NECOS Project: Towards Lightweight Slicing of Cloud Federated Infrastructures

Felipe S. Dantas Silva[1], Marcilio O. O. Lemos[1], Alisson Medeiros[1], Augusto Venâncio Neto[1], Rafael Pasquini[2]
David Moura[3], Christian Rothenberg[2], Lefteris Mamatas[4], Sand Luz Correa[5], Kleber Vieira Cardoso[5]
Cesar Marcondes[6], Antonio ABelem[7], Marcelo Nascimento[8], Alex Galis[9], Luis Contreras[10], Joan Serrat[11]
Panagiotis Papadimitriou[4]

[1]Federal University of Rio Grande do Norte (UFRN), Natal, Brazil
Email: felipe.dantas@ifrn.edu.br, marcilio.cc.lemos@gmail.com, alisson@ppgsc.ufrn.br, augusto@dimap.ufrn.br
[2]Federal University of Uberlândia (UFU), Uberlândia, Brazil - Email: rafael.pasquini@ufu.br
[3]University of Campinas (UNICAMP), Campinas, Brazil - Email: dfcmoura@gmail.com, chesteve@dca.fee.unicamp.br
[4]University of Macedonia (UOM), Thessaloniki, Greece - Email: {emamatas,papadimitriou}@uom.edu.gr
[5]Federal University of Goiás (UFG), Goiânia, Goiás - Email: {sand,kleber}@inf.ufg.br
[6]Federal University of São Carlos (UFSCAR), São Carlos, Brazil - Email: marcondes@dc.ufscar.br
[7]Federal University of Pará (UFPA), Belém, Brazil - Email: abelem@ufpa.br
[8]CPqD Telecom R&D Center (CPqD), Campinas, Brazil - Email: marcelon@cpqd.br
[9]University College London (UCL), London, England - Email: a.galis@ucl.ac.uk
[10]Telefónica Investigación y Desarrollo (TID), Spain - Email: luismiguel.contrerasmurillo@telefonica.com
[11]Universitat Politècnica de Catalunya (UPC), Spain - Email: serrat@tsc.upc.edu

Abstract—The *Novel Enablers for Cloud Slicing* (NECOS) project addresses the limitations of current cloud computing infrastructures to respond to the demand for new services, as presented in two use-cases, that will drive the whole execution of the project. The first use-case is focused on Telco service provider and is oriented towards the adoption of cloud computing in their large networks. The second use-case is targeting the use of edge clouds to support devices with low computation and storage capacity. The envisaged solution is based on a new concept, the *Lightweight Slice Defined Cloud* (LSDC), as an approach that extends the virtualization to all the resources in the involved networks and data centers and provides uniform management with a high-level of orchestration. In this position paper, we discuss the motivation, objectives, architecture, research challenges (and how to overcome them) and initial efforts for the NECOS project.

Index Terms—Cloud Federation, Slicing as a Service, Management and Orchestration

I. INTRODUCTION

The proliferation, increasing popularity and maturity of *Cloud Computing* (CC), virtualization and *Data Center* (DC) technologies come with tremendous attention from the research community, the practitioners, and the *Information and Communication Technology* (ICT) industry. There are hundreds of providers that are offering services in any of three service models: (*i*) *Software as a Service* (SaaS), (*ii*) *Platform as a Service* (PaaS), and (*iii*) *Infrastructure as a Service* (IaaS). This variety of existing cloud services creates a new challenge for service providers that are using separately managed computing, connectivity, and storage resources to easily deploy new services, as well as enforcing reasonable *Service Level Agreements* (SLAs).

Over the last ten years, CCs/DCs have established themselves as a highly economical way of providing high-performance computation and storage services of all kinds and continue to have a significant impact on service architectures and implementations. Economies of scale can be achieved through the concentration of operations and highly improved different abstractions available at the CC/DC level. Recent research and development have focused on improving application and service deployment and management, energy efficiency, the versatility, the security, and the availability of CCs/DCs. In parallel to the CCs/DCs evolution, network technologies have evolved tremendously. Nevertheless, these network technologies are currently based on routing devices, middleboxes and data servers tighten together to the specific service that they are supporting. This means that modifying a service or deploying a new one requires installation and reconfiguration of many of those network devices. This is creating several significant difficulties for *Service Providers* (SP) and *Network Operators* (NO) in provisioning innovative network functionalities, new services, which need to benefit from the increasing dynamicity of the ICT markets. Launching new services is still time-consuming and requires extensive and expensive efforts, thus preventing any rapid roll-out of new businesses models and limiting opportunities.

Having identified a problem that needs to be addressed, in this position paper we introduce the *Novel Enablers for Cloud Slicing* (NECOS), which employs a new paradigm to fulfill the following three requirements. The first systemic requirement is making the innovation and integration cycles of applications, networking, and services features faster and simpler given that there is tighter integration of CCs/DCs and networking systems. The service provider infrastructure is transitioning from the traditional fixed linked model of inter-connected DCs

towards a model of integration and software-defined storage, computation and communication capabilities represented by virtual services. Paradigms such as *Software-Defined Networking* (SDN) [1], *Network Function Virtualization* (NFV) [2], *Mobile Edge Computing* (MEC) [3] and Network Slicing [4], if properly designed and deployed, could help in fulfilling a deeper integration of networking, cloud computing, and storage resources. Consequently, the second systemic requirement is to develop better-integrated management systems (currently mainly developed and carried out separately for CCs/DCs and Edge Clouds).

The reduction in OPEX can be achieved by reducing or limiting human intervention and by reducing inaccurate manual operations. This process combined with flexible and optimal provisioning of network and server functions and services, across domains, can also mitigate systems costs. As such, the third systemic requirement is the energy efficiency and interoperability across multiple domains. With the right abstractions, heterogeneous multi-tenant cloud infrastructures can offer greater control over the separate management of multiple types of resources. They would be able to provide differentiated services from the physical to the application level, leveraging both the horizontal (DC interconnection over multiple domains) and vertical (through the multiple layers of the physical infrastructure and its diverse abstractions) integration. This would allow huge savings and the acquisition of higher flexibility in service and application provisioning.

One technique that is a clear enabler for a system fulfilling the above three requirements is virtualization. Virtualization of physical resources, as seem in cloud computing environments, for computation, storage or networking, gives the ability to split a physical resource into many resources that behave like they were independent resources, or the capability to aggregate different physical resources to behave as though they were only one resource.

In response to the above systemic requirements and challenges, we propose a novel solution based on resource virtualization, the *Lightweight Slice Defined Cloud* (LSDC), an approach for automating the process of optimal cloud configuration by extending the virtualization concept to all resources in a data center and by providing a uniform management with a high-level of autonomicity for the currently separated computing, connectivity, and storage resources. The key differentiating novel characteristics of an LSDC are:

1) It presents a new service model, the *Slice as a Service*: The enhanced management for such a slice enabled infrastructure creates slices on-demand, and (re)configures them as appropriate to provide the end-to-end service. The slice management takes over the control of all the service components, virtualized network functions and system programmability functions assigned to the slice.

2) It enables the configuration of slices across the physical resources in the cloud networking infrastructure, to better accommodate the various service demands. Adaptions and reconfigurations are done at a per slice level, rather than for a whole cloud. This configuration

is achieved through using specially designed software that can describe and manage the various aspects that comprise slices within the cloud environment.

3) It allows each aspect that comprises the cloud environment – from the networking between virtual machines to the SLAs of the hosted applications – to be managed via software. This reduces the complexity related to configuring and operating the infrastructure, which in turn eases the management of the cloud infrastructure. Such infrastructure tends to be very large scale, commonly composed of thousands of servers and network elements, supporting tens of thousands of virtual machines, virtual networks, and applications.

4) It utilizes lightweight and uniform management and virtualization systems, with small footprint components, deployable on a large number of small servers and cloud systems both at the core and the edges of the network. These lightweight elements enable the integration of core data centers and mobile edge into cloud networks.

The NECOS proposal aims to realize the above mentioned systemic requirements and the four characteristics through the development of an advanced cloud infrastructure – the NECOS platform. The goal of NECOS is to show the feasibility of the LSDC concept as a means to fulfill the three systemic requirements mentioned above. To achieve this goal, we propose making a reference implementation as a proof of concept of *Slice as a Service*, by means of two use cases, that will leverage existing cloud platforms and tools employing the design, implementation, and integration of the new software components that, all together, will constitute the NECOS solution to the problems faced by both use cases. The resulting solution will be general enough to be applicable to a multitude of scenarios that utilize cloud technologies as *Slice as a Service*.

To the best of our knowledge, the LSDC concept described above is not yet developed or feasible with current technology. Although virtualization is now common, current proposals cloud management tools overlook the service and management aspects of slicing when considering a non-federated or a federated cloud environment.

The remainder of this paper is structured as follows: Section II presents NECOS, exposing its architecture and main features. Section III discusses some research challenges. Section IV conducts a study of related works to underline the value of our contribution. Finally, Section V examines the outcomes.

II. NECOS PROPOSAL

The potential of Cloud Slicing through the *Lightweight Slice Defined Cloud* (LSDC) is massive, as it focuses on creating a *Slice as a Service* mechanism which spans multiple cloud infrastructures, from the data center to the edge. One of the main tasks of the NECOS proposal is to formalize the requirements and to provide the integrated management and orchestration mechanisms that need to be used to achieve this goal. As service providers offer more applications on the cloud and as consumer's demand increases at a non-linear rate,

eventually even the largest cloud infrastructure will run out of computing, networking, and storage resources to supply. While clouds can be over-provisioned to handle spikes in demand, this is hugely expensive and not scalable.

Within NECOS, we believe the deployment of a Cloud Slicing mechanism is the best approach to address this. We will verify this under the guidance and demonstration of two use-cases, namely: the Telco Cloud use-case and the Mobile Edge Computing use-case.

The following subsections introduce the LSDC approach, indicating its role in the NECOS proposal, and highlighting the NECOS platform as an enabler for Cloud Slicing through the two mentioned use cases.

A. Lightweight Slice Defined Cloud (LSDC) – A Cloud Slicing platform

Integration of network and cloud computing domains and their related management operations will allow for massive savings and ability for greater flexibility in service provisioning. These operations are mainly carried out separately, so combining them is a goal of NECOS. Maturing paradigms such as SDN and network virtualization, for instance with NFV, when properly designed and deployed [5], can help in fulfilling the requirement of making such an innovation for services more reliable, faster, and simpler.

The NECOS concept on Cloud Slicing builds upon already available cloud platform features and functions. Slicing is associated with the partition of resources, being able to create and redefine these partitions as needed. A slice is a grouping of physical or virtual (network, compute and storage) resources which can act as a seemingly independent sub-cloud, sub-network and can accommodate service components. Key characteristics of Cloud Slicing include:

- The concurrent deployment of multiple logical, self-contained and independent, shared or partitioned slices on a common infrastructure platform.
- Dynamic multi-service support, multi-tenancy, and the integration mean for vertical market players.
- The separation of functions, simplifying the provisioning of services, the manageability of networks, and integration and operational challenges especially for supporting communication services.
- Network operators, ISP and Cloud infrastructure owners can exploit slicing for: reducing operations expenditure, allowing programmability and innovation necessary to enrich the offered services, for offering tailored services, and allowing network programmability to OTT providers and other market players without changing the physical infrastructure.

The LSDC represents a novel approach for automating the process of optimal cloud configuration by creating the Cloud Slice concept across all resources in a set of federated data centers, as well as providing uniform management of currently separated computing, connectivity, and storage resources.

B. NECOS Architecture

The NECOS architecture is adopting, as a starting, point elements from the SONATA 5G PPP[1] architecture, the 5GEx 5G PPP[2] architecture, and the architecture from RESERVOIR[3] project, with the aim of building a unified environment integrating connectivity, computation, and storage in order to create the *Slice as a Service* model. This requires an enhancement of the current control and management planes of state of the art platforms and the introduction of orchestration capabilities for slices.

NECOS aims to define the overall system architecture for LSDC to (*i*) support the functionality of the selected use cases; (*ii*) to enable cloud federation and distributed application deployment through slices; and (*iii*) to efficiently integrate the orchestration, cloud management, and monitoring components for slices. The fundamental characteristics of this enhanced LSDC cloud architecture are shown in Figure 1.

NECOS aims to realize an integrated platform encompassing cloud and network management, service orchestration, and distributed resources monitoring to support Cloud Slicing. The platform exposes interfaces for resource allocation and service deployment. This platform is the instantiation of the LSDC concept and will include, as shown in Figure 1, the Cloud Manager, the Network Manager, the Control Element for VMs and the Service Orchestration within a single deployable and distributed environment. Executing the platform components in integrated demonstrators will exemplify how the NECOS solution can be put in practice for demonstrating inter-connection and inter-operation of multiple clouds into a federated cloud environment. It will also describe how *Slice as a Service* brings about an improved cloud environment. The NECOS platform components and its environment are planned as open source releases.

C. Uniform Management Framework for LSDC

A unified management framework for both networking and computational domains is needed to support LSDC. This includes the main intelligence for federated cloud infrastructures: the management and the orchestration components. When accepting requests for cloud slices, the NECOS LSDC must determine the most efficient way of creating the federated virtual network, configure it, and then be able to analyze its correct functioning. Moreover, the LSDC must be able to dynamically decide which end-user devices are pertinent to the task, contact them via the NECOS API to perform the required tasks on those devices, and then report their performance back to the platform or the user requesting those services. This management framework will decide how to perform and execute efficiently these tasks across federated domains and the components responsible for ensuring that everything works according to what was planned without conflicts. New algorithms for the orchestration of slice resources, as well as their implementation for orchestrating innovative

---

[1]http://www.sonata.eu
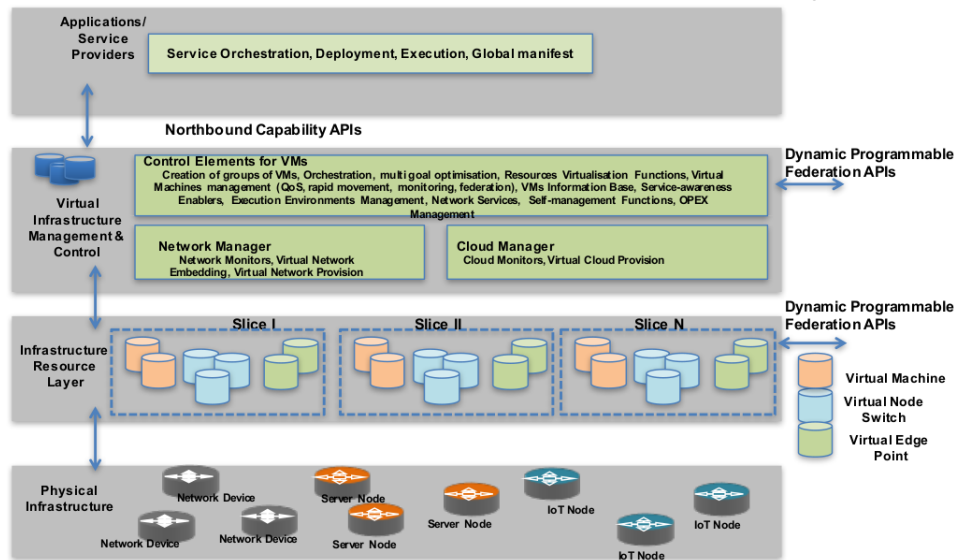[2]http://www.5gex.eu
[3]http://www.reservoir-fp7.eu

Figure 1. LSDC Architecture

overlay cloud virtualization frameworks, are needed. The main novel functional elements involved in the Management and Orchestration operations are: the *Intelligent Orchestrator* – responsible for deciding where and how to place virtual elements over the federated slice; the *Cloud Manager* – responsible for executing the decision of the orchestrator efficiently; and the *Monitoring Manager* – responsible for checking that the decisions and the commitments made by all the individual components are met, thus allowing to take corrective actions if they are not.

D. Use Cases for LSDC

The overarching challenge for the NECOS initiative is to research and develop the technologies to enable slice enabled Cloud Slicing infrastructure providers to provide value for themselves and to their customers – typically, application suppliers. To this end, NECOS picked two use cases from its industrial and academic partners to derive requirements for the design of the architecture and also to test the developed systems and to demonstrate the validity of the NECOS solution. These two use-cases are described below.

1) The Telco Cloud – Use Case 1: The traditional Telco business is evolving towards offering a richer set of services beyond basic connectivity services. Future telco networks are expected to support the needs of a hyper-connected society, which is continuously demanding very high data rate access, independence from the technology of attachment to the network, and an increasing number of almost permanently connected devices. Conventional ways of engineering services, based on monolithic devices statically located in the network is not valid anymore. Evolution in time, location, and requirements of the workloads generated by the end-users advocates a flexible infrastructure able to allocate resources that can be instantiated and removed, scaled-up and down, and being

made closer to the user, according to real needs of the overall services, in real time.

A versatile execution environment is required, capable of running different workloads. This is only economically sustainable when using commodity hardware, and defining simple and automated lifecycle processes associated with both the infrastructure and the workloads. Cloud computing and virtualization are the key technologies to allow for this change. However, traditional cloud computing could not be enough to address this challenge due to some essential differences concerning a telco cloud (e.g., performance bound to CPU *vs.* performance bound to I/O; node-centric *vs.* network-centric; many and small VMs *vs.* few and large VMs; etc.). A key concept of a telco cloud environment will be represented by the evolution of the existing Central Offices for integrating both virtualized and non-virtualized equipment.

As shown in Figure 2, NECOS will address the provision of a distributed but integrated Telco cloud architecture for the hosting and deployment of virtualized components enabling those new services in an automated fashion, meeting telco-grade requirements for security, availability, reliability, and performance. This integrated Telco cloud leverages a common control and management framework for a variety (in terms of size, infrastructure and management capabilities) of Data Centers dispersed across the network, in coordination with the control of the network for optimal usage of resources.

2) Mobile Edge Computing – Use Case 2: Wearable devices can often create situations where hundreds of devices are present in a few square meters getting data to be processed and transformed in appropriate stimulus to their owners to manage their digital lives seamlessly.

As an example, we may foresee a Metropolitan Tourist Center responsible for offering tourist information facilities to a significant number of visitors each year (high-profile
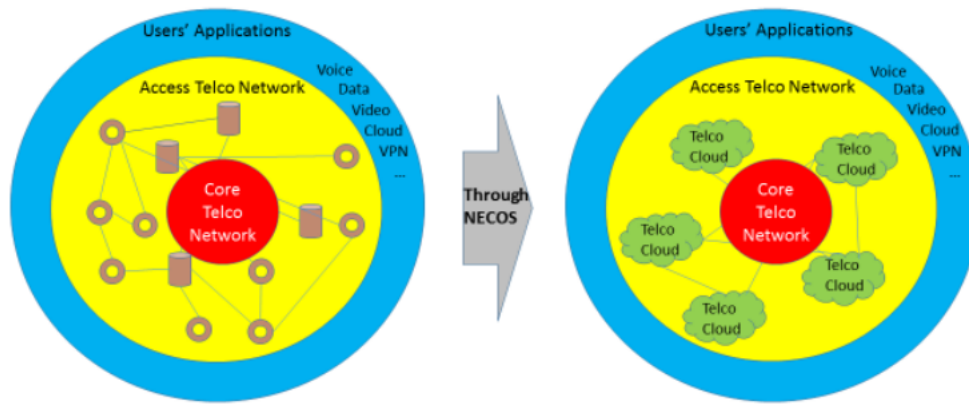
Figure 2. The evolution of the Telco's network fostered by the NECOS concept of Telco-Cloud

Metropolitan area). The *Metropolitan Tourist Center* (MTC) aims at enriching the visiting experience of its visitors, by offering state-of-the-art location-aware cultural content available in the city. These services will be offered: (*i*) within public transport vehicles (buses, subway, etc.), as they move throughout the city; and (*ii*) to visitors at various city locations (e.g., squares, museums) possibly taking advantage of public Wi-Fi infrastructure.

MTC will utilize (mobile) edge clouds deployed near crowded areas, museums, and other sights, as well as larger cloud infrastructure, potentially operated by a telco which is geographically present in the specific city. The telco's cloud infrastructure will maintain all content required by the services offered to the visitors and will further provide data processing (e.g., video transcoding) for personalized application delivery. Edge clouds will be mainly used for caching to reduce the latency while accessing data by the visitors.

Additionally, we may envision a *Command, Control, and Communications* (C3) *Public Safety Center* (PSC) scenario, responsible for enhancing incident management and resolution to a significant number of first responders (e.g., police department, fire rescue services, emergency medical services, and public works), as well as their interaction with civilian population in a metropolitan area. The C3 PSC aims at leveraging the information flow level between citizens, responders, and agencies, by quickly offering the ability to receive, correlate, and share information.

In such environment, there shall be services to be offered (*i*) within public transport vehicles (buses, subway, trains, etc.), as they move throughout the city, (*ii*) to citizens and visitors at various public buildings (e.g., city hall, hospitals, public schools, and museums) possibly taking advantage of public Wi-Fi infrastructure, or (*iii*) by storage, computing and infrastructure resources deployed by C3 PSC agencies (for instance, public agency mobile and/or stationary data centers, C3 applications, and tactical radio terminals).

Similarly, C3 PSC use case may utilize either (mobile) edge clouds deployed near crowded areas, museums, and other public buildings, in an opportunistic manner or based on agency-owned infrastructure. Furthermore, it may use larger cloud telco-operated infrastructure, maintaining all content (e.g., weather channel, traffic updates, and emergency alerts) required by the services offered to the civilian population, as well as data processing (e.g., intelligence reports, suspects identification) for on-demand secure application delivery to authorized personnel. On the other hand, edge clouds caching will reduce latency to access data for stakeholders (civilians and first responders) when deployed on the field.

In scenarios like this, where we can have high-density distributions of devices in wide geographical areas, a scalable approach consists of extending the capabilities of wearable devices using the massive resources offered by cloud data centers. However, instead of the classical cloud servers situated on the Internet, the edge cloud can be used as computation and storage resources. Although the capacity of the edge cloud may be lower than the classical cloud, its proximity to end-users and dense geographical distribution can lead to reduced latency, improved QoS and better mobility support for the above-applications and services. Figure 3 presents the functional elements of NECOS to be developed to support this scenario.

NECOS will have to provide an ecosystem of lightweight edge clouds to be in support of the mobile user applications. These users will interact with these clouds though the User-to-Cloud API of the NECOS platform. This API will present a homogeneous cloud environment constituted by the corresponding virtual machines and their controlling functions. The mapping of the VMs within their slice, to the physical substrate network, will be managed by the NECOS platform. This will adopt decisions based on algorithms to control the flows, the topology and the balancing of both the load on the servers and the network. The NECOS capabilities are extended by means of an orchestrator to take the highest level of decisions. Through this scenario, NECOS will be facing the following key challenges to pave the way of LSDC, namely QoS and optimal resource allocation, scarcity of resources, lack of good connectivity and users' mobility.

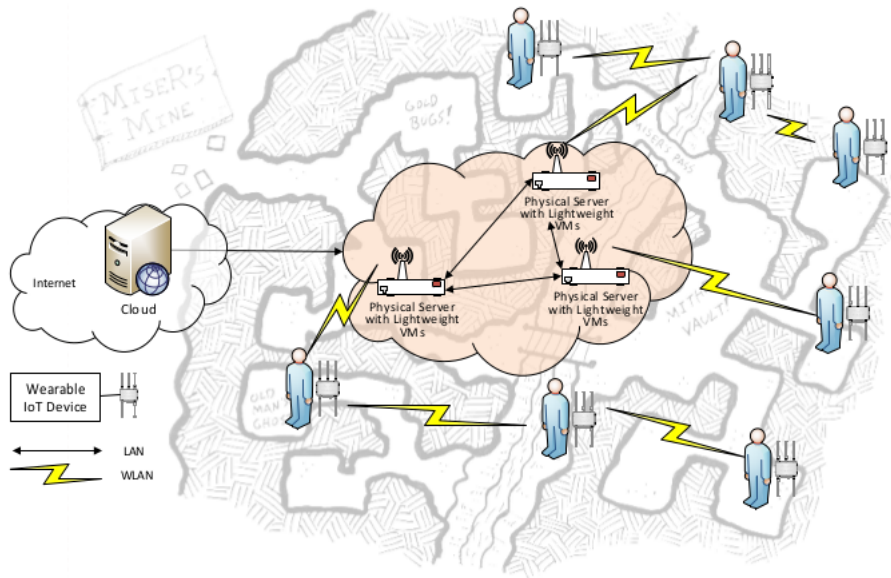As a by-product, this configuration results in multiple levels

Figure 3. Application point of view of Use Case 2

of data storage (local, edge cloud, remote and federated cloud). The criteria to decide where to store particular data would be a combination of several factors, among which we can mention the minimization of energy consumption, the network load balancing, as well as the cloud load balancing, and the fulfillment of the relevant regulations of personal data protection. These decisions need to be made transparently to the user, who would use her/his data as if they were in local storage. This scenario leverages network and computing/storage resources for a long-term and global data accessibility as well as to perform processing and analysis of these data.

III. RESEARCH CHALLENGES

While NECOS *Slice as a Service* approach brings many benefits to future clouds and networks, there are also some challenges that need to be addressed to deploy it properly.

A. Virtualization of Physical Infrastructures

Physical infrastructures that can be virtualized possess the flexibility that is demanded, however, as an example, once these physical infrastructures become virtualized applications and services have to be adapted to the new virtualized scenario, and the service model has to be adopted to make effective use of virtual resources.

B. Smart Computing Elasticity

Another challenge is the elasticity [6] that the assigned virtual resources to a given service have to offer. In its lifetime span a given service may need more or fewer resources, and hence these resources should be provided. Consequently, the virtualized environment has to be able to provide more or less resource as needed by the service. In this way, there is a need to provide autonomous management of the available resources for all slices, giving optimization in the use of the virtual resources.

Given this, the NECOS proposed approach will require advanced mechanisms capable of providing intelligent cloud management allowing for on-demand setup and tear down of network resources. In this context, the available resources for distribution among all slices should operate employing elastic properties insofar as resources are granted or withdrawn from each slice according to demand.

To overcome the problem of virtual resources elasticity are necessary advanced features able to support the slice orchestrator in the management operations appropriately. The intelligent computational elasticity features should be based on a resource monitoring system as well as provide an API so that the orchestrator can make use of all its functionalities efficiently.

C. Heterogeneity

Another challenge is the complexity that can arise as a consequence of diversification. The number of independent entities may become several orders of magnitude higher by using virtualization rather than physical elements. The question is how to face this growing complexity. Finally, the overhead associated with a system based on resource virtualization should be not too high. Otherwise, it could not be deployed on relatively resource-scarce environments.

D. Security in Slicing Management

It is of utmost importance that the NECOS platform provides a security context that can meet the different availability, security, and reliability requirements of hosted services over the lifetime of each cloud slice.

Also, it must implement secure end-to-end isolation of slices in the sense that incidents in a slice (e.g., network failures, overhead, cyber attacks, etc.) do not affect the operation of the others. However, this is a complex task because the sharing

of physical resources is not naturally well-designed to support the multi-tenant model required for cloud computing [7].

Another significant concern about cybersecurity is that each cloud slice should be protected based on security requirements that may vary according to the type of service hosted on the slice. Therefore, each slice must have independent security functions. Last but not least, access to orchestration mechanisms and management interfaces must be strictly controlled and monitored, and orchestration components must be protected [8].

E. Mobility Management

The vertiginous increase in the number of mobile users has required the development of new technologies able to provide guarantees of communication efficiently [9]. However, the emergence of these new technologies will potentially increase the costs of network providers due to the increasing complexity of infrastructure management and operation. In this context, the *Wireless Network Virtualization* (WNV) [10] paradigm, which enables the sharing of common network infrastructure, including radio resources between different virtual networks, has emerged as an alternative to solve the main associated efficiency problems. However, due to the particularities of wireless communications, ensuring the isolation between slices has become a very complicated task, especially concerning QoS and SLA. In this case, advanced techniques are required to cope with performance deterioration problems (e.g., wireless channel conditions) from other slices or the intrinsic complexity of wireless communications such as user interference and need for mobility (inappropriately depleting the resources on the destination network).

Given such issues, technologies such as SDN (crucial to achieving the flexibility and programmability needed to manage a sliced-network) and NFV (capable of providing slice management more smoothly if some functions can be decoupled from proprietary hardware, virtualized and run centrally) appear to be capable of offering alternative solutions. When it comes to mobile QoS slicing, it is necessary to meet specific requirements, such as the delivery of explicit QoS guarantees within slices (e.g., slice for services that serve a particular group of devices with the same requirements or slices for certain types of applications, etc.).

IV. RELATED WORK

So far there have been many successful efforts in designing either clean-slate or evolutionary clouds. The major categories of NECOS evolution are the following: (*i*) Cloud Management Platforms; (*ii*) Mobile Cloud, Mobile Edge and Fog Computing; (*iii*) Slice as a Service; (*iv*) Orchestrating Slices Across Federated Domains; and (*v*) Federation of Cloud Services and Resources. The next subsections provide an overview of each of this categories.

A. Cloud Management Platforms

The component that manages both physical and virtual resource in a cloud infrastructure is called the *Cloud Management Platform* (CMP). Management of cloud and network resources is indeed one of the main functionalities of the NECOS platform. It executes different networking operations necessary to realize the virtual networks calculated by the Orchestrator (e.g., bandwidth reservation, subnetworks configuration, the definition of routes in the network elements) and realizes the VM management operations on the physical infrastructure (creation, destruction, and migration of virtual machines, as well as their reconfiguration).

There are several actors in the cloud management field, and there are several CMP implementations that we can identify as forming part of two sets: proprietary CMPs and open source CMPs. In the first group, we can include solutions like VMware vCloud[4], IBM smart cloud[5], HP Cloud[6], OpenStack[7]) proprietary distributions like for instance Piston[8] software, etc. These solutions are mature and stable, but they lack the openness needed to perform modifications and thus extend their functionality. In the open source CMPs we find four major players worth highlighting, each of them with their particularities: Eucalyptus[9] is designed to build Amazon Web Services-compatible private, and hybrid cloud computing environments, OpenStack follows a similar approach paying more attention to the underlying physical resources, while CloudStack[10] and OpenNebula[11] started with a bottom-up approach, focusing first on data center virtualization and building cloud interfaces afterwards. Anyway, current state-of-the-art CMPs offer limited support for efficient resource partitioning, isolation, and elastic management. Therefore NECOS will perform a deep state-of-the-art analysis to select the reference CMP for the project, focusing on the requirement of a flexible architecture that eases the integration with necessary third-party components, extensions, and improvements.

B. Mobile Cloud, Mobile Edge, and Fog Computing

The concepts of Mobile Cloud, Mobile Edge, and Fog Computing are embedded in the LSDC approach of NECOS, i.e., bringing efficiently together with the lightweight clouds with the mobile resources, to enable novel network edge services. This concept is the center of Use case 2, as currently foreseen in NECOS.

[11] deals with the problem of enhancing the capacity of mobile devices by offloading resource-intensive applications to components external to the device itself. Such external devices can be cloud servers, other mobile devices, or cloudlets (powerful servers that are spread around the mobile's coverage area to provide extra capacity or to offload workload to the cloud). NECOS will exploit this approach of offloading tasks of mobile users to a cloud data center or a cloudlet (or micro-data center) in the edge network. The motivation for offloading CPU-intensive tasks is to enable mobile devices to perform

---

[4]https://www.vmware.com/products/vcloud-suite

[5]com/cloud-computing/eg/en/private-cloud.html

[6]http://www8.hp.com/us/en/business-solutions/solution.html

[7]https://www.openstack.org

[8]http://www.pistoncloud.com/openstack-cloud-software/

[9]https://eucalyptus.com

[10]http://cloudstack.apache.org/

[11]http://opennebula.org/

activities that otherwise would not be possible because of the limitations in resources on the device (e.g., energy, CPU, network bandwidth). Honeybee [12] is a framework to enable mobile devices to offload tasks, utilize resources from devices, and perform human-aided computations. [13] proposed an architecture to offload computation to nearby devices using P2P techniques. [14] proposed an approach for mobile cloud by exploring a middleware between the devices and the cloud. [15] proposed an approach to offload parts of a computation to the cloud to reduce energy consumption on the device.

C. Slice as a Service

Nowadays, all slicing proposals consider each slice as some kind of virtual network, including computation and storage services. In the past recent years, considerable attention has been paid to alternatives for the virtualization of networks. Among these works, there are some achievements in optimal VM placement inside data centers having in consideration traffic features. There is one implementation and deployment of these ideas in the GENI [16] network automatically co-schedules and provisions heterogeneous networked resources. In the context of NECOS, the scenario is more complex because, in addition to the networking infrastructures, we have to consider both data center compute infrastructure and storage infrastructure as elements to be included on the slice. The most critical requirements for virtual infrastructures can be summarized as: (*i*) the possibility of co-existence of different slices on the same substrate infrastructure, (*ii*) the necessity to manage and program the slices individually, and finally (*iii*) the isolation of the slices to avoid conflicts between co-existing ones.

D. Orchestrating Slices Across Federated Domains

The Intelligent Orchestrator is responsible for deciding where and how to place virtual elements over the federated cloud network. It constitutes one of the main building blocks of the NECOS platform.

[17] presents a work on orchestrated arrangements of heterogeneous resources (compute, storage, networks, content, scientific instruments etc.), through a single interface consisting of a meta control architecture designed independently of any a priori substrate assumptions and capable of driving multiple heterogeneous substrates using their native control/management interfaces and creating orchestrated arrangements of these resources acquired from multiple independent providers. [18] presents a scenario for realizing sophisticated cloud services requirements, for which a cloud control framework that can orchestrate cloud resource provisioning, configuration, utilization and decommissioning across a distributed set of physical resources must be in place. Following this approach, cloud resources are modeled as structured data that can be queried by a declarative language, and updated with well-defined transactional semantics. [19] describes one of the outcomes of the FP7 OFELIA project on the joint orchestration of IT (computing, storage) and networking resources in the context of an SDN based network. The cloud orchestration requires the interworking of an SDN

controller (e.g. OpenDaylight [20]) with a Cloud Computing Controller (e.g. OpenStack. OpenStack offers a networking module with a plugin mechanism to connect different network controllers. [21] presents a proof of concept of an idea coined in the FP7 UNIFY project [22], which aims at unifying cloud and carrier network resources in a common orchestration framework. This consists of a *Service Function Chain* (SFC) control plane which can integrate any infrastructure domains including different network function execution environments, SDN networks or legacy data centers. This solution supports automated, dynamic service creation and multi-domain NFV orchestration. Central to this design is the definition of a joint virtualization and resource programming interface unifying cloud and network resources. RESERVOIR [23] has taken service orchestration a step further by allowing for automatic allocation and deployment of resources to service as demand fluctuates.

E. Federation of cloud services and resources

Outsourcing the demand to the federated cloud has been proposed as a mean to overcome the local demand spikes without having to rely on over-dimensioning or just rejecting service requests at peak demand times. Nevertheless, the decision as to which and how much data should be outsourced is not trivial. In fact, it directly impacts the revenues of all service providers of the federated ecosystem. The existence of a platform adopting autonomous decisions in that respect widens the market of cloud providers and stimulates new entrants to enter in the game. The federation of cloud services and resource is inherent to the concept of NECOS because it considers and architecture based on a Cloud-to-Cloud API in support of federations.

[24] considers a general cloud architecture to highlight the limitations of cloud federation and proposes some enhancements which add new federation capabilities. To address such concerns, he proposes a solution based on the Cross-Cloud Federation Manager, a new component to be placed within the Cloud Manager layer of the cloud architectural stack, allowing a cloud to establish the federation with other clouds according to a three-phase model: discovery, match-making, and authentication. [25] shows how a layered Cloud service model of software (SaaS), platform (PaaS), and infrastructure (IaaS) leverages multiple independent Clouds by creating a federation among the providers. [26] presents a characterization of provider's federation in the Cloud, including decision equations to outsource resources to other providers, rent-free resources to other providers or shut down unused nodes to save power, and he characterizes these decisions as a function of several parameters. [27] discusses a reference architecture that enables new service models by horizontal and vertical integration.

NECOS platform will enable inter-cloud federation with the following embedded capabilities: (*i*) a Physical Infrastructure Provider to publish compute, storage, and connectivity resources without exposing the actual topology of the physical substrate; (*ii*) the publication and composition of offered and

required SLA constraints; (*iii*) the abstraction of the device and cloudlet resources in mobile edge networks, bringing the cloud beyond the network edge; (*iv*) *Monitoring as a Service* – the allocation of resources and services must be always coupled with trustable monitoring instances reporting the actual compliance with the required service; and (*v*) *Slice as a Service* – the orchestration of resource to enable end-to-end network services across different segments of the federated platform in NECOS.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we present the NECOS proposal towards addressing the limitations of current cloud infrastructures by proposing the *Lightweight Slice Defined Cloud* (LSDC) approach. LSDC acts as an enabler of Cloud Slicing through the *Slice as a Service* concept as a unified management framework for both networking and computational domains, providing intelligent orchestration for federated cloud infrastructures. In future work, we will present results from the platform implementation work and the experimental validation of use cases to demonstrate the feasibility and potential benefits of the proposed concepts.

REFERENCES

[1] M. Boucadairm and C. Jacquenet, "Software-Defined Networking: A Perspective from within a Service Provider Environment," Internet Requests for Comments (RFC), Internet Engineering Task Force Request for Comments (RFC 7149, Mar 2014. [Online]. Available: http://www.rfc-editor.org/rfc/rfc7149.txt

[2] M. Chiosi, D. Clarke, P. Willis, A. Reid, J. Feger, M. Bugenhagen, W. Khan, M. Fargano, C. Cui, H. Deng, J. Benitez, U. Michel, H. Damker, K. Ogaki, T. Matsuzaki, M. Fukui, K. Shimano, D. Delisle, Q. Loudier, C. Kolias, I. Guardini, E. Demaria, R. Minerva, A. Manzalini, D. Lopez, F. Salguero, F. Ruhl, and P. Sen, "Network functions virtualisation," in *white paper at the SDN and OpenFlow World Congress , ETSI, Tech. Rep.*, 2012. [Online]. Available: https://portal.etsi.org/nfv/nfv_white_paper.pdf

[3] M. Patel, Y. Hu, P. Hede, J. Joubert, C. Thornton, B. Naughton, J. R. Ramos, C. Chan, V. Young, S. J. Tan, D. Lynch, N. Sprecher, T. Musiol, C. Manzanares, U. Rauschenbach, S. Abeta, L. Chen, K. Shimizu, A. Neal, P. Cosimini, A. Pollard, and G. Klas, "Mobile-edge computing - introductory technical white paper," in *ETSI, Tech. Rep.*, 2014. [Online]. Available: https://portal.etsi.org/portals/0/tbpages/mec/docs/mobile-edge_computing_-_introductory_technical_white_paper_v1%2018-09-14.pdf

[4] H. Zhang, N. Liu, X. Chu, K. Long, A. H. Aghvami, and V. C. M. Leung, "Network slicing based 5g and future mobile networks: Mobility, resource management, and challenges," *IEEE Communications Magazine*, vol. 55, no. 8, pp. 138–145, 2017.

[5] A. Manzalini, C. Buyukkoc, P. CHEMOUIL, S. Kuklinski, F. Callegati, A. Galis, M.-P. Odini, C.-L. I, J. Huang, M. Bursell, N. Crespi, E. Healy, and S. Sharrock, "Towards 5G Software-Defined Ecosystems: Technical Challenges, Business Sustainability and Policy Issues," Research Report, Jul. 2016. [Online]. Available: https://hal.archives-ouvertes.fr/hal-01419600

[6] R. da Rosa Righi, V. F. Rodrigues, G. Rostirolla, C. A. da Costa, E. Roloff, and P. O. A. Navaux, "A lightweight plug-and-play elasticity service for self-organizing resource provisioning on parallel applications," *Future Generation Computer Systems*, vol. 78, pp. 176 – 190, 2018. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167739X17302339

[7] M. Pattaranantakul, R. He, A. Meddahi, and Z. Zhang, "Secmano: Towards network functions virtualization (NFV) based security management and orchestration," in *2016 IEEE Trustcom/BigDataSE/ISPA, Tianjin, China, August 23-26, 2016*, 2016, pp. 598–605. [Online]. Available: https://doi.org/10.1109/TrustCom.2016.0115

[8] NGMN, "5g security recommendations package 2: Network slicing," Frankfurt, Germany, Tech. Rep., 2016. [Online]. Available: https://www.ngmn.org/fileadmin/user_upload/160429_NGMN_5G_Security_Network_Slicing_v1_0.pdf

[9] F. S. D. Silva, A. V. Neto, D. Maciel, J. Castillo-Lema, F. Silva, P. Frosi, and E. Cerqueira, "An innovative software-defined winemo architecture for advanced qos-guaranteed mobile service transport," *Computer Networks*, vol. 107, pp. 270 – 291, 2016, mobile Wireless Networks. [Online]. Available: https://doi.org/10.1016/j.comnet.2016.04.019

[10] E. J. Kitindi, S. Fu, Y. Jia, A. Kabir, and Y. Wang, "Wireless network virtualization with sdn and c-ran for 5g networks: Requirements, opportunities, and challenges," *IEEE Access*, vol. 5, pp. 19 099–19 115, 2017.

[11] M. Satyanarayanan, "A brief history of cloud offload: A personal journey from odyssey through cyber foraging to cloudlets," *GetMobile*, vol. 18, no. 4, pp. 19–23, 2014. [Online]. Available: http://doi.acm.org/10.1145/2721914.2721921

[12] N. Fernando, S. W. Loke, and W. Rahayu, "Honeybee: A programming framework for mobile crowd computing," in *Mobile and Ubiquitous Systems: Computing, Networking, and Services*, K. Zheng, M. Li, and H. Jiang, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 224–236.

[13] G. Huerta-Canepa and D. Lee, "A virtual cloud computing provider for mobile devices," in *Proceedings of the 1st ACM Workshop on Mobile Cloud Computing & Services: Social Networks and Beyond*, ser. MCS '10. New York, NY, USA: ACM, 2010, pp. 6:1–6:5. [Online]. Available: http://doi.acm.org/10.1145/1810931.1810937

[14] H. Flores and S. N. Srirama, "Mobile cloud middleware," *Journal of Systems and Software*, vol. 92, pp. 82 – 94, 2014. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0164121213002318

[15] B.-G. Chun, S. Ihm, P. Maniatis, M. Naik, and A. Patti, "Clonecloud: Elastic execution between mobile device and cloud," in *Proceedings of the Sixth Conference on Computer Systems*, ser. EuroSys '11. New York, NY, USA: ACM, 2011, pp. 301–314. [Online]. Available: http://doi.acm.org/10.1145/1966445.1966473

[16] GENI, "GENI - Global Environment for Network Innovations," 2018. [Online]. Available: http://www.geni.net/

[17] I. Baldine, Y. Xin, A. Mandal, C. H. Renci, U. C. J. Chase, V. Marupadi, A. Yumerefendi, and D. Irwin, "Networked cloud orchestration: A geni perspective," in *2010 IEEE Globecom Workshops*, Dec 2010, pp. 573–578.

[18] A. Hindle, "Cloudorch: A portable soundcard in the cloud," in *NIME*, 2014.

[19] "Design and implementation of the ofelia fp7 facility: The european openflow testbed," *Computer Networks*, vol. 61, pp. 132 – 150, 2014, special issue on Future Internet Testbeds – Part I.

[20] OpenDayLight, "OpenDayLight Project," 2018. [Online]. Available: https://www.opendaylight.org/

[21] B. Sonkoly, R. Szabo, D. Jocha, J. Czentye, M. Kind, and F. J. Westphal, "Unifying cloud and carrier network resources: An architectural view," in *2015 IEEE Global Communications Conference (GLOBECOM)*, Dec 2015, pp. 1–7.

[22] FP7-UNIFY, "UNIFY Project," 2018. [Online]. Available: https://www.fp7-unify.eu/

[23] RESERVOIR, "RESERVOIR FP7 Integrated Project," 2018. [Online]. Available: https://www.reservoir-fp7.eu/

[24] A. Celesti, F. Tusa, M. Villari, and A. Puliafito, "Three-phase cross-cloud federation model: The cloud sso authentication," in *2010 Second International Conference on Advances in Future Internet*, July 2010, pp. 94–101.

[25] D. Villegas, N. Bobroff, I. Rodero, J. Delgado, Y. Liu, A. Devarakonda, L. Fong, S. M. Sadjadi, and M. Parashar, "Cloud federation in a layered service model," *Journal of Computer and System Sciences*, vol. 78, no. 5, pp. 1330 – 1344, 2012, jCSS Special Issue: Cloud Computing 2011. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0022000011001620

[26] I. Goiri, J. Guitart, and J. Torres, "Characterizing cloud federation for enhancing providers' profit," in *2010 IEEE 3rd International Conference on Cloud Computing*, July 2010, pp. 123–130.

[27] T. Kurze, M. Klems, D. Bermbach, A. Lenk, S. Tai, and M. Kunze, "Cloud Federation," in *Proceedings of the 2nd International Conference on Cloud Computing, GRIDs, and Virtualization (CLOUD COMPUTING 2011)*. IARIA, Sep. 2011, Inproceedings. [Online]. Available: http://www.aifb.kit.edu/images/0/02/Cloud_Federation.pdf