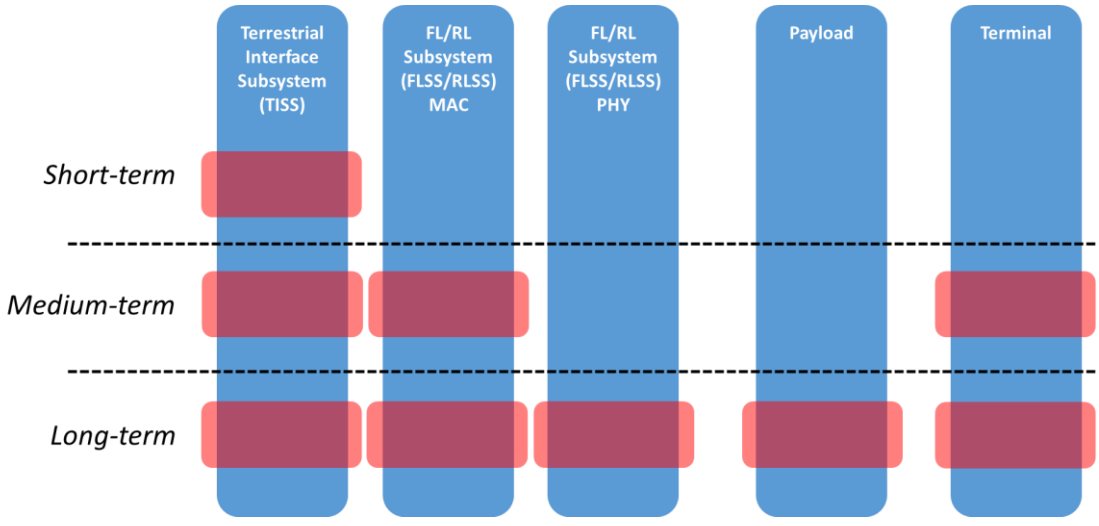

Figure 275. Proposed SP migration roadmap to NFV

In the short term, the foreseen upgrade would be to virtualize the functions at the interface with the terrestrial network (routers, firewalls, NATs, filters, accelerators

(only at hub side) as well as customer functions). This approach requires minimal intervention in the satellite hub.

At a later stage, virtualization could expand to the MAC-layer functions of the FLSS/RLSS e.g. scheduling, multiplexing, RL resource allocation, access control etc. These should be considered core SP functions but could also be offered in an SVNO customer. Generally, these functions are not directly under the control of SPs, but only controlled at high level (e.g. abstract QoS control and resource management).

Also, VNF-capable terminals could be considered for the medium-term.

For the long term scenario, we could envisage a fully NFV capable infrastructure, allowing virtualization also at the PHY front-end (CloudRAN scenario) and also maybe at the payload.

SDN

Benefits and opportunities

As with NFV, most of the satcom service providers' benefits associated with SDN adoption can be derived from the benefits referred to in Sec 0. These are:

- Unified, vendor-agnostic management of all network components with per-flow granularity *(direct benefit)*. This is possibly the most important advantage, also facilitating the use of components from multiple vendors in the same infrastructure as well as integration with terrestrial/5G.
- Better QoS and flow handling support with low response time, also an NFV enabler thanks to traffic steering *(direct benefit)*
- Better support for hybrid (satellite/terrestrial) delivery and also multi-access (handovers between LEO/MEO and GEO/HTS systems) *(direct benefit)*

Uncertainties, weaknesses, risks and threats

As with NFV, an important risk in SDN adoption lies in the constant evolution of the SDN technology per se and the technical immaturity of SDN controllers (which are well proven in lab environments but not yet widely used in production networks).

Another weakness lies at the current inability of SDN to handle radio resource management and PHY parameters, which are essential for satcom network management. For this purpose, an SDN controller is currently unable (even in principle) to perform end-to-end satcom management by controlling all the parameters of the satcom network chain.

Security issues associated with SDN are generally not very applicable to satcom, since SDN control is not assumed to be exposed to multiple stakeholders. However, we see a possible issue when it comes to the responsibilities of operating the SDN functions that ultimately control how data are basically handled, forwarded and switched *for multi-service provider deployments.* Indeed, an SDN-capable infrastructure cannot be operated simultaneously by the access network provider (to address its own needs) and by the Service Providers themselves, unless a strong cooperation is expected. This

level is certainly beyond what is currently done today with clear separation of domains – and responsibilities - bordered at interconnection/exchange point(s).

Migration roadmap

Similarly to the NFV case, we can propose an evolutionary roadmap towards SDN, as depicted in Figure 276.

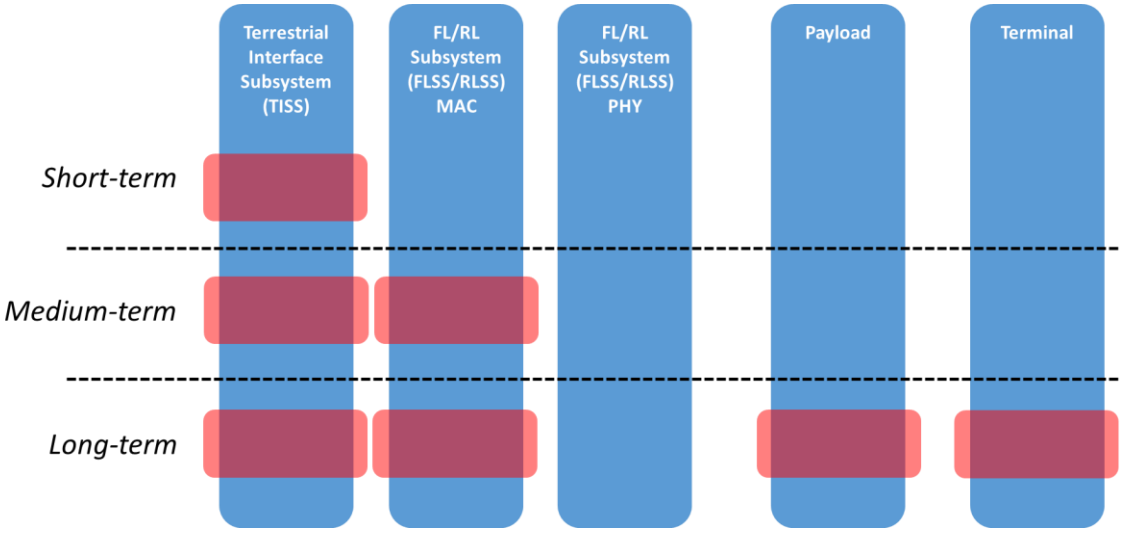

Figure 276. Proposed SP migration roadmap to SDN

In the short term, SDN capabilities could be inserted at the boundary with the terrestrial network e.g. replacing some TISS components with SDN-capable ones. This would allow some basic set of functionalities e.g. integration with terrestrial (up to a certain degree) and traffic steering for NFV services at the terrestrial interface.

Later on, SDN could be expanded to embrace FLSS/RLSS functionalities in order to better control resource allocation and QoS. However, this depends on whether future SDN protocol evolutions can handle RRM requirements.

In the long term, SDN could be supported at payload and also at the terminal. The latter case (terminal SDN) is foreseen for the long term not because of technology readiness purposes, but mostly because of the limited range of use cases which such a feature can serve.

SDN support in the radio front-ends is not foreseen, given the current management scope of SDN (which excludes PHY control).

Conclusions

Playing a central role in the value chain, satcom Service Providers are directly impacted by the adoption of software technologies.

Regarding NFV, it would be considered quite beneficial – and also safe from both a technical and business point of view – to adopt in the short term some virtualisation strategies, especially at the edges of the network (terrestrial interface and later also

at the terminals). Longer term evolutions should be carefully planned given also the evolution and adoption of the NFV technology in general.

Regarding SDN, introducing SDN capable devices at the TISS could be beneficial in the short term, enabling control by SDN controllers and basic integration with terrestrial. Further expansion of SDN support should bring significant additional benefits since it would allow end-to-end management, but it would essentially depend on the evolution of the SDN capabilities per se.

8.1.2.4. Customers

From the customers' point of view, the adoption of SDN and NFV would bring significant direct benefits, such as:

- Access to a **wide range of added-value services**, mostly thanks to NFV.
- Fine-grained **QoS assurance and control**.
- Ability to **pay-per-use** and employ elastic resource consumption (better OPEX control).
- Ability to **offload hardware functions to software appliances** (CAPEX reduction), which are also easier to maintain and upgrade.

No major risks and/or weaknesses are associated with the use of network softwarisation for most customers. Some concerns however could apply to specific customer groups:

- Customers with strict security requirements should ensure that VNFs off-loaded to the SP present adequate level of security and privacy with regard to traffic handling, compared with on-premises hardware appliances.
- Customers with requirements for high availability should ensure that software network services fulfil their requirements and/or (if needed) request virtual standby/failover units.

8.2. Recommendations and roadmap

8.2.1. What is/could be needed

8.2.1.1. SDN/NFV at the Ground segment

We currently see the investigated Cloud Networking technologies (mainly SDN & NFV) as enablers in the two domains of **network service deployments and orchestration**, and virtual **routing and forwarding**. Those two categories already span over a large number of functions to be supported in a satcom system.

SDN/NFV evolutions requirements

We firstly wonder whether any evolution requirements for SDN and NFV would be required or recommended to ease their integration in satcom or whether it could simply be neutral. Although a definitive answer is hard to provide, **we have no clear hints stating that SDN / NFV technologies should be changed to take into account the core satcom specificities**:

- SDN: for systems based on single-hop links (current GEO and MEO/LEO without ISLs), interconnection with external networks or devices is done with standard switching/routing appliances. Since with SDN a variety of rules and actions can be implemented from any combination of header fields (Ethernet addresses, VLANs, IP address, TCP ports, …) recognized in OpenFlow, we see no real interest for a evolution of the OpenFlow protocol.

  For systems involving on-board routing (LEO with OBP / ISL and GEO/MEO with OBP/router) there could be an interest to extend current OpenFlow capabilities based e.g. on specific fields used for the switching process (such as MPEG-2TS / ATM / ULE / GSE fields in DVB architectures) when no direct mapping to the previous headers can be made. An idea is that system would be that the Satcom provider could configure the on-board switch just as any other switch in the system, hiding (at some point) the specificity of the satellite connectivity within the network.

  Also, what can be observed is that the temporal granularity of this "reprogrammability" is inline is what is supported in current SDN, and based on "events" that have limited occurrence (e.g. few times a day, maximum). Those questions are important since GW sites and consequently the "resource hypervisor" are typically distant of hundreds of kilometers and traffic exchanges on such distances imply additional costs. In such deployments, the resource hypervisor would be co-located to the SDN controller, if it exists. We note that this is also a main trend in terrestrial networks where the MME (*Mobility Management Entity*) is responsible of resource allocation for several eNBs.

- NFV: current NFV solutions offer rich features as well (and more are planned for the future), and there does not seem to be a fundamental new requirement for NFV from a satcom point of view regarding its applicability to the ground segment.

However, although no mandatory changes are required in the SDN/NFV technology core in order to allow basic integration with satcom, there are some technical aspects which, if addressed, would greatly increase the added-value of softwarising the satellite infrastructure. These aspects are discussed below.

Technical issues to address for applicability

We identify two main categories of aspects (control and management plane) which system designers could tackle to maximise the benefits of virtualisation technologies for satcoms.

- Control Plane

As far as the control plane is concerned, we believe a possible next step could be the **virtualization of access resources (allocation issue),** similarly to the concepts introduced with the CloudRAN approach in cellular networks (see Chapter 2), but with some main differences, regarding notably the degree of centralization. Whether the benefits are not evident for the current generation of GEO systems we know today, we can infer two main use cases for a virtualized management of access resources:

- In multi-GW HTS systems: some resources may be shared among GW for variety of reasons: rain fading diversity issues with high-band system (Ka and above); load balancing between GW to overcome the HW, frequency, or processing limits at each site – when such constraints exist. In this case centralized processing or resource "hypervision" will be needed, borrowing exactly the same idea of orchestrating IT and network capacity resources in SDN.
- LEO/MEO systems when potentially the density of GW/access point is high.

To support this, a kind of extension of OpenFlow could be studied and proposed, extending the role of the SDN controller to a centralized resource management allocator, and the GWs the actual agents (clients) enforcing the radio resource allocation rules. Such extension could be seen as an additional layer to map generic resource descriptors (bandwidth, in kbps or in packet per second) into RF resource (e.g. MHz). Such units can be used when introducing traffic metering conditions and related actions, in order to configure a virtual network. The extension could take the form of an additional processing, both at SDN controller and vSwitch, able to transform generic requests/queries into satellite-aware request/queries, without changing the rest of the processing.

An ultimate goal of this new component could be – in many years - an applicability and a similarity of the whole forwarding configuration of the system **for any type of satcom topologies, including GEO/nGEO-based.** Although each topology (and actual systems) do have their specific requirements, a minimal level of abstraction could be defined with some level of reuse, thanks to a generic API implemented for example in a multi-purpose / multi-mission SW hypervisor. The controlled parameters could encompass information such as frequency plan, carriers, breakdown of frequency/time resources, polarizations, available power and/power constraints, uplink/downlink allocation, etc. This list shall be extensible as desired.

Finally, considering the CloudRAN perspective, one could notice that the proposed approach to be applied to satcom would be limited to the **centralization** of some resource allocation mechanisms, instead of actually preparing and sending the baseband signal remotely. Once again a main reason not to do this is related to the extra cost of sending much more data over longer distance that would involve much higher OPEX for the supporting network infrastructure the SP has to lease.

Finally, a side benefit would be to ease the integration with terrestrial systems such that operating tools and methods between equipments would be equivalent. Such

feature could be important driver for emergence of hybrid satellite/terrestrial systems, including 5G at short-term.

- Interaction with Management Plane and OSS/BSS solutions

There are also a couple of aspects associated with the management plane, and its interactions with OSS/BSS. These interactions take place over different types of interfaces, also identified in several system models such as e.g. the ETSI NFV model.

For the satellite community perspective, we observe that service configuration, deployments, activations are probably the cornerstones of programmability and flexibility capabilities of future satcom systems. Heavy operational procedures remain the norm in this area. Even a simple reconfiguration (e.g. SLA change, addition/change of routes or connectivity, addition/suppression of Virtual operator, etc.) often require at some time the involvement of several (human) operators at different levels of the overall operational & management chain. Automation of those processes would encompass not only generic IT OSS/BSS software and integrated solutions, but also the availability and the full usage of interfaces between the OSS/BSS and the satcom network management system.

In this respect, the satcom industry could have major interests on the new possibilities offered by “SDN/NFV-ready” OSS/BSS solutions.

8.2.1.2. Payload virtualisation

The concept of virtualizing payload OS needs to be developed further and investigated for long-term applicability. Also we note that other candidate technologies, such as Docker (and generally the virtual Container) approach, shall not be excluded because they achieve a satisfactory trade-off among resource isolation and HW resource saving. This kind of approach should probably gain interest if the need to isolate on-board service and applications –i.e. implementation of the NFV concept at the payload – was considered feasible and an important feature for next-gen systems.

8.2.2. How to do it

This section identifies several possible interactions with the software network and satcom community to be pursued in order to facilitate the foreseen technical achievements in a most efficient and effective manner. Such actions would involve:

- Interaction with relevant standardisation initiatives
- Interaction with collaborative open-source projects
- Participation in co-funded research projects
- Activities to engage and motivate stakeholders

8.2.2.1. Tracking of and contributions to standardizations bodies

ETSI NFV ISG

ETSI NFV ISG, described in Chapter 2 has lately proceeded to Phase 2, underway with over 30 new Work Items, also including normative specifications. ETSI NFV phase 2 extends the NFV charter, which is now mainly targeted at technology adoption and addressing areas such as testing/validation, performance/assurance, security, stability, interoperability, reliability, availability and maintainability. Collaboration with external bodies is also a key priority for NFV phase 2.

From a satcom perspective, the working groups IFA and EVE deserve special attention, as described below.

IFA (Interfaces and Architecture): The responsibilities of this working group include the delivery of information models and information flows to support interoperability at reference points and the refinement of the architecture and interfaces leading to the production of detailed specifications.

Tracking: IFA activities should be tracked so that satcom vendors seeking ETSI NFV compliance can adjust the management end data interfaces of their (physical or virtual) appliances in concordance with IFA specs.

Contributions: Satcom-specific contributions could be foreseen to the specific work items:

- IFA013: Os-Ma-Nfvo reference point - Interface and Information Model Specification (Interfacing with satellite OSS/BSS systems)
- IFA005: Or-Vi reference point - Interface and Information Model Specification (Communication with satcom-specific virtualised infrastructure managers)

EVE (Evolution and Ecosystem): According to the terms of reference, this working group is responsible to develop feasibility studies and requirements in relation to a) new NFV use cases and associated technical features, b) new technologies for NFV and c) relationship of NFV with other technologies. It should also maintain an overall view of NFV-related work performed elsewhere (e.g. SDOs, industry groups, open source communities) and develop gap analysis on industry standards in areas relevant to NFV.

Tracking: EVE activities should be monitored in order to observe the correlation of NFV with other network technologies (such as SDN). This would also affect the SDN/NFV migration roadmaps laid out in 8.1.2.3. ., in case either technology is dependent on the other.

Contribution: Satcom-specific contribution could be foreseen to the specific work items:

- EVE006: Report on NFV Industry Roadmap (Feedback from satcom vendors and SPs regarding their plans)
- EVE003: Report on NFVI Node Physical Architecture Guidelines for Multi-Vendor Environment (Specialisation for satcom NFV-enabled equipment (payload/terminal) and also satellite-terrestrial multi-domain cases)

Open Networking Foundation (ONF)

Open Networking Foundation (ONF) [ONF] is a user-driven organization dedicated to the promotion and adoption of Software-Defined Networking (SDN) through open standards development. ONF is the SDO maintaining the OpenFlow specification.

ONF Technical Communities continue to analyse SDN requirements, evolve the OpenFlow spec to address the needs of commercial deployments, and research new standards to expand SDN benefits.

Tracking: OpenFlow is currently the dominant SDN protocol. Following the ONF advances and, more specifically, the evolution of the OpenFlow spec, is vital for any interested stakeholder – including the satcom community – so that they can keep up with latest OF capabilities and develop OF-compatible equipment.

Contributions: Among the ONF Technical Communities, the following two can be identified as most relevant for contributions:

- The *Open Datapath* project is maintaining and evolving the OpenFlow protocol and associated datapath modelling technologies. Candidate extensions to OF to serve satcom-specific requirements could be submitted there – even though they would be more relevant in the frame of Open Transport project (see below).
- The *Open Transport project* addresses SDN and OpenFlow Standard-based control capabilities for transport technologies of different types, including optical and wireless transport. Up to now, the project was focusing on optical networks, yet it is expanding to the wireless domain. Satellite-driven use cases could be contributed, as well as candidate extensions of OF to suit satellite transport (such as the OF-S extensions for optical transport).

ITU IMT-2020

In early 2012, ITU-R embarked on a programme to develop "IMT for 2020 and beyond", setting the stage for 5G research activities that are emerging around the world. Through the leading role of Working Party 5D [IMT2020], ITU's Radiocommunication Sector (ITU-R) has finalized its view of a timeline towards IMT-2020.

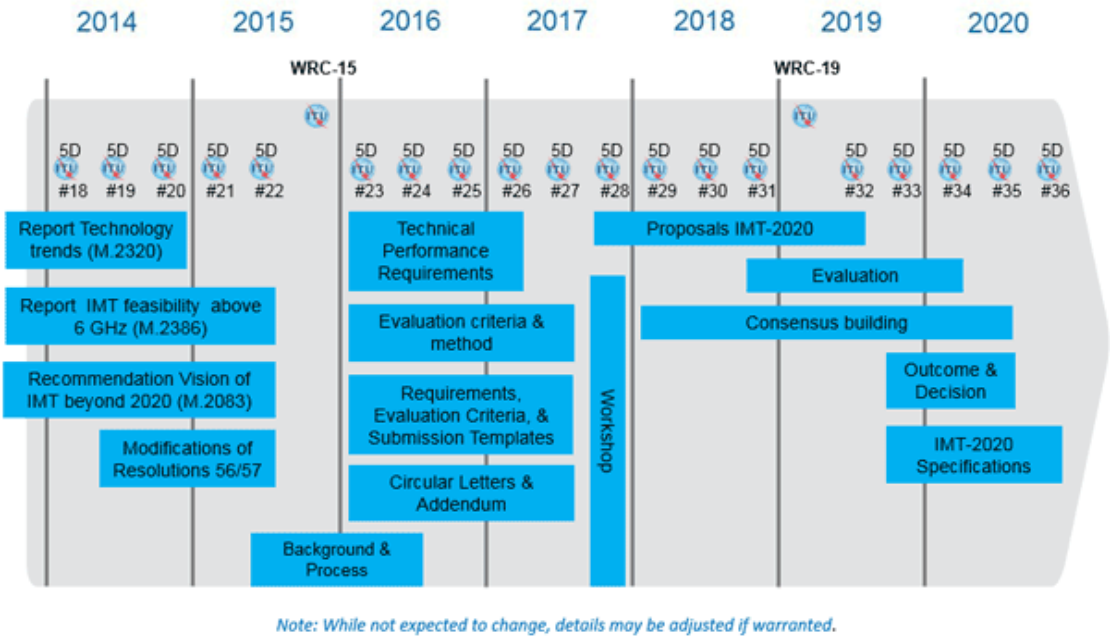

Figure 277. Detailed timeline & process for IMT-2020 in ITU-R [IMT2020]

In this framework, ITU-R has very recently (September 2015) released Rec. M2083-0 [ITUM2083]which presents the Framework and overall objectives of the future development of IMT for 2020 and beyond. This Recommendation, apart from addressing radio issues, which lie within the main scope of ITU-R, explicitly states that software network technologies are an essential component of the 5G vision:

"Future IMT will require more flexible network nodes which are configurable based on the Software-Defined Networking (SDN) architecture and network function virtualization (NFV) for optimal processing the node functions and improving the operational efficiency of network. Featuring centralized and collaborative system operation, the cloud RAN (C-RAN) encompasses the baseband and higher layer processing resources to form a pool so that these resources can be managed and allocated dynamically on demand, while the radio units and antenna are deployed in a distributed manner." (Sec. 2.3.2)

Tracking: IMT-2020 activities, still at a very early stage, need to be tracked in order to follow the ITU-R vision for 5G, especially given that ITU-R is expected to pay specific attention to the interplay of network softwarisation with radio technologies and spectrum issues, which are of particular interest for satcom stakeholders. This interplay is not expected to be well covered by other SDOs (ETSI NFV, TMF, IETF).

Contributions: Rec. ITUM2083 also states that *"interworking will be necessary among various access technologies, which might include a combination of different fixed, terrestrial and satellite networks. Each component should fulfill its own role, but also should be integrated or interoperable with other components to provide ubiquitous seamless coverage."* (Sec. 6.1.2) In this context, the contribution of the satcom community to the ITU activities is essential to promote the requirements of the satcom community within the 5G vision. Although specific technical contributions cannot be foreseen at this early stage, it would be probably relevant within ITU-R to address the relationship between network softwarisation and satcom radio aspects.

8.2.2.2. Collaborative open-Source projects

OpenStack

OpenStack has become the de facto standard in open-source cloud management. Behind this project is the OpenStack Foundation, counting more than 28.000 individual members from 140 countries around the world. More details on OpenStack can be found in Chapter 2. OpenStack was used as the cloud controller platform during CloudSat experimentation.

OpenStack has already become quite mature during the last years. The current release is named Liberty, while the next one (Mitaka) is expected to become available April 2016.

Tracking: The role of OpenStack in providing virtualised network services is indispensable, also being the dominant open platform for cloud management and also a key component of OPNFV platform. Of essential interest to satcom are its developments in the compute (nova), networking (neutron) and telemetry (ceilometer) components.

Contributions: Nova extensions for managing lightweight compute nodes with restricted resources could be relevant to satcom in order to enable onboard or terminal-side VNFs; also adaptations to Neutron to match the specific requirements of satellite networks, as well as Ceilometer enhancements in order to integrate metrics for satcom component monitoring.

OpenDaylight

OpenDaylight is becoming the dominant SDN controller platform, with a wide and ever-increasing set of features, developed as a Linux Collaborative Project. More details on ODL can be found in Chapter 2. ODL was used as the network controller platform during CloudSat experimentation.

OpenDaylight is currently backed by a wide development community, in which engineers from most networking vendors actively participate. The Lithium release has recently been available, featuring critical enhancements for NFV enablement, such as Service Function Chaining (SFC). The next release (Beryllium) is planned for February 2016.

Tracking: OpenDaylight could be the likely candidate for satcom SDN management also. It is also an essential component of the OPNFV platform. Therefore, its overall progress (in terms of protocol support, plug-ins and built-in network services) needs to be closely monitored.

Contributions: ODL support for network devices is based on management plug-ins which interface with its southbound Service Abstraction Layer and targets at both SDN and also non-SDN network elements. A very promising contribution would be the development of specific plugins so that satcom-specific elements (e.g. hub components) can be managed via OpenDaylight.

OPNFV

OPNFV is the Linux Foundation Collaborative Project which aims at providing an open-source management and execution platform for NFV services. More details can be found in Chapter 2.

To date, OPNFV has already delivered its first release (Arno) in January 2015, focusing on the VIM and Infrastructure layers. This first release is far from mature, yet it is quickly evolving. The second release (Brahmaputra) is planned for February 2016.

Tracking: Tracking of OPNFV activities is essential for anyone interested in NFV evolution, since it is most likely that OPNFV will constitute the prevailing open-source NFV platform in short/mid- term.

Contributions: Satcom-related contributions to OPNFV would very likely be related to virtualization enablers on resource-restricted hardware. The current OPNFV releases require a considerable volume of computing and storage resources in order to be deployed.

8.2.2.3. Activities in Research Programs

H2020 and 5G – Relevance of Phase 1 projects

Horizon 2020 (H2020) is the R&D programme of the European Commission for the period 2014 – 2020. With a total budget of over €80 billion, it aims to be the flagship funding instrument for supporting innovation in the EU. H2020 activities are categorised into three main funding pillars: Excellent Science (which targets at supporting basic and long-term research), Industrial Leadership (which aims at reinforcing the European industries and SMEs) and Societal Challenges (which focuses on inter-disciplinary research to meet the needs of the society).

Within the Industrial Leadership pillar and in the frame of the ICT (Information and Communication Technologies) theme, the EC has launched a specific activity which targets at the research and development of a 5G communication system within the H2020 programme. This activity is implemented under a Public – Private Partnership (PPP), still using competitive calls. The 5G PPP Phase 1 call was closed November 2014 and the first wave of approved projects are already running, targeting at various technological domains which will form the basis of the future 5G system. Table 19 lists the 5G PPP Phase 1 projects in the software networks area and identifies the most important outcomes which would be relevant to satcom. Most of them started July 2015 and will last for either 2.5 or 3 years.

The 5G PPP Phase 2 call has already been issued as a part of the ICT 2016-2017 Workprogramme [ICTWP] –currently in last draft version- and will close November 2016.

Table 19. Relevant H2020 5G PPP Phase 1 projects on software networks

Project	Main project scope	Outcome/Knowledge relevant for satcom
5G NORMA (5G NOvel Radio Multiservice adaptive network Architecture)	5G NORMA follows the concept of adaptive (de)composition and allocation of mobile network functions, which flexibly decomposes the mobile network functions and places the resulting functions in the most appropriate location.	Strategies and technologies for placement of mobile network functions (e.g. EPC) within the satcom infrastructure
5GEx (5G Exchange)	5G Ex develops an open platform enabling cross-domain orchestration of services over these multiple domains, with a set of open source software tools and extensions.	Tools for satellite/ terrestrial integration and unified management
SESAME (Small cEllS coordinAtion for Multi-tenancy and Edge services)	SESAME introduces the Cloud-Enabled Small Cell (CESC) concept, a new multi-operator enabled Small Cell that integrates a virtualised execution platform ("LightDC") for deploying VNFs.	Techniques for edge NFV and generally NFV on devices with constrained resources (terminals, satellite payload)
SONATA (Service Programing and Orchestration for Virtualized Software Networks)	SONATA provides a novel SDK for NFV service development as well as a NFV Service Platform to manage service execution.	Orchestration platforms for NFV service management
SUPERFLUIDITY (A super-fluid, cloud-native, converged edge system)	SUPERFLUIDITY decomposes network services into reusable primitives to be deployed in a native, converged cloud-based architecture.	Strategies and technologies for placement of mobile network functions (e.g. EPC) within the satcom infrastructure

In addition to the above, the VITAL project (VIrtualized hybrid satellite-TerrestriAl systems for resilient and fLexible future networks) [VITAL], funded under the H2020 ICT theme but not specifically belonging to the 5G projects cluster, is investigating SDN/NFV integration into satcom towards satellite/terrestrial interconnection and federated management. The outcomes of the VITAL project, in conjunction with CloudSat findings, could constitute an excellent foundation for advancing a step further and experimenting with the actual implementation of software technologies in an operational satellite network.

Last but not least, the outcomes of the last relevant FP7 projects on network softwarisation which were identified in Chapter 2 (4WARD, SAIL, ALICANTE, MCN, T-NOVA, XIFI, FIWARE, ALIEN, OFELIA, iJOIN, CROWD) should also be taken into account in a future "software satcom" research activity.

Opportunities within the 5G initiative

In general, it is true that 5G is a very wide term whose scope is not strictly defined. 5G technologies are expected to go beyond air interface and MAC specifications and also embrace unified management and control of heterogeneous network infrastructures. In other words, 5G is an umbrella term embracing most future communication technologies, rather than specific to a single domain (e.g. cellular/mobile communications).

In this converged landscape, it is recognised that satellite communications have a clear role to play, enabling global and truly ubiquitous coverage and supporting use cases where terrestrial networks are proved inadequate. Satcom, as infrastructure, is considered an essential building block of the heterogeneous 5G infrastructure layer.

At the same time, network softwarisation (including SDN/NFV) has been identified as a key enabling concept towards the 5G vision. Apart from the other advantages it brings (service agility, rapid reconfiguration, novel added-value services etc.), network softwarisation is seen as the major "gluing factor" towards the unification of heterogeneous infrastructures.

Concluding, it can be said that the EC 5G initiative, implemented within the H2020 programme, presents an excellent opportunity for attracting funding for further research on the interplay between satcom and software networks.

ESA ARTES activities

The ESA Advanced Research in Telecommunications Systems (ARTES) programme supports R&D activities in the field of satcom.

Different opportunities could be identified in the various elements of the ARTES programme:

Future Preparations (previously **ARTES 1)** is dedicated to strategic analysis, market analysis, technology and system feasibility studies. It is the programme where CloudSat belongs. Further activities could be additional studies focusing on e.g.:

- the feasibility of applying SDN/NFV technologies to specific satcom components (e.g. satellite payload or various processing stages within the satellite hub)
- detailed specification of integrated satcom software networks (exact architectures, interfaces, protocols etc.), probably after a few years when the terrestrial SDN and especially NFV landscape will be clearer
- detailed techno-economic studies via thorough market assessment, probably after a few years where the potential of the SDN/NFV market can be more accurately predicted
- exploitation of terrestrial 5G R&D achievements for satcom

- business and technical investigation of one or more specific ETSI NFV use case(s) (such as vCDN)

Competitiveness & Growth (previously **ARTES 3-4)** is dedicated to the development, qualification and demonstration of products. This element could support projects pursuing outcomes feasible in the short term such as e.g.:

- SDN- and NFV-capable hub elements
- Virtualisation-capable terminals and/or Virtual appliances (virtual Home Gateways or any Virtual appliances operated from customer remote location).

Advanced Technology (previously **ARTES 5.1/5.2)** is dedicated to long term technological development of the satcom industry. Possible projects could involve e.g.:

- SDN for inter-satellite links for LEO/MEO constellations
- SDN extensions for radio resource management
- Software-defined satellite payloads (including their qualification)
- Business and technical investigation of one or more specific ETSI NFV use case(s) (such as vCDN)

8.2.2.4. Involvement of stakeholders

As for any other technology, it is essential that partnerships, feedback, and exchanges between the different satcom stakeholders can take place, not only to precisely capture the community requirements, but also to plan the most important activities, and identify the related organization to support them (in terms of budget, programs, planning, etc.). In particular **feedback from customers and satcom Service Providers** are of utter importance for them, and because they are the **primary users** of the technologies that will operate the systems day-to-day.

Various examples of activities could support the possible involvement

- Follow-up of virtualization technologies and developments by the different parties. In particular, the emergence and/or growing success of solutions shall be tracked.
- Organisation of workshops focused on Virtualisation technologies.
- Involvement into non-satellite events and exchanges with terrestrial community and operators.
- Prototyping and testing of different kinds of SDN / NFV-enabled solutions and platforms.

9. CONCLUSIONS

Cloud and Virtualization Networking technologies will have important impacts on future satcom systems. High cost, low resource availability, and conservative architectures that predominate today in the satellite landscape, certainly constitute major obstacles to cross for this family of technologies. On the other hand, many applications could be targeted, and the interests and requirements on cloud and virtualisation networking expressed from all stakeholders do justify additional works in the spatial area. Tangible results such as prototyping/pre-development and/or proof of concepts shall be considered as the main next step to achieve.

SDN and NFV, the two main concepts investigated in this work, have different kinds of implication for satcoms. As shown via the results of this study, SDN is mainly intended to be implemented at the border of the satcom telecom system, possibly without any impact for the development of its core service in mid-term application, still needing to be integrated with the satcom NMS and OSS/BSS. NFV could have shorter-term applications, related to the operations and management of specific features, wherever they are implemented. For long-term, SDN could also be supported more in depth in satcom. With the advent of projects aiming at developing low-cost LEO constellation composed of many small satellites, the opportunities to develop and operate on-board SDN-compatible routers could become reality.

10. REFERENCES

[4WARD] FP7 4WARD project, http://www.4ward-project.eu

[5GEx] Project 5GEx (5G Exchange), http://www.5gex.eu/

[5GNORMA] Project 5G NORMA (5G NOvel Radio Multiservice adaptive network Architecture), https://5gnorma.5g-ppp.eu/

[802.1ad] 802.1ad-2005 – "IEEE Standard for Local and Metropolitan Area Networks---Virtual Bridged Local Area Networks---Amendment 4: Provider Bridges", May 2006

[802.1ah] 802.1ah-2008 – "IEEE Standard for Local and metropolitan area networks -- Virtual Bridged Local Area Networks Amendment 7: Provider Backbone Bridges", June 2008

[Abarca13] E. Abarca, J. Grassler, G. Schaffrath, S. Schmid, "A Federated CloudNet Architecture: The PIP and the VNP Role", arXiv:1303.6753v1 [cs.NI], March 2013

[ABNO] D. King, A. Farrel, "A PCE-based Architecture for Application-based Network Operations", IETF draft draft-farrkingel-pce-abno-architecture-13 (work in progress)

[ACG] ACG Research study (2013) "NFV Promises Cost Savings of Nearly 70%"

[ADVA10] ADVA Optical Networking white paper, M. Ritter, "Virtualized Optical Networks for Sustainable Cloud Services," Jan. 2010.

[ALC] Alcatel-Lucent's white paper (2014) "Business Case for Moving DNS to the Cloud"

[ALICANTE] FP7 ALICANTE project, http://www.ict-alicante.eu

[ALIEN] FP7 ALIEN project, http://www.fp7-alien.eu/

[AMS] Analysys Mason (2014), http://www.analysysmason.com/Research/Content/Reports/SCN-sizing-forecast-Jun2014-RMA16/

[AQOS14] UC Interoperability Forum, Automating QoS: UC SDN Use Case, v.1.2, February 27, 2014, http://ucif.org/Portals/0/documents/2014_02_27_Use_Case.pdf

[Baldine12] Ilia Baldine, Yufeng Xin, Anirban Mandal, Paul Ruth, Chris Heerman, and Jeff Chase, "ExoGENI: A Multi-Domain Infrastructure-as-a-Service Testbed", in Proc. TridentCom 2012

[Bao14] J. Bao, et al., "OpenSAN: A Software-defined Satellite Network Architecture", Proc. SIGCOMM '14, August 17-22, Chicago, Illinois, USA, pp. 347-348

[Bari13] Md. Faizul Bari, Raouf Boutaba, Rafael Esteves, Lisandro Zambenedetti Granville, Maxim Podlesny, Md Golam Rabbani, Qi Zhang, and

Mohamed Faten Zhani, "Data Center Network Virtualization: A Survey", IEEE Communications Surveys & Tutorials, Vol.15, No. 2, Second Quarter, 2013.

[BATS] Broadband Access Terrestrial Satellite (BATS) project webpage, http://www.batsproject.eu/

[Beacon] What is Beacon? https://openflow.stanford.edu/display/Beacon/Home

[Bernardos14] C.J. Bernardos, A. de la Oliva, P. Serrano, A. Banchs, L.M. Contreras, H. Jin, J.C. Zúñiga, "An Architecture for Software Defined Wireless Networking", IEEE Wireless Communication Magazine, Vol. 21, No. 3, pp. 52-61, 2014.

[Bertaux15] L. Bertaux, et al., "Software Defined Networking and Virtualization for Broadband Satellite Networks", to appear in IEEE Communications Magazine, Special Section on Satellite Communications and Networking: Emerging Techniques and New Applications, 2015

[Bitar13] N. Bitar, S. Gringeri, T.J. Xia, "Technologies and Protocols for Data Center and Cloud Networking", IEEE Communications Magazine, September 2013

[BMC] Osterwalder, Pigneur (2010), The Business Model Canvas

[BOO] Booz&Co, "Why satellites matter", http://www.esoa.net/news-info-30.htm

[BOS] Blue ocean strategy and value innovation, http://www.harbott.com/2011/05/17/blue-ocean-strategy-and-value-innovation/

[Broad14] Broadcom Announces World's First Single-chip Hybrid Direct Broadcast Satellite Terrestrial and IP Devices for Set-top Boxes with Integrated HEVC and MoCA 2.0, Press release, http://www.broadcom.com/press/release.php?id=s869811

[Carapinha09] J. Carapinha, J. Jimenez, "Network Virtualization – a View from the Bottom", in Proc. ACM SIGCOMM VISA'2009 Workshop, Barcelona, 17 August 2009

[CDI] Mindlin D. (2013) Present Values, Investment Returns and Discount Rates, http://www.cdiadvisors.com/papers/ CDIDiscountRate.pdf

[Chang12] Y.-J. Chang et al., "Scalable and Elastic Telecommuncation Services in the Cloud", Bell Labs Technical Journal, vol. 17, no. 2, pp. 81-96, 2012.

[Cisco09] Cisco VN-Link: Virtualization-Aware Networking, 2009.

[CloudNFV] CloudNFV, http://www.cloudnfv.com.

[Cloudstack] The Apache Software Foundation. Apache CloudStack - Open Source Cloud Computing, http://cloudstack.apache.org/.

[CONTENT] Anna Tzanakaki (ed.) et al, "D2.3 Overall System Architecture Definition and Specifications", FP7 CONTENT Project, November 2013

[Contreras12] L. M. Contreras, V. López, O. González De Dios, A. Tovar, F. Muñoz, A. Azañón, J. P. Fernandez-Palacios, J. Folgueira, “Toward cloud-ready transport networks”, IEEE Communications Magazine, Vol. 50, No. 9, pp. 48-55, 2012

[CRAN11] China Mobile Research Institute, “C-RAN: The Road Towards Green RAN”, White Paper, v.2.5, October 2011

[CROWD] FP7 CROWD Project, http://www.ict-crowd.eu/

[CRU] Carnevale Chuck (2013), What Is The Correct Discount Rate To Use?, http://www.fastgraphs.com/research_articles/2013-09-26-chuck-carnevale-what-is-the-correct-discount-part2b

[Davie12] B. Davie and J. Gross, "A Stateless Transport Tunnelling Protocol for Network Virtualization", 2012.

[Docker] Docker - Build, Ship and Run any app, anywhere, https://www.docker.com/

[DPDK] Data Plane Development Kit, http://dpdk.org/

[DSI15] Bitcoin Pioneer Inks Contract for Satellite Constellation, Deep Space Industries, http://deepspaceindustries.com/bitcoin-pioneer-inks-contract-for-satellite-constellation/

[DVBRCS2] “Digital Video Broadcasting (DVB); Second Generation DVB Interactive Satellite System (DVB-RCS2); Part 1: Overview and System Level specification”, ETSI TS 101 545-1, V1.1.1, May 2012

[DVBRCS2] Digital Video Broadcasting (DVB); Second Generation DVB Interactive Satellite System (DVB-RCS2); Part 1: Overview and System Level specification

[DynQoS] C. Baudoin, M. Gineste, E. Chaput, P. Gelard, J. Bernard, “Dynamic satellite system QoS architecture integrated with IP Multimedia Subsystem core network”, in International Journal of Satellite Communications and Networking (IJSCN), July 2014

[ETSIMAN] ETSI GS NFV-MAN 001, “Network Functions Virtualisation (NFV); Management and Orchestration”, v.1.1.1, December 2014

[ETSIMEC] ETSI Mobile Edge Computing White Paper. http://portal.etsi.org/Portals/0/TBpages/MEC/Docs/Mobile-edge_Computing_-_Introductory_Technical_White_Paper_V1%2018-09-14.pdf

[ETSINFV] ETSI GS NFV 002 “Network Functions Virtualisation (NFV); Architectural Framework”, v1.1.1, 2013

[ETSIREQ] ETSI GS NFV 004 “Network Functions Virtualisation (NFV); Virtualisation Requirements”, v1.1.1 , 2013

[Eucalyptus] Eucalyptus Systems, Inc. Eucalyptus. http://www.eucalyptus.com.

[ExoGENI] ExoGENI project, http://www.exogeni.net/

[Ferguson98] P. Ferguson, G. Huston, “What is a VPN?”, Tech. Rep., Cisco Systems (1998)

[FIWARE] FP7 FI-WARE project, http://www.fi-ware.org/

[Floodlight] Project Floodlight - Open Source Software for Building Software-Defined Networks Project Floodlight, www.projectfloodlight.org/floodlight

[Flowvisor] ON.LAB. Flowvisor. http://onlab.us/flowvisor.html

[FMG] Prasanna Chandra (2011), Financial Management, McGrawHill

[FORCES] IETF Forwarding and Control Element Separation Working Group, https://datatracker.ietf.org/wg/forces/documents/

[Gall13] A. Gall, “Leveraging SDN to provide Layer-2 VPNs as a Service on a MPLS-free core”, Proc. TERENA 2014 Networking Conference, Dublin, Ireland, 19-22 May 2014

[Garimella07] P Garimella, Y.W.E Sung, N. Zhang and S. Rao, "Characterizing VLAN usage in an operational network", Proc. 2007 SIGCOMM workshop on Internet network management, pp. 305-306.

[GENI] GENI project, https://www.geni.net/

[Greenberg08] A. Greenberg, J. Hamilton, D. Maltz, and P. Patel, "The cost of a cloud: research problems in data center networks", ACM SIGCOMM Computer Communication Review, vol. 39, pp. 68-73, 2008.

[Greenberg11] A. Greenberg, J. Hamilton, N. Jain, S. Kandula, C. Kim, P. Lahiri, D. Maltz, P. Patel, and S. Sengupta, "VL2: a scalable and flexible data center network", Commun. ACM, vol. 54, pp. 95-104, 2011.

[Horizon] Openstack Horizon Dashboard, https://wiki.openstack.org/wiki/Horizon

[HWP] Hewlett-Packard white paper (2014) “The Reality of Cost Reduction”

[ICTWP] Draft Horizon 2020 Work Programme 2016 – 2017 in the area of Information and Communication Technologies, https://ec.europa.eu/programmes/horizon2020/sites/horizon2020/files/05i.%20LEIT-ICT_2016-2017_pre-publication.pdf

[IEEE] IEEE. IEEE Guide for Developing System Requirements Specifications. s.l.: ETSI, 1998. IEEE Std 1233

[IEEE802] IEEE Standard for Local and Metropolitan Area Networks – Virtual Bridged Local Area Networks, IEEE Std 802.1Q-2005 (May 2006)

[IEEE802bg] IEEE standard 802.1Qbg - Edge Virtual Bridging, 2011.

[iJoin] FP7 iJOIN project, http://www.ict-ijoin.eu/

[IMT2020] ITU towards “IMT for 2020 and beyond”, http://www.itu.int/en/ITU-R/study-groups/rsg5/rwp5d/imt-2020/Pages/default.aspx

[INTEL] White Paper, Realising the Benefits of Network Functions Virtualisation in Telecoms Networks, February 2014.

[IrNEXT] Iridium NEXT: The Bold Future of Satellite Communications, https://www.iridium.com/About/IridiumNEXT.aspx

[IRR] Infonetics Research Report (2014), http://www.infonetics.com/pr/2014/Carrier-SDN-NFV-Market-Highlights.asp

[iSATVNO] iSAT Virtual Network Operator , http://www.isatafrica.com/services/carrier-class-services/satellite-virtual-network

[ITUM2083] IMT Vision – Framework and overall objectives of the future development of IMT for 2020 and beyond, ITU-R Rec. M.2083-0, September 2015

[ITUY3011] ITU-T Recommendation Y.3011, "Framework of network virtualization for future networks", January 2012.

[Jain13] S. Jain, et al., "B4: Experience with a Globally-Deployed Software-Defined WAN", in Proc. ACM SIGCOMM '13, August 12-16, 2013

[Jinno09] M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, and S. Matsuoka, "Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies", IEEE Commun. Mag., vol. 47, no. 11, Nov. 2009.

[Kim11] C. Kim, M. Caesar and J. Rexford, "SEATTLE: A Scalable Ethernet Architecture for Large Enterprises", ACM Trans. Comput. Syst., vol. 29, 2011.

[Kreeger12] L. Kreeger, D. Dutt, T. Narten, D. Black, and M. Sridharan, "Network Virtualization Overlay Control Protocol Requirements", 2012.

[Krishnan11] B. Krishnan, H. Amur, A. Gavrilovska, K. Schwan, "VM power metering: Feasibility and challenges", Performance Evaluation Review, vol 38, No. 3, pp. 56-60, 2011.

[KVM] Kernel-Based Virtual Machine, http://www.linux-kvm.org/page/Main_Page

[Laserre12] M. Lasserre, F. Balus, T. Morin, N. Bitar, Y. Rekhter, and Y. Ikejiri, "Framework for DC Network Virtualization", 2012.

[LMG] LEO, MEO and GEO Satellite Systems: A Comparison, http://www.durofy.com/leo-meo-geo-satellite-systems/

[LPMS] Mallory Caroline, Pak Hei Wu (2012), Lean Production Management System, Shanghai University

[Maestro] Maestro Platform -A scalable control platform written in Java which supports, http://code.google.com/p/maestro-platform

[Maha12] M. Mahalingam, D. Dutt, K. Duda, P. Agarwal, L. Kreeger, T. Sridharand, M. Bursell, and C. Wright, "VXLAN: A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", Internet Draft, https://tools.ietf.org/html/draft-mahalingam-dutt-dcops-vxlan-00

[MARK] Kotle Philip(2013), Principles of Marketing, 6th Edition, Pearson

[MCA] MCAfeeWorks, http://www.mcafeeworks.com

[McKeown08] N. McKeown, et al, “Openflow: enabling innovation in campus networks” ACM SIGCOMM Computer Communication Review, vol. 38 (2), April 2008, pp. 69-74

[MCN] FP7 Mobile Cloud Networking project, http://www.mobile-cloud-networking.eu/

[Mudigonda11] J. Mudigonda, P. Yalagandula, J. Mogul, B. Stiekes, and Y. Pouffary, "NetLord: a scalable multitenant network architecture for virtualized", ACM SIGCOMM, p. 62-73, 2011.

[Mysore09] R.N. Mysore, A. Pamboris, F. Nathan, N. Huang, P. Miri, S. Radhakrishnan, V. Subramanya, and A. Vahdat, "PortLand: a scalable fault-tolerant layer 2 data center network fabric", ACM SIGCOMM, pp. 39-50, 2009.

[Nag10] A. Nag, M. Tornatore, and B. Mukherjee, “Optical Network Design with Mixed Line Rates and Multiple Modulation Formats,” IEEE/OSA J. Lightw. Technol., vol. 28, no. 4, pp. 466-475, Feb. 2010.

[Naous08] J. Naous, D. Erikson. G. Covington, G. Appenzeller, N. McKeown, “Implementing an OpenFlow Switch on the NetFPGA platform”, in Proc. ANCS ’08, November 6-7, 2008, San Jose, CA, USA

[Narten12] T. Narten, M. Sridharan, D. Dutt, D. Black, and L. Kreeger, "Problem Statement: Overlays for Network Virtualization”, RFC 7364, October 2014

[NETCONF] R. Enns (ed.) et al, Network Configuration Protocol (NETCONF), RFC 6241, https://tools.ietf.org/html/rfc6241

[Neutron] OpenStack Neutron, http://wiki.openstack.org/wiki/Neutron

[NFV1] ETSI GS NFV. Network Function Virtualisation; Architectural Use Cases, V1.1.1.

[NFV2] ETSI GS NFV. Network Functions Virtualisation (NFV);Virtualisation Requirements, V1.1.1.

[NFV3] ETSI GS NFV. Virtual Network Function Architectural Framework, V1.1.1.

[NFV4] ETSI GS NFV. Network Functions Virtualisation (NFV);Terminology for Main Concepts in NFV, V1.1.1.

[NFV5] ETSI GS NFV. Network Functions Virtualisation (NFV);Proof of Concepts; Framework, V1.1.1. http: //www.etsi.org/deliver/etsi_gs/NFV-PER/001_099/002/01.01.01_ 60/gs_NFV-PER002v010101p.pdf

[NFVArch] ETSI GS NFV.?Virtual Network Function Architectural Framework, V1.1.1.

[NFVMAN] ETSI GS NFV-MAN 001, “Network Function Virtualisation (NFV) Management and Orchestration”, v.0.6.1 (July 2014)

[NFVMAN] ETSI GS NFV-MAN 001, “Network Function Virtualisation (NFV) Management and Orchestration”, v.0.6.1 (July 2014)

[NFVPerform] ETSI GS NFV; NFV Performance & Portability Best Practices. http://www.etsi.org/deliver/etsi_gs/NFV-PER/001_099/001/01.01.01_60/gs_NFV-PER001v010101p.pdf

[NFVUC] ETSI GS NFV. Network Function Virtualisation; Architectural Use Cases, V1.1.1.

[NIST] P. Mell and T. Grance. “The NIST denition of cloud computing” Technical report, National Institute of Standards and Technology, 2011

[NOX] NOX Openflow Controller, http://www.noxrepo.org

[NPD] NPD DisplaySearch Quarterly Smart TV Shipment and Forecast Report, 2014

[NSR] NSR's report Wireless Backhaul via Satellite, 9th Edition

[O3B] O3b Networks, http://www.o3bnetworks.com/

[ODL] OpenDaylight network controller, http://www.opendaylight.org

[OF14] Open Networking Foundation, OpenFlow Switch Specification, v.1.4, October 14, 2013, ONF TS-012

[OFELIA] FP7 OFELIA project, OpenFlow in Europe: Linking Infrastructure and Applications, http://www.fp7-ofelia.eu/

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

[ONF-W&M] ONF Solution Brief, OpenFlow™-Enabled Mobile and Wireless Networks, September 30, 2013

[OpenDOVE] OpenDOVE - Distributed Overlay Virtual Network. https://wiki.opendaylight.org/view/Project_Proposals:Open_DOVE .

[OpenFlow] Open Networking Foundation, https://www.opennetworking.org/sdn-resources/openflow

[OpenMANO] Telefonica OpenMANO, https://github.com/nfvlabs/openmano

[OpenNaaS] OpenNaaS – Open Platform for Network-as-a-Service, http://opennaas.org

[OpenNFV] HP OpenNFV, http://www8.hp.com/us/en/cloud/nfv-architecture.html

[OpenSAND] OpenSAND, http://opensand.org/

[Openstack] OpenStack cloud controller, http://www.openstack.org

[OpenVirteX] OpenVirtex. Programmable virtual networks. http://ovx.onlab.us/

[OPN] OPNFV, The Open Platform for NFV, http://www.opnfv.org

[OPNFV] Open Platform for NFV, https://www.opnfv.org/

[OStackArch] http://docs.openstack.org/training-guides/content/associate-getting-started.html#associate-openstack-architecture

[OStackCeil] http://docs.openstack.org/developer/ceilometer/architecture.html

[OStackCinder] http://docs.openstack.org/training-guides/content/associate-computer-node.html#associate-block-storage

[OStackDash] http://docs.openstack.org/user-guide/content/ch_dashboard.html

[OStackKey] http://docs.openstack.org/training-guides/content/associate-controller-node.html#associate-keystone-arch

[OVS] Open vSwitch, Production Quality, Multilayer Open Virtual Switch, http://openvswitch.org/

[OWB] OneWeb, http://www.oneweb.net/

[Peng13] Shuping Peng, Reza Nejabati, and Dimitra Simeonidou, "Impairment-aware Optical Network Virtualization in Single-Line-Rate and Mixed-Line-Rate WDM Networks," Journal of Optical Communications and Networking, Vol. 5, No. 3, March 2013.

[Pica8] Pica8 Open Networking, www.pica8.com

[POX] POX Openflow Controller, http://www.noxrepo.org/pox/about-pox

[QSAT] Andrew D. Santangelo, "QuickSAT/Xen: An Open Source Space Hypervisor", AIAA SciTech, 13-17 January 2014, National Harbor, Maryland, 52nd Aerospace Sciences Meeting

[Riera14] J. Ferrer Riera, J., Batallé, E., Escalona, "OpenNaaS: An enabler to deploy Virtual Network Functions". Proc. TERENA Networking Conference, 2014.

[Rosen06] E. Rosen, Y. Rekhter, BGP/MPLS IP Virtual Private Networks (VPNs), RFC 4364 (February 2006).

[SAIL] FP7 SAIL project, http://www.sail-project.eu

[Sandvine] Sandvine Global Internet Phenomena report, 2H 2013, https://www.sandvine.com/downloads/general/global-internet-phenomena/2013/2h-2013-global-internet-phenomena-report.pdf

[SESAME] Project SESAME (Small cEllS coordinAtion for Multi-tenancy and Edge services), http://www.sesame-h2020-5g-ppp.eu/

[Smith05] J. Smith, R. Nair, "Virtual machines: Versatile platforms for systems and processes", Morgan Kaufmann 2005.

[SONATA] Project SONATA (Service Programing and Orchestration for Virtualized Software Networks), http://www.sonata-nfv.eu/

[Spring07] Giovanni Giambene, Resource Management in Satellite Networks Optimization and Cross-Layer Design, Springer, 2007

[Sridharan11] M. Sridharan, K. Duda, I. Ganga, A. Greenberg, G. Lin, M. Pearson, P. Thaler, C. Tumuluri, N. Venkataramiah, and Y. Wang, "NVGRE: Network Virtualization using Generic Routing Encapsulation", Internet Draft, https://tools.ietf.org/html/draft-sridharan-virtualization-nvgre-00

[SRIOV] Single-Root I/O Virtualisation and Sharing Specification, Rev. 1.1, PCI SIG, January 2010

[STC] Satcom Services, http://www.satcom-services.com/broadband2.htm

[Sun10] X. Sun, Y.W Sung, S. Krothapalli and S Rao, "A systematic approach for evolving VLAN designs", IEEE INFOCOM 2010, pp. 1451-1459.

[SUPERFLUIDITY] Project SUPERFLUIDITY (A super-fluid, cloud-native, converged edge system), http://superfluidity.eu/

[SVNOTT] SVNO – A unique customer management tool, http://www.globaltt.com/files/slide/svno/

[TND21] J. Carapinha (Ed.), "System Use Cases and Requirements", FP7/ICT T-NOVA Deliverable D2.1,", http://www.t-nova.eu/results/

[TND221] G. Xilouris (Ed.), "Overall System Architecture and Interfaces", T-NOVA Deliverable D2.21, July 2014, http://www.t-nova.eu/wp-content/uploads/2014/01/TNOVA_D2.21_Overall_System_Architecture_and_Interfaces_v.1.0.pdf

[TNOVA] FP7 T-NOVA project, Network Functions as-a-Service over Virtualised Infrastructures, http://www.t-nova.eu

[TNOVAD241] A. Gamelas (ed.) et al, "Specification of the Infrastructure Virtualisation, Management and Orchestration – Interim", T-NOVA Deliverable D2.31, September 2014, http://www.t-nova.eu/wp-content/uploads/2014/12/TNOVA_D2.31_Spec_of_IVM_and_Orchestrator_I.pdf

[Touch09] J. Touch and R. Perlman, "Transparent Interconnection of Lots of Links (TRILL): Problem and Applicability Statement", IETF RFC 5556, 2009.

[UNIFY] Mario Kind (ed.) et al, "D2.1: Use Cases and Initial Architecture", FP7 UNIFY Project, August 2013

[vCloud] vmware. vCloud Suite, http://www.vmware.com/products/vcloud-suite

[Verchere11] D. Verchere, "Cloud Computing over Telecom Network", Proc. OFC 2011.

[VillasenorDC] A. Villasenor, "Digital Cinema: Using Satellite CDN as Delivery and Transport Platform", White Paper, Globecomm Systems, Inc., http://www.globecommsystems.com/pdf/wp-digital-cinema-gsi-rev-2.pdf

[VIPS] Technical answer to AO 726212/NL/US, "Validation of complex payloads and associated systems and networks", October 2011

[VITAL] Project VITAL (VIrtualized hybrid satellite-TerrestriAl systems for resilient and fLexible future networks), http://www.ict-vital.eu/

[VTN] OpenDaylight Virtual Tenant Network (VTN) https://wiki.opendaylight.org/view/OpenDaylight_Virtual_Tenant_Network_%28VTN%29:Main.

[WLC] Choice of Discount Rates, http://www.wlcf.org.uk/page32.html

[XIFI] FP7 XIFI Project, https://fi-xifi.eu/home.html

[XTRAT] M. Masmano, I. Ripoll, A. Crespo and S. Peiro, "XtratuM for LEON3: an Open Source Hypervisor for High Integrity Systems", In European Conference on Embedded Real Time Software and Systems. ERTS2 2010., Toulouse (France), 19-21 May 2010.

[Yang04] I. Yang et al, "Forwarding and Control Element Separation (ForCES) Framework", IETF RFC 3746, 2004

[Yang13] Mao Yang, Yong Li, Depeng Jin, Li Su, Shaowu Ma, Lieguang Zeng, "OpenRAN: A Software-defined RAN Architecture Via Virtualization", Proc. SIGCOMM 2013.

11. List of Acronyms

Acronym	Explanation
ABNO	Application Based Operations Architecture
API	Application Programming Interface
BEP	Breakeven Point
BEY	Breakeven Year
BGP	Border Gateway Protocol
BMC	Business Model Canvas
BoD	Bandwidth-on-Demand
BSS	Billing Support System
BSS	Broadcast Satellite Services
CAPEX	Capital Expenditure
CBA	Cost Benefit Analysis
CDN	Content Delivery Network
CEA	Cost Effectiveness Analysis
CF	Cash Flow
CFA	Cash Flow Analysis
CFM	Cash Flow Model
COTS	Commercial-off-the-shelf
CPE	Customer Premises Equipment
CR	Customer Relationships
C-RAN	Centralised RAN or Cloud RAN
DC	Data Centre
DCF	Discounted Cash Flow
DoS	Denial of Service
DP	Discounted Payback
DPDK	Data Plane Development Kit
EPC	Evolved Packet Core
ETSI	European Telecommunications Standards Institute
EU	End User
FA	Financial Analysis

FE	(Customer) front-end
FM	Federated Manager
FSS	Fixed Satellite Service
GEO	Geostationary Earth Orbit
GRE	Generic Routing Encapsulation
GUI	Graphical User Interface
HTS	High Throughput Satellite
HW	Hardware
IaaS	Infrastructure as a Service
ICMP	Internet Control Message Protocol
InPs	Infrastructure Providers
IR	Interest Rate
IRR	Internal Rate of Return
IT	Information Technology
KVM	Kernel-based Virtual Machine
LEO	Low Earth Orbit
LTE	Long-Term Evolution
M2M	Machine-to-Machine
MACRS	Modified Accelerated Cost Recovery System
MANO	Management and Orchestration
MEC	Mobile Edge Computing
MEO	Medium Earth Orbit
MNO	Mobile Network Operator
MPLS	Multi Protocol Label Switching
MSP	Media Service Provider
MSS	Mobile Satellite Services
NaaS	Network-as-a-Service
NFV	Network Functions Virtualization
NFVI	Network Functions Virtualisation Infrastructure
NFVI-PoP	NFVI Point-of-Presence
NGSO	Non-Geostationary
NMS	Network Management System

NPV	Net Present Value
NVGRE	Network Virtualisation using Generic Routing Encapsulation
OBP	On-Board Processing
OF	Openflow
ONF	Open Networking Foundation
OPEX	Operating Expenditures
OPNFV	Open Platform for NFV
OS	Operating System
OSS	Operation Support System
OTT	Over-the-top
PaaS	Platform as a Service
PI	Profitability Index
QoE	Quality of Experience
QoS	Quality of Service
RAN	Radio Access Network
REST	Representational State Transfer
ROI	Return on Investment
RRM	Radio Resource Management
RTT	Round-Trip Time
SaaS	Software as a Service
SDK	Software Development Kit
SDN	Software defined Networking
SDO	Standards Developing Organisation
SDWN	Software-Defined Wireless Network
SFC	Service Function Chaining
SLA	Service Level Agreement
SOHO	Small Office Home Office
SP	Simple Payback
SRIOV	Single-Root I/O virtualisation
STT	Stateless Transport Tunneling
SVNO	Satellite Virtual Network Operator
SW	Software

TTM	Time-to Market
UC&C	Unified Communications & Collaboration
VAS	Value Added Service
VC	Value Chain
VIM	Virtualised Infrastructure Management
VLAN	Virtual Local Area Network
VM	Virtual Machine
VNF	Virtual Network Function
VNFaaS	Virtual Network Function as-a-Service
VNO	Virtual Network Operator
VNSP	Virtual Network Service Provider
VPN	Virtual Private Network
VXLAN	Virtual Extensible LAN
WAN	Wide-Area Network
WIP	Work In Progress

12. APPENDIX I: DETAILED DESCRIPTION OF OPENSTACK COMPONENTS

12.1. Horizon (Dashboard)

The OpenStack dashboard [OStackDash] enables the user to provision their own services within the limits set by administrators. The user can create different types and sizes of server instances. Horizon provides a Web UI to send requests to the APIs provided by the rest of elements in OpenStack Architecture.

12.2. Keystone

The Keystone identity service performs these functions:

- User management: Tracks users and their permissions. The main components of Identity user management are:
  - Users
  - Tenants
  - Roles
- Service catalog: Provides a catalog of available services with their API endpoints. The Identity Service provides the following service management functions:
  - Services
  - Endpoints

The Identity Service acts as a common authentication system across the cloud operating system and can integrate with existing backend directory services like LDAP.

More details can be found in [OStackKey].

12.3. Glance

OpenStack Image Service (Glance) provides discovery, registration and delivery services for disk and server images. Stored images can be used as a template. They can also be used to store and catalog an unlimited number of backups. The Image Service can store disk and server images in a variety of back-ends, including OpenStack Object Storage. The Image Service API provides a standard REST interface for querying information about disk images and lets clients stream the images to new servers.

Capabilities of the Image Service include:

- Administrators can create base templates from which their users can start new compute instances
- Users can choose from available images, or create their own from existing servers
- Snapshots can also be stored in the Image Service so that virtual machines can be backed up quickly
- A multi-format image registry, the image service allows uploads of private and public images in a variety of formats.

Glance serves a central role to the overall IaaS picture. It accepts API requests for images (or image metadata) from end users or Nova components and can store its disk files.

12.4. Nova

Nova is the most complicated and distributed component of OpenStack. A large number of processes cooperate to turn end user API requests into running virtual machines. Below is a list of these processes and their functions:

- nova-api accepts and responds to end user compute API calls. It supports OpenStack Compute API, Amazon's EC2 API and a special Admin API (for privileged users to perform administrative actions). It also initiates most of the orchestration activities (such as running an instance) as well as enforces some policy (mostly quota checks).
- The nova-compute process is primarily a worker daemon that creates and terminates virtual machine instances. The process by which it does so is fairly complex but the basics are simple: accept actions from the queue and then perform a series of system commands (like launching an instance) to carry them out while updating state in the database.
- The nova-scheduler process is conceptually the simplest piece of code in OpenStack Nova: it takes a virtual machine instance request from the queue and determines where it should run (specifically, which compute server host it should run on).
- Nova also provides console services to allow end users to access their virtual instance's console through a proxy. This involves several daemons (nova-console, nova-novncproxy and nova-consoleauth).

Nova interacts with many other OpenStack services: Keystone for authentication, Glance for images and Horizon for web interface. The Glance interactions are central. The API process can upload and query Glance while nova-compute will download images for use in launching images.

The nova-conductor service enables OpenStack to function without compute nodes accessing the database. Methods exposed by nova-conductor are relatively simple methods used by nova-compute to offload its database operations. The conductor service implements long running complex operations, ensuring forward progress and

graceful error handling. This will be especially beneficial for operations that cross multiple compute nodes, such as migrations or resizes.

12.5. Neutron

Neutron provides "network connectivity as a service" between interface devices managed by other OpenStack services (most likely Nova). The service allows users to create their own networks and then attach interfaces to them. Like many of the OpenStack services, Neutron is highly configurable due to its plug-in architecture. These plug-ins accommodate different networking equipment and software.

neutron-server accepts API requests and then routes them to the appropriate Neutron plug-in for action. Neutron plug-ins and agents perform the actual actions such as plugging and unplugging ports, creating networks or subnets and IP addressing. These plug-ins and agents differ depending on the vendor and technologies used in the particular cloud. Neutron ships with plug-ins and agents for: Cisco virtual and physical switches, NEC OpenFlow products, Open vSwitch, Linux bridging, the Ryu Network Operating System, and VMware NSX.The common agents are L3 (layer 3), DHCP (dynamic host IP addressing) and the specific plug-in agent.

12.6. Heat

Heat implements an orchestration engine to launch multiple composite cloud applications based on templates in the form of text files that can be treated like code. A native Heat template format is evolving, but Heat also endeavours to provide compatibility with the AWS CloudFormation template format, so that many existing CloudFormation templates can be launched on OpenStack. Heat provides both an OpenStack-native ReST API and a CloudFormation-compatible Query API.

Heat works as follows:

- A Heat template describes the infrastructure for a cloud application in a text file that is readable and writable by humans, and can be checked into version control.
- Infrastructure resources that can be described include: servers, floating IPs, volumes, security groups, users, etc.
- Heat also provides an autoscaling service that integrates with Ceilometer, so one can include scaling rules in a template.
- Templates can also specify the relationships between resources (e.g. volume A is connected to server B). This enables Heat to call out to the OpenStack APIs to create the infrastructure in the correct order to.
- Heat manages the whole lifecycle of the application - when there is a need to change the infrastructure, the template can be modified and used to update the existing stack.

12.7. Cinder

OpenStack Block Storage (Cinder) [OStackCinder] provides persistent block level storage devices for use with OpenStack compute instances. The block storage system manages the creation, attaching and detaching of the block devices to servers. Block storage volumes are fully integrated into OpenStack Compute and the Dashboard allowing for cloud users to manage their own storage needs. Block storage is appropriate for performance sensitive scenarios such as database storage, expandable file systems, or providing a server with access to raw block level storage. Snapshot management provides powerful functionality for backing up data stored on block storage volumes. Snapshots can be restored or used to create a new block storage volume.

The current services available in OpenStack Block Storage are:

- cinder-api - The cinder-api authenticates and routes requests throughout the Block Storage system.
- cinder-scheduler - The cinder-scheduler is responsible for scheduling/routing requests to the appropriate volume service.
- cinder-volume - The cinder-volume service is responsible for managing Block Storage devices, specifically the back-end devices themselves.

12.8. Ceilometer

The Openstack Telemetry module, also called Ceilometer [OStackCeil] collects the metering data about resource utilisation. It also collects data by monitoring notifications sent from services or by polling the infrastructure. Ceilometer configures the type of collected data to meet various operating requirements.

Ceilometer consists of the following basic components:

- A compute agent (ceilometer-agent-compute). Runs on each compute node and polls for resource utilization statistics.
- A central agent (ceilometer-agent-central). Runs on a central management server to poll for resource utilization statistics for resources not tied to instances or compute nodes.
- A collector (ceilometer-collector). Runs on one or more central management servers to monitor the message queues (for notifications and for metering data coming from the agent). Notification messages are processed and turned into metering messages and sent back out onto the message bus using the appropriate topic. Telemetry messages are written to the data store without modification.
- An alarm notifier (ceilometer-alarm-notifier). Runs on one or more central management servers to allow setting alarms based on threshold evaluation for a collection of samples.

- A data store. A database capable of handling concurrent writes (from one or more collector instances) and reads (from the API server).
- An API server (ceilometer-api). Runs on one or more central management servers to provide access to the data from the data store. These services communicate using the standard OpenStack messaging bus. Only the collector and API server have access to the data store.

13. APPENDIX II: THE BUSINESS MODEL CANVAS

The Business Model Canvas (BMC) was developed by Alex Osterwalder and Yves Pigneur, and co-created with an array of 470 practitioners from around the world.

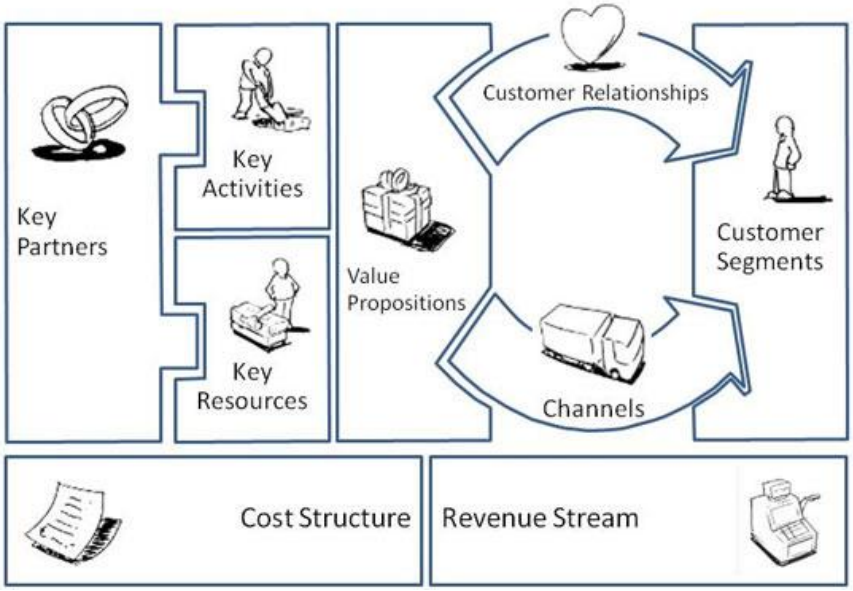

Figure 278. Business Model Canvas (PictorialView as per Osterwalder and Pigneur's) [BMC]

It offers a visual, one-page canvas providing a way of composing a business model with nine building blocks [BMC]:

- Customer Segments
- Value Propositions
- Channels
- Customer Relationships
- Revenue Streams
- Key Resources
- Key Activities
- Key Partnerships
- Cost Structure

Unlike many of the frameworks, the BMC was built out of careful research and it has also been tested and enhanced through the input of many practitioners.

Below depicts the Business Model canvas into a flow chart approach, where all the blocks are organized around the value proposition block, which reflects the business objective of value to be delivered and acknowledged and respectively the belief from

the customer that value will be delivered and experienced by the proposed innovation/business.

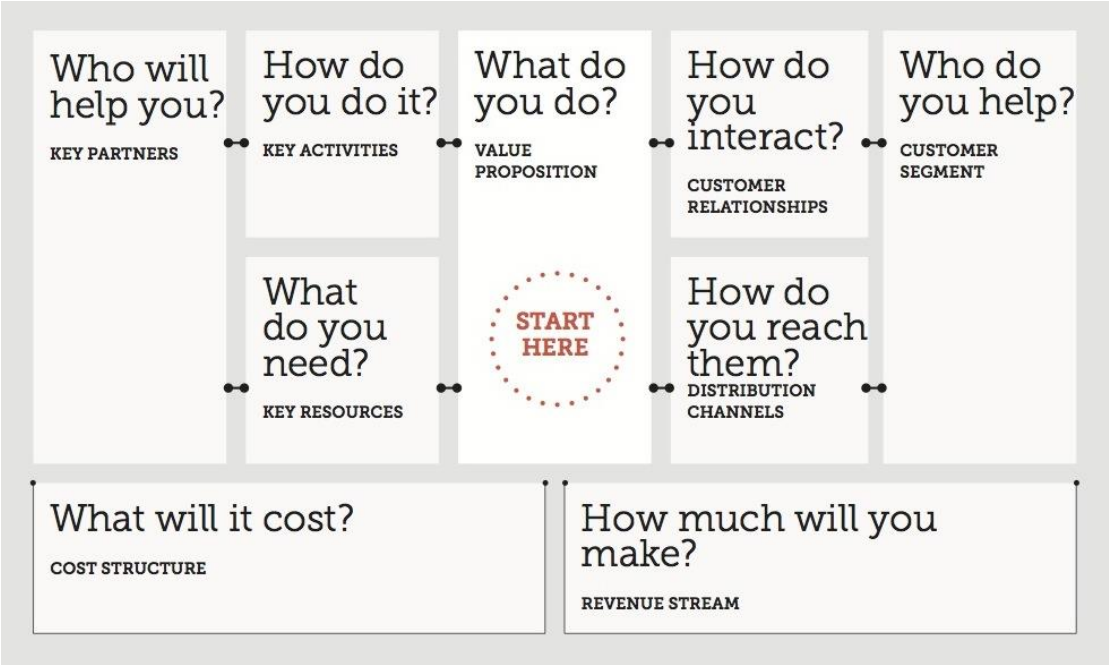

Figure 279. Business Model Canvas key questions [BMC]

As above figure depicts, in our case the value proposition is strongly bounded with the NFV/SDN applicability on the satellite network components, considering the value gained by the cloudification of the involved –till today- HW units. Through this cloudification, the respective value proposition can apply to the entire Satellite Service Provider, although parts of the value can be considered to the provided products or the services or even the customer. Thus around the value proposition are gathered all the key questions relative to the business model definition of the cloud networking model, ranging from the key partnerships to customers segments and from cost structure to revenue streams.

These nine blocks, together with the key questions, map on the four main business areas of a business concept, which are:

- Customers
- Offer
- Infrastructure
- Financial viability

14. APPENDIX III: FINANCIAL ANALYSIS METHODOLOGY AND TERMS

Innovations require financial investments either for developing the new technology and infrastructure or for upgrading an existing one so as the business entity willing to offer such a service to be able to provide it sufficiently and effectively. The different financial models, methodologies and frameworks that might be developed and used depend on the type of investment, current market conditions, revenue sources, as well as the accompanying business and pricing models that the business has adapted.

Following the CloudSat business and market framework of analysis, in this appendix we describe the methodology and the tools that were used for the development of the CloudSat financial framework.

The accuracy of the financial analysis can be further reinforced by estimating the financial viability of the new investment by applying the discounted Cash Flow Model (CFM) based on the CAPEX, OPEX, Revenues and Contribution Margin (CM) estimated figures. CFM is a 5-step approach which is applied in the rest sections of this financial framework. In specific:

1. Business and Market analysis
2. Estimate Costs, Revenues, Contribution Margin (CM)
    a. Initial investment cost, FC - CAPEX
    b. Operating/Running cost, VC - OPEX
    c. Revenues
    d. Net Selling Price (NSP), Variable Manufacturing/Provision Cost (VMC), Variable Cost (VC), Cost of Goods Sold (COGS), altogether as expressed through Contribution Margin (CM= NSP -VC) where CM represents the portion of sales revenue that is not consumed by variable costs and so contributes to the coverage of fixed costs.
3. Breakeven Analysis (EBITDA, Taxes, Cash Flows, Discounted Cash Flows)
    a. EBITDA (Earnings Before Interest, Taxes, Depreciation and Amortization). EBITDA, is a popular equity evaluation metric for analyzing companies in the telecommunications sector mainly because of what the metric excludes, such as depreciation.
    b. Breakeven year
    c. Cash Flows, Discounted Cash Flows
4. High Level Financial Analysis/Financial Ratios
5. Cost/Benefit Analysis

So, following the above discounted Cash Flow Model (CFM), the structure of this section, presenting the financial framework of analysis, is summarized as per below:

- Identify the general reasonable assumptions of the financial analysis
- Define the benchmarking cases (as per GEO, MEO, LEO satellite systems) and estimate corresponding CAPEX costs and Cost reduction rates
- Summarize findings of benchmarking cases in a CAPEX Cost-Benefit Analysis (CBA)
- Define the three CloudSat case scenarios to be analyzed

and perform the following steps of financial analysis per CloudSat case/scenario:

- Estimate the initial Investment cost
    - Per component
    - FC/CAPEX
- Estimate the Operating/Running cost
    - Per component
    - VC/OPEX
- Perform high level financial analysis
    - Revenues
    - Cash Flow Model (CFM) and discounted Cash Flow
    - Financial Ratios
- In addition, for the three CloudSat cases, three different business environment scenarios of evaluations are used with their specific high level financial analysis variables/parameters
    - Optimistic (Blue Ocean)
    - Normal
    - Pessimistic (Red Ocean)
- Summarize findings of all CloudSat cases in a Cost-Benefit Analysis (CBA)

14.1. Identification of reasonable assumptions

Since a financial analysis relies on many assumptions, it is important to document all of them, and, if possible, justify them on the basis of prior experiences or actual data originating for the internal or external business, market and financial environment.

Initially, financial business environment assumptions must be set and analyzed such as restrictions, conditions and hypotheses concerning variables, rates and taxes. In addition, specific to the business sector assumptions must be made like productivity, cash flow, costs, annual depreciation rates, trends. Finally, assumptions about external environmental factors that may impact the business such as political, social

and legal must be taken into consideration e.g. possible risks, possible impact (quantitative approach) etc.

14.2. Estimation of initial investment cost

The accurate estimation of the initial investment cost is fundamental for any financial analysis. Initial investment is the amount required to start a business or a project. It is also called initial investment outlay or simply initial outlay.

Capital budgeting or initial investment decisions involve careful estimation of the initial investment outlay and future cash flows of a project. Correct estimation of these inputs helps in taking decisions that increase shareholders wealth.

Initial investment equals the amount needed for capital expenditures, such as machinery, tools, shipment and installation, etc.; plus any increase in working capital, minus any after tax cash flows from disposal of any old assets. Sunk costs are ignored because they are irrelevant.

Initial Investment = Capital Expenditures + Increase in Working Capital − Disposal Inflows

Capital expenditure, or CAPEX, is funds used by a company to acquire or upgrade physical assets such as property, industrial buildings or equipment in order to undertake new projects or investments. This type of outlay is also made by companies to maintain or increase the scope of their operations.

In terms of accounting, an expense is considered to be a capital expenditure when the asset is a newly purchased capital asset or an investment that improves the useful life of an existing capital asset. If an expense is a capital expenditure, it needs to be capitalized. This requires the company to spread the cost of the expenditure (the fixed cost) over the useful life of the asset. If, however, the expense is one that maintains the asset at its current condition, the cost is deducted fully in the year of the expense.

The amount of capital expenditures a company is likely to have depends on the industry it occupies. Some of the most capital intensive industries have the highest levels of capital expenditures including oil exploration and production, telecom, manufacturing and utilities.

Capital expenditure should not be confused with revenue expenditure or operating expenses (OPEX). Revenue expenses are shorter-term expenses required to meet the ongoing operational costs of running a business, and therefore they are essentially identical to operating expenses. Unlike capital expenditures, revenue expenses can be can be fully tax-deducted in the same year in which the expenses occur.

In order to have an accurate initial investment cost and CAPEX estimate, we break the initial investment to its components and we provide an accurate estimate of its FC (per component).

The major financial ratio that can be used to evaluate the initial investment ability of a company is the cash flow to capital expenditure ratio, or CF/CapEX, which is related to a company's ability to acquire long term assets using cash flow.

Profitability Index = CF/CapEx =

Cash Flow From Operations / Capital Expenditures

14.3. Estimation of operating/running cost

An operating expense, operating expenditure, operational expense, operational expenditure or OPEX is an ongoing cost for running a product, business, or system. It is a category of expenditure that a business incurs as a result of performing its normal business operations. Considering the economy's changing dynamics, it has become essential for companies (public as well as private) to convert capital expenditure (CAPEX) to operational expenditure (OPEX). In specific, there is a direct correlation between OPEX and the value of the enterprise, in that when the OPEX decreases, while maintaining the same level of production and quality, the overall value of the enterprise increases.

Both business and accounting wise, operating expenses include:

- Leasing equipment and service expenses
- maintenance and repairs
- advertising/promotion
- supplies
- utilities
- accounting expenses
- license fees
- insurance
- fees
- property management
- unforeseen expenses and
- taxes (analyzed on separate section of a financial analysis)

In order to have an accurate operating/running cost and OPEX estimate, we break the initial investment to its components and we provide an accurate estimate of its CAPEX (per component) and then based on it we estimate the operating and variable expenses. This analysis will lead as to the breakeven year (year on which Cash Flows turn positive), Payback period and the calculation of a positive total EBITDA for the analyzed period.

14.4. High level financial analysis

Over the years, investors and analysts have developed numerous analytical tools, concepts and techniques to compare the relative strengths and weaknesses of

companies and investments. These tools, concepts and techniques form the basis of fundamental financial analysis.

Ratio analysis is a tool that was developed to perform quantitative analysis on numbers found on financial statements. Ratios help link the three financial statements together and offer figures that are comparable between companies and across industries and sectors. Ratio analysis is one of the most widely used fundamental financial analysis techniques.

For the scope of TN3.3 analysis, breakeven analysis (using EBITDA, which is a popular equity evaluation metric for analyzing companies in the telecommunications sector mainly because of what the metric excludes, such as depreciation) and four financial ratios/metrics have been selected for the first part of the high level financial analysis and evaluation of the CloudSat architecture along with CFM method. In addition, CloudSat high level financial analysis and its financial efficiency in terms of revenues and profitability will be evaluated through the use of 3 business environment scenarios: Optimistic (Blue Ocean), Normal and Pessimistic (Red Ocean). Under those three scenarios conditions, the four financial ratios/metrics, along with the CAPEX and OPEX estimates and CFM results, will be recalculated and interpreted.

1. Return on Investment (ROI) is a financial accounting measurement for determining the value of making a specific investment. ROI is a ratio of the net benefits to the total cost of an investment for the same specific period. The two principle concerns with ROI are that the calculations do not account for the time value of money and the calculations assume a consistent annual rate of return. ROI is a useful measure when comparing alternatives using the same cost and benefit criteria for the same period. The formula for calculating ROI is:

   ROI% = (Net Benefit / Cost of investment) X 100

   where Net Benefit = Benefits from investment – Cost of investment

2. The Net Present Value (NPV) of an investment is the present (discounted) value of future cash inflows minus the present value of the investment and any associated future cash outflows. By considering the time value of money, it allows consideration of such things as cost of capital, interest rates, and investment opportunity costs.
   NPV is important because without using the net present value of benefits and cost the comparisons drawn between solutions in the out years are not accurate. This metric recognizes that money has different real value over time and makes the values of money constant by discounting costs and benefits over a specific period of time—an asset's life cycle or any selected period of analysis.

   NPV allows managers and investors to compare, on purely financial factors, investment alternatives with widely disparate cash flows. NPV facilitates objective evaluation of projects regardless of scale differences or the existence of capital rationing, and can be used to compare independent or mutually exclusive projects. For each year of the analysis period, cash inflows (benefits)

and cash outflows (costs) are totaled and then summed to arrive at the net impact on cash. The net cash flow is then multiplied by an appropriate discount factor to arrive at a discounted cash flow for each year. NPV is the total of these discounted cash flows over the period of analysis.

Generating a meaningful NPV requires sound estimates of the costs and benefits of a project, use of the appropriate discount rate, and the identification of the timing of cash receipts and disbursements. NPV focuses on an investment's impact on cash flow rather than net profit, or savings in the case of non-revenue generating entities.

Given the (period, cash flow) pairs (t, R_t) where N is the total number of periods, the net present value NPV is given by:

$$NPV(i,N) = \sum_{t=0}^{N} \frac{R_t}{(1+i)^t}$$

where

t – the time of the cash flow

i – the discount rate (the rate of return that could be earned on an investment in the financial markets with similar risk.); the opportunity cost of capital

R_t – the net cash flow i.e. cash inflow – cash outflow, at time t. R_0 is commonly placed to the left of the sum to emphasize its role as (minus) the investment.

Fundamentally, the discount rate, used in NPV calculation, reflects the opportunity cost of capital, i.e. by investing in one project we sacrifice the return from investing in another project. The discount rate also refers to the interest rate used in discounted cash flow (DCF) analysis to determine the present value of future cash flows. The discount rate in DCF analysis takes into account not just the time value of money, but also the risk or uncertainty of future cash flows; the greater the uncertainty of future cash flows, the higher the discount rate.

- A lower bound for the discount factor is therefore a risk-free, social discount rate, assuming a perfectly functioning financial market (a range up to 2%). This value may be country specific, since it is supposed to reflect the long-term rate of growth in the economy [CRU].
- A higher bound can then be obtained by incorporating the risk related to the financial assets, the financial portfolio of the specific actor, and the risk related to the underlying project (above 8%). Note that a discount rate can also include inflation (nominal rate) or not (real rate) [CRU].

Nowadays, it is recommended economic development organizations to use a discount rate of 4% to 5% when the financial markets are stable, economies not in risk and degree of uncertainty is not high (conditions that apply to developed and most of the developing countries). Based on the above, the value of the discount rate to be used in the TN3.3 financial analysis should lay

between a lower and a higher bound (eg a discount rate of 5%) and adjusted as required by the conditions of three 3 business environment scenarios of CloudSat analysis: Optimistic (Blue Ocean), Normal and Pessimistic (Red Ocean) [CDI], [WLC], [CRU].

3. The **Internal Rate of Return** (**IRR**) or economic rate of return (ERR) is the rate of return used in capital budgeting to measure and compare the profitability of investments. The IRR on an investment or project is the "annualized effective compounded return rate" or rate of return that makes the net present value of all cash flows (both positive and negative) from a particular investment equal to zero. It can also be defined as the discount rate at which the present value of all future cash flow is equal to the initial investment or in other words the rate at which an investment breaks even. Given the (period, cash flow) pairs (n, C_n) where n is a positive integer, the total number of periods N, and the net present value NPV, the internal rate of return is given by r in:

$$NPV = \sum_{n=0}^{N} \frac{C_n}{(1+r)^n} = 0$$

4. The **Payback Period method** determines the time necessary for a new investment to pay for itself. Payback does not measure profitability, but cash recoverability. Payback tends to show the risk factor by pointing out the recovery time of an investment. Its primary advantage is its simplicity - it is quick to calculate and easy to understand. Its limitations include:
    - Does not consider the benefit of net results after the investment has been repaid
    - Does not take into account the time value of money (as NPV does).

    Payback period is usually expressed in years. Start by calculating Net Cash Flow for each year: Net Cash Flow Year 1 = Cash Inflow Year 1 - Cash Outflow Year 1. Then Cumulative Cash Flow = (Net Cash Flow Year 1 + Net Cash Flow Year 2 + Net Cash Flow Year 3, etc.) Accumulate by year until Cumulative Cash Flow is a positive number: that year is the payback year.

In addition, as result of the application of the Cash Flow Model (CFM) in our analysis, the Profitability Index will be also calculated for evaluating the initial investment ability of a company to acquire long term assets using expected future cash flows.

In the second part, CloudSat high level financial analysis and its financial efficiency in terms of revenues and profitability is evaluated through the use of 3 business environment scenarios: Optimistic (Blue Ocean), Normal and Pessimistic (Red Ocean). This evaluation will apply to the three CloudSat scenarios (and not to the two benchmarking cases)

As per [FMG], Scenario analysis is designed to allow improved decision-making by allowing consideration of outcomes and their implications. Consequently, a scope of possible future outcomes is observable. In specific, not only are the outcomes observable, but also the development paths leading to the outcomes become

noticeable. Although an infinite set of different assumptions or conditions exists, scenario analysis includes the three basic aforementioned scenarios.

In the Optimistic scenario, the financial analysis assumes that the outcomes of some variables are better than the normal (most likely) values e.g operating revenues, rates, costs. In the Pessimistic scenario, the outcomes of some variables are worse than the normal values. For instance, the initial outlay, the taxation or the interest rate could be higher than expected and the investment payback period longer than anticipated. These two case scenarios represent the two extremes of the Normal (most likely) scenario in which all variables have their expected-normal values [WLC].

Different strategies being applied by businesses as part of their business plan for achieving their goals may lead to completely different business environment scenarios. Within this context, a red ocean strategy is more likely to create the conditions of a pessimistic scenario, while a blue ocean strategy may bring results better than expected and setup the conditions of an optimistic scenario.

Red oceans represent all the industries in existence today. In red oceans industry boundaries are defined and companies try to outperform their rivals to gain a greater market share. As the space gets more and more crowded, profits are reduced and products turn into commodities, and increasing competition turns the water bloody [BOS].

Blue oceans denote all the industries not in existence today – the unknown market space with no competition. In blue oceans, demand is created rather than fought over. There is ample opportunity for growth that is both profitable and rapid.

Creating a blue ocean will allow rapid growth and high profits but eventually the space will invite competitors and erode profitability so a blue ocean strategy requires that a company continually search for new ways to break away from the crowd.

From a financial perspective, blue ocean strategy argues that the simultaneous pursuit of differentiation and low cost is achievable. A blue ocean is created in the region where a company's actions favorably affect both its cost structure and its value proposition to buyers. Cost savings are made from eliminating and reducing the factors an industry competes on. Buyer value is lifted by raising and creating elements the industry has never offered. Over time, costs are reduced further as scale economies appear, due to the high sales volumes that superior value generates [BOS].

The figure below summarizes the characteristics of Red and Blue Ocean strategies.

Figure 280. Red and Blue Ocean Strategy characteristics [BOS]

14.5. Cost-Benefit Analysis (CBA)

The last step of this financial analysis framework will provide a summarization of findings in a combined Cost-Benefit Analysis (CBA) in which the results in terms of economic, technical and business-social gains, tangible and intangible benefits versus all type of costs will be evaluated.

Identifying these gains-benefits will usually require an understanding of the business processes of the project, the business and its customers. Some benefits realized by the business are flexibility, organizational strategy, risk management and control, organizational changes, and staffing impacts. These benefits are often measured in terms of productivity gains, staffing changes, and improved business effectiveness. Possible benefits to customers include improvements to the current IT services and the addition of new services. These benefits can be measured in terms of productivity gains and cost savings, but the customers must be the ones to identify and determine how to measure and evaluate the benefits. Customer surveys are often needed to identify these benefits.

Benefits might be tangible or intangible. Tangible benefits originate from increased revenue, cost reduction, and cost avoidance. They measure, in monetary savings, the impact of an alternative on society, people, businesses, equipment, time, facilities, and support materials. Intangible benefits are subjective issues that can exert strong influences on the entire process, but can seldom be measured in monetary terms. Some intangible benefits are: better and/or timelier decision-making, more accurate information, better reporting, political response, goodwill in the community, personnel morale etc.

In specific, the CloudSat CBA will present and evaluate the:

- Direct and indirect benefits
- Positive and negative impacts
- Economical gains and limitations/constraints
- Evaluation of social affect

Conclusions on the financial viability and suitability of CloudSat