Open Cloud eXchange (OCX): A Pivot for Intercloud Services Federation in Multi-provider Cloud Market Environment

Yuri Demchenko, Cosmin Dumitru, Ralph Koining, Cees de Laat
University of Amsterdam
{y.demchenko, C.T.A.M.deLaat}@uva.nl

Taras Matselyukh
Opt-Net BV
tmatsely@opt-net.eu

Sonja Filiposka
Ss. Cyril and Methodius University in Skopje
sonja.filiposka@finki.ukim.mk

Migiel de Vos
SURFnet
migiel.devos@surfnet.nl

Daniel Arbel
IUCC
dani@noc.ilan.net.il

Damir Regvart
CARNET
damir.regvart@carnet.hr

Tasos Karaliotas
GRNET
karaliot@noc.grnet.gr

Kurt Baumann
SWITCH
kurt.baumann@switch.ch

***Abstract*—This paper presents results of the ongoing development of the Open Cloud eXchange (OCX) that has been proposed in the framework of the GN3plus project. Its aim is to provide cloud aware network infrastructure to power and support modern data intensive research at European universities and research organisations. The paper describes the OCX concept, architecture, design and implementation options. OCX includes 3 major components: distributed L0-L2 (optionally L3) network infrastructure that includes OCX points of presence (OCXP) interconnected with GEANT backbone; the Trusted Third Party (TTP) for building dynamic trust federations; and the marketplace to enable publishing and discovery of cloud services. OCX intends to be neutral to actual cloud services provisioning and limits its services to Layer 0 through Layer 2 in order to remain transparent to current cloud services model. The recent developments include an architectural update, API definition, integration with higher-level applications and workflow control, signaling and intercloud topology modelling and visualization. The paper reports about results and experiences learnt from the recent OCX demonstrations at the SC14 Exhibition in November 2014 that demonstrated the benefits of an OCX enabled Intercloud infrastructure for running data intensive real-time cloud applications on top of the advanced GEANT multi-gigabit network. The implemented OCX functionality allowed applications to control the network path for data transfer and service delivery connectivity between multiple Cloud Service Providers (CSPs). It was used in combination with a multi-cloud workflow management and planning application (Vampire) that enables data processing performance monitoring and migration of VMs and processes to an alternative location based on performance predictions.**

Keywords- Open Cloud eXchange (OCX); Intercloud Architecture Framework (ICAF); Intercloud Federations Framework; Vampire Application Workflow Management/Optimisation System; Intercloud Network Infrastructure Monitoring and Visualisation; Cloud Services Marketplace.

I. INTRODUCTION

The increasing complexity of scientific research conducted at universities and research organisations requires continuously increasing power of computing resources and storage volume that in most cases are required for limited period of time (e.g. for collected data processing and reporting) and to be elastically scaled. This motivates research and university community to use cloud based services [1, 2] and cloud enabled High Performance Computing (HPC) and Big Data technologies [3, 4]. However, large volumes of data used in modern research as well as required to support future Data Science education programs create new challenges for the National Research and Education Networks (NRENs) that provide connectivity for research and education institutions. When using cloud services, the typical NREN user will turn to the commodity Internet and best effort service. However, there are cases when the use of best effort Internet can no longer satisfy the demands of certain user groups (hereafter referred to as power users) and tasks that require high performance network with guaranteed Quality of Service (QoS) to ensure secure and reliable connections to potentially multiple Cloud Service Providers (CSP). Such power users need to manage complex cloud services that involve large data volumes transfers to and/or from the CSPs combined with augmented need for computing and storage resources, i.e. Big Data, analytic tools working with large data volumes [5].

At European level, the European Research and Education backbone GEANT network works to provide high performance network connectivity and advanced cloud services access and delivery infrastructure to satisfy the power users requirements such as high speed network, guaranteed QoS, security, federated access control integrated with organizational and national (at NREN level) federated Identity Management.

The Open Cloud eXchange (OCX) has been proposed by the GN3plus project JRA1 activity in 2013 [6] as a new conceptual and functional component of the general inter-cloud service delivery infrastructure with the intent to bridge the gap between the global CSP's infrastructure and NRENs and campus infrastructure, in particular, to solve the "last mile" problem in delivering cloud services to customer locations and individual (end-) users.

The paper refers to the general Intercloud Architecture Framework (ICAF) [7, 8] proposed in the earlier authors' work as a result of cooperative efforts in a number of EU funded projects and currently being submitted as an Internet-

Draft to IETF [9]. The proposed OCX solution is positioned as an important infrastructure component of the Intercloud Federation Framework (ICFF) [10].

The remainder of the paper is organized as follows. Section II describes the OCX concept and section III provides details about the OCX Architecture and major functional components. Section IV provides suggestions about the OCX design and implementation, while API design and monitoring tools are discussed in section V. Section VI describes the recent OCX demo at SC14 exhibition and summarises lessons learnt. Section VII discusses possible OCX extension with cloud services marketplace what is demanded by the GEANT community and appreciated by Cloud Services Providers. Related works are discussed in section VIII, and the paper concludes with remarks on future development in section IX.

II. Open Cloud Exchange (OCX) Concept

In order to address the currently existing problems in delivering cloud services to organizational/enterprise customers and end users, in this paper we propose the Open Cloud eXchange (OCX). Having in mind the power cloud users, the main goal of the gOCX architecture is to provide dedicated infrastructure that will bring together the CSPs and users in an efficient, fast, reliable and cost effective manner facilitating intercloud computing federations.

The proposed OCX concept is based on and extends the Internet eXchange Points (IXP) [11] and GLIF Optical Lightpath Exchange (GOLE) [12] service models with additional functionalities to allow ad hoc dynamic Intercloud federation establishment and non- restricted peering between cloud providers, customers, and also local infrastructure providers, in case cloud services delivery requires involvement of such entities.

Additionally to providing physical location for (network) interconnecting of all involved actors, the OCX declares two basic principles that simplify and facilitate services delivery:

- No value-added third party services (i.e. service composition, integration or operation). In this way, OCX will not be involved in the business dealings related to the actual cloud services provisioning and delivery;
- Trusted Third Party (TTP) services for ad-hoc/dynamic federations establishment: OCX may provide the directory service, trusted repository of provider certificates operating under supervision of the community (representatives), which can act as a policy authority for security and operational practices.

The proposed OCX role as a TTP will facilitate creation of dynamic federations and establishment of dynamic trust relation between CSPs and customers.

Referring to the generic Cloud Services Model (CSM) defined in [7, 8] as a part of the Intercloud Architecture Framework (ICAF), the OCX functionality can be related to the Intercloud Access and Delivery Infrastructure (ICADI) layer where the main goal is to deliver cloud based services to organizational customers and end users. Structurally ICADI includes all infrastructure components between the CSP, the final consumer and other entities involved into cloud services delivery and operation. However, to allow easy integration into existing cloud infrastructures and remain transparent to current service models, the OCX limits its services to Layer 0 through Layer 2 transport networks. OCX can be similarly defined as a place for inter-connection and peering between CSPs and customers. Thus, it may also benefit from being collocated with the service provider, NREN exchange points or regional data centers servicing the regional/national research community.

The introduction of the OCX TTP service intends to support the federated cloud and inter-cloud service provisioning that simplifies heterogeneous multi-provider services integration and operation. In this respect OCX will provide important functionality related to the Intercloud Federation Framework (ICFF) [10].

III. OCX Architecture and Components

Architecturally and functionally, the OCX includes the following services and functional components (see Figure 1):

- Physical Point of Presence (PoP) or OCX Access Points (OCXP) for providers and customers
- L0-L2 network interconnection facility (optionally also connectivity with dedicated optical links)
  - The associated service should allow customer related topology information exchange (such as related to virtual private clouds) between providers and customers in a secure and consistent way (this is of extreme importance since topology information in most cases is considered as commercial or restricted information)
- Trusted Third Party (TTP) services for support of dynamic peering, business/service and trust relations establishment between OCX members; the specific services may include:
  - Trusted Certificates repository and associated Trusted Introducer service to allow dynamic trust associations and/or federations establishment
  - Additionally Trust Broker service can be provided and supported by either or both Trusted Introducer and privacy/data security policy Registry or clearinghouse.

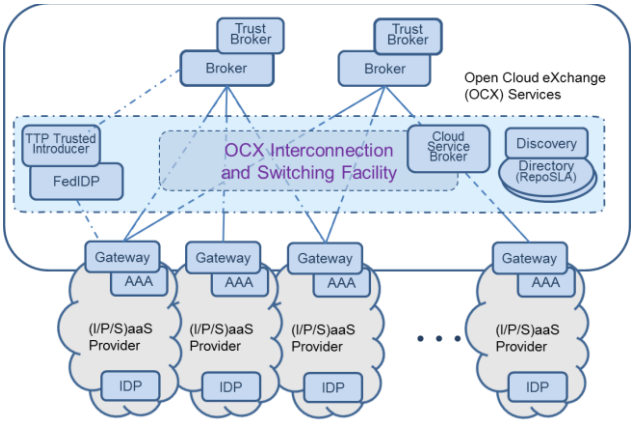

Figure 1. OCX functional component (as part of the Intercloud Federation Infrastructure).

Additionally, OCX may include services that should enable creation of the cloud services marketplace for GEANT community:

- Publish/subscribe Services Directory and Discovery; additionally the SLA Clearinghouse service can be provided.
- Optionally, Cloud Service Broker to provide service advice and integration for contracted community.

IV. OCX Design and Implementation

OCX is currently at the stage of functional design and pilot implementation of the major services to demonstrate benefits of the OCX enabled cloud services delivery infrastructure. As a conceptually new component of the inter-cloud infrastructure, OCX will require functional definition of new services, control and management interfaces that should be integrated with the current cloud management services and network providers infrastructure (also known as cloud carriers [2]). This opens a possibility to use the generic Software Defined Network (SDN) design principles [13, 14] for its implementation, i.e. the separation of the data plane that applies the rules defined by the SDN controller that implements the advanced control plane policies.

The OCX design team will also look into a possibility to use the Network Service Interface (NSI) [15] to control OCX connectivity services that is already used in the GEANT network. On the other hand, OCX management capabilities should allow their interaction with the emerging industry standard for cloud infrastructure services interoperability such as OASIS TOSCA (Topology and Orchestration Specification for Cloud Applications) [16].

The following provides suggestions for design and implementation of the functionalities described in the OCX Architecture section.

A. OCX interconnection network and peering design

Topologically OCX should allow any-to-any interconnection at Layer 0, Layer 1 and Layer 2. This can be implemented using corresponding L0-L2 optical switches. Figure 2 illustrates the switching topology of OCX, together with the TTP services operation model.

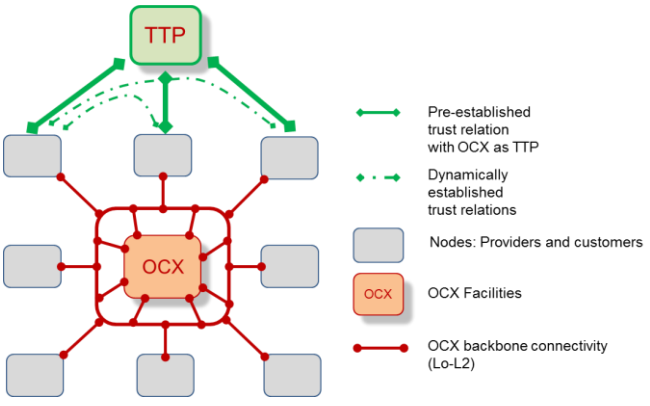

Figure 2. OCX interconnection capability and TTP role in establishing dynamic trust relations between OCX members

When deployed on top of the GÉANT/NREN network, the OCX can be seen as a hierarchical distributed system with OCX instances in multiple NRENs, and one, or several, OCX instances on the GÉANT level, which are used not only for connecting CSPs, but also for orchestration and performance optimization or load balancing purposes.

Each OCX instance is connected to the rest via backbone links engineered and dimensioned in such a way that the requested performance metrics could be guaranteed. Downstream, OCX instances connect users with CSPs that intends to offer cloud services. Finally, upon user's request, OCX will be able to provide connectivity between any two or more OCX Access Points in a secure and isolated manner, e.g. via E-LAN services. OCX access ports can multiplex various services on one port using VLANs that maintain logical traffic separation.

B. OCX SDN-based Design

The OCX operation and characteristics require fast decision-making and policy enforcement mechanisms for transparent coordination of the cloud service transactions traffic. OCX can benefit from a SDN architecture by adopting its main design principle, the separation of the control and data planes. This way, the data plane can be optimized for efficiently applying forwarding rules at any layer (L0-L2) while the SDN controller will implement features such as routing, data filtering, policy enforcement, TTP services, etc. In an SDN driven OCX architecture, the process of setting up direct connectivity between customers and CSPs can be seen as setting up slices of virtual networks connecting the client(s) and provider(s). Each CSP can setup and manage its own virtual controller that will run on top of a network virtualization layer. A modular implementation of the SDN controller such as the one offered by Floodlight [17] or even an implementation using the Network as a Service (NaaS) concepts developed by frameworks like OpenNaaS [18] provide the flexibility and extensibility that allow an easy adaptation of interconnectivity requirements.

C. OCX Trusted Third Party services

Figure 2 illustrates how OCX can operate as a TTP to establish direct/dynamic trust relations between OCX members. These trust relationships can be used for establishing identity management federations among OCX members.

The OCX trust model can contain a TTP for all members, storing their trust anchors like a trusted certificated repository in TACAR [19]. Relationships between unknown members can depend on the trust threshold values determined from other existing relationships as proposed in [20]. It is recommended that members have trust policies that define such criteria.

V. OCX API and Monitoring Tools

A. General suggestions and requirmenets

In order to efficiently use the potentially heterogeneous resources from multiple CSPs connected to OCX, customers

should have programmable access to the services offered by the participating or selected CSPs. This raises several problems. The first one is related to the fact that CSPs usually offer APIs, which are specific to their own services. This implies that users that request resources from different CSPs must express them using the respective API. The second problem is related to the way CSPs describe the types of resources available to users. For example, some CSPs offer fixed hardware templates that users can instantiate and then use, while others allow full customisation of the requested resource. Another issue is the semantics used by CSPs for describing conceptually similar concepts. Thus, users must express their request using disk templates for one CSP, while they need to use machine images for other CSPs. Although the end result is similar, i.e. a virtual block device used by the virtual machine, the provisioning process differs slightly from provider to provider. It is unrealistic to expect CSPs to adopt common semantics (or vocabulary) and APIs, despite the existing standardisation efforts to define common interfaces such as OCCI [21]. OCX must at least provide repository of available cloud resources and corresponding APIs. A more advanced OCX functionality would be to provide mapping between CSP's APIs or cloud services brokering service which would abstract the common functionality and differences between the APIs of the inter-connected CSPs. Such functionality can be built leveraging known efforts jClouds [22] or libcloud [23].

B. *Main OCX API Functions*

In order to respond to the demand for elasticity and dynamicity that cloud services present, OCX should impose a Networking API to all participating parties (primarily to CSPs) that should implement a northbound interface to their networking components. This would provide the necessary tools to facilitate OCX service setup and termination to appropriate OCX Access Points depending on user's request. The target is to provide via API an almost automated provisioning procedure to OCX members that would result with instant connectivity between the cloud customer and CSP over OCX. The API should incorporate basic authentication and authorization functionality, as part of OCX's Trusted Third Party function.

Working together with CSPs, OCX should define and implement a standard set of solutions to expose (or extend) datacenter's virtual networks to the Internet/WAN and multiplex different virtual networks belonging to different service instances over the same port(s). One of the solutions is to use different VLANs, but there are certain scenarios that ask for other solutions. This networking interface should be exposed and managed over the mentioned API.

When OCX runs as a service on top of the GÉANT/NRENs infrastructure, multiple already implemented technologies can be used such as L2 P-2-P MPLS VPNS, VPLS, Bandwidth on Demand (BoD), MD-VPN, and SDN), all of which need to be coordinated with the individual CSPs.

C. *Intercloud Network Infrastructure Monitoring and Visualisation*

The Next Gen Cloud Management System (NG-CloudMS) has been developed by Opt/Net [24] to provide inter-cloud infrastructure monitoring and visualisation for OCX. Its prototype version provides near real-time visibility of the networks and cloud infrastructures and interconnected computing and data storage resources that are part of the OCX enabled Intercloud architectures.

NG-CloudMS is connected to the OCX controller and automatically maps and monitors addition and deletion of CSPs, network links and nodes to the managed architecture. It collects the most complete information about the OCX Cloud network's inventory, topology, mapping of IP address space and also provides analysis of syslog events and SNMP alarms both in near real time and from the historical archives. Figure 3 provides example of the OCX enabled inter-cloud topology visualization.

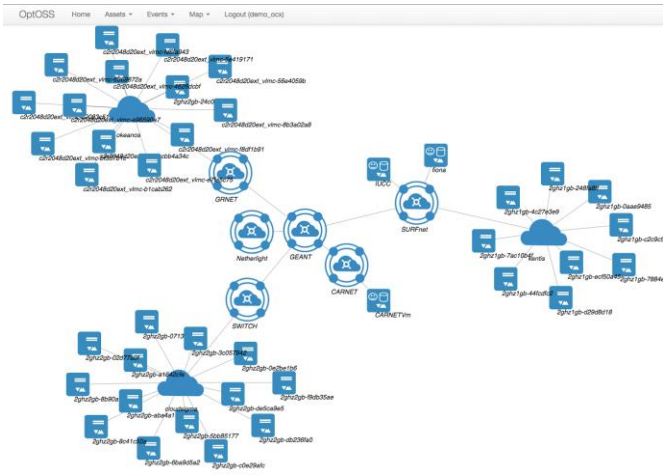

Figure 3. Real-time map of the OCX enabled Intercloud network used during SC14 demo

NG-CloudMS operates autonomously once it has been initiated and configured by the OCX operator. It consists of the following modules: network audit and host polling modules, central database, event collector modules and device specific plugin modules, Web GUI. The system should be configured with authentication methods for the entire managed domain. All inter-cloud components are configured to send syslog messages to the NG-CloudMS via UDP or TCP network protocols, where these messages are processed and profiled by event collector modules in near real-time. All VM templates must include such configuration and authentication settings.

In order to discover and conduct complete inventory of the inter-cloud infrastructure, NG-CloudMS needs network topology information, which in our OCX testbed and demo implementation is created and maintained by the Vampire Application Workflow Management/Optimization System (VAWMOS) [25]. When new VMs are dynamically spawned by the VAWMOS, the infrastructure update cycle is triggered.

The information is provided in JSON table format and is retrieved by NG-CloudMS automatically, as soon as it receives a notification message that the inter-cloud infrastructure has changed. Following this, NG-CloudMS initiates the network audit cycle and rediscovers the entire cloud topology. The network discovery may be partial or complete, depending on the type of change that took place.

All discovered nodes are monitored continuously and periodic inventory of the complete network takes place. Events from managed nodes are continuously received by collectors and stored in the central database for archival purposes and analysis. The cumulative severity of the received messages may be plotted using the Web GUI. This graphical representation facilitates interpretation of the network and device activity and simplifies search for the causes of different events. The historical view may be used for data analytics and search for the root causes of different events. Often, it helps with troubleshooting of different problems on the network and associated cloud services.

NG-CloudMS is an Open Source project hosted on SourceForge [24] and distributed under GPL3.0 license. This guarantees that NG-CloudMS will benefit free software and research and education communities.

VI. OCX Demonstration

In order to highlight the benefits of using OCX and evaluate the best options on setting up the OCX infrastructure, a multiuser test scenario has been defined and presented live at the SC14 Exhibition [26]. The users are opting to use OCX for the purposes of obtaining a joint service from multiple CSPs. In order to make a qualitative and quantitative comparison of both approaches, the scenario is compared to the traditional commodity Internet approach.

A. SC14 Demo topology

The SC14 demo topology is illustrated in Figure 4. It involved GEANT multi-gigabit backbone facilities, NRENs SURFnet, GRNET, SWITCH, CARNET hosting OCX Access Points interconnected via GEANT backbone, cloud providers involved into the main demo scenario Okeanos, CloudSigma, Kentis and Amazon Web services and Microsoft Azure connected via regular Internet. Participating universities include University of Amsterdam, St Cyril and Methodius University in Skopje, Inter-University Computation Center (Israel). Opt/Net provided their NG-CloudMS application for OCX infrastructure monitoring and visualization.

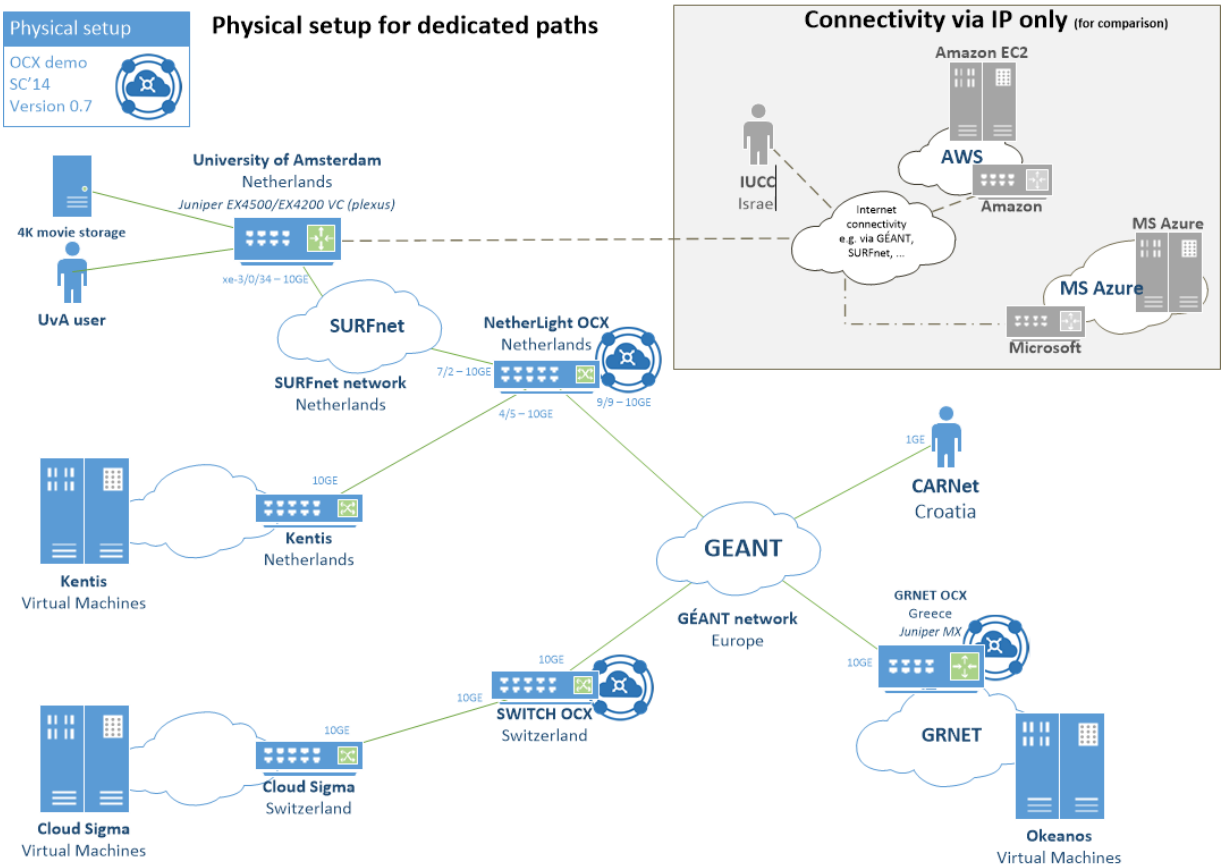

Figure 4. OCX demo scenario and topology at SC14.

The demonstration scenario is based on one of the main use cases that would benefit from dedicated connections to CSPs such as HD video streaming and real-time compression. The demo scenario is defined as follows: two or multiple institutions would like to collaborate for editing, compressing and then viewing HD video content. The video editing and compression is done using IaaS providers (in our case, Okeanos connected via GRNET, CloudSigma connected via SWITCH, and Kentis connected via SURFnet and Lighthouse GOLE) connected to OCX. The result of the video manipulation is then sent to one of the participating institutions. The application workflow, resource reservation and application execution is handled by Vampires [36] a multi-cloud scheduler and batch processing framework.

In a traditional approach using the public Internet for interconnecting to the cloud providers, after the institutions send the videos to the cloud providers, the resulting video has to be sent back over the Internet to one of the institutions that will subsequently view it. All of the data transfers are done over the public Internet using the best effort approach without any QoS guarantees or special traffic isolation.

The customer networks connect to their local OCX Access Points, which enable them to get direct access to the necessary cloud resources via high performance dedicated network links. Additionally, on-demand network connectivity using OGF NSI [15] connection service has been demonstrated as it is implemented by Kentis and SURFnet.

The NG-CloudMS is running on Ubuntu 14.04 LTS and can run on just one core in a virtual machine requiring only 1GByte of RAM (at the very minimum) as was shown during SC14 demo. The larger networks require more computing power and storage. It may be deployed on the real server hardware or virtualized into the cloud as one of the services.

B. Lessons learnt

The major benefit of the OCX infrastructure is the use of dedicated links towards the cloud providers that will substantially improve the data transfer performances between the users and the CSPs by completely mitigating the public Internet. The use of multi-domain VLAN services ensures traffic isolation. Also, all already available layer 2 connections can be reused for future cloud service delivery. After the initial setup, the customers can use the provided cloud service transparently. Furthermore, by enabling end-to-end connectivity, a number of additional performance enhancements (compared to the traditional best effort Internet approach) can be implemented, including jumbo frames, bypassing firewalls/policies, using private addresses to Cloud VMs/networks, etc. Moreover, the performances of the system are not influenced by different MTU sizes or number of firewalls that are traversed since all links are established below L3.

The presented scenario has highlighted the benefits of direct connectivity, but also underlined the tedious task of setting up all direct connection links needed to interconnect the parties involved. Additionally, when performing the connections setup manually, there is lots of room for human errors and prolongations due to poor synchronization and misunderstandings. Thus, the test has also revealed that the true benefits of the OCX architecture will be available, usable and transparent to the end-users only if the connection setup process is done automatically, preferably by virtualizing all OCX instances and using a web service based platform.

Automation using an OCX inter/intra connectivity management web portal will allow NREN customers to choose among different CSP offers, setup their preferred connections, and manage and monitor the demanded service, thus setting up a complete autonomous marketplace. Additionally, by employing automated process CSPs can dynamically setup their offerings on top and monitor their subscribers' activity, too.

VII. OCX AS A MARKETPLACE

The implemented OCX concept can be used as a basis for a cloud services marketplace for the European Research and Education community that can offer cloud service directory where the connected CSPs publish their services (including SLA, API and X.509 certificates), "connectivity as a service" where customers can setup on-demand connectivity to CSPs, SLA Repository and Clearing house.

The concept of a generic marketplace needs to provide the following features:

- The cloud services catalogue, which should include registered cloud services that can be accessed via federated OCX infrastructure; the catalogue should also register a new services deployed by OCX customers; optionally also cloud brokering service;
- Possibility to integrate third party, where the machine interface needs to enable automated access to the market place for cloud services aggregation;
- The customer access to the marketplace and CSP services should be provided both through a web portal and command line interface (CLI);
- The northbound OCX APIs should be open, possibly using commonly accepted standards (e.g. REST, OCCI) covering CSPs capabilities and solving user requirements;
- A trusted Single Sign On (SSO) to allow customers use their Federated Identity Management infrastructure.

Important for a successfully implementation is that the market place is technology agnostic and flexible enough to cope with the most requirements/features of the end-users/CSPs, otherwise they would prefer direct peering to aggregate and provide cloud services with their chosen CSPs. Automation of the OCX enabled inter-cloud services provisioning will be a key factor in marketplace operation and usability.

VIII. RELATED DEVELOPMENTS AND SERVICES

The proposed OCX architecture and service model is built upon successful services like Internet Exchange (IXP) [11]

for general Internet traffic exchange and GOLE (GLIF Open Lightpath Exchange) [12] that provides lightpath interconnection service. In the following we provide a short reference to the GOLE and review some works related to other OCX functionality.

The OCX intends to fill the gap between global CSP's infrastructure with limited number of regional access points and existing GEANT and NREN and campus infrastructure that provide network access to the active cloud services user community. In this respect, OCX targets to (inter)connect to CSP's dedicated connectivity services such as AWS DirectConnect [27] and Azure ExpressRoute [28] that are offered for customers that require high bandwidth and dedicated connectivity.

Using these services the customer can increase the bandwidth throughput and have more consistent network experience compared to the Internet-based connections, while reducing the network costs. However, these services are generically Layer 3, while offering partitioning of the direct connection into multiple virtual interfaces using IEEE 802.1q VLANs.

The Equinix Cloud Exchange [29] is an interconnection solution that enables on-demand, direct access to multiple clouds from multiple networks in a given number of Equinix datacenter locations around the world that hosts multiple CSPs. It provides virtualized, private connections with high performances that can be created and managed via the Equinix Cloud Exchange portal. While this solution seems to follow the same idea and motivation as OCX, the Equinix Cloud Exchange services are limited to a datacenter location. The project team is in discussion with Equinix about possible collaboration to extend reach of OCX/GEANT community to CSPs worldwide.

A number of research and developments are available in NREN and GEANT community to operate federated access control services where the community has recognized leadership. The Moonshot Project [30] develops a single unifying technology for extending the benefits of federated identity to a broad range of non-Web services, including Cloud infrastructures, High Performance Computing & Grid infrastructures and other commonly deployed services like mail, file store, remote access and instant messaging. The project implements the technology developed by the IETF Working Group Application Bridging for Federated Access Beyond web (ABFAB) [31].

The OpenStack KeyStone project [32] provides Identity, Token, Catalog and Policy services for use specifically by projects in the OpenStack family. We consider it as a candidate platform for OCX TTP implementation that can also integrate solutions proposed in the research work by the University of Kent [20] and also integrate the authors' earlier works on trust establishment trust based policy evaluation [33] and trust bootstrapping protocol [34].

IX. Conclusion and Future Development

This paper presents an on-going research and development of the Joint Research Activity JRA1 in the GN3plus project conducted by a group of cooperating universities and NRENs to develop the Open Cloud eXchange (OCX) – a new service and component of the Intercloud Architecture that addresses problems with multi-domain heterogeneous cloud based applications integration and inter-provider and inter-platform interoperability.

The current stage of development concludes the OCX architecture, functional design and pilot implementation has been successfully fulfilled. The next stage will include detailed services and API design and validation (in particular, modeling the OCX network infrastructure, monitoring and operational models) what will be in the framework of future GEANT4 projects that will start from April 2015. The future development will also address security and Federated Identity Management issues in integrated (multi-)provider and campus cloud infrastructure.

OCX intends to provide a basis to support cloud based collaborative infrastructure for emerging new applications, in particular Big Data infrastructure for universities and research organizations that should support data intensive research domains and applications like: particle physics LHC (Large Hadron Collider) experiments, SKA (Square Kilometer Array) astronomy observations, genomics, climate research, etc.

The proposed approach and definitions are intended to provide an input to standardization activities in the area of Intercloud architecture and services.

Acknowledgement

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7 2007–2013) under Grant Agreement No. 238875 (GÉANT). Intercloud federation and trust management research and development are partly supported by the Horizon2020 project CYCLONE.

References

[1] NIST SP 800-145, "A NIST definition of cloud computing", [online] http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

[2] NIST SP 500-292, Cloud Computing Reference Architecture, v1.0. [Online] http://collaborate.nist.gov/twiki-cloud-computing/pub/CloudComputing/ReferenceArchitectureTaxonomy/NIST_SP_500-292_-_090611.pdf

[3] NIST Big Data Working Group Documents [online] http://bigdatawg.nist.gov/V1_output_docs.php

[4] Demchenko, Yuri, Peter Membrey, Cees de Laat, Defining Architecture Components of the Big Data Ecosystem. Second International Symposium on Big Data and Data Analytics in Collaboration (BDDAC 2014). Part of The 2014 International Conference on Collaboration Technologies and Systems (CTS 2014), May 19-23, 2014, Minneapolis, USA

[5] The Adoption of Cloud Services. TERENA. ASPIRE Report, September 2012. http://www.terena.org/publications/files/ASPIRE%20-%20The%20Adoption%20of%20Cloud%20Services.pdf

[6] Demchenko, Yuri, Jeroen van der Ham, Canh Ngo, Taras Matselyukh, Sonja Filiposka, Cees de Laat, Open Cloud eXchange (OCX): Architecture and Functional Components. Proc. "The 3rd workshop on

Network Infrastructure Services as part of Cloud Computing (NetCloud 2013)", in conjunction with The 5th IEEE International Conference and Workshops on Cloud Computing Technology and Science (CloudCom2013), 2-5 December 2013, Bristol, UK

[7] Demchenko, Y., C.Ngo, M.Makkes, R.Strijkers, C. de Laat, Intercloud Architecture for Interoperability and Integration. Proc. The 4th IEEE Conf. on Cloud Computing Technologies and Science (CloudCom2012), 3 - 6 December 2012, Taipei, Taiwan. IEEE Catalog Number: CFP12CLU-USB. ISBN: 978-1-4673-4509-5

[8] Demchenko, Y., C.Ngo, C. de Laat, J.A.Garcia-Espin, S.Figuerola, J.Rodriguez, L.Contreras, G.Landi, N.Ciulli, Intercloud Architecture Framework for Heterogeneous Cloud based Infrastructure Services Provisioning On-Demand. The 2nd Intl Workshop on inter-Clouds and Collective Intelligence (iCCI-2013). The 27th IEEE International Conference on Advanced Information Networking and Applications (AINA2013). 25-28 March 2013. ISBN-13: 978-0-7695-4953-8.

[9] Cloud Reference Framework. Internet-Draft, version 0.6, January 4, 2013. [online] http://www.ietf.org/id/draft-khasnabish-cloud-reference-framework-06.txt

[10] Makkes, M., C.Ngo, Y.Demchenko, R.Strijkers, R.Meijer, C. de Laat, Defining Intercloud Federation Framework for Multi-provider Cloud Services Integration, The Fourth International Conference on Cloud Computing, GRIDs, and Virtualization (CLOUD COMPUTING 2013), May 27 - June 1, 2013,Valencia, Spain.

[11] Promoting the Use of Internet Exchange Points: A Guide to Policy, Management, and Technical Issues, FInternet Society Report. 14 May 2009 [online] http://www.internetsociety.org/promoting-use-internet-exchange-points-guide-policy-management-and-technical-issues

[12] The GLIF "Automated GOLE Pilot" Project. http://staff.science.uva.nl/~delaat/sc/sc10/GLIFAutomatedGOLEPilot.SC.pdf

[13] Software-Defined Networking: The New Norm for Networks, ONF White Paper. April 13, 2012 https://www.opennetworking.org/images/stories/downloads/sdn-resources/white-papers/wp-sdn-newnorm.pdf

[14] OFELIA and GÉANT Cooperation on OpenFlow Experimental Facilities. Posted 22 August 2013. [online] http://www.fp7-ofelia.eu/news-and-events/press-releases/ofelia-and-gEant-cooperation-on-openflow-experimental-facilities/

[15] GFD.173 Network Services Framework v1.0, OGF Standard [online] http://www.gridforum.org/documents/GFD.173.pdf

[16] Topology and Orchestration Specification for Cloud Applications, Version 1.0. Candidate OASIS Standard. 11 June 2013. [online] http://docs.oasis-open.org/tosca/TOSCA/v1.0/TOSCA-v1.0.html

[17] Floofligh OpenFlow SDN Controller [online] http://www.projectfloodlight.org/floodlight/

[18] OpenNaaS: Open platform for Network as a Service resources [online] http://www.opennaas.org/

[19] TERENA Academic Certification Authority Repository. [online] https://www.tacar.org/

[20] Chadwick, D., M.Hibbert, Towards Automated Trust Establishment in Federated Identity Management. Proc. The 7th IFIP WG 11 International Conference on Trust Management (2013), Malaga, Spain.

[21] Open Cloud Computing Interface - OCCI [Online]

[22] Apache Libcloud [online] https://libcloud.apache.org/

[23] Apache Jclouds [online] https://jclouds.apache.org/

[24] Opt/Net NG-CloudMS [online] https://sourceforge.net/projects/ngcms/

[25] Dumitru, C. , Oppresscu, AM., Živković, M. , R van der Mei, Grosso P. , de Laat, C. A queueing theory approach to Pareto optimal bags-of-tasks scheduling on clouds, Euro-Par 2014 Parallel Processing, pp. 162-173.

[26] Demchenko, Y., C. Dumitru, R. Koning, C. de Laat, M. de Vos, D. Regvart, T. Karaliotas, K. Baumann, D. Arbel, S. Filiposka, T. Matselyukh, GEANT Open Cloud eXchange (gOCX): Architecture, Components, and Demo Scenario, SC14, New Orleans, USA, November 2014

[27] Amazon Direct Connect service [online] http://aws.amazon.com/directconnect

[28] Azure ExpressRoute [online] http://azure.microsoft.com/nl-nl/services/expressroute/

[29] Equinix Cloud Exchange [online] http://www.equinix.com/services/interconnection-connectivity/cloud-exchange/

[30] Moonshot Project [online] https://community.ja.net/groups/moonshot

[31] IETF Application Bridging for Federated Access Beyond web (Active WG) [online] http://tools.ietf.org/wg/abfab/

[32] Keystone, the OpenStack Identity Service! [online] http://docs.openstack.org/developer/keystone/

[33] Ngo, C., Y.Demchenko, C. de Laat, Toward a Dynamic Trust Establishment Approach for Multi-provider Intercloud EnvironmentThe 4th IEEE Conf. on Cloud Computing Technologies and Science (CloudCom2012), 3 - 6 December 2012, Taipei, Taiwan

[34] Membrey, P., K.C.C.Chan, C.Ngo, Y.Demchenko, C. de Laat, Trusted Virtual Infrastructure Bootstrapping for On Demand Services. The 7th International Conference on Availability, Reliability and Security (AReS 2012), 20-24 August 2012, Prague.