The FELIX Architecture for Testbed Federation

Kostas Pentikousis*, Matthew Broadbent*, Radek Krzywania†, Gino Carrozzo‡, Albert Vico-Oton§,
Tomohiro Kudoh¶, Carolina Fernandez§, Atsuko Takefusa¶, Bart Puype**, Jin Tanaka‖, Tom Rothe*
* EICT, † PSNC, ‡ Nextworks s.r.l., § i2CAT, ¶ AIST, ** iMinds, ‖ KDDI
Corresponding Author: k.pentikousis@eict.de

***Abstract*—Network programmability emerges as a key ingredient of Future Internet testbeds, which are established through international collaborations across country boundaries throughout Europe and around the world. Software Defined Networking challenges long-established invariants of the Internet architecture and fosters a paradigm shift in the way we design and operate networks, and enables the establishment of larger and more diverse advanced programmable network testbeds with a global scale. This paper surveys several programmable testbed networks in Europe (FIRE) and Japan, from which we derive the requirements for federating SDN testbeds between the two regions.**

I. Introduction

Large-scale testbeds have emerged as a vital part of both academic and industrial research. They often facilitate multinational experiments that would otherwise be impossible to conduct by a single entity or organization. Over the last years, significant efforts by several projects in Europe, Asia and the US have taken on the task to create experimental research infrastructures that are reusable and can incorporate the latest network technologies as they emerge. Yet, despite said investments, cross-border, open testbed federation is still an inexact science. In this paper, we review relevant testbed architectures and identify common design considerations which will be used to guide an architecture design for a testbed federation spanning Europe and Japan.

Such an experimental platform, which enables Future Internet research, is a necessity to enable network researchers to validate their designs and applications under near real-world conditions. In general, it is advantageous for each partner to host their own specialised set of hardware and equipment, and offer it as a complete facility to the community. This diversity often draws experimenters to the testbed, but also poses significant technical challenges: creating a platform which provides an authoritative method for experimenters to create an experiment, whilst still maintaining the flexibility needed for experiments, is not an easy task. The FELIX project (see www.ict-felix.eu) aims to address this.

II. Related Work

In this section, we briefly survey a number of relevant testbeds, with particular emphasis on the testbed architectural elements that we consider to be valuable in generalising a conclusive testbed design.

The **OFELIA** project [1] established a pan-European experimental network facility which enables researchers to experiment with real OpenFlow-enabled network equipment and to control and extend the network itself in a precise and dynamic fashion. The OFELIA facility uses the OpenFlow protocol (and related tools) to support virtualization and control of the network environment through secure and standardised interfaces. Ten interconnected sites form a diverse OpenFlow-based infrastructure that allows experimentation on multi-layer and multi-technology networks. A key objective is to provide experimental facilities which allow for the flexible integration of test and production traffic by isolating the corresponding traffic domains inside the OpenFlow-enabled network equipment. This creates realistic test scenarios and facilitates the seamless deployment of successfully tested technology into the real-world.

OFELIA consists of two layers. The *physical layer* is comprised of the computing resources (servers, processors) and network resources (routers, switches, links, wireless devices and optical components). Resources are managed by the OFELIA Control Framework (OCF). Furthermore, the *control framework layer* contains components which manage and monitor the applications and devices in the physical layer. Aggregate Managers and Resource Managers are crucial components of this layer, which can be seen as the combination of three components: *Expedient* is the GUI and allows the connection and federation with different Aggregate Managers via its plugins; *Aggregate Managers* (AMs) enable experimenters to create both compute and network resources via the VT AM and OF AM respectively; *Resource Managers* directly interact with the physical layer, provisioning compute resources (OFELIA Xen Agent) or flow rules to establish the topology (FlowVisor).

The **FIBRE** project [2] federates testbeds distributed across Europe and Brazil. The FIBRE-EU system builds on top of the OFELIA OCF and incorporates several wireless nodes based on commercial Wi-Fi cards and Linux open source drivers. On the other hand, the FIBRE-BR testbed includes nine Brazilian partners interconnected using private L2 channels. The VLAN-based L2 physical link between Europe and Brazil is provided by GÉANT, Internet2 and RedCLARA.

Unlike OFELIA, the FIBRE infrastructure is managed by different types of control and monitoring frameworks (CMFs). Indeed, FIBRE includes and enhances testbeds from other projects like OFELIA, OMF and ProtoGENI, which have been modified with the necessary software components to align their northbound interface to a common specification. The FIBRE project has opted to have two top-domain authorities, one in Brazil and one in Europe, to manage and own resources in the respective continents. These inter-connected authorities interoperate to allow the federation of BR and EU testbeds. The FIBRE architecture is composed of several multi-layer building blocks which are briefly summarized next. The SFA Registry is a database able to store the information related to users and projects, and to manage the certificates provided by the authorities. The MySLICE tool is used as the graphical (web) user interface for administration and experiment management. The SFA gateway is designed to translate the user's requests to a common API and provide slice management functions. FIBRE reuses and enhances the Aggregate Managers (AM) previously developed in OFELIA related to OpenFlow (OPTIN AM) and Xen-based (VT AM) resources and introduces a new AM to manage optical switches (ROADM) devices.

In Japan, **GridARS** [3] provides a reference implementation of the Open Grid Forum (OGF) Network Services Interfaces Connection Service (NSI-CS) protocol standard [4]. NSI is a web service-based interface for reserving, provisioning, releasing and terminating a network service, such as an end-to-end connection, via a two-phase commit protocol. GridARS can coordinate multiple resources (services), such as a network connection, virtual machines and storage spaces, via the NSI-CS protocol. It provides experimenters a virtual infrastructure, which spans several cloud resources, realised by multiple management domains including commercial solutions.

GridARS consists of three main components. First, the *Resource Management Service* (RMS) is based on NSI-CS and consists of Global Resource Coordinators (GRCs) and Resource Managers (RMs) for Computers (CRM), Networks (NRM), and Storage (SRM). Coordinating with GRCs and RMs, RMS can coordinate heterogeneous virtual resources on multiple cloud environments. GRC has a co-allocation planning capability, which determines a suitable resource allocation plan. Second, the *Distributed Monitoring Service* (DMS) allows experimenters to monitor the virtual environment allocated to them. DMS does not have a central database, but gathers distributed monitoring information, tracking the hierarchical RMS reservation tree using the reservation ID, automatically. DMS consists of Aggregators (DMS/A) and Collectors (DMS/C). Each DMS/A gathers monitoring information from related DMS/As or DMS/Cs distributed over multiple domains, and provides the information to the requester. Each DMS/C monitors the reserved resources periodically, filters the monitoring information according to the domain policy, and provides the requester with the authorized information. Finally, the *Resource Discovery Service* (RDS) collects static resource information items from each resource domain and provides the aggregated information. The RDS implementation is based on the Catalog Service Web (CSW), defined by Open Geospatial Consortium (OGC) as an online XML-based database. Each resource domain can POST its static resource information, such as network topology, number of VMs, and storage spaces.

Our review indicates that several testbed federations have been either based on the OCF or the NSI framework. The joint EU-Japan collaboration project FELIX aims to bridge this gap and develop for the first time a testbed federation that successfully spans across these technological and geographical boundaries and incorporates both approaches in a recursive and scalable model.

III. Common Design Considerations

As can be surmised from the previous section, certain commonalities are present in several testbed architectures. In this section we discuss a number of elements that, in order for the testbed to be successful, should be addressed in the final FELIX architecture design.

Resource Orchestration – Orchestration of various virtualised resources, not only computers, but also network and storage, provided from multiple domains, is required.

Domain Resource Management – Coordination of various resources provided by heterogeneous resource management systems within a domain.

Resource Allocation Planning – It is essential to create a suitable resource allocation plan for both computing and network resources. This should take into consideration user and resource administrator aspects, such as, for instance cost, energy consumption and load balancing.

Provisioning – It is important to provide applications with a virtual flat environment, just like a dedicated cluster, using dynamic resource information, such as IP addresses.

Authentication and Authorization – It is vital to ensure that any action is proven to be performed (only) by authenticated actors and ensure that those actors are who they claim to be, as well as keeping track of any authenticated action.

Monitoring – Each domain has to provide monitoring information for resources (which are often virtual rather than physical). Such monitoring information from multiple domains has to be coordinated and provided to the user.

User Access/GUI – A friendly interface is important to allow experimenters to manage their experiments in a faster and easier way.

Many of the testbed implementations described in Section II have the concept of a slice. A **slice** is an abstraction built upon a set a of physical resources, such as compute and networking, using virtualization technologies. When a slice is provisioned it establishes a virtual infrastructure upon which the experiments have full control of the resources. A slice is often dynamic in nature, as the resources will only be provisioned when needed. This is in contrast to having them allocated statically and prior to use. The only requirement to create a user's slice is that there is still available capacity within the physical substrate. Importantly, a slice also offers isolation, ensuring that there is no interference across separate experiments running in parallel on the same physical layer.

IV. Conclusion

This paper presents a number of design considerations to be used in the FELIX architecture for testbed federation. These are driven by our our analysis of the current status and limitations of several projects also working on SDN infrastructures for experimentation and specifically three European FIRE projects (OFELIA and FIBRE) and one Japanese Project (GridARS). We will be using these as influence to form a conclusive FELIX architecture design.

Acknowledgment

This work was conducted within the framework of the FELIX project, which is partially funded by the Commission of the European Union FP7 program under Grant Agreement no. 608638, and National Institute of Information and Communications Technology (NICT), Japan. Study sponsors had no role in writing this report. The views expressed do not necessarily represent the views of the authors' employers, the FELIX project, or the Commission of the European Union.

References

[1] A. Köpsel and H. Woesner, "OFELIA: Pan-european Test Facility for Openflow Experimentation," in *Proceedings of the 4th European Conference on Towards a Service-based Internet*, ser. ServiceWave'11, 2011, pp. 311–312.

[2] S. Sallent et al., "FIBRE Project: Brazil and Europe Unite Forces and Testbeds for the Internet of the Future," in *Testbeds and Research Infrastructure. Development of Networks and Communities*, ser. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, 2012, vol. 44, pp. 372–372.

[3] A. Takefusa et al., "GridARS: An Advance Reservation-based Grid Co-allocation Framework for Distributed Computing and Network Resources," in *Job Scheduling Strategies for Parallel Processing*, vol. 4942/2008: 152-168, April 2008.

[4] G. Roberts et al., "NSI Connection Service, Working Draft, Version 2.0."

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.