Federating SDN-enabled islands with an extended NSI Framework

Bartosz Belter[1], Radosław Krzywania[1], Tomohiro Kudoh[2], Gerben van Malenstein[3]

[1] Poznan Supercomputing and Networking Center, ul. Z. Noskowskiego 12/14, 61-704 Poznań, Poland,
Email: bartosz.belter@man.poznan.pl; radek.krzywania@man.poznan.pl

[2] National Institute of Advanced Industrial Science and Technology (AIST), 1-1-1 Umezono Tsukuba 305-8568 Japan
Email: t.kudoh@aist.go.jp

[3] SURFnet, P.O. Box 19035, 3501 DA Utrecht, The Netherlands,
Email: Gerben.vanMalenstein@surfnet.nl

Keywords

Software Defined Networking, OpenFlow, Network Services Interface, Connection Service, Federation

1. Introduction

Software Defined Networking (SDN) [1] is an emerging architecture for computer networking. SDN separates the control plane from the data plane in switches and routers. Under SDN, the whole logic and control is moved to the control plane, which is implemented as a software layer and separated from the network equipment. The network equipment remains responsible for packet/flow switching on the data plane. OpenFlow [2][3] is a defined standard communication protocol used between the control and data planes in an SDN architecture.

This paper describes a novel framework capable of provisioning end-to-end network services between resources attached to SDN clouds (e.g. data centres) controlled by OpenFlow or similar protocol. This paper proposes extensions in the form of a new NSI service, which will integrate SDN and NSI-CS messaging and functionality, delivering adequate mechanism for dynamic management of remote slices and corresponding interconnectivity between data centres. The framework extends the existing NSI model with a new service enabling the multi-domain service provisioning for SDN networks, in consequence providing a real end-to-end network service spanning multiple administrative domains. This framework also defines a new service within the NSI framework, enabling the multi-domain network provisioning between SDN clouds, not only restricted to providing connectivity but also able to manage other resources, e.g. computation resources or storage. The extensions defined within the framework will be proposed for further standardization through the OGF standardization procedure.

2. Project architecture

The ultimate goal of this project is to allow control and management of both IT and network resources seamlessly, taking into account: i) actual capabilities of the existing SDN platforms running in data centres and local networks and ii) new extensions to the NSI Framework to enable multi-domain processing of end-to-end users' requests (which are explained later in this paper).

The outcomes of the project will be a key enabler for new experiments between distant SDN clouds requiring dynamic network connectivity and bandwidth elasticity. These new experiments will take advantage of on-demand path control toolkits deployed in European Research and Education Networks (GÉANT & NRENs), GLIF and Japan (JGN-X).

Figure 1 details the proposed project architecture. The key elements introduced by the project to the existing NSI framework include:

- The NSI SDN Service located in NSI Agents
- The Multi-domain SDN Lookup Service
- OpenFlow Controllers for virtualization of Data Centers' resources

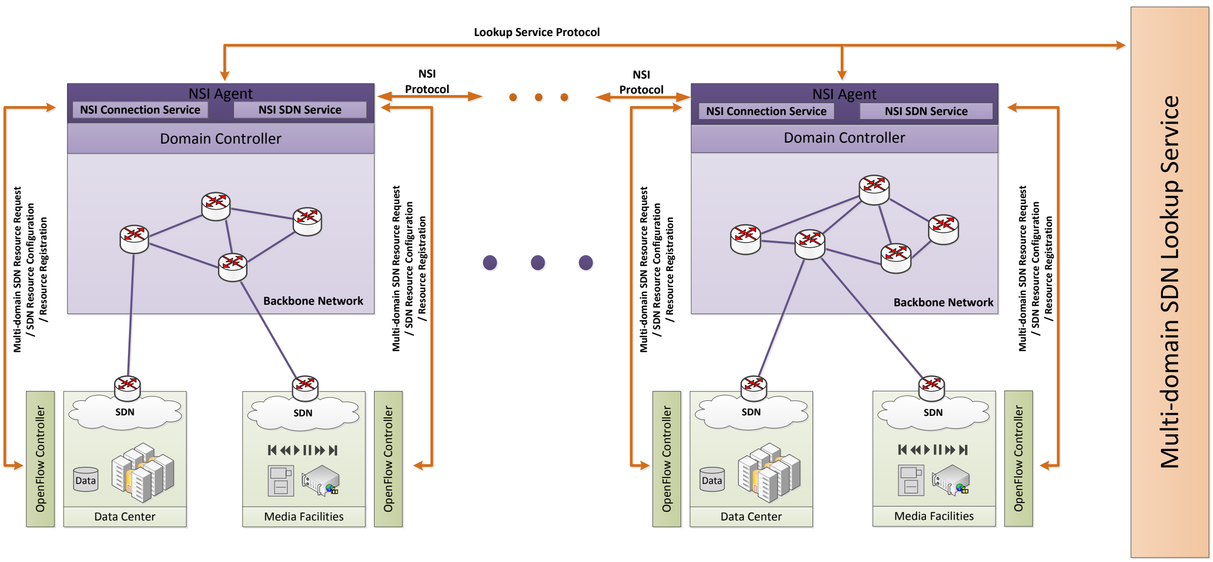

Figure 1: Multi-domain end-to-end connectivity in an extended NSI framework

The NSI Agent functionality is extended with an additional functional component, the NSI SDN Service. The NSI SDN Service is responsible for: i) handling multi-domain SDN resource requests issued from other SDN-enabled domains attached to the NSI enabled backbone network; and ii) requesting last mile configurations of the end-to-end service in local SDNs.

The new service will be assisted by Multi-domain SDN Lookup Service, that stores information about network resources available in a federation for experiments. Within the project, resource registration protocol and procedures for SDNs will be defined and proposed for standardization. Data centres may join the service federation by subscribing the facility in the SDN Lookup Service and registering selected resources for further use in multi-domain experiments.

In order to utilize information available in the SDN Lookup Service the Lookup Service Protocol for NSI Agents will be defined and proposed for standardization. The Lookup Service Protocol adapters will be implemented in NSI Agents, enabling resource discovery in the multi-domain environment and correlation of NSI and SDN resources.

The SDN Lookup Service may be implemented either as a centralized, fully distributed or hierarchical service. The project within its lifetime will investigate which model fits into the proposed architecture and perform simulation studies to evaluate the solution.

3. The reference use case scenario

The following reference use case scenario is considered in the project:

1. The three backbone networks participate to the use case: D_A, D_B and D_C.
2. Each backbone network runs its own Bandwidth on Demand system with specialized NSI interfaces for multi-domain NSI message exchange: NSI_A, NSI_B and NSI_C.
3. Three data centres are connected to the backbone networks, DC_A to D_A, DC_B to D_B and DC_C to D_C through the Service Termination Points, respectively STP_A, STP_B and STP_C.

Figure 2 presents the topology demonstrated in the reference use case scenario.

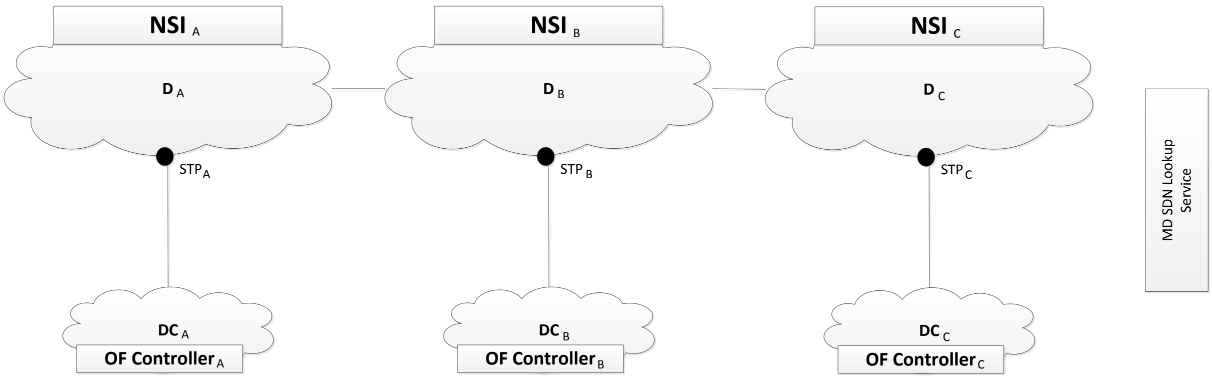

Figure 2: The reference use case scenario

Currently, the NSI framework supports the termination of network services at Service Termination Points only. This implies information about resources in clouds (e.g. data centres) attached to the STPs is not taken into account while setting up connectivity services. The service is terminated at the ingress/egress interface of the NSI domain, in consequence the end-to-end path is not delivered to end users.

The proposed extensions to the existing NSI framework will enable an establishment of connectivity services between specific resources available in data centres (e.g. storage services, multimedia facilities or computational resources) and networks managed by OpenFlow controllers.

The following steps present an extended functionality proposed by the project providing new features for end users.

4. Each data centre registers its resources in the Multi-domain SDN Lookup Service using Resource Registration Protocol.
5. A cloud manager (scheduler/middleware) of a data center DC_A requests outside resources since there are not sufficient resources for demand in its own data center. Therefore, the controller located in DC_A requests a service from NSI_A via a new NSI interface.
6. NSI_A issues the request to the SDN Lookup Service to localize the domain where the counterpart SDN may be attached.
7. Taking into account the response from the SDN Lookup Service NSI_A calculates the location of the remote part of the virtual infrastructure – in this case the data centre DC_C.
8. NSI_A triggers the connectivity setup between STP_A and STP_C, using the state-of-the-art procedure according to the NSI v2.0 specification.
9. When the multi-domain path is set up, the two edge domain controllers, NSI_A and NSI_C issue the request towards OpenFlow controllers to configure the corresponding virtual infrastructures.
10. When the two virtual infrastructures are created, confirmation is sent to the component initiating the whole procedure (in this case to the OpenFlow controller in DC_A).
11. OpenFlow controller in DC_A confirms to the user the availability of the virtual infrastructure spanning multiple domains.

4. Experiments enabled by new developments

Disaster Recovery Scenario

From the earthquake and the tsunami disaster at Mar. 11, 2011, Japan have learned that business continuity is sometimes more important than data integrity. That is, even if the data is slightly obsolete (up to hours or days) when they are recovered, it will be highly beneficial if the service can be continued immediately at a remote site. In addition, according to a survey done after the disaster, most of data centres and wide area networks in the earthquake affected area had been operational for a few hours after the disaster, using backup power supply. If services can be migrated to remote sites while equipment is operational, services can be continued without a long time suspension.

These days, virtualisation is commonly used in cloud computing systems. Virtual machines are used to run a tenant program on top of physical machines in data centres. One of the advantages of virtualisation is that snapshot of executing program can be taken easily, and execution of the program can be resumed using different physical machines. If virtual machine snapshots are taken periodically and are stored at a remote site, the program can be easily resumed at the remote site. Such snapshot does not require any change of running program itself. So, this method can be easily deployed.

On the other hand, when snapshot of virtual machines are taken, the size of the snapshot data tend to be very large. In addition, encryption of the data may be required to ensure required level of security. To meet these requirement, simultaneous allocation of storage to which snapshot will be stored, computers which will do encryption and QoS guaranteed network connection between local and remote sites will be required. If data backup after a disaster is going to be supported, such backup should be given a higher priority than other requests.

To support these requirements, in the Disaster Recovery Scenario a simultaneous allocation of storage, network and computing resources will be realized according to request, with priority given to urgent requirements. Then, using the allocated data, virtual machines in data centres will be migrated, demonstrating how services can be resumed. The experiment will dynamically allocate QoS guaranteed storage system Papio, computers which will be used for encryption/decryption, network between data centres.

High Quality (4K and beyond) Media Transmission over long-distance networks

4K refers to variety of resolutions with approximately 4 thousand of points horizontally. Single image consists of about 8-9 megapixels. The resolution often depends on the screen aspect ratio, but the most common resolutions are DCI (Digital Cinema Initiatives) compliant which is 4096x2160 and Quad HD which is 3840x2160.

Technology for automatic adjusting connection paths and parameters in SDN networks could act as a powerful test engine for advanced multimedia purposes. High resolution visualisation solutions, especially streaming technologies, could be tested in order to pinpoint the sensitivity for network issues changes. High resolution multimedia streams are very demanding and require high quality of transmission, therefore can act as very sensitive test engines for transmission problems.

The experiment would help to determine the behaviour of the transmission mechanisms in case of streaming high resolution multimedia content at a very long distance. Required bandwidth for 4K connections differs from 1Gb/s to 20Gb/s, depending on compression aspects. Stereoscopic transmission doubles these values, since two streams have to be delivered separately for left and right eye. Synchronisation of these streams is extremely important from quality point of view and requires even more reliable transmission engines.

The key aspect of such a kind of transmission is to take heed on sensitivity for negative effects in network (as jitter). Large data stream are very demanding from network point of view. Therefore data transmission parameters have to be configured respectively in order to reduce risk of transmission disruption.

Long distance network connections would be most factual test case for advanced content distribution systems. Extensive and properly planned tests could define possibility of future development and usage delivery systems regardless of the distance of transmission. Many aspects of this sensitive communication could be identified and qualified in order to define sufficient set of requirements for seamless streaming. Additionally, QoS issues could be examined in order to find correlation between network problems and their visual reflection.

4K and 8K resolutions are defined as a part of standards for broadcasting technologies. Therefore it is extremely important to have a full view of the dependencies between quality of multimedia stream and transmission parameters. Tests of long distance multipoint connections would give important and interesting overview of reasonable transmission possibilities and identify new issues. The way of behaving of advanced network engines should be also verified in case of transmission of many parallel streams.

The broad-based initiative would be also a great opportunity to test the important streaming issues by simulating different distances and network parameters. Extensive tests with usage of different paths, parameters and on-demand adjusting would give invaluable results, which would be hardly possible to carry through without advanced test-bed. All tests use cases, as well as experiments would be done in order to address future usage and application.

5. Conclusions

This article introduces the concept of a new service in the NSI Framework – the NSI-SDN service. Recent advances in Software Defined Networking prove the technology behind this concept is becoming mature enough to be adopted in research networking to run advanced networking and computational facilities. The existing NSI standard does not fully reflect all state-of-the-art SDN technologies, although it may be considered itself as an element of SDN – it logically separates the control plane from the data plane, understanding each single administrative domain under control of the NSI Agent acts as a virtual network element. The Connection Service (NSI CS) proposed under the NSI umbrella is limited in scope to an end-to-end network connections delivery, which is insufficient for some groups of researchers interested in creating cross-test-bed slices to run computationally-extensive tasks in a distant and distributed networking environment.

References

[1] http://en.wikipedia.org/wiki/Software_Defined_Networking
[2] OpenFlow consortium website: URL http://www.openflow.org
[3] http://www.openflow.org/documents/openflow-wp-latest.pdf

Biographies

Bartosz Belter received the M.Sc. degree in Computer Science from the Poznan University of Technology in 2002. He works in Poznan Supercomputing and Networking Center as a Senior Network Engineer. He participated in several FP6 IST projects: 6NET (IST-2001-32063), PHOSPHORUS (IST034115) and GN2 (IST511082). He also participated in a number of national initiatives funded by Polish Ministry of Science and Higher Education (Clusterix, Polish LDAP and others). Currently he is involved in several EU funded projects. His main research activities concern the architectural aspects of Control and Management Planes in optical networks and Quality of Service in next generation packet networks. He is author or co-author of papers in professional journals and conference proceedings.

Radosław Krzywania received the M.Sc. degree in Computer Science – Software Engineering from the Poznan University of Technology in 2003. He is working in Poznan Supercomputing and Networking Centre as a senior network engineer. He participated in several FP6 IST projects: 6NET (IST-2001-32063), PHOSPHORUS (IST034115) and GN2 (IST511082). He also participated in a number of national initiatives funded by Polish Ministry of Science and Higher Education (e.g. Clusterix). Currently he is involved in the national project "Engineering of Future Internet" and FP7 project GN3 (Project no. 238875). The main experience is Bandwidth on Demand services, network control planes, and network management. He is author or co-author of papers in professional journals and conference proceedings.

Tomohiro Kudoh received his Ph.D. degree from Keio University in Japan in 1992. He joined National Institute of Advanced Industrial Science and Technology (AIST) in 2002. He currently serves as a deputy director of Information Technology Research Institute, AIST. In the past few years his research has focused on network as a

Grid and Cloud infrastructure. His recent work also includes the G-lambda project which target is to define an interface to manage network as an IT resource. He is a co-chair of the Open Grid Forum NSI working group.

Gerben van Malenstein received his M.Sc. degree in System and Network Engineering from the University of Amsterdam in 2007. Gerben is working as Network Manager within SURFnet's Network Services department. In his roles as GigaPort3 project leader and Technical Product Manager, his focus is on developing NetherLight, SURFnet's GLIF Open Lightpath Exchange (GOLE) in Amsterdam. Alongside this work, he is involved in Bandwidth On Demand (automated network provisioning) both within SURFnet and internationally through NetherLight.

The scientific/academic work is financed from financial resources for science in the years 2013 - 2016 granted for the realization of the international project co-financed by Polish Ministry of Science and Higher Education.