Computer Networks 61 (2014) 176–183

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

The FEDERICA infrastructure and experience

M. Campanella *, F. Farina

Consortium GARR, Via dei Tizii 6, 00185 Roma, Italy

ARTICLE INFO

Article history:
Received 30 August 2012
Received in revised form 11 September 2013
Accepted 20 December 2013
Available online 3 January 2014

Keywords:
Future Internet
Network architectures
NREN
Virtualization
Distributed systems
Infrastructure as a Service

ABSTRACT

The European Commission co-funded project FEDERICA started in 2008 with the objective to support Future Internet research and experimentation. The project created a Europe-wide infrastructure based on virtualization in wired networks and computing elements, offering fully configurable and controllable virtual testbeds as a service to researchers. This article reviews the architecture, its deployment and current active status, usage experience, including virtual resource reproducibility and elaborates on challenges for Future Internet testbed support facilities.

1. Introduction

The FEDERICA project (Federated E-Infrastructure Dedicated to European Researchers Innovating in Computing Network Architectures) [2–6] has been designed and operated to support research and experimentation on current and Future Internet. The project consortium comprised twenty stakeholders from Research Institutes, Universities, National Research and Education Networks (NREN) and industrial partners. The project started in January 2008 as an EC 7th Framework Programme project (RI-213107) co-funded by the e-Infrastructure Unit, and formally ended on November 2010. Since then, the NREN partners maintained the facility, which is now partially supported by the GN3 and GN3plus FP7 projects (GÉANT) [10]. The facility is currently serving other EC projects in the FIRE EC Unit [9] (Future Internet Research and Experimentation, now Unit E4 "Experimental Platforms").

The article is organized as follows: Section 2 provides a brief overview of the FEDERICA architecture, implementation, results on reproducibility and current status. Section 3 provides the experience and lessons learned in providing its services. Section 4 elaborates on the issues in operating the infrastructure and in its evolution. The article concludes with Section 5 on infrastructure perspective for the future.

2. The facility

To support research and experimental validation on new networking technologies and applications, the facility has been engineered using computing and network physical resources, implementing virtualization capabilities in all its resources. The choice poses the minimum number of technological constraints to researchers and virtual testbeds.

Virtualization is defined here as the capability to create a virtual version of a physical resource. The virtual resources (e.g., a virtual network circuit, disk partition, virtual computer) are typically created by segmenting a physical resource. Virtualization creates un-configured, clean, virtual resources, e.g., an image of the hardware of a computing element on which (almost) any operating system can be installed, a point-to-point network circuit, a portion of disk space. The virtual resources can be then tai-

* Corresponding author. Tel.: +39 0649622000.
E-mail addresses: mauro.campanella@garr.it (M. Campanella), fabio.farina@garr.it (F. Farina).

http://dx.doi.org/10.1016/j.bjp.2013.12.029
1389-1286/

lored and configured to users' needs and even moved from one virtualization-enabled physical platform to another.

The use of virtualization has been the key enabler of the novel architecture, decoupling the physical substrate from the testbeds. Virtualization allows the facility to provide to the user complete control and configuration capabilities within the assigned virtual resources, allows disruptive testing and permits serving more than one user at the same time, maintaining separation and privacy, and optimizing the use of the physical infrastructure. Some additional requirements for the facility have been taken into account:

- Ensure reproducibility of experiments when requested. Given the same initial conditions and hardware models, the behavior of a single virtual resource should be the same, as a basic principle to obtain the same experimental results.
- Be capable to interconnect or federate with other e-Infrastructures and Internet.
- Favour testing of functionalities, protocols and new ideas rather than very high performance capabilities.

The framework for such an infrastructure is based on two distinct layers. The lower is the virtualization substrate, i.e. the physical environment which contains all the hardware and software to instantiate the virtual resources. On top of it there is the virtual infrastructures layer, containing all the virtual sets of resources (named slices). In the case of FEDERICA some network physical resources are provided by GÉANT and NRENs [16] infrastructures, which are the physical layer at the bottom. Fig. 1 provides a pictorial view of the architecture.

2.1. Service description

The service architecture of FEDERICA follows the Infrastructure as a Service (IaaS) paradigm. The main FEDERICA service is to provide testbeds in a wired large scale environment, in the form of "slices". A slice is a set of virtual computing and network resources in a network topology and with characteristics according to users' requirements. The researchers may select from a collection of preconfigured virtual host, usually Linux based, and request to the network operation centre (NOC see 2.2.3) an almost fully configured slice, or choose to receive the resources un-configured (i.e. with no operating systems, applications, IP addresses and routing). The user can access the service once he/she has agreed on an Acceptable user Policy and has been granted a set of credentials to access the resources in his/her slice. The slice creation process has been initially based on human interaction.

Reproducibility for a virtual host and a Quality of Service (QoS) for a network circuit can be requested. The QoS cannot be requested for an aggregate of resources, but only for single ones. It's the responsibility of the user to ensure a global QoS in his/her experiment.

The service includes constant monitoring of virtual resources' main characteristics (e.g., traffic load on circuits, CPU, RAM usage).

The current FEDERICA service is open to all research entities, including the private sector and there is no fee associated for its use. If a user requests non-standard hardware or a direct physical connection, incurred costs are supported by the user.

The facility may evolve its services to include "cloud" services, like Platform as a Service (where various pre-configured system images and network scenarios can be provided from a repository), and Application as a Service (where a set of applications can be preselected by users).

2.2. Implementation

The FEDERICA physical topology (substrate) is depicted in Fig. 2, it is composed of circuits, network equipment and computing physical resources. At present it is composed of 10 distributed physical points of presence. The topology is a mesh of one Gigabit Ethernet circuits, which offer resiliency and load balancing to the network substrate. The capacity has been chosen as a compromise between the cost of wide area circuits and the total capacity. The choice has been adequate for the users' requirements. The circuit capacity can be sliced, still providing high-speed links, and although expensive, it is contributed by the participating NRENs without requesting a users' contribution. The circuits have been provided initially as Ethernet over SDH and will be, after August 2012, provided using IP/MPLS as

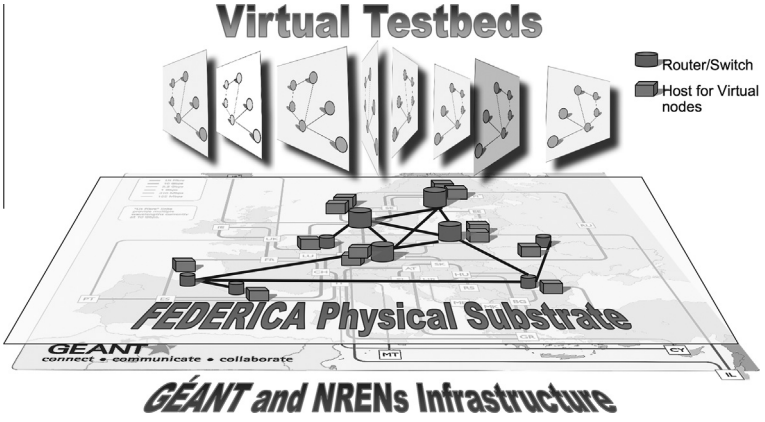

Fig. 1. Pictorial view of the FEDERICA facility.