

The NECOS Approach to End-to-End Cloud-Network Slicing as a Service

Stuart Clayman, Augusto Neto, Fábio Verdi, Sand Correa, Silvio Sampaio, Ilias Sakelariou, Lefteris Mamas, Rafael Pasquini, Kleber Cardoso, Francesco Tusa, Christian Rothenberg, and Joan Serrat

ABSTRACT

Cloud-network slicing is a promising approach to serve vertical industries delivering their services over multiple administrative and technological domains. However, there are numerous open challenges to provide end-to-end slices due to complex business and engineering requirements from service and resource providers. This article presents a reference architecture for the cloud-network slicing concept and the practical realization of the slice-as-a-service paradigm, which are key results from the Novel Enablers in Cloud Slicing (NECOS) project. The NECOS platform has been designed to consider modularity, separation of concerns, and multi-domain dynamic operation as prime attributes. The architecture comprises a set of interworking components to automatically create, manage, and decommission end-to-end cloud-network slice instances in a lightweight manner. NECOS orchestrates slices at runtime, spanning across core/edge data centers and wired/wireless network infrastructures. The novelties of the multi-domain NECOS platform are validated through three proof-of-concept experiments: (i) a touristic content delivery service slice deployment featuring on-demand virtual infrastructure management across three countries on different continents to meet particular slice requirements; (ii) intelligent slice elasticity driven by machine learning techniques; and (iii) marketplace-based resource discovery capabilities.

INTRODUCTION

Modern Internet services are characterized by large-scale deployments of distributed software nodes to serve myriads of service consumers. Increasingly, many of these services need to be adaptable to challenging and stringent application requirements, such as ultra-low delay. Such demanding services must coexist with “standard” applications, a situation that introduces conflicting requirements over shared infrastructures. The availability of distributed computing, storage, and networking resources, along with diverse service requirements and varying workloads, differ from traditional deployments on large-scale central data centers (DCs).

To embrace new environments of geographically spread cloud capabilities, multi-access edge

computing (MEC), telco clouds, and *network slicing* are emerging, as network operators and service providers seek to extend their service platform footprints. Network slicing provides the ability to divide a single physical network infrastructure into multiple coexisting and isolated logical networks, with *each slice* being tailored to cope with and address the networking requirements of a particular service. This logical grouping of networking resources can cover the complete end-to-end path, including both radio/network equipment and physical cloud server resources, and may cross several administrative boundaries.

An evolutionary view of the implementation of network slicing [1] harnesses the network design and operation principles of the combined utilization of network functions virtualization (NFV) with software-defined networking (SDN). Surveying the current state of affairs [2, 3] reveals that:

- There is still little evidence of works [4] that have attempted holistic end-to-end slicing approaches spanning across multiple administrative domains of network and cloud deployments.
- Those available are not sufficient to deliver the broadest and most ambitious visions around network slicing.

The Novel Enablers in Cloud Slicing (NECOS) project [5] advances the state of the art by offering slice-as-a-service capabilities through a novel modular, multi-domain, end-to-end cloud-network slicing architecture with dynamic operation and a separation of concerns in the design. A NECOS-provisioned cloud-network slice instance represents a group of managed elements: a set of infrastructure facilities (network, edge/cloud) plus networking functions from providers; a set of manageable resources (i.e., connectivity, compute, and storage); and the service applications that feature operations and attributes specially devised for the requirements of specific industry verticals. Each activated cloud-network slice instance holds a set of computing and networking resources across a number of domains, whereby tenants can deploy services and/or applications running in a virtualized form. To ensure the appropriate deployment, slices that span across a shared multi-domain infrastructure are orchestrated in a way that matches the dynamic end-to-end quality requirements of the slice, while still being

Cloud-network slicing is a promising approach to serve vertical industries delivering their services over multiple administrative and technological domains. However, there are numerous open challenges to provide end-to-end slices due to complex business and engineering requirements from service and resource providers. The authors present a reference architecture for the cloud-network slicing concept and the practical realization of the slice-as-a-service paradigm, which are key results from the Novel Enablers in Cloud Slicing (NECOS) project.

Stuart Clayman and Francesco Tusa are with University College London; Augusto Neto and Silvio Sampaio are with Federal University of Rio Grande do Norte; Fabio Verdi is with Federal University of São Carlos; Sand Correa and Kleber Cardoso are with Federal University of Goiás; Ilias Sakelariou and Lefteris Mamas are with the University of Macedonia; Rafael Pasquini is with Federal University of Uberlândia; Christian Rothenberg is with the University of Campinas; Joan Serrat is with Universitat Politècnica de Catalunya.

Modern Internet services are characterized by large-scale deployments of distributed software nodes to serve myriads of service consumers. Increasingly, many of these services need to be adaptable to challenging and stringent application requirements, such as ultra-low delay. Such demanding services must coexist with “standard” applications, a situation that introduces conflicting requirements over shared infrastructures.

independent of each other. Toward this direction, NECOS introduces the following novel features:

- F1. The **slice-as-a-service** model, which fully addresses the life cycle (creating, reconfiguring, and decommissioning) of cloud-network slice instances on demand and at runtime, allowing end-to-end multi-tenancy highly isolated service provisioning. The tenant specifies slice requirements through appropriate service descriptors and application programming interfaces (APIs) such that the NECOS platform creates a new end-to-end cloud-network slice instance and associated slice operation-specific management facilities.
- F2. The **lightweight slice defined cloud (LSDC)** architecture, for tackling the complexity of large-scale virtualized infrastructure environments, and leveraging intelligent management and orchestration functions. Instead of overloading heavyweight management and operation (MANO) systems with extra slicing capabilities, LSDC follows a modular, lightweight, microservices architectural approach.
- F3. **Multi-domain cloud-network slicing** enables cooperative business models and large-scale slice deployments over both network and cloud providers.
- F4. **Intelligent elasticity** capabilities support both horizontal and vertical slice elasticity, harnessing machine learning (ML) techniques.
- F5. An **infrastructure manager on-demand** model allocates dedicated management and control points, providing the ability to allocate a new virtual infrastructure manager (VIM) on demand for each DC slice [6], and a new wide-area infrastructure manager (WIM) on demand for each network slice part [7].
- F6. A **resource marketplace** performs dynamic discovery of cloud and network resources for slice parts across multiple geographic domains, via an online resource marketplace, rather than having predetermined providers set previously configured in a federation [8].

The main contributions of this article are three-fold:

- An overview of the capabilities required to support end-to-end cloud-network slicing over multiple cloud and network resource providers
- The final design of the novel NECOS LSDC architecture, offering those capabilities via the slice-as-a-service model
- An evaluation of the NECOS modular architecture and the above six NECOS features through the testbed experimentation of three different scenarios built upon proof-of-concept software prototypes

The work described in [9] introduces the initial proposal in the context of the project, while [10] covers an earlier feasibility study of NECOS. This article extends both previous publications by exercising the architecture and system platform (fully detailed in project deliverables [11, 12]) in three proof-of-concept experiments on the final implementation. Finally, it should be noted that the NECOS platform has been released as open source artifacts, featuring its novel capabilities through five different demonstrators [5].

RELATED WORK

In our previous work [3], we surveyed, from an architectural perspective, work from six standardization bodies — International Telecommunication Union — Telecommunication Standards Sector (ITU-T), Next Generation Mobile Network alliance (NGMN), Internet Engineering Task Force (IETF), Third Generation Partnership Project (3GPP), European Telecommunications Standards Institute (ETSI), and Open Network Forum (ONF) — related to 5G slicing, and from eight relevant project initiatives considered to be the most closely related to NECOS. We identified 14 key capabilities for slicing, including connectivity resource slicing, connectivity service slicing, network cloud slicing, end-to-end slicing, multi-domain slicing, uniform life cycle management, tenant slice management, slice as a service, VIM/WIM on demand, marketplace for slices, network elasticity, cloud elasticity, and full monitoring. Then the six standards and eight project initiatives were analyzed from the point of view of the fulfillment of these key capabilities, and, as a result, none of them fulfilled all of the capabilities. Furthermore, we observed that those standards presented divergences in their approach to where slicing was done and what their scope was. Our survey results suggest that those approaches do not fully encompass all of the needs that network slicing could bring about.

Literature surveys on slicing [1, 2] confirm that the concept of slicing is broadening from multiple perspectives, and that SDN and NFV are part of the solution [4]. Much of the related work on slicing is limited to the 5G domain, and only a few consider joint slicing of cloud and network domains or fully address heterogeneous service demands.

Recent efforts at the ETSI Zero-Touch network and Service Management Industry Specification Group (ZSM ISG) [13] have resulted in a ZSM framework to deliver fully automated network and service management in multi-domain environments for 5G and beyond based on a model-driven architecture. The proposed approach also has an end-to-end slice design, which can be recursively built via combining multiple slice subnets in different segments of the infrastructure. NECOS relies on a similar model, which, although not recursive, also considers the compute slice parts in the possible end-to-end slice building blocks. In the NECOS model, compute and network slices can be directly accessed, and securely attached to an end-to-end slice, by the provider that initiated the slice creation workflow. This can overcome the need for peer-to-peer orchestrator interactions after a slice has been put in place, increasing resource isolation and performance.

NECOS PLATFORM ACTORS AND BUSINESS CASE

The main actors and architectural tenets of NECOS are illustrated in Fig. 1. Details about each actor, their top-level interactions, and an exemplary business case are presented next.

THE TENANT

A NECOS tenant is an organization requiring cloud-network slices for their own services to run. A tenant will initiate the slice creation process by presenting a request specifying a slice descriptor

covering different scopes (e.g., commercial, geographical, contractual commitments).

NECOS SLICE PROVIDER

The NECOS slice provider is a management entity responsible for creating end-to-end cloud-network slice instances from a set of constituent network and DC-like slice parts, with each NECOS created slice instance being mapped to a set of federated resources. When a slice is requested, the NECOS slice provider initiates a search via the resource marketplace for slice parts. Upon the receipt of answers from the marketplace, in the form of alternative resource offers for each part, the slice provider selects and combines the slice parts into a single aggregated slice. The slice provider is in charge of orchestrating end-to-end cloud-network slices, as well as their life cycle management at runtime. Furthermore, it is also responsible for instantiating both virtual machines (VMs) and virtual links for the services running inside the resource domains.

RESOURCE MARKETPLACE

The resource marketplace implements a dynamic resource discovery approach, initiated by the NECOS slice provider. The design was influenced by existing online marketplaces for flight pricing and hotel reservations, which provided a working operational model to follow. Rather than configuring pre-determined provider sets in a federation approach, NECOS applies a more flexible runtime model, where an auctioneer component announces a *call for proposals* for resource offers on a slice part, aggregating alternative offers by bidders. The aggregator can be a subsystem of any resource provider participating in NECOS, or in a more complex setting, it can be an external aggregator that provides either compound offers or offers at a price point. The results of this auction process are collected and communicated back to the slice provider, who selects the actual slice resources.

The primary reason for these design choices was to provide a mechanism that reaches and interacts with resource providers in multiple geographic locations in order to provision end-to-end cloud-network slices. Existing federation approaches often lack coverage to these kinds of networks. The second reason is that, in general, slices are created in a more dynamic fashion than federation agreements; therefore, a highly-dynamic, market-inspired provider searching mechanism is required. In contrast to delegating resource management to the respective domains, our approach trades a potential performance increase at the domain-level resource adaptations for greater flexibility, since it permits changes in the slice graph, including allocating services to new providers, service reallocation to other slice parts, and so on. A detailed description of the marketplace can be found in [8].

RESOURCE PROVIDERS

Resource providers are those organizations in charge of providing the set of required resources to provision slice parts. Other types of resources can be manifested by organizations featuring MEC, networked sensors, or wireless networks. The allocation of a new VIM (for DC slice parts), or a new WIM (for network slice parts) is also the

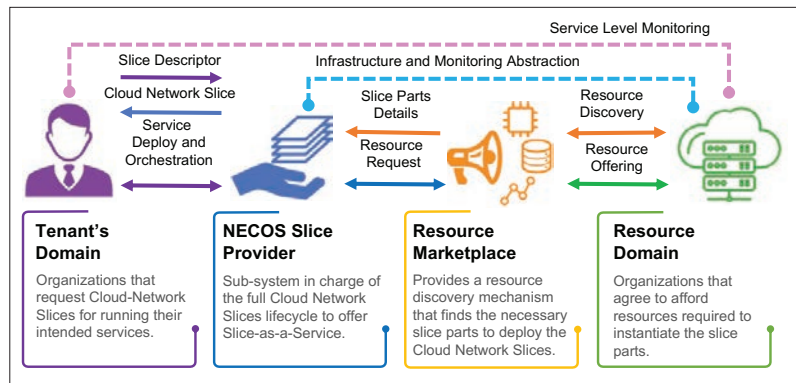


FIGURE 1. The four main actors participating in the NECOS architecture.

responsibility of the resource provider, allowing the slice resource orchestrator to remotely manage and utilize the resources.

BUSINESS CASE

NECOS relies on the slice-as-a-service model to set up network-cloud slices for a wide range of business-oriented applications [14]. Here, we briefly present a business case built around the NECOS platform, by aligning with the business model canvas approach for creating slices that utilize a logistics Internet of Things (IoT) scenario.

Value Proposition: The business case relates to the provisioning of network-cloud slices accommodating services for international logistic companies, allowing tracking of assets during worldwide delivery, with the customer being interested in the geographical location, temperature, humidity, and so on.

Customer Segments: The end customer, acting as a NECOS tenant, can directly be any international logistics, delivery, or transportation company, or can be a provider of Internet-based services to the logistics companies.

Key Resources: The business case requires NECOS platforms deployed in a number of geographical locations, owned by unique or different organizations. These can be network operators, cloud providers, or new business entities, acting as NECOS slice, resource marketplace, cloud, network, or IoT infrastructure providers.

Key Activities: The key activities concern the creation, activation, efficient runtime operation, and teardown of cloud-network slice instances and service elements, tailored to the evolution of the customers' needs.

Key Partnerships: Slice providers receiving slice creation requests may transfer part of the requests to one or more partners in the case of resource shortage. A novelty of the NECOS approach enables partnerships of existing or new types of business entities, such as sensor network providers, participating in NECOS roles.

Revenue and Cost Streams: Both revenue from and cost of a slice can be based on cost models being proportional to the amount of DC and network resources needed, the utilization of the sensor network, and the slice duration.

NECOS FUNCTIONAL BLOCKS

We now describe the NECOS architecture functional blocks, which are presented in Fig. 2 (with more details in [11]).

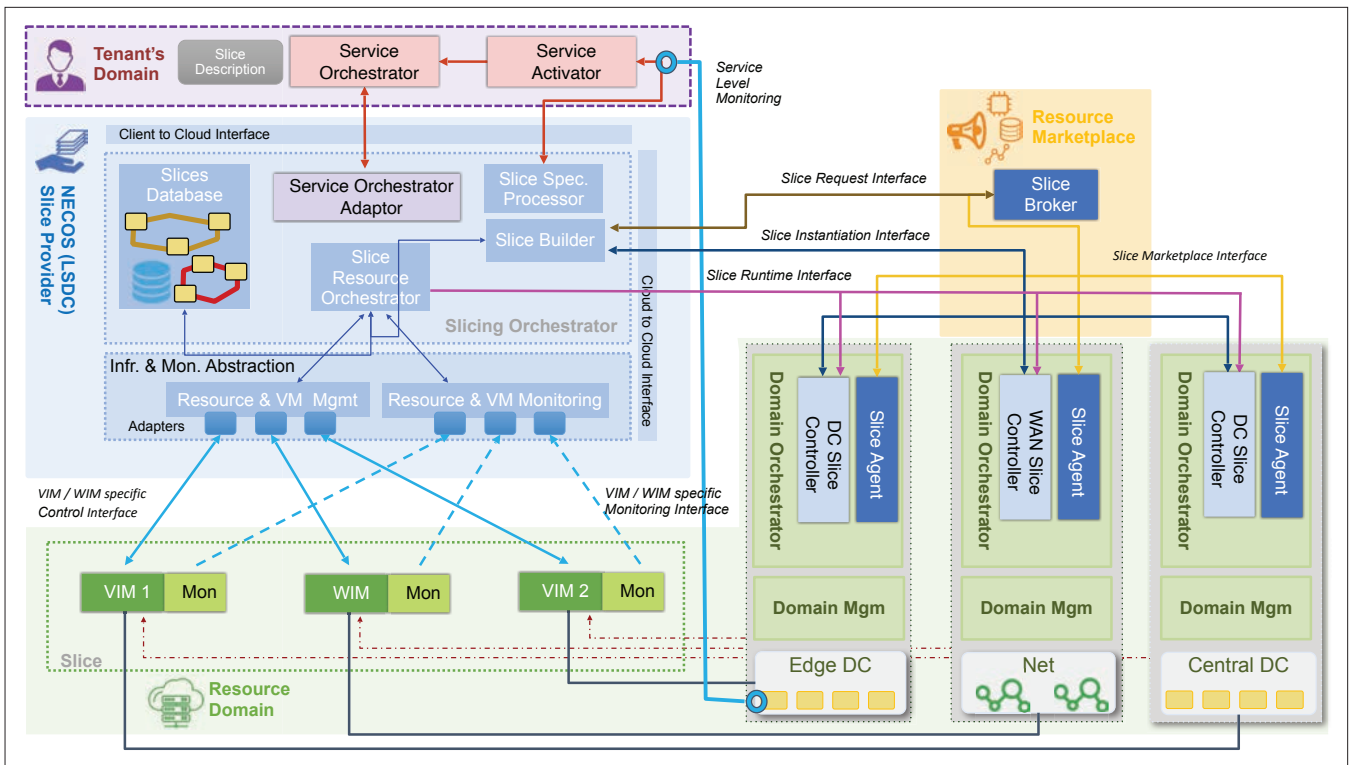


FIGURE 2. Functional view of the NECOS architecture.

Service Orchestrator: Residing in the tenant's domain, this component is responsible for deploying and managing tenant services on the allocated slice by interacting with the slice resource orchestrator.

Service Activator: This component is responsible for two tasks:

1. Passing the tenant's slice description to the NECOS slice provider
2. Handling the response, which informs the tenant's service orchestrator about the slice status

Slice Specification Processor: This component transforms the tenant's slice descriptor, which only requires a service specification, into an information-rich *slice creation request* for the slice builder, which includes resource options, geographical information, decomposition for slice parts, and so on.

Slice Builder: This component builds a full end-to-end cloud-network slice from a set of constituent slice parts. It is responsible for selecting, among the alternative offers, each slice part that the slice broker reports back as output from the discovery process via the resource marketplace.

Slice Resource Orchestrator (SRO): This component joins together the slice parts into a single aggregated slice to be used for service deployment, orchestration, and runtime life cycle management, including the placement and embedding, into the resource domains, of the VMs and virtual links for the services.

Infrastructure and Monitoring Abstraction (IMA): In order to interact with various VIM/WIM instances and the monitoring subsystems, the NECOS slice provider relies on IMA mechanisms that abstract away technology-specific details through plug-in API adaptors. The IMA subsystem provisions the tenant services and also monitors

the underlying resources, collecting additional telemetry data that is not available via the service-level monitoring interface.

DC Slice Controller: This component creates DC slices so that a given slice part is configured with compute and storage resources, as well as returning a handle to the respective VIM manager. This VIM can be either an instance deployed on demand, complying with the tenant's specification, or an existing running VIM in that resource domain, offering an isolated interaction point for the tenant.

WAN Slice Controller: For each network domain, a WAN slice controller instantiates each new network slice to interconnect two DC slices. Like the VIM, instances of the WIM can be deployed on demand, per tenant, or shared through an existing WIM using an isolation shim.

Slice Broker: This component manages the resource discovery process of the marketplace by decomposing a full slice request into a per-slice-part call for proposals. This triggers a bidding process involving the slice agents in each provider domain. The broker then aggregates the alternative responses, each annotated with a set of costs, and sends them to the slice builder.

Slice Agent: Running in each DC/WAN provider domain, this component's role is to receive slice part requests and to match the detailed slice part requirements to available resources reported by its own DC/WAN controller. If a call for proposals match is successful, it then returns an offer for the slice part.

EXPERIMENTAL EVALUATION

In order to validate the NECOS LSDC platform presented earlier, different software components were developed and evaluated in a set of experiments. These experiments are grouped into three

categories according to their functional domain and capabilities:

- Marketplace-based resource discovery
- Intelligent vertical elasticity functionalities
- Life cycle management of end-to-end slices and services

Additional experiments are provided through demonstrators,¹ including on end-to-end slicing for both logistics IoT and touristic content delivery network (CDN) services.

Although the prototype artifacts were not engineered for performance but rather for a proof-of-concept demonstration of the main NECOS novelties, when using lightweight components as in [7], the processing steps, from sending a slice request to having a handle returned to an executing manager, take around 6 seconds, depending on the capacity and load of the server.

The experiments were executed on a geographically distributed testbed deployed on six sites: four located in Brazil (Pará, Goiás, Campinas, and São Carlos), and two in Europe (Greece and Spain). In addition, 14 resource providers were emulated and used during the tests to better assess the experimental behavior on a larger scale. Radio access network (RAN)-level slicing is not part of the experiments since the testbed lacks wireless domains with virtualization and softwarization substrates.

MARKETPLACE

The novelties of the NECOS platform with respect to the slice-as-a-service model F1, the lightweight architecture F2, the multi-domain operation F3, and the marketplace's dynamic discovery of distributed resources F6, were evaluated in a use case involving up to 20 DC resource providers and 3 network resource providers. A slice request generator was used to introduce Tenant slice request workloads.

The size of each slice request is characterized by a 3-tuple, indicating the different numbers of DC slice parts, network slice parts, and virtual functions hosted, presented respectively as (*dc*, *net*, *vf*) in column 1 of Table 1. The required attributes of the complete slice request [8] were populated with random values. The experimental results presented in Table 1 are values for each slice request averaged over five runs. Columns 2 and 3 present the number of DC and network slice part alternatives collected, column 4 has the alternatives of end-to-end slices considered (with a maximum value of 50,000), and finally, column 5 presents the average time since the initial request to make the final slice selection. From these results and similar experiments, we conclude that the marketplace is:

- Fully functional
- Scales well for a wide range of parameters, including the number of slice parts, virtual functions, attributes' range, and diversity of resource providers

VERTICAL ELASTICITY ENABLED BY ML

One of the core functions of the NECOS slice resource orchestrator (SRO) is the runtime management of the slice life cycle, which includes the capability to carry out dynamic upgrades and downgrades of the resources allocated to the slice, to respond to demand changes. This

Slice request 3-tuples	Alternative DC parts	Alternative net parts	Slice alternatives	Wall time (s)
(2, 1, 4)	11	27	27	11.3
(3, 2, 6)	14.2	39.6	320.4	18.8
(4, 5, 8)	18.2	256.2	26,924.4	91.4
(6, 7, 24)	55.4	1755	50,000	605.8
(8, 9, 32)	73	2205	50,000	763.1

TABLE 1. Performance results of the marketplace.

capability, which is called elasticity [12], consists of two models: *vertical elasticity*, entailing the change of the amount of resources of the same slice parts; and *horizontal elasticity*, referring to the change of resources through the addition or removal of slice parts. This experiment summarizes the tests conducted to confer vertical elasticity to the SRO by leveraging ML.

Although this experiment also supports the validation of NECOS novel features F1 and F2, it is specifically focused on demonstrating feature F4, the intelligent elasticity, through an SRO implementing an ML algorithm that receives slice key performance indicator (KPI) metrics, and infers KPIs associated to the slice service level agreement (SLA) metrics. Through the time evolution of these KPIs, it is possible to predict the need to trigger upscale or downscale vertical elasticity processes to avoid SLA violations, while providing resource efficiency.

The experiment consists of a service using the Cassandra database, with the KPI of interest being the read-write time operations. The service is deployed on a slice over the distributed testbed and is characterized by infrastructure parameters like CPU, memory, network traffic, and others. As the number of service users accessing the service increases or decreases (via a load generator), the monitored resources are affected, and the read-write time service KPIs are impacted.

The SLA metric defined with a threshold of "response time below 180 ms" means that whenever the orchestrator predicts an SLA violation in the next 30 seconds, the elasticity process is started. The timings collected for a vertical elasticity process are plotted in Fig. 3. At around 225 seconds after the start of the time window, the orchestrator foresees an SLA violation (green star) and proactively triggers the vertical elasticity operations to modify the slice with additional resources. The elasticity process takes a couple of seconds, and this fixes the foreseen violation, as can be seen by the response times that are mainly maintained under the specified SLA, even as the number of users keeps growing.

FULL SLICE LIFE CYCLE IN MULTI-SITE DISTRIBUTED ENVIRONMENT

We evaluate the full process of end-to-end, multi-domain slice-as-a-service delivery through the deployment of a touristic CDN service on a slice that includes resource domains in Brazil and Europe (Fig. 4). This experiment uses all the components of the NECOS architecture of Fig. 2, and is devised to validate a rich set of the novel features of NECOS.

The touristic CDN service, highlighted via a blue dotted line in Fig. 4, operates on top of a cloud-network slice and delivers touristic con-

¹ Pointers to videos, source code, and reproducibility instructions of all NECOS demonstrators are available at <http://www.h2020-necos.eu/demonstrators/>

tent to users based on their geographic location. For the sake of simplicity, the figure only shows those providers that met the requirements of the requested slice, that is, resource providers along their slice controller and slice agent instances located in Campinas, São Carlos, Spain, and

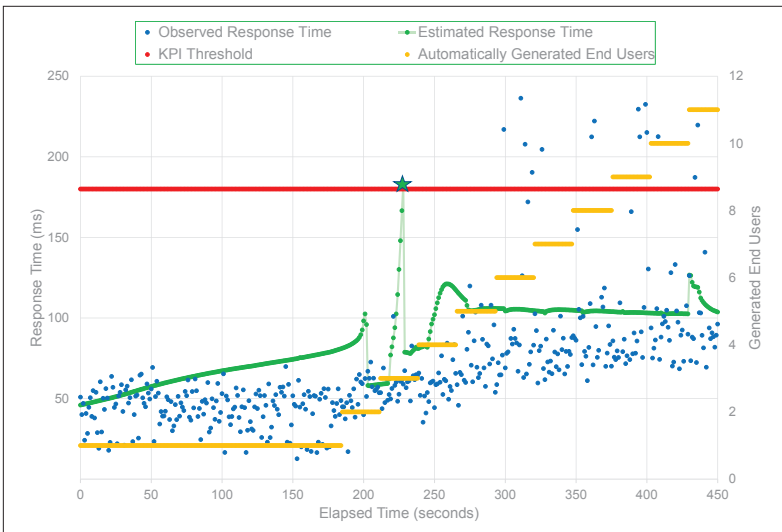


FIGURE 3. Estimated response time vs. observed response time as a function of the number of Cassandra users.

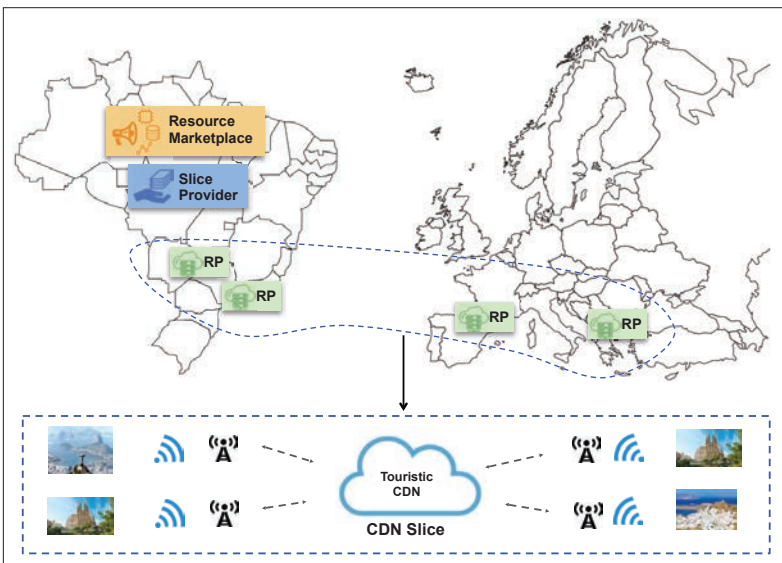


FIGURE 4. Touristic CDN service slice deployment in resource provider facilities in Brazil and Europe.

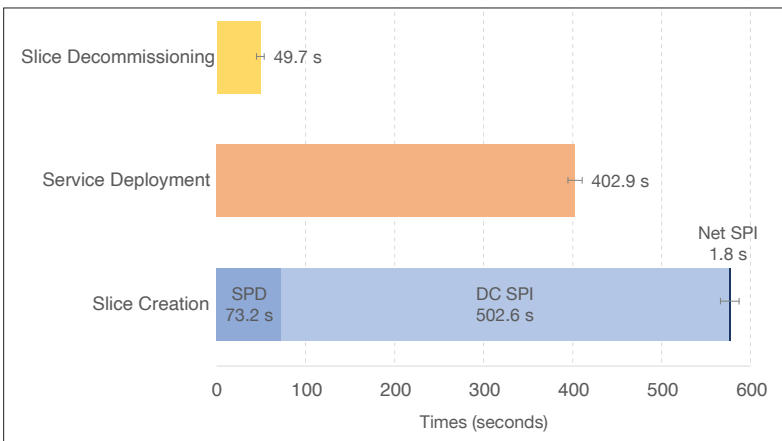


FIGURE 5. Slice life cycle elapsed times of the touristic CDN trial.

Greece. The process workflow is executed multiple times, and for each run, data is sampled to gather the times to create the slice, deploy the service, and decommission the slice.

In Fig. 5, we can observe that, on average, the process workflow takes almost 10 minutes (≈ 578 seconds) to create the slice, and nearly 6 minutes (≈ 403 seconds) to deploy the touristic CDN service, resulting in a total of 16 minutes to fully provision the slice with its service. We now look into the duration of the different operations, namely, the time spent looking for candidate slice parts, both DC and network, discovering the slice parts (SPD), instantiating the DC slice parts (DC SPI) in the selected resource providers, and setting up the network slice parts to form the end-to-end cloud network slice (Net SPI). In our experiments, it takes on average less than 2 minutes to find 7 slice parts (4 DC slice parts and 3 Net slice parts) among all the combinations involving 20 resource providers distributed across Europe, Brazil (SPD in Fig. 5) and the 14 emulated resource providers. On average, it takes less than 9 minutes to instantiate a DC slice part (DC SPI segment of slice creation) containing a fully configured Kubernetes cluster and less than 2 s to set up a VxLAN network slice part for the touristic CDN slice (Net SPI fraction of slice creation).

These times should be viewed with a perspective of our prototype establishing a trans-oceanic slice – an entity that may last for days, weeks, or even years – atop the physical and virtual resources of the testbed infrastructure. Therefore, these numbers should be considered as acceptable overheads for a relatively long lasting slice.

CONCLUSIONS AND OUTLOOK

The NECOS cloud-network slicing concept presents a set of innovative characteristics that allows for slicing using the slice-as-a-service model, in which service components are dynamically mapped into an end-to-end slice based on resources from different providers. The NECOS platform offers participating resource providers the ability to federate their infrastructures through a marketplace approach. This enables the execution of tenant services using virtualized resources to span across the cloud/network infrastructures offered by providers presenting different configurations. Such a concept is not purely technical, as it also considers business, cultural, and geographical relationships among domains. A number of experiments over instances of the NECOS architecture have served to prove its core concepts and novelties.

There are still many open challenges around the broader vision of slicing, for example:

- Full slice isolation, that is, highly efficient creation of isolation-guaranteed slices for each of the data, control, management, and service planes, which demands the development of enablers for safe, secure, and efficient slice multi-tenancy
- Slicing service mapping, that is, setting an efficient service mapping model bound over network cloud slicing, specifying policies and methods to honor service requirements without imposing infrastructure re-engineering

We expect the evolution of slicing to keep embracing heterogeneous resources at all hardware/software layers in order to realize service-tailored end-to-end architectures. Using innovative business and operational models, governed by the economies of sharing, where multi-tenancy at multiple realms becomes the norm, slicing will take the networking principle of resource multiplexing to new levels.

ACKNOWLEDGMENTS

Work supported by the H2020 EU-Brazil collaborative call, grant no. 777067 (NECOS), funded by the European Commission and the Brazilian Ministry of Science, Technology, Innovation, and Communication (MCTIC) through CTIC/RNP. The authors would like to thank the Editor and the anonymous reviewers for their valuable comments and suggestions as well as the whole NECOS team and stakeholders.

REFERENCES

- [1] A. A. Barakabitze *et al.*, "5G Network Slicing Using SDN and NFV: A Survey of Taxonomy, Architectures and Future Challenges," *Computer Networks*, vol. 167, 2020, p. 106,984.
- [2] L. U. Khan *et al.*, "Network Slicing: Recent Advances, Taxonomy, Requirements, and Open Research Challenges," *IEEE Access*, vol. 8, 2020, pp. 36,009–28.
- [3] A. Galis *et al.*, "Slicing 5G Networks: An Architectural Survey," *Wiley 5G Ref: The Essential 5G Reference Online*, Wiley, 2020, pp. 1–41; <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781119471509.w5GRef095>.
- [4] T. Taleb *et al.*, "On Multidomain Network Slicing Orchestration Architecture and Federated Resource Control," *IEEE Network*, vol. 33, no. 5, Sept./Oct. 2019, pp. 242–52.
- [5] NECOS, "EU-Brazil – Novel Enablers for Cloud Slicing," 2017; <http://www.h2020-necos.eu>, accessed Nov. 2020.
- [6] S. Clayman, F. Tusa, and A. Galis, "Extending Slices into Data Centers: The VIM On-Demand Model," *IEEE 9th Int'l. Conf. Network of the Future*, Poznań, Poland, 19–21 Nov. 2018.
- [7] S. Clayman *et al.*, "WIM On-Demand – A Modular Approach for Managing Network Slices," *IEEE Conf. Network Softwarization*, Ghent, Belgium, 29 June–3 July 2020.
- [8] P. D. Maciel *et al.*, "A Marketplace-Based Approach to Cloud Network Slice Composition Across Multiple Domains," *2nd Wksp. Advances in Slicing for Softwarized Infrastructures*, co-hosted at the 5th IEEE NetSoft, Paris, France, 24–28 June 2019.
- [9] F. S. D. Silva *et al.*, "NECOS Project: Towards Lightweight Slicing of Cloud Federated Infrastructures," *4th IEEE Conf. Network Softwarization and Wksp.*, 25–29 June 2018.
- [10] P. Valsamas *et al.*, "Multi-PoP Network Slice Deployment: A Feasibility Study," *IEEE CloudNet*, 4–6 Nov. 2019.

- [11] NECOS project, "D3.2: NECOS System Architecture and Platform Specification, V2," Tech. Rep., Apr. 2019; http://www.maps.upc.edu/public/necos_d3.2.v4.11_final_web.pdf, accessed Nov. 2020.
- [12] NECOS Project, "D5.2: Intelligent Management and Orchestration," Tech. Rep., Oct. 2019; <http://www.maps.upc.edu/public/D5.2%20final.pdf>, accessed Nov. 2020.
- [13] ETSI, "Zero-Touch Network and Service Management (ZSM) Industry Specification Group (ISG)," 2019; <https://www.etsi.org/committee/zsm>, accessed Nov. 2020.
- [14] NECOS Project, "D2.2: Consolidated Definition of Use Cases, Business Models and Requirements Analysis," Tech. Rep., Dec. 2018; http://www.maps.upc.edu/public/d2.2_final_v2.0.pdf, accessed Jan. 2021.

BIOGRAPHIES

STUART CLAYMAN (s.clayman@ucl.ac.uk) is a principal research fellow at the Electronic and Electrical Engineering Department, University College London, United Kingdom.

AUGUSTO NETO (augusto@dimap.ufrn.br) is an associate professor in the Informatics and Applied Mathematics Department, Federal University of Rio Grande do Norte, Brazil.

FÁBIO VERDI (verdi@ufscar.br) is an associate professor in the Computing Department, Federal University of São Carlos, Brazil.

SAND CORREA (sand@inf.ufg.br) is an associate professor in the Instituto de Informática, Universidade Federal de Goiás, Brazil.

SILVIO SAMPAIO (silviocs@imd.ufrn.br) is an assistant professor at the Digital Metropolis Institute, Federal University of Rio Grande do Norte.

ILIAS SAKELLARIOU (iliass@uom.edu.gr) is an assistant professor in the Department of Applied Informatics, University of Macedonia, Greece.

LEFTERIS MAMATAS (emamatas@uom.edu.gr) is an assistant professor in the Department of Applied Informatics, University of Macedonia.

RAFAEL PASQUINI (rafael.pasquini@ufu.br) is an associate professor, in the Faculty of Computing at Universidade Federal de Uberlândia, Brazil.

KLEBER CARDOSO (kleber@inf.ufg.br) is an associate professor in the Instituto de Informática, Universidade Federal de Goiás.

FRANCESCO TUSA (francesco.tusa@ucl.ac.uk) is a research associate in the Electronic and Electrical Engineering Department, University College London.

CHRISTIAN ROTHENBERG (chesteve@dca.fee.unicamp.br) is an assistant professor at the University of Campinas, Brazil.

JOAN SERRAT (serrat@tsc.upc.edu) is a full professor at Universitat Politècnica de Catalunya, Spain.

We expect the evolution of slicing to keep embracing heterogeneous resources at all hardware/software layers in order to realize service-tailored end-to-end architectures. Using innovative business and operational models, governed by the economies of sharing, where multi-tenancy at multiple realms becomes the norm, slicing will take the networking principle of resource multiplexing to new levels.