DISTRIBUTED CLOUD SERVICES BASED ON PROGRAMMABLE AGILE NETWORKS

Sonja Filiposka
Faculty of Computer Science and Engineering, Ss. Cyril and Methodius University, Rugjer Boshkovikj 16, 1000 Skopje, Macedonia
e-mail: sonja.filiposka@finki.ukim.mk
Yuri Demchenko
University of Amsterdam, Science Park 904, 1066HC Amsterdam, Netherlands
e-mail: y.demchenko@uva.nl
Tasos Karaliotas
Greek Research & Technology Network S.A. GRNET, Leof. Mesogeion 56, Athina, Greece
e-mail: karaliot@noc.grnet.gr
Migiel de Vos
SURFnet, PO Box 19035, 3501 DA Utrecht, The Netherlands
e-mail: migiel.devos@surfnet.nl
Damir Regvart
Croatian Academic and Reseach Network CARNet, Josipa Marohnića 5 10000 Zagreb, Croatia
e-mail: damir.regvart@carnet.hr

Paper type

Research paper.

Abstract

In order to allow researchers to focus on their work, all newly designed applications and services must not only offer features needed to tackle Big Data problems, but need to work seamlessly and intuitively, while efficiently hiding all non-necessary technical and networking details from the end user. Thus, today's requirements of the research and education community demand a holistic converged approach in the design of new generation intelligent cloud aware networks that need to work in concert with the distributed application components. The goal of this paper is to contribute toward setting up the building blocks of an open multi-cloud ecosystem based on a programmable cloud service delivery infrastructure. For these purposes we examine the possibility of introducing a third-party multi-cloud marketplace placed on top of a cloud aware agile network that can be reconfigured based on the defined multi-cloud application workflow. We discuss the overall architecture of the ecosystem and define the requirements for the underlying network infrastructure. The described proof of concept demonstrates the benefits of the proposed solution providing an example of how the concepts of network function virtualization and software defined networks can be leveraged in order to obtain an agile programmable network that will respond to the changes in the application traffic flow. With self-provisioning being acknowledged as a crucial component among service providers, we discuss all steps needed in order to translate user requests in a fully automated end-to-end service delivery and monitoring especially in a multi-domain environment.

Keywords

Agile programmable network infrastructure, marketplace, multi-cloud services, service function chaining, zero touch provisioning.

1. Introduction

High Performance Computing (HPC) traditionally used for storing and processing large amounts of data remains rigid and difficult to use from both programming and data management point of view. This is especially emphasized with the latest trends in modern research that became increasingly data driven and deal with Big Data problems (Chen 2014) requiring processing of huge amount of distributed data on-demand in an easy to use manner. The majority of todays' scientific data intensive tasks can easily be broken down into a set of independent tasks that can be processed in parallel (e.g. using cloud platform and not requiring HPC facilities), while the problem of distributed data management, storing and fast transfer remains unsolved.

In order to focus on their research, users need to be able to perform application-specific data analysis and processing in an intuitive and straightforward way. They don't need to understand the underlying building blocks of the infrastructure that should solve the problems of distributed computing, storing and interconnection seamlessly (Katal, 2013). Use case examples that highlight these challenges can be found in almost every scientific branch like bioinformatics, geoscience, high quality video streaming and real-time processing, or simple collaborative efforts of a large group of geographically distributed scientists.

Cloud computing in all of its available flavours like IaaS, PaaS and SaaS have come to play a major role in this attempt to alleviate the researchers from the burden of knowing the details of the underlying infrastructure and its management in order to be able to use it for advanced data intensive purposes. By providing resource abstraction and simple automation tools, the modern cloud platforms simplify the majority of routing tasks such as setup, maintenance, backup, security, etc. In this way, cloud based applications have become an essential tool for modern researchers. Even more, today they tend to be the most favourite way of dealing with Big Data problems (Hashem, 2015).

An omnipresent component in all of the discussed scenarios is the simple fact that the quality of the cloud services as it is being perceived by the client is crucial for the continuous usage and operations of the cloud-based applications (Toosi, 2014). However, Big Data researchers can under no circumstances be treated as “typical” cloud service clients. Their data intensive tasks require high bandwidth, low latency, and high anytime reliability when moving the data to/from the client which can be quite challenging in the best effort Internet environment. This ongoing trend has put huge stress on the network infrastructure that needs to provide high performing connectivity between the researchers as clients and the multitude of cloud service providers (CSPs) that offer the required set of resources.

Another commonly recognised problem is the delivery of cloud services from cloud data centres to the end used location either for fast data transfer or data visualisation, defined as the “last mile” problem. Current major model for cloud services access and delivering is based on the best effort open Internet access that doesn’t guarantee the quality of services (QoS) which are critical for many scientific applications. To address this problem, the GÉANT project proposed the Open Cloud eXchange (OCX) (Demchenko, 2014) as a component and central hub of the GÉANT based cloud-aware backbone infrastructure. OCX is envisioned as a hierarchical set of points where CSPs can directly connect to the GÉANT network via high bandwidth, low latency special dedicated links. The GÉANT and NRENs clients (researchers) can then connect to the CSPs of their choice by establishing on-demand links from the client via one or several OCX devices to the chosen CSP. In this way, the ground has been set for providing dynamic dedicated links between any user in the GÉANT community and a connected CSP. The OCX concept of bringing cloud providers and R&E users closer together has been successfully demonstrated on several occasions (Demchenko at al, 2014) showing the multitude of benefits of the proposed approach especially in terms of collaborative efforts and solving Big Data problems.

The proposed cloud services delivery infrastructure architecture has put forth the means for further development in the area of supporting collaborative Big Data research via cloud services by providing the foundation for further developing two equally important components: (1) a third-party cloud service marketplace and (2) an intelligent agile network that can respond to the dynamic nature of the service usage.

By providing a single access point for delivery of services offered by multiple cloud providers, the OCX enabled Cloud Services Delivery Infrastructure (CSDI) provides a basis for a multi cloud marketplace where all providers and clients can go to in order to publish, browse and subscribe to one or multiple services. With the ever more prominent use and development of multi-cloud services, the marketplace has the potential to become a central point for the multi-cloud applications deployment, operation and management as it is discussed in the rest of this paper.

The OCX demonstrations have also shown that for this approach to be widely adopted and successful as a solution in the research community, it needs to be fully transparent. In other words, it is of great importance that the underlying network infrastructure is in fact empowering the cloud-based elastic and dynamic experience while working with big data sets and high computing demands (Pathan, 2013). This means that the on-demand connections between the entities need to be established in a fast, reliable way provisioned via a fully automated process with dynamicity, elasticity, and scalability that extend the generic properties of a single cloud to multi-cloud and inter-cloud environment. The first steps towards this intelligent agile network infrastructure between the providers and the clients, able to respond in real time while supporting federated multi-cloud services is another topic of this paper.

In summary, the main goal of this paper is to bring together the aspects of multi-cloud applications and a third-party cloud marketplace, and analyse the possibilities for their provisioning using a fully automated agile programmable network model.

The remainder of the paper is organized as follows: In Section 2 we discuss the architecture of an open multi-cloud R&E ecosystem providing a marketplace that will bring together the providers and users in the multi-cloud environment and enable the execution of complex workflows that span multiple clouds. Section 3 focuses on the requirements from the underlying network, especially agility and programmability, which need to be provided in order to support the defined complex workflows of data intensive multi-cloud applications. As a proof of concept a demo scenario is presented that demonstrates the benefits of the approach while providing a seamless interface to the end user. The final section concludes the paper reflecting on related attempts and sets the ground for future efforts in the field of open science cloud environments.

2. Open multi-cloud R&E ecosystem

Typical cloud services provided by individual CSPs today are offered via dedicated portals that are developed and run by each individual CSP. This puts a strain on users that are interested in different CSPs for different parts of their research process, e.g. storage at one CSP, processing at two other CSPs. The process of selecting the best CSP for a given purpose is troublesome and tedious without a comprehensive catalogue that can unify all existing offers and their characteristics. As a first step towards addressing this issue, a service activity within GÉANT has created the GÉANT Cloud Catalogue (GÉANT, 2016) that is a form of a structured list of all available CSPs that have committed to collaboration with the R&E community. The catalogue can be used in order to make a comparison of the different CSPs based on different characteristics defined as important cloud requirements by the R&E community such as AAI, type of user provisioning, billing, data protection, backup and restore, user availability, national and international regulations, data ownership and similar. The catalogue is a great starting place for choosing CSPs but does not help in the rest of the process starting from subscription to the chosen services. For the next steps, users are needed to use the CSP provided dashboards or other means of service subscription and provisioning defined by the particular CSP.

In addition, the focus on cloud powered service development today turns to the distributed complex applications that are multi-cloud based (Huang, 2014). This means that ever more frequently researchers today are interested in developing and using multi-cloud applications or services that are actually composed by combining together a number of interlinked cloud services. The problem gets really interesting in the case when these different cloud services that are to be composed together are offered by different CSPs, thus creating the "multi-"cloud aspect.

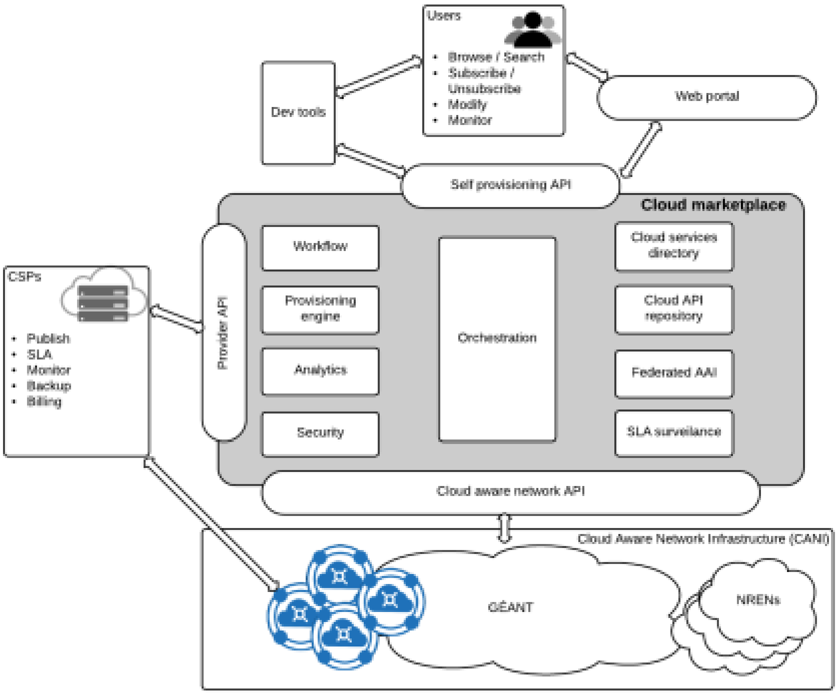

Figure 1 OCX enabled multi-Cloud Services Delivery Infrastructure.

This is one of the major points where an CSDI based marketplace for the R&E community can play a crucial role in order to alleviate the burden of multi-cloud application developers, see Fig. 1. On top of the envisioned OCX enabled CSDI a self-service user portal can be developed that will act as a trusted third-part marketplace for all interested parties (the CSPs that have established a direct link to a GÉANT gOCX instance and the R&E community served by the GÉANT and NRENs network). This marketplace will provide a detailed service portfolio and catalogue built by enabling all CSPs to publish their offered services. The CSDI based marketplace features can be seen as an extension of the existing cloud catalogue today: it will not just list the offers and characteristics of each CSP, but it will also act as a single point of entry (sort of a multi-cloud enabled dashboard) where the users can not just browse the offers, but also subscribe to the offered services, monitor the complete set of used resources across all clouds and manage and combine resources from different providers.

Another seamless integration in this proposed marketplace environment is the possibility for using the hybrid cloud model, which is a common request within the R&E community. By having a single point of entry for both public and private clouds, the marketplace user can view the available resources from the private (research/university) providers that usually offer specialized services to their researchers and the public providers that augment the amount of storage and computing resources available.

With all of these capabilities, eventually the marketplace is not envisioned to be used by the end users of the multi-cloud applications, but by the multi-cloud application developers that can benefit the most from the true values of the marketplace: single platform where they can locate the providers of needed resources and then decide on the final set of development tools that are compatible with their CSPs choice.

Of course, the goal is to develop a multi-cloud application that will provide seamless usage to the end users and they must under no circumstances be burdened with resource provisioning, usage monitoring, setup, backup and similar management and maintenance tasks. The workings of the multi-cloud application need to be fully automated and for these purposes a highly capable service orchestration and application deployment tools are needed such as some of the newest cloud based tools like Chef, Puppet or Ansible (Venezia, 2013). All of these solutions are based on using recipes or cookbooks to describe machine configurations in a declarative language, bring them to a desired state, and keep them there through automation. For a multi-cloud scenario where many application components spawned in different providers exist, the complete application topology and the components interrelationship needs to be described using a language like OASIS TOSCA (Binz, 2014).

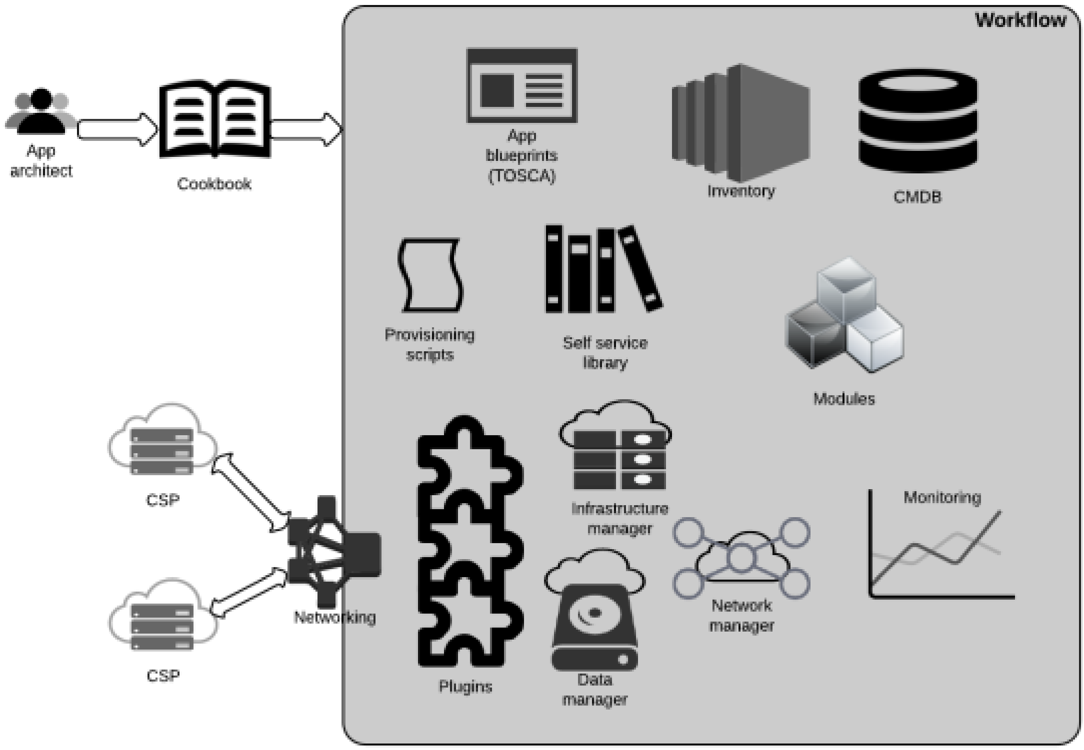

Figure 2 Multi-cloud application workflow components.

Multi-cloud services are characterized with complex workflows, see Fig. 2, that define the schedule and timing of the invocations of the different cloud based service components. The workflow also defines the data path during the application lifecycle, the original input to the application, how the output of once service component becomes an input for another one, and how the final output of the multi-cloud application is going to be provided to the user. A well-defined workflow can provide streamlined cloud management that can enable gauging the application response time by modelling the full multi-cloud application as an end-to-end process. The streamlined cloud management process is actually a starting point for the possibility of including the DevOps philosophy (Cukier, 2013) in the CSDI based marketplace. The final goal here is to choose and use a single DevOps tool that will be able to integrate the different environments and in this way will simplify the complete lifecycle management for the application developers.

A possible extension of the proposed marketplace model is the ability to allow users/developers to publish and offer their created composite cloud services to the rest of the community. In this way the marketplace can also transform into a repository of useful composite tools and services that are of interest for the GÉANT community thus minimizing the repeated effort among research projects while providing the application developers with different bundled components they can use for their envisioned solution. Similar vision of this composite tools repository intended for the bioinformatics researchers is being developed within the Cyclone project (Cyclone, 2015) and is part of the Slipstream cloud automation tool (SixSq, 2015).

3. Agile programmable cloud aware network infrastructure

The described R&E cloud ecosystem can be seen as a specialized, yet versatile and sophisticated cloud management platform that provides a very much needed single point of abstraction for the researchers that rely on cloud computing to tackle their challenging problems. For providing full automation of the process thus enabling complete transparency and high usability to the end user, the marketplace can be seen as a self-service portal that offers a massive collection of network accessible services/APIs. Thus, in order to ensure tangible value and provide the needed high performances, the network setup underneath must be intricately intertwined in all steps of the services lifetime from setting up, during their elastic usage, till tear down. Of course, this marriage between the cloud management automation tools and the network infrastructure underneath must be fully automated by exposing partial control to the network via a well-defined set of APIs that will enable the creation of on-demand high performing dynamic connections between all cloud application components with the possibility of a fine grain user control of the full virtual network encompassing all application entities.

In this way, the growth in complex cloud implementations has increased the intra- and inter-domain network requirements. Yet, one must concede that network performance is one of the key issues when implementing multi-cloud solutions. This leads to treating network governance and management as a core concern; it being an integral part needed to provide integrated security and application performance.

This translates in the requirements that the cloud aware network infrastructure must become highly agile, dynamically responding to the changing requests as the complex multi-cloud application workflow is being implemented. To achieve this the network must be fully automated which leads not only to lower cost for provisioning of new infrastructure, but, most importantly, allows for self-provisioning. Self-provisioning, on the other hand, means that the network becomes service oriented, provides automated governance with adaptable degrees of security and control.

All of this leads towards highly improved user experience where by using APIs the network becomes an agile programmable environment that works in concert with and responds to the requirements of the cloud application layer. Thus, in the process of orchestration of the operations and management of the overall system discussed in the previous section, network orchestration must be included. The main goals of the network orchestration process are to deliver the dynamic behaviour of the network that can be fully aligned to the client requirements by self-adapting and enhanced agility.

We can summarize this discussion with the following requirements from the cloud aware network:

- **Intelligent network connectivity** - the network should respond to the application self-provisioning requests enabling auto-provisioned network capacity, intelligent traffic routing and highly scalable IP service, possibilities for traffic prioritization, and services that will enable network optimization;
- **Integrated network security** – enabling partial client "control" of the network can only be done in an environment that implements integrated policy-driven security with intrusion detection and proactive protections, full real-time monitoring and management, and cloud-aware security services;
- **Network and applications performance management** – the overall performances that are delivered to the end user are to be obtained by correlating the application and network performance log analysis, while the client must be able to have an end-to-end visibility of his/her portion of the network via a single unified console.

In this way the GÉANT network will turn into an intelligent WAN providing a series of stitched bandwidth on-demand services between the gOCX instances becoming an always on, burstable, virtual private network for the multi-cloud ecosystem with increased security owing to the use of existing, private connections instead of the public Internet or physically separate direct connections that can not scale to the level of the R&E community.

This holistic integrated approach for federated inter-cloud services will expose the underlying CSDI via the proposed cloud marketplace. Essentially, this means that the marketplace will be augmented with a network as a service API. In this way, cloud application developers can integrate network requirements into their custom defined multi-cloud workflow so that the complete data flow between the component services can be done via requested (long and short-term) virtual circuits that are automatically set up. In other words, the network will be dynamically programmed so as to ensure that the data flow from one service to another will be transported under the conditions as requested by the controlling application.

The provisioning of the seamless, Zero-Touch based, implementation of this convergence of networking and cloud services is far from trivial. For these purposes we have proposed the Zero-Touch Provisioning, Operations and Management (ZTPOM) architecture (Demchenko, 2016). In general, the ZTPOM infrastructure includes a ZTPOM Server (or Engine) for each application domain, all coordinated by a central ZTPOM server, and ZTP clients. The server holds the full information about application or infrastructure to be provisioned and initiates clients to run the application or perform infrastructure provisioning locally on the hosting cloud. Within a typical ZTP scenario, the ZTP client may need to discover the ZTP server and download necessary information and images before staring the deployment process or workflow based on the application blueprint or recipe. It is anticipated that the ZTP process may need to adjust the provided general recipe to the local environment and cloud platform. Thus, the ZTP process is to be implemented on several levels: device-element configuration management, network wide synchronization, multi-domain network orchestration, and top level central management that should be an integral part of the marketplace.

For high performing multi-cloud applications this task would include dynamic set up and tear down of a number of virtual connections, wherein the optimal path between a given set of gOCX instances needs to be computed according to the current complete CSDI topology and network status. The problem complexity increases exponentially when considering the fact that the gOCX instances are placed in different network domains, and that setup of one dynamic virtual link translates into a multi-domain network wide orchestration process.

The exposed network as a service API will enable partial user control of the virtual network with a full range of operation support services. The currently available DevOps and NetOps tools do not include this level of programmability options for the underlying network. Thus, to increase the agility of the network they must be extended with additional features that will enable not only creation of a custom virtual network overlay but will provide full dynamic setup and tear down of individual connections and in-depth management and monitoring of their characteristics. In this way, the newly developed complex multi-cloud applications are going to become composite network-cloud services solutions that enable flexible, dynamic network service composition across heterogeneous networking systems. Considering the complexity of the task at hand, the solution is only feasible if we turn to complete network functions abstraction and network resource virtualizations provided by Network Function Virtualization - NFV (Han 2015) and Software Defined Networks - SDN (Li 2013), as these are the necessary tools that will enable the agility of the underlying network infrastructure via programmability, while providing combined control and management.

3.1 Multi-cloud demo and testbed

With the aim to investigate the suitability of NFV and SDN as the basis for providing the programmable requirements for the CSDI's ecosystem, a multi-cloud application for real-time UHD video editing was developed and demonstrated at the SC15 conference in Austin, Texas (Suerink, 2015). Several specifically targeted main goals were aimed with the demonstration:

- **Full transparency for the end-user**: provided with an intuitive graphical user interface that successfully hides all technical issues and allows the user to focus on the task at hand, i.e. choosing from a number of editing options and deciding on the order of implementation.
- **Highly dynamic features**: the number of CSPs involved is dynamically ordained according to the user defined workflow that is translated into a dynamic setup and tear down of links between the parties with on-demand spawning and continuous monitoring of the VM instances involved.
- **Orchestration of a programmable network**: all user initiated changes in the workflow are sent via an API to an orchestrator that needs to accordingly decide on changes in the links between the distributed application components so as to make sure that the UHD real-time data flow will follow the user specified path and the QoS parameters will be satisfied at all times.

In order to achieve the laid out goals the demonstration was developed using NFV combined with Service Function Chaining (SFC) (Halpern, 2015) all taking place over the programmable network designed using an SDN controller. NFV was used to move the functions traditionally performed by specialized equipment to the cloud virtualized environment which provides high flexibility and scalability. This makes it possible to manage network traffic through several functions, creating a chain of network functions. Adding network functions to active traffic with full transparency and in real-time requires a programmable network, and this was enabled by creating SDN based gOCX instances. With a programmable network, one can add and remove network functions at any time, in any order which is one of the main requirements of the demo.

The decided demo scenario was laid out as follows, see Fig. 3: an UHD camera captures the user image in real time; the real time data stream undergoes through a custom defined series of video effects; using an intuitive drag and drop web based interface, the end-user defines the number and order of different video editing functions, which effectively defines the SFC workflow; the composite result of the editing process is presented to the user on an UHD display.

The gOCX programmable instance was implemented as an OpenDayLight controller (Medved, 2014) located in Amsterdam. This instance receives API calls from the SFC workflow and orchestrates the directions of the traffic flow between the different CSPs. Each of the CSPs (SURFnet, Cloud Sigma, SURFSara, Okeanos) is considered as a provider of one type of video editing function: rotate, text add-on, grayscale, mirror, respectively. The total generated traffic flow was 3.2 Gbit/s. Using mostly 100G links, with no video functions activated, the traffic came straight back to Austin from Amsterdam with a delay of around 240ms. When video functions were introduced, the increase in delay depended on the currently chosen workflow by the user. However, in every single instance the end users were provided with a very high quality of experience characterised with no glitches in the streamed video and non-noticeable latency.

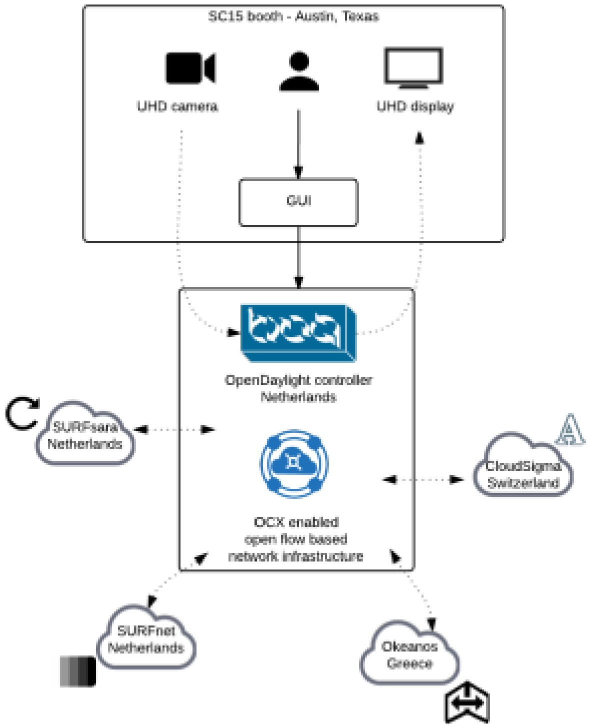

Figure 3 SC15 demonstration scenario.

SFC provides the ability to define an ordered list of network services. These services are then "stitched" together in the network to create a service chain. SFC is conceptually related to Policy Based Routing in physical networks but is typically thought of as SDN technology. Fundamentally, SFC causes network packet flows to route through a network via a path other than the one that would be chosen by routing table lookups on the packet's destination IP address. The SFC functions were implemented using VLAN stitching and MAC rewriting so that different paths and functions are kept separate. With the experience gathered during the demo development, the SFC option has proven to be quite complex and not as flexible as firstly imagined. However, the possibilities for SFC are still in development, and a new protocol for identifying service function paths in the network called Network Services Headers (NSH) is now underway.

The SC15 demo has successfully proven the idea of integration of network and cloud services in a federated and distributed infrastructure. The demo setup and lessons learned should be used as the basis for defining the CSDI's network as a service API that should enable GÉANT to grow into a cloud aware agile programmable network that will respond to the requirements of the developed cloud applications.

The demo setup is thus planned to grow into a testbed that will provide means for live testing of the zero-touch based network as a service combined with different multi-cloud based applications. The first order testbed setup will also become an excellent example of the possible future application functions solutions that can find their place in the CSDI's marketplace. Based on the NFV paradigm, a large number of additional services can be developed and then offered on the marketplace, thus creating a pool of virtualized functions that the end user, using methodologies like SFC, can combine in a desired and innovative way creating custom multi-cloud solutions that will enrich the R&E community.

4. Conclusion and future development

The main goal of this paper is to contribute towards establishing GÉANT as a cloud aware intelligent network by introducing the idea of the CSDI enabled marketplace envisioned to become the focal point of interoperable and integrated cloud services that dynamically interact with the agile programmable network infrastructure for high performance delivery of targeted data flows. Similar marketplace concept in the CSP community is viewed as an integral part of the new generation OSS/BSS solutions (Chen, 2015) where the idea is translated into a self-service user portal as an entry point that invokes the automated behaviour of the business intelligence tools and the operation support components related to provisioning, monitoring and billing. It is defined as a part of the digital service reference architecture by the TeleManagement Forum (TM Forum, 2015) and it is widely accepted and supported by the MEF community (MEF, 2015). However, these proposals are all mainly focused on solving a single provider problem, while just recently the commercial service providers community has started to delve into the issue of partnership and providing common composite services. A simple form of this partnership is seen in the Equinix Cloud Exchange and Marketplace (Equinix, 2015) that aims to act as a gateway to a set of CSPs that are directly reachable via the distributed Equinix data centres.

In the R&E community the effort to bring together multiple CSPs that will enable richer experience and opportunities for collaboration to distributed researchers is on the rise. Having in mind the latest European Cloud Initiative (ECI, 2016) being one of the 16 initiatives in the Digital Single Market strategy (EU Commission, 2015), it is straightforward to conclude that the GÉANT network infrastructure needs and should play a crucial role in the delivery of the European cloud services to the R&E community. The aim to establish a European Open Science Cloud (EOSC) as a trusted environment for hosting and processing research data must be supported by solutions that will provide seamless usage of service, such as implementations of the proposed CSDI architecture in conjunction with provisioning, operations and management automation tools.

The ultimate goal of these efforts should enable a user transparent end-to-end automation of the overall process of design, configuration and deployment, and operations and management of multi-provider services in a next-generation converged environment. In order to achieve this goal, all necessary components must be well aligned and fully support automation. This includes tight interaction between all involved components starting from an intelligent agile network infrastructure exposed with an open API that is dynamically reconfigured from a fully automated distributed set of responsive services that combined together support the complete portfolio of management and operations. In this way the process of provisioning services requested by the client can be entirely automated, where the marketplace on top will present the clients with an integrated view and control of the requested services augmented with real-time monitoring information. Thus, the user requests and modifications can be analysed by the operations services, translated into a set of necessary configuration adjustments accordingly, and then automatically propagated to the infrastructure for an almost instant response of the complete system architecture.

Acknowledgement

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 691567 (GÉANT 4, Phase 1). The authors would like to thank Ronald van der Pol, Hans Trompert and Gerben van Malenstein from SURFnet and the JRA1T1 team for their contribution to the development of the SC15 demonstration, especially Dani Arbel and Kurt Baumann for their constructive comments and discussions.

References

Binz, T., Breitenbücher, U., Kopp, O. and Leymann, F., 2014. TOSCA: portable automated deployment and management of cloud applications. In *Advanced Web Services* (pp. 527-549). Springer New York.

Chen, M., Mao, S. and Liu, Y., 2014. Big data: A survey. *Mobile Networks and Applications*, 19(2), pp.171-209.

Chen, Y., Qin, Y., Lambe, M. and Chu, W., 2015. Realizing network function virtualization management and orchestration with model based open architecture. In *Network and Service Management (CNSM), 2015 11th International Conference on* (pp. 410-418). IEEE.

Cukier, D., 2013, October. Devops patterns to scale web applications using cloud services. In *Proceedings of the 2013 companion publication for conference on Systems, programming, & applications: software for humanity* (pp. 143-152). ACM.

Cyclone Horizon 2020 project, 2015. Available at: < http://www.cyclone-project.eu/ > (Accessed 16 April 2016)

Demchenko, Y., van der Ham, J., de Laat, C., Matselyukh, T., Escalona, E., de Vos, M., Filiposka, S., Karaliotas, T., Mavrin, A., Regvart, D., Baumann, K., Arbel, D. and Breach, T., 2014. Open Cloud eXchange (OCX): Bringing Cloud Services to NRENs and Universities. *Terena TNC 2014*.

Demchenko, Y. at al., 2014. GÉANT Open Cloud eXchange (gOCX): Architecture, Components, and Demo Scenario. *Supercomputing conference 2014*.

Demchenko, Y., Grosso, P., de Laat, C., Filiposka, S., Regvart, D., de Vos, M., Karaliotas, T., 2016. Zero-Touch Provisioning Model and Infrastructure Components for Multi-Provider Cloud Services Provisioning. *Intercloud workshop, IEEE International Conference on Cloud Engineering (IC2E)*.

Equinix, 2015. The Home of the Interconnected Cloud: Where Service Providers and Enterprises Meet to Grow and Innovate Together. Available at <http://www.equinix.de/resources/whitepapers/home-interconnected-cloud/> (Accessed

on 16 April 2016)
EU Commission, 2015. A Digital Single Market Strategy for Europe – COM (2015) 192 final. Available at <https://ec.europa.eu/digital-single-market/en/news/digital-single-market-strategy-europe-com2015-192-final> (Accessed on 16 April 2016)
European Cloud Initiative, 2016. Building a competitive data and knowledge economy in Europe, COM(2016) 178 final, Brussels. Available at <https://ec.europa.eu/digital-single-market/en/news/communication-european-cloud-initiative-building-competitive-data-and-knowledge-economy-europe> (Accessed 19 April 2016)
GÉANT Cloud Catalogue, 2016. Available at: <https://catalogue.clouds.GÉANT.net> (Accessed 16 April 2016)
Halpern, J. and Pignataro, C., 2015. Service Function Chaining (SFC) Architecture (No. RFC 7665).
Han, B., Gopalakrishnan, V., Ji, L. and Lee, S., 2015. Network function virtualization: Challenges and opportunities for innovations. *Communications Magazine, IEEE*, 53(2), pp.90-97.
Hashem, I.A.T., Yaqoob, I., Anuar, N.B., Mokhtar, S., Gani, A. and Khan, S.U., 2015. The rise of "big data" on cloud computing: Review and open research issues. *Information Systems*, 47, pp.98-115.
Huang, J., Liu, G. and Duan, Q., 2014. On modeling and optimization for composite network–Cloud service provisioning. *Journal of Network and Computer Applications*, 45, pp.35-43.
Katal, A., Wazid, M. and Goudar, R.H., 2013. Big data: issues, challenges, tools and good practices. In *Contemporary Computing (IC3), 2013 Sixth International Conference on* (pp. 404-409). IEEE.
Li, C.S. and Liao, W., 2013. Software defined networks. *IEEE Communications Magazine*, 51(2), pp.113-113.
Medved, J., Varga, R., Tkacik, A. and Gray, K., 2014, June. Opendaylight: Towards a model-driven SDN controller architecture. In *2014 IEEE 15th International Symposium on* (pp. 1-6). IEEE.
MEF, 2015. The Third Network: Lifecycle Service Orchestration Vision. The Metro Ethernet Forum 2015.
Pathan, A.S.K., Monowar, M.M. and Fadlullah, Z.M. eds., 2013. *Building Next-generation Converged Networks: Theory and Practice*. CRC Press.
SixSq, 2015. SaaS, Slipstream and Helix Nebula Marketplace. *SixSq showcase at Cloud Expo Europe*. Available at <http://www.cloudexpoeurope.com/Exhibitor> (Accessed 16 April 2016)
Suerink, T., Karaliotas, T., Regvart D., Demchenko, Y., Baumann, K., Kooman, F., Mufty, W., Trompert, H., van der Pol, R. van Malenstein, G., de Vos, M., 2015. Real-Time Worldwide Service Function Chaining Using a Programmable Network with Network Function Virtualisation. *SuperComputing15 conference*.
TM Forum, 2015. Digital Service Reference Architecture Guide R15.0.0. Available at <http://www.tmforum.org/resources/standard/digital-services-reference-architecture-guide-r15-0-0/> (Accessed on 16 April 2016)

Toosi, A.N., Calheiros, R.N. and Buyya, R., 2014. Interconnected cloud computing environments: Challenges, taxonomy, and survey. *ACM Computing Surveys* (CSUR), 47(1), p.7.
Venezia, P., 2013. Review: Puppet vs. Chef vs. Ansible vs. salt. *InfoWorld*.

Biographies

Sonja Filiposka received her Ph.D. in 2009 at the Faculty of Electrical Engineering and Information Technologies, Ss. Cyril and Methodius University in Skopje. Currently, she is an associate professor at the Faculty of Computer Science and Engineering at the university. Her main research areas and interests include network design and performances, network science and complex networks theory, simulation methodologies and analysis. She is a co-author of over 100 scientific papers and is involved in a number of research and application projects on international and national level, as well as Women in engineering and Girls in ICT activist.

Yuri Demchenko received his M.Sc. degree and later Ph.D. from the Kiev Polytechnic Institute, National Technical University of Ukraine. He is a Senior Researcher with the System and Network Engineering (SNE) Research group at the University of Amsterdam. His main research areas include Cloud and Big Data technologies, security architectures and distributed authorization service infrastructure, complex resource provisioning and manageable security services. He is coordinating H2020 CSA project EDISON and involved in two European projects CYCLONE and GÉANT4 where he develops components of the federated Intercloud infrastructure. Yuri is actively contributing to standardization activity at RDA, NIST, OGF, IETF.

Migiel de Vos received his M.Sc. degree in Computer Science and Engineering from the Technical University of Eindhoven in the Netherlands in 2009. Since 2010 he is working on Research and Development at the department network services of SURFnet. He is involved in the operations and development of the SURFnet network and the NetherLight exchange point. This includes dedicated international connectivity, On-Demand networking and connecting cloud service providers. As of 2013 Migiel is also participating in the GN3plus project where he actively contributes to the concept of the Open Cloud eXchange.

Tasos Karaliotas received his diploma in Computer Engineering and Informatics from the University of Patras in 2002, and since has been working as a professional network engineer first in the Research academic computer technology institute, and then in the Mediterranean Nautilus and the Greek Academic and Research Network GRNET. He has participated in a number of European research and development project, and as such has also been a part of the team that manages the academic IaaS CSP Okeanos.

Damir Regvart is a CARNet (Croatian Academic and Research Network) employee working as Head of Network Operation Center. Damir received his master's degree in Electrical Engineering from the University of Zagreb, Faculty of Electrical Engineering and Computing in 2002. His research interests are in the field of secure network communication, implementation of new network protocols, architectures and WAN technology. Currently he is involved in GN4 JRA1 project as task leader for Network Architectures for Cloud Services for Horizon 2020.