Federation as part of Infrastructure as a Service in Transport Networks

David Stezenbach

University of Vienna

Chair of Future Communication (Endowed by Telekom Austria)

Universitätsstrasse 10/T11, A-1090 Vienna, Austria.

david.stezenbach@univie.ac.at

Abstract—This work transfers the approach of federating resources, such as computation power and connectivity, from connecting and reusing test facilities like Planetlab [1] or GpENI [2] to combining network resources (virtual links, nodes, etc.) for transport networks. These new transport networks might facilitate both, *a)* new technical types of core or backhaul systems in provider networks but also *b)* new ways of doing core network business and new operational procedures for core networks. In particular, the new operational procedures should allow more dynamic in the topology and traffic management of future core networks.

We show some options how application providers or Internet Service Providers could benefit from using federation with using an Infrastructure as a Service approach covering all network entities included in new transport services like *Transport Virtualization* [3].

Keywords – Federation, Network Virtualization, Infrastructure as a Service.

I. Introduction

In this work we present a new approach of traffic handling using virtual topologies based on federation. While federation is very common in test facilities it not yet covered by Infrastructure as a Service (IaaS) approaches. Current infrastructure installation is very inflexible regarding both computation power and connectivity. As a result of that, new upcoming services are not published due to financial risks. We propose the possibility of dynamical resources leasing like in cloud services, but extend it to connectivity in order to offer more control over the application traffic. In the first section we present the general idea of IaaS. The second section motivates the use of federation in productive networks. The last section points out two possible scenarios that use virtual topologies to support traffic handling of specific applications.

II. Infrastructure as a Service

Nowadays, common network services are operated on a relatively fixed infrastructure, regarding both the network and the servers. Thus, changing the service mostly comes with a change in the infrastructure. Neither offering the service to larger group of endusers nor adapting Quality of Service requirements, such as maximum delay or bandwidth, can be transposed without transferring the application to a new system. This leads to inflexible services, which can not be easily adapted to changing requirements as changing the infrastructure causes high extra costs.

In order to counter these problems currently a new business field is emerging. Cloud services offer the possibility to pay only for resources which were really used by the service. Not only that this is reducing the financial risks, also the flexibility is improved. Most cloud services allow the operator to adapt the resources used by his service dynamically. Thus, one has the possibility to scale the service if it is accepted by the user community without moving it to a new hardware. Current cloud services offer flexibility regarding resources such as computation time and storage.

The flexibility described above makes it necessary to virtualize the resources in order to share them among different services. If we now think ahead, this sharing needs not to be limited on storage and computation time, it could also be extended to network infrastructure. Thus, new degrees of freedom are added in how the different server locations are connected which offers new optimization possibilities in terms of traffic management.

As the network between different server locations is generally not managed by a single carrier a request and reservation mechanism is required in order to manage the resources shared among different services. In the next section we will describe how such a federation approach, that is very common in test facilities, works and how it can be transferred to an IaaS system.

III. The Federation Approach

Currently federation is generally know in the context of test facilities. The basic idea followed by projects like Planetlab, GENI, and many others is to share resource located at different sites among a wide researcher community. Therefor, a federated testbed provides a central point of trust, called clearinghouse. Each participant trusts this central point, which offers the option to trust other participants without contracts between each partner. As a result of that, a researcher is able to use the infrastructure provided by others in order to use it for its own purposes. The GpENI project, which is part of the GENI project, in addition offers the possibility to control the network topology connecting different sites. This is one step towards a full IaaS, but it is missing the strict isolation between different user flows.Furthermore, not all network entities connecting the sites are under the control of the framework.

Fig. 1. Virtual Topologies Maintained by an ISP

Fig. 2. Virtual Topologies Maintained by an Application Provider

If we take this approach, we think it is possible to transfer it to a productive network and thus, offer the possibility of a full IaaS covering nodes and connectivity. As the ISPs have full control over their resources, they could offer connectivity between two points with a certain level of QoS. Clearinghouses present in each ISP network could offer application providers the possibility to dynamically book resources within their networks.

IV. Virtual Topologies in Networks

The IaaS service as described above offers the possibility to dynamically adapt the server capabilities and the infrastructure connecting them. In the following, we will describe two scenarios how this flexibility can be used in order to manage different services using IaaS.

The first scenario shown in Figure 1 considers an Internet Service Provider (ISP) that offers two different topologies for applications A and B. The ISP therefor uses its physical infrastructure and provide virtualized topologies on top of it for each application. Each node can be member of multiple virtual topologies, which requires the nodes to be virtualized in order to guarantee an isolated traffic handling. As pointed out the two topologies can consist of different nodes adjusted to the requirements of each application. The topology then corresponds to the network connecting the end-users to the application. Furthermore, a different view of an application can be the aggregation of multiple user connections to a single service class. For example an ISP could summarize all Voice over IP (VoIP) in a single virtual topology in order to optimize this traffic in terms of minimum delay. The motivation of an ISP to offer such a service can be various. We will only give two possibilities without considering legal or economy purposes since this is out of scope of this work. First, the ISP could optimize the utilization of its infrastructure and use the virtual topologies as a method for traffic engineering. Second, a provider could offer premium service to its customers in order to increase users Quality of Experience (QoE). We think that several more fields of application will arise if the technology is present.

The second scenario presented in Figure 2 shows virtual topologies created by application providers in order to optimize the network to support their applications best. In this scenario a virtual topology includes network entities of different ISPs. This requires that the ISPs act as an IaaS provider. One possibility to achieve this is to instantiate one clearinghouse in each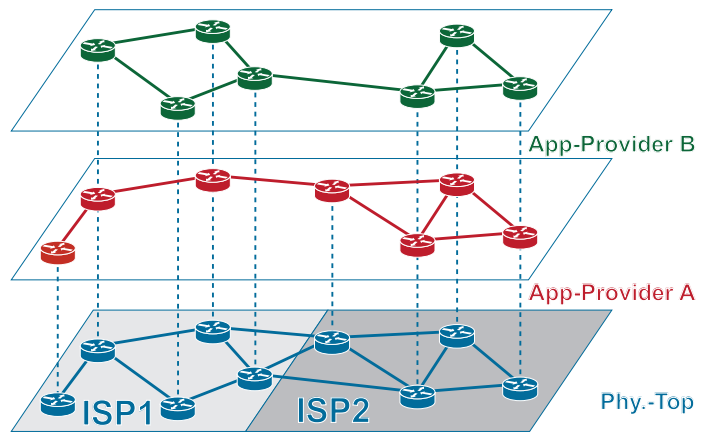 ISP network, which acts as a local resource reservation point. We currently leave out of scope how to locate this clearinghouse within each ISP, but one possibility could be to attach this to existing services like BGP. The virtual topology can either consist of the points of present of the application and the network connecting them, or aggregate the end-user traffic in the topology when it first reaches a virtualized node controlled by the application provider. This offers the application provider to handle the traffic on its own in order to optimize the data transmission. Thus, application providers gain more control over the data transmitted between the end-systems running their application. As a result of that they can guarantee better QoE to their users as they do not have to rely on basic IP routing, which possible transmit the traffic over high utilized links.

In this section we presented two possible fields of application for virtual topologies. One resides within the network of a single ISP, while the second offers new possibilities in traffic handling to an application provider.

V. Outlook and Intended Benefit from the Summer School

In the following, I will point out shortly how my future research regarding the WWTF project OptFI [4] will benefit from the Ofelia summer school. First, I am interested in the concept and the configuration of OpenFlow-based networks for researching operational and optimization procedures. Second, how OpenFlow technology could be used to instantiate network cooperation in core networks. This includes questions on scalability and dynamics of OpenFlow-based virtual networks. Therefor, models of OpenFlow nodes are in the focus of my interests, also regarding on optimization procedures.

References

[1] PlanetLab, "PlanetLab - An open platform for developing, deploying, and accessing planetary-scale services," [Online; accessed 22-July-2011].
[2] GpENI, "GpENI - Great Plains Environment for Network Innovations," [Online; accessed 22-July-2011].
[3] T. Zinner, K. Tutschku, A. Nakao, and P. Tran-Gia, "Using Concurrent Multipath Transmission for Transport Virtualization: Analyzing Path Selection," in *Proceedings of the 22nd International Teletraffic Congress (ITC)*, Amsterdam, Netherlands, Sep. 2010.
[4] OptFI, "OptFI - Optimization Challenges in the Operation of the Future, Federated Internet," [Online; accessed 22-July-2011].