Networking solutions in the federation of clouds

An approach implemented by the NOVI project and dynamic circuit systems as the extensions for inter-cloud connectivity

Roman Łapacz, Błażej Pietrzak
Poznań Supercomputing and Networking Center
ul. Wieniawskiego 10, 61-704 Poznań, Poland
{romradz, blazej.pietrzak}@man.poznan.pl

***Abstract*—This paper presents the cloud federation concept implemented by the NOVI project and ideas of possible extensions. The authors pay much attention to the networking aspect and available bandwidth provisioning systems which offer reliable network transfer services to the cloud systems and their federations.**

Keywords—cloud; federation; infrastructures, networking; bandwidth on demand; circuit provisioning

I. INTRODUCTION

The growing interest in cloud computing results in an increasing number of cloud infrastructures offering various resources and services to their users. At the same time the applications are getting more demanding and expect high levels of resource parameters. It may be predicted that resources of single clouds will not suffice to run the complex tasks. In such a case automated and standardised communication between infrastructures is required to share the resources. To address these trends and offer new solutions for building innovative cloud-based services there are initiatives working on cloud federation aspects, including inter-cloud communication. This paper presents an implementation of a federation solution for independent virtualized resource infrastructures, proposes an architecture approach and possible extensions which enhance inter-cloud communication functionality on the networking level.

II. FEDERATION OF CLOUDS

The Future Internet (FI) emerges as a complex ecosystem, offering services to users over shared federated environments. FI infrastructures should provide dynamic provisioning and interconnection of resources belonging to independent heterogeneous infrastructures [1]. Early efforts concentrated on creating infrastructures for running disruptive experiments like Virtual Network Infrastructure (VINI) [2], PlanetLab [3]. Emulab [4] and ProtoGENI [5] in the US GENI NSF initiative [6], and in Europe in FIRE EC initiative [7] FEDERICA [8], PanLab [9], and OneLab [10].

The current efforts are focused on developing tools and services to federate virtualized infrastructures, including Slice-based Federation Architecture (SFA) [11] adopted by GENI in US and the OneLab/OneLab2 EC project as well as Teagle architecture [12] employed within the European PanLab infrastructure.

Federation of infrastructures requires federated Authorization and Authentication Infrastructure that can scale. Centralized approaches may not be accepted due to technical and security problems of sharing crucial information belonging to the underlying substrates. If the policy-based management approach [13] is applied then sharing the policies among separate domains must be addressed by federated policy management services. Another crucial element of the federated environment is the ability to provision resources belonging to various infrastructures. There is a need for a distributed database of resources. Resources may need to interact with other resources belonging to remote domains. Federated infrastructure should provide this ability at the data plane level.

III. THE NOVI ARCHITECTURE

The EC FP7/FIRE STREP project NOVI [14] – Networking innovations Over Virtualized Infrastructures – provides tools and services for an efficient approach to compose virtualized e-infrastructures managed by separate yet interworking providers. These infrastructures may offer heterogeneous resources belonging to various levels such as networking, storage and processing. NOVI offers services allowing users to request for resources, control them, and run disruptive experiments without affecting other users' experiments.

The high level overview of NOVI architecture reported in [15] is shown in Fig. 1. It consists of three layers:

1. The Infrastructure Layer that contains virtualized resources ready to be allocated to users' requests.
2. The Integration Framework that offers facilities employing strong modularity and allows integrating components written in various programming languages.
3. The NOVI Service Layer that provides Control & Management (C&M) services offering advanced functionalities to the federation. Data plane connectivity within NOVI is achieved using the NSwitch software component.

Infrastructures offer provisioning of their resources via publicly available application programming interfaces (APIs) encapsulating the underlying implementation details. NOVI minimizes the footprint that is required to provide federated environment and exploits the facilities APIs to reserve resources.

The Integration Framework offers common facilities allowing components written possibly in different programming languages to cooperate with each other. The Services Layer runs within the Java Runtime Environment (JRE) [16] process, although some components are written in the Python language as well. The Integration Framework offers a Python interpreter that allows for running components written in Python inside the Java process. The logging facility provides a uniform service for logging important system information which makes an administrator's life easier. The components communicate synchronously as well as asynchronously depending on their needs. The Integration Framework offers both methods providing a dedicated event queue for the asynchronous communication. Orchestration of all components is also one of the responsibilities offered by The Integration Framework. The components only specify what they need and the process of wiring is performed automatically by the framework.

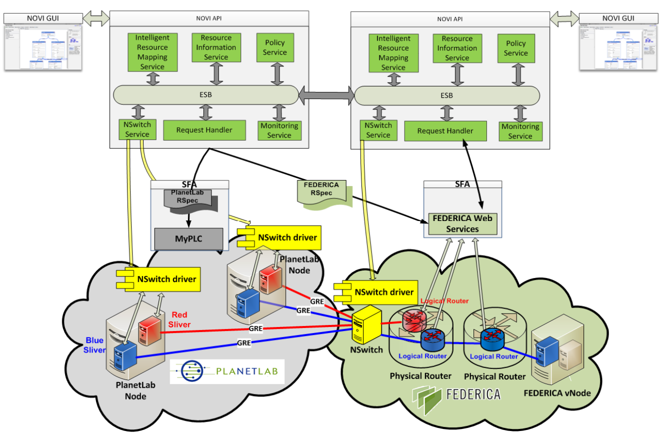

Fig. 1 : The NOVI architecture [15].

NOVI API provides a facade to all C&M services. It accepts requests incoming from authenticated users, delegates them to appropriate NOVI components and provides a feedback to users on how their request is being handled. Users can follow the process of creating the requested topologies in real time. API may receive requests from NOVI GUI or from the clients using the Representational State Transfer (REST) API interface.

The Resource Information Service (RIS) is responsible for providing information about all resources available for users. It acts as a single point of contact for the whole NOVI service layer to acquire information about the status of virtual and substrate resources. The resource discovery locates and retrieves information across all federated substrate networks in a distributed way allowing for a scalable query process.

The Intelligent Resource Mapping (IRM) service maps the requested topology into underlying virtualized substrate resources. The inter-domain Virtual Network Embedding (VNE) problem can be divided into the following sub-problems [17]:

- Splitting the request between infrastructure providers
- Solving the intra-domain VNE problem.

Intra-domain mapping algorithms tailor-made for the specific underlying substrate allow to optimize resource utilization. The IRM service uses the RIS service to find available resources for mapping. Then it uses an infrastructure specific embedding algorithm to choose map optimal resources. Each infrastructure provider must extend the NOVI service layer with an appropriate implementation of the embedding algorithm.

The policy service (PS) manages services by control policies, following the policy-based management approach [18]. It supports access control policies that specify the authorization rights to resources and event-condition-action policies triggering actions when the event under a given condition happens like a resource failure. Inter-domain policies are defined in the form of mission policies. The authentication service is testbed-specific and must be implemented by the infrastructure provider.

The monitoring service provides information about substrate and slice monitoring. The heterogeneous federated environment consists of a diverse combination of monitoring tools deployed on different infrastructures. This service provides a uniform way to access them. It provides hooks and configuration files that must be updated by infrastructure providers so that the monitoring service can be applied on the substrate.

The Request Handler (RH) service is an intermediate component between the NOVI services and the underlying substrate resources exposed via public infrastructure APIs. The responsibility of RH is to translate NOVI requests into infrastructure public API and analyze the responses received from the testbeds. NOVI implemented SFA RH to communicate with PlanetLab and FEDERICA testbeds that provide their functionalities via SFA API [19].

The NOVI Distributed Virtual Switch (NSwitch) provides a unified way of interaction between heterogeneous domains at the data plane level. It considers concurrence, isolation and programmability aspects, and allows to connect virtual entities belonging to different domains at layer 2. Each platform provider must implement the NSwitch service dedicated to their platform.

As a proof of concept NSwitch was developed for FEDERICA and PlanetLab infrastructures. These two testbeds use dissimilar communication protocols and virtual machines hypervisors. PlanetLab does not provide data-plane connectivity options. It relies on legacy Internet protocols like IP/BGP that cannot be overridden by the user, while FEDERICA offers Juniper logical routers and Ethernet connectivity based on L2 VLAN technology. In order to federate PlanetLab with FEDERICA resources Ethernet over the GRE (EGRE) [20] protocol was used to provide point-to-point virtual network functionality to a virtual resource over the Internet. Performance evaluation measurements indicated that this stitching approach provides minimal performance degradation compared to physical (substrate) entities [21].

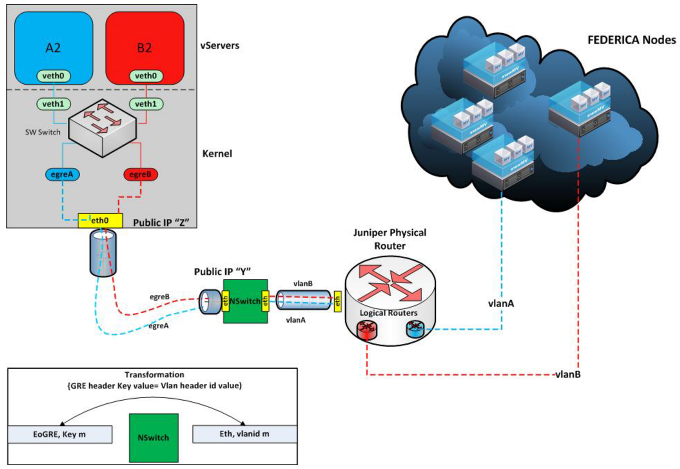

Fig. 2 : L2 federation example [13].

Due to Juniper MX 480 router limitations deployed at the FEDERICA side, GRE tunnel interfaces could not be tagged with VLAN IDs. Therefore, an additional NSwitch node was deployed on FEDERICA substrate to translate GRE tunnels from PlanetLab into VLAN IDs at the FEDERICA side and vice versa as depicted on Fig. 2 .

IV. DYNAMIC CIRCUIT SYSTEMS AND THEIR POTENTIAL

In the recent years one can observe fast progress of works on advanced middleware systems which automate establishing network circuits in multi-domain environments. Those efforts were mainly driven by two factors: 1) presence of applications that transfer high volumes of data and require a guaranteed minimum bandwidth and a minimized level of packet loss (e.g., The Large Hadron Collider, LHC, producing on average 15 PBytes of data a year, with data distribution around the world), and 2) integration of independent systems within international collaboration initiatives (e.g., Open Grid Forum - OGF [22]) in order to provide multi-domain seamless end-to-end connection services in heterogeneous networks. For the moment the NOVI solution for connecting independent infrastructures is based on basic tunneling configurations but the architecture of the NSwitch networking component allows to implement an interface to one of the existing circuit provisioning systems (Fig. 3). Such an interface would be an abstract layer with control mechanisms to manage an underlying network infrastructure that is composed of devices of different vendors, implemented in different technologies and deployed in different network domains.

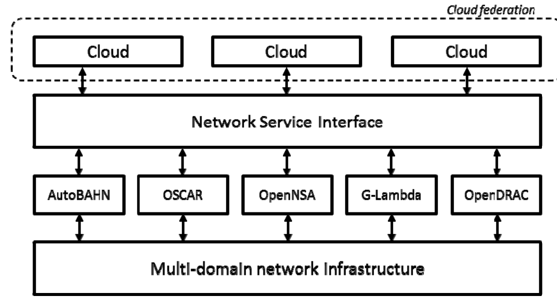

Fig. 3 : Circuit provisioning systems between clouds and multi-domain network infrastructure

The networking component of NOVI does not have to communicate with provisioning systems directly. The standardization community OGF offers a solution called Network Service Interface (NSI) [23] that defines standardized communication. In general, the main objective of NSI is to provide a unified communication method enabling independent single domain resource management tools to collaborate at the global scale providing multi-domain services in heterogeneous environments [24]. For the moment, NSI is supported by a set of provisioning systems, i.e. AutoBAHN [25], OpenDRAC [26], DynamicKL, G-Lambda [27], OpenNSA [28], and OSCARS [29].

The benefits of provisioning tools that can be offered to the clouds are also known to BonFIRE [30]. This is an EU-funded project that provides facilities for large scale experimentally-driven cloud research (see Fig. 4). One of its networking scenarios is based on connections between resources of independent clouds created by a provisioning tool. The scenario is implemented by AutoBAHN, a GÉANT [31] system for establishing dynamic circuits over the global research and education (R&E) network infrastructures.

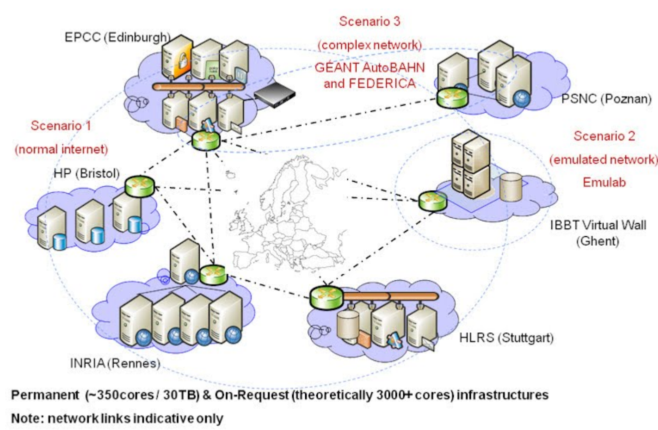

Fig. 4 : The BonFIRE infrastructure [30]

The work on integrating highly distributed and complex systems, like clouds, with provisioning tools shows the potential in this research scope. Also new ideas and efforts [32] integrating Software-defined networking (SDN)/OpenFlow [33][34] with NSI may bring valuable results that will be used in the future to improve networking in federated clouds.

V. SUMMARY

In this paper we have briefly identified federation challenges and possible solutions. As an example we have presented the architecture of the EC FP7/FIRE STREP project NOVI. It provides services for managing topologies created from virtualized heterogeneous resources belonging to separate yet interworking providers.

One of the important aspects of the federated environment is network stitching. NOVI addresses this aspect by introducing the NOVI Distributed Virtual Switch (NSwitch). It provides data plane connectivity for resources belonging to different domains. As a proof of concept PlanetLab and FEDERICA resources were federated. These two testbeds use different communication protocols and virtual machines hypervisors. PlanetLab relies on legacy internet protocols like IP/BGP that cannot be overridden by the user, while FEDERICA networking resources and Ethernet connectivity based on L2 VLAN technology. PlanetLab and FEDERICA resources were federated using Ethernet over GRE (EGRE) protocol. Due to limitations on the Juniper MX 480 router provided in FEDERICA a dedicated NSwitch node had to be deployed to translate GRE tunnels from PlanetLab into VLAN IDs in FEDERICA to solve the problem of tagging tunnel interfaces.

The architecture of NOVI allows to adopt the existing dynamic circuit-based systems. These systems are an interesting possibility that should be taken into account when facing the problem of creating the data-plane level for the federated architecture. They may require additional work with installation, configuration and possibly extension for the specific infrastructure.

ACKNOWLEDGMENT

The NOVI project was funded by the EC in the 7th framework programme, Grant Agreement No. 257867.

The work on NSI and AutoBAHN performed partially in the GN3 project, was funded by the EC in the 7th framework programme, Grant Agreement No. 238875.

REFERENCES

[1] Lymberopoulos, L., Grosso, P., Papagianni, C., Kalogeras, D., Androulidakis, G., van der Ham, J., de Laat, C., Maglaris, V.: Managing federations of virtualized infrastructures: A semantic-aware policy based approach. In Proc. of Integrated Network Management Symposium (IM) 2011. pp.1235-1242

[2] VINI – A VIrtual Network Infrastructure, http://www.vini-veritas.net

[3] PlanetLab, http://www.planet-lab.org

[4] Emulab, http://www.emulab.net

[5] ProtoGENI - Prototype implementation and deployment of Global Environment for Network Innovations, http://www.protogeni.net

[6] GENI - Global Environment for Network Innovations, http://www.geni.net

[7] FIRE – Future Internet Research & Experimentation, http://cordis.europa.eu/fp7/ict/fire

[8] FEDERICA – Federated E-infrastructure Dedicated to European Researchers Innovating in Computing network Architectures, http://www.fp7-federica.eu

[9] PanLab, http://www.panlab.net

[10] OneLab, http://www.onelab.eu

[11] Slice Federation Architecture (SFA), http://groups.geni.net/geni/wiki/SliceFedArch

[12] Teagle, http://www.fire-teagle.org

[13] Sloman,M.S. , Policy Driven Management for Distributed Systems, Journal of Network and Systems Management, 1994, Vol:2, Pages:333-360, ISSN:1064-7570

[14] NOVI – Network Innovation over Virtualized Infrastructures FP7 STREP Project, http://www.fp7-novi.eu

[15] L. Lymberopoulos, M. Grammatikou, M. Potts, P. Grosso, A. Fekete, B. Belter, M. Campanella and V. Maglaris, "NOVI Tools and Algorithms for Federating Virtualized Infrastructures", in Future Internet – From Technological Prom-ises to Reality, Springer Lecture Notes in Computer Science, pp. 213-224, 2012.

[16] Java SE Overview, http://www.oracle.com/technetwork/java/javase/overview/index.html

[17] I. Houidi, W. Louati, W. B. Ameur, D. Zeghlache, "Virtual network provisioning across multiple substrate network" ELSEVIER Computer Networks, vol.55, no. 2, pp. 1011-1023, 2011, dx.doi.org/10.1016/j.comnet.2010.12.011.

[18] Sloman,M.S. , Policy Driven Management for Distributed Systems, Journal of Network and Systems Management, 1994, Vol:2, Pages:333-360, ISSN:1064-7570

[19] Generic SFA Wrapper, http://sfawrap.info/

[20] RFC 2784, Generic Routing Encapsulation (GRE). D. Farinacci, T. Li, S. Hanks, D. Meyer, P. Traina. March 2000.

[21] C. Argyropoulos, G. Androulidakis, D. Kalogeras, B. Pietrzak, B. Belter, L. Lymberopoulos, V. Maglaris, "Network Virtualization over Heterogeneous Federated Infrastructures: Data Plane Connectivity", IFIP/IEEE Integrated Network Management Symposium (IEEE IM 2013), Ghent, Belgium, May 2013.

[22] Open Grid Forum (OGF), http://www.ogf.org/

[23] Network Service Interface (NSI), http://redmine.ogf.org/tab/show/nsi-wg

[24] R. Krzywania, at al., "Network Service Interface – gateway for future network services", Terena document, Feb 2012.

[25] M. Büchli, at al., "Deliverable DJ.3.3.1:GÉANT2 Bandwidth on Demand Framework and General Architecture", GÉANT, 2005

[26] OpenDRAC, https://www.opendrac.org

[27] G-Lambda, http:// www.g-lambda.net

[28] OpenNSA, https://github.com/jeroenh/OpenNSA

[29] OSCAR, http://www.es.net/services/virtual-circuits-oscars/

[30] BonFIRE, http://www.bonfire-project.eu/

[31] GEANT, http://www.geant.net/

[32] Belter B., T. Kudoh, R. Krzywania, G. van Malenstein , "Federating SDN-enabled islands with an extanded NSI Framework", TNC 2013

[33] Open Networking Fundation, https://www.opennetworking.org/

[34] OpenFlow, http://www.openflow.org/