



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.

© 2014 Elsevier B.V. All rights reserved.

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).

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

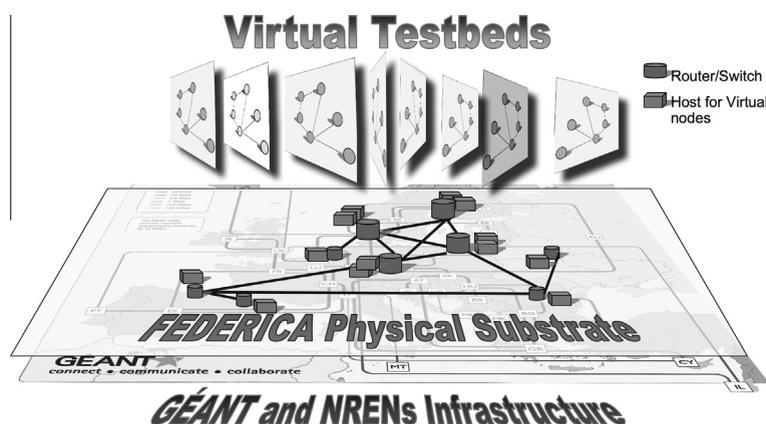


Fig. 1. Pictorial view of the FEDERICA facility.

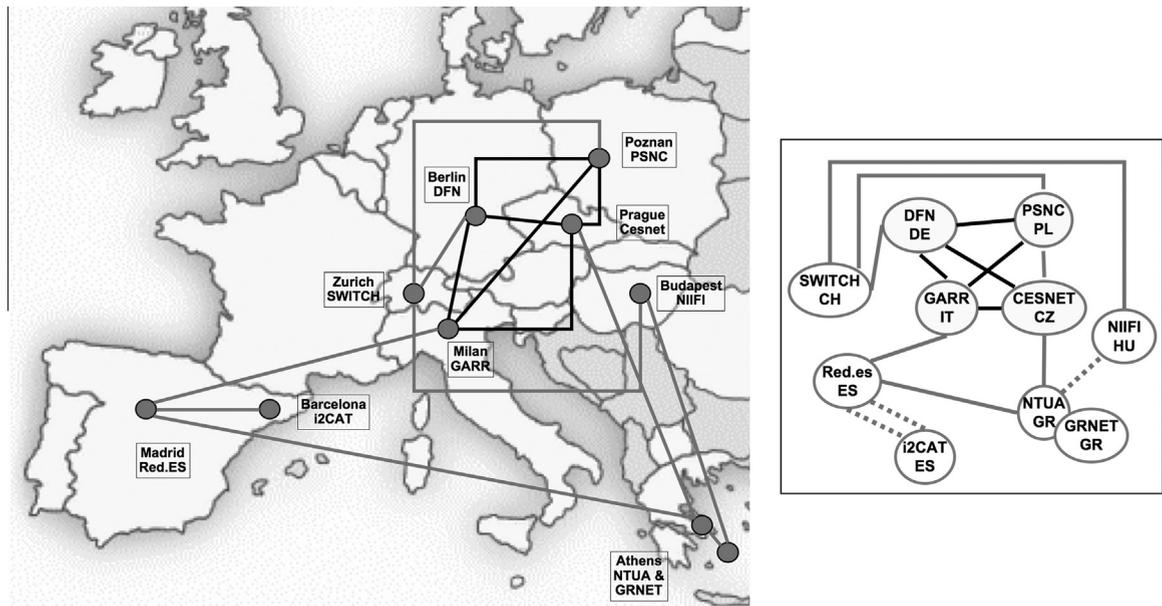


Fig. 2. FEDERICA network topology (June 2012).

point to point virtual circuits carrying Ethernet frames. Each PoP hosts a rack with network and computing elements, Fig. 3 shows a typical installation. The four PoPs in DE, IT, PL, CZ are fully meshed and each hosts a programmable high-end router/switch from Juniper Networks, an MX480 model with dual CPU and 1 line card with 32 ports at 1 Gb Ethernet (8 optical and 24 copper). The MX functionalities include virtual and logical routing, Multiprotocol Label Switching, VLANs, IPv4, IPv6. The MX480 in Germany and Poland are equipped with Ethernet linecards with hardware QoS capabilities capable of a very precise

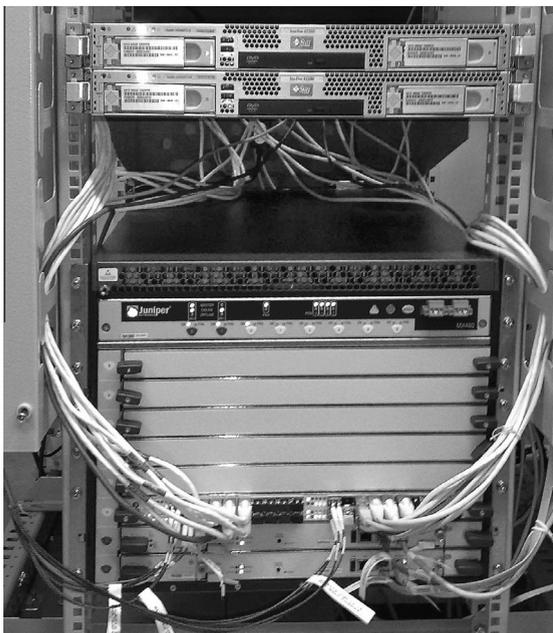


Fig. 3. FEDERICA rack in the Milan PoP.

capacity slicing. These provide up to 1000 virtual circuits with strict capacity guarantees, when requested. Smaller multi-protocol switch/routers (Juniper EX series) are installed in non-core PoPs. The computing equipment (called V-Node in the following) is based on off the shelf PC hardware, running virtualization software. Each PC contains 2 quad-core CPU running at 2 GHz, 32 GB RAM, 8 Ethernet 10/100/1000 Mbps network interfaces and 2 SATA disks of 500 GB each. All the interfaces of the V-Nodes are connected to the Juniper equipment. Figs. 3 and 4 display the typical configuration. There are typically two V-Nodes in each core PoP and one in the other PoPs.

The substrate is configured to be a single IPv4 and IPv6 public Autonomous System. The AS has a full BGP peering in the four core nodes with hosting NRENS, which announce it to the Internet. A slice may receive public or private IP addresses and its resources are reachable from Internet for remote management through a user access portal.

The choice of the virtualization software for the V-nodes is the free version of the VMware [25] hypervisor (ESXi). This choice has been made initially after a review of available virtualization software. In particular the completeness and structure of the application programming interface, the availability of usage examples, available expertise, free tools to import and convert node images were evaluated. The capabilities and performance of the free version have been adequate for the requirements.

Each PoP is ready to host user's equipment, according to space availability, and to physically interconnect user's sites.

2.2.1. Reproducibility and quality of service

Determining the reproducibility of a combined set of many resources is still an open research area. The interaction among all the elements in a testbed makes their behavior complex to predict and to regulate. The addition

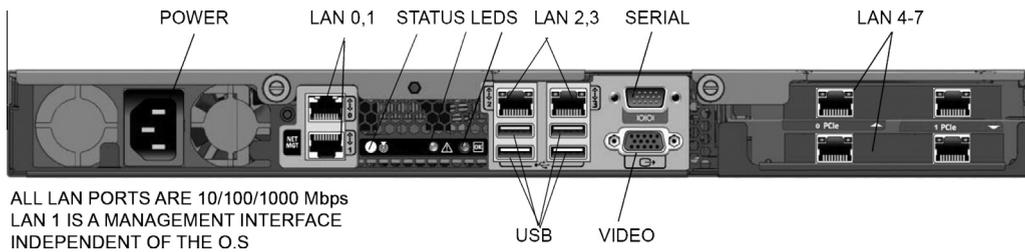


Fig. 4. Back view of a V-Node.

of the virtualization software introduces an additional element of variability that may alter the stability of the behavior and performance of each resource, see e.g. [23].

The project performed an internal validation of reproducibility capabilities of its elementary building blocks: virtual nodes (V-Nodes) and point-to-point virtual circuits.

Measurements have been collected using a dedicated slice made of four virtual computing nodes connected by virtual circuits. The nodes were configured in pairs: two “small” virtual nodes (0.9 GHz CPU, 1 GB RAM, 10 GB local storage, 1 Ethernet interface) and two “large” (1.8 GHz CPU, 2 GB RAM, 100 GB local storage, 1 Ethernet interface). The small VMs are used primarily to test CPU and memory, the large are used also for disk access evaluation. The networking performance and its reproducibility have been tested generating 100 Mbps full-duplex traffic between pairs of VMs, hosted in different PoPs.

The characteristics that have been sampled during the reproducibility tests are: the use of the CPU, the memory bandwidth, the read/write bandwidth of the disk channel and the use of network capacity. Standard benchmarking tools have been adapted to exercise the components. Following an approach similar to the ones proposed in [13] [12], High Performance Linpack (HPL) and STREAM benchmarks from the High-Performance Computing Cluster benchmarking suite have been used to sample CPU and RAM bandwidth respectively. The high-performance computing systems benchmarks results, expressed in flops, are used to compare performance with and without the hypervisor on a physical machine. The IOzone tool has been used to measure disk I/O performance with different disk access patterns (e.g., sequential read and write, random read and write). Netperf has been adopted to measure the circuits’ bandwidth.

The results indicate that reproducibility can be provided configuring in the hypervisor guaranteed, fixed values for CPU cycles and memory size of the virtual machine. The performance remains stable even in case of physical resource contention, with a slightly lower stability of the bandwidth to disk storage, mainly due to the simple hardware in the testbed. The variability of a virtual resource due to virtualization itself can be limited to below 1% for CPU cycles and less than 0.5% for random access to memory. The result is in line with reproducibility and performance measurement in clouds [22,24]. The configuration of a virtual circuit to guarantee capacity is also achievable and it performs exactly as configured when hardware assistance in the router is used.

Descriptive statistics and statistical tests on mean, deviation and on the distributions profile have been used to ana-

lyze the data collected and validate the correlations. The end-user license for VMware restricts the publication of benchmarking and results are a function of hardware and software release used. However, the statistical indicators show that benchmarks on the VMs are stable over time and that variability can be reduced, enforcing reproducibility. The minimum bias introduced by virtualization and by the substrate network has been measured to calibrate the results. Virtualization, when no specific configuration is used, implies a loss of 5 MHz with standard deviation of 1.8 MHz on the tested CPU (2 GHz), therefore it can be assumed that the virtualization represents a negligible offset and it does not interfere in the tests of the researchers. FEDERICA active network monitoring has been used to confirm that the network circuits in the substrate are stable and errorless in the typical timescale of tens of minutes of each test. The results of HPL and STREAM indicate stable behavior also when contention was introduced. When CPU affinity (the static binding of a virtual CPU on a specific physical core) is activated, the deviation of the observations is further reduced by 60%. An additional reduction of deviation has been observed for memory intensive tasks when both CPU and RAM affinity are active. The effect of the activation of both affinities on the STREAM benchmark running over a small instance is shown in Figs. 5 and 6. The bandwidth sustained by the virtual network interfaces has been assessed both for half and full-duplex transfers. The netperf benchmarking tool has been adopted to gauge the rates for periods of 120 s and the measurements have been repeated. The observed sustained rates are reported in Fig. 7. The VMs are capable of using 100 Mbps full-duplex connections even in case of contention. Higher

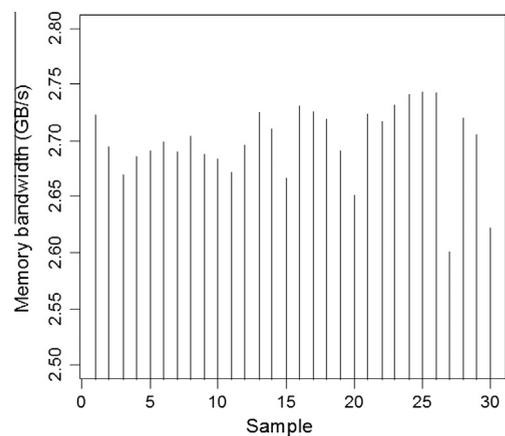


Fig. 5. Memory bandwidth rate during STREAM benchmark execution for a Large instance with default parameters.

bandwidths experiments require a specific mapping over the substrate resources, balancing the trade-off between resources underutilization and reproducibility.

Measurements on storage have shown a reduced level of reproducibility mainly given by the V-Nodes' local storage. Every V-Node is equipped with two SATA disks: one entirely dedicated to store virtual machines images, and one shared between the images and VMware installation. The IOzone tool has been used for sequential and random tests under different contention patterns. When the storage is accessed exclusively, measurements indicate that the behavior of the experiments is repeatable. On the contrary, when storage is under contention, the performance can fluctuate up to the 20% of the observed bandwidth. Based on the V-Nodes storage setup, it can be concluded that slices with intensive disk activity need stronger isolation to minimize the interactions with other experiments.

As a result of testing, the project has defined a set of default configuration rules for virtual instances of type small, medium and large, which provide a default level of reproducibility. The "medium" configuration allocates a CPU core to each image, a fixed block 2 GB of RAM, 100 GB of local disk space and a physical network Ethernet interface. With this configuration a single V-Node can host up to 7 images. Reproducibility can be further enforced using specific capabilities of the virtualization software, like fixed allocation and "affinity".

2.2.2. Monitoring

The substrate and virtual resources in each slice (if the user agrees) are constantly monitored. The monitoring system is based on the G3 system developed by Cesnet [7] and it includes the Nagios [14] system for monitoring the computing elements. The system uses passive monitoring, using simple network management protocol (SNMP), and includes proprietary extensions, e.g. to be able to monitor VLAN packets inside a VLAN. In the Cesnet PoP a flow monitoring probe is also available.

The monitoring system is currently monitoring about 40.000 physical and virtual entities (circuits, CPU cores, RAM, temperature, interfaces). Due to the large amount of entities to monitor, the timers and the monitoring pro-

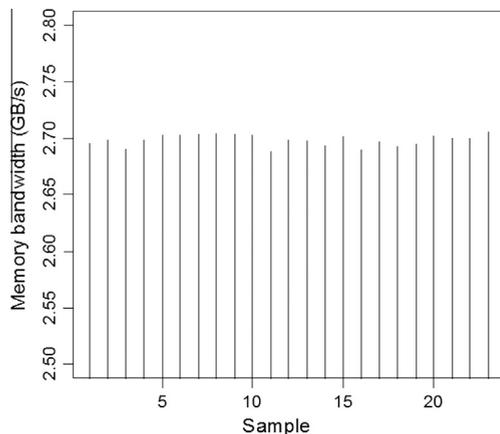


Fig. 6. Repetition of STREAM memory benchmark with NUMA and CPU affinity activated.

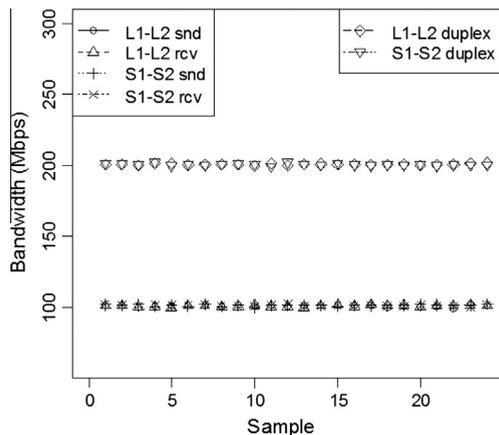


Fig. 7. Network bandwidth: half and full duplex measurements.

cedures have been tuned to reduce and spread evenly in time the frequencies of queries to avoid generating peaks of network traffic and loading CPU of the resources. The system analyses the information in real time and is capable of signalling faults. Each slice monitoring information is available to the respective owner via the web in the user portal and raw data can be downloaded in rrdtool [21] export xml format.

Fig. 8 shows the active map of all monitored entities, the shades of gray represent the load of the element. As an example at the right is shown the traffic statistics of a V-Node ethernet interface in Cesnet, where test traffic at about 1 Gb/s is generated to a node in GARR, the data shown are averaged over 300 s each.

2.2.3. User access and the network operating centre

The use of the infrastructure is accessible, free of charge, to public and private researchers on the current and Future Internet. The access to the infrastructure is supervised by a User Policy Board (UPB). The UPB examines with the users their requests, grants access and defines a priority according to the analysis of technical feasibility, and current resources availability. The decision is not based on the scientific merit of the proposal. The motivation for mandating a controlled access is related to the need of an initial strong support to discuss the capabilities of FEDERICA and the user requirements. The user should also accept a policy document for the use of FEDERICA. More information on user's forms and templates is available in the web site [2] in the users' section.

The infrastructure is centrally managed and monitored by a Network Operation Centre. The NOC has also the task to initially create the slices using a partially automated procedure. The detailed up-to date knowledge by the NOC of facility use and resource allocation permits to ensure guarantees of the service in the overall network, avoiding congestion due to overbooking and enforcing single resource reproducibility (quality of service) when requested. The user can access using a web browser a User Portal to access its personal account, request a slice, upload files, request slice information and access information on its slice.

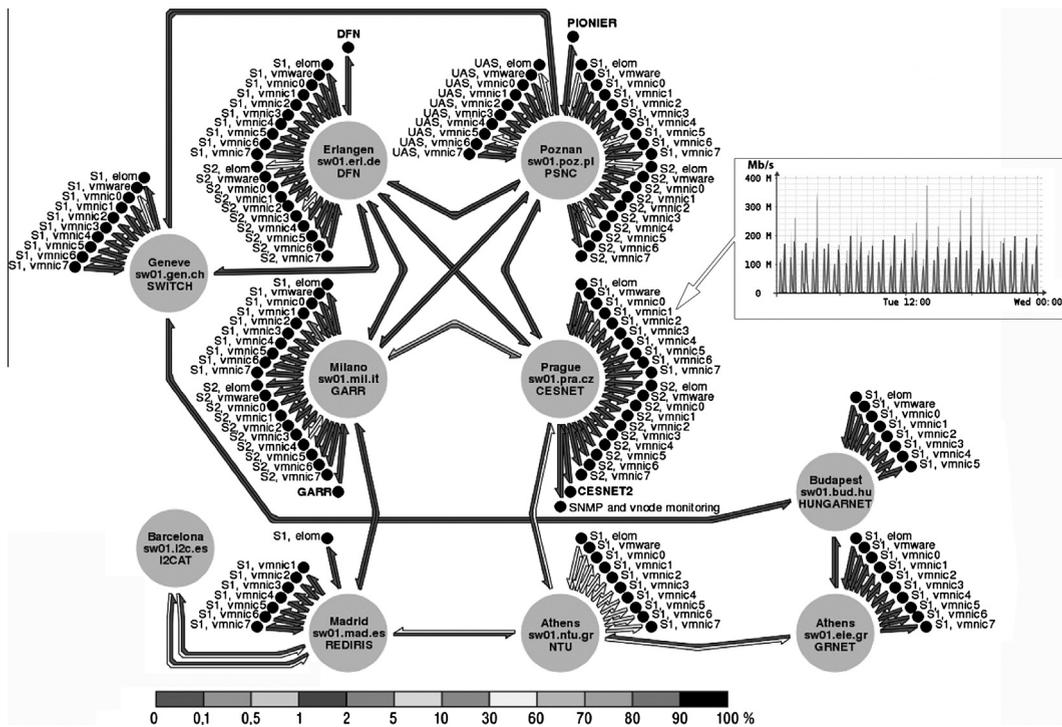


Fig. 8. FEDERICA facility active monitoring map (June 2012).

2.2.4. Authentication and authorization, security

The system is based on credentials assigned to each user and managed through an LDAP server. The core routers implement strict access firewall policies for the traffic to and from Internet, permitting traffic generated only by addresses related to users or the NOC.

2.3. Current status

The facility is active and its operation is partially supported by the EC FP7 project GN3 [10] and it is currently supporting FIRE projects. The development activities have been focused to increase the automatization of slice provisioning and to introduce a resource specification description with is compatible with the Slice-based Federation Architecture (SFA) [20], extending the associated resource specification (rspec) version 2 to represent the facility. SFA is widely adopted by future Internet Testbed facilities, notably by many GENI [11] projects and it is used by the federation of PlanetLab [19] based domains.

The facility plans to enable OpenFlow [17] in its network hardware in the next months, as a service. It is already possible to support Openflow in slices, using virtual resources and software implementation of nodes.

3. Service offering experience

During the operational lifetime of the project, since January 2009, the consortium partners approached and consulted more than 50 potential user groups all over Europe that resulted in 20 slice requests for specific user experiments, with 5 slices active at the same time on aver-

age. The article [5] provides a comprehensive overview of the research projects that have been using the virtual infrastructure slices of FEDERICA until 2011 to validate their research concepts, even being disruptive to their testbed slice, to obtain results in realistic network environments. The projects may be grouped according to their main objective; validation of virtual infrastructure features, evaluation of multilayer network architectures, and design of novel data and control plane protocols at various layers.

After the project ended in 2010, the facility continued operation and at present it is supporting the EC co-funded projects NOVI [15] on approaches to compose virtualized e-Infrastructures, CONFINE [8] on the development of a unified access to an open testbed where to deploy, run, monitor and experiment with services, protocols and applications on real-world community IP networks, and BonFIRE [1] to support experimentation and testing of innovative scenarios from the Internet of Services research community, specifically focused on the convergence of services and networks. In the case of NOVI, in addition to a set of slices, the collaboration extends to SFA compliance, rspec extension to describe the FEDERICA substrate and API to automatically configure resources. The research and development of federation is also ongoing with the BonFIRE project. FEDERICA is providing to each project a slice, connecting also their local testbeds sites using Layer 2 Ethernet virtual circuits.

The duration of the experiments ranged between 3 and 8 months. The slice set-up phase initially required between few days and a couple of weeks, as a function of slice complexity and user requirements. The manual procedure to create the slice has been progressively automated using web services to offer a graphical user interface and an

application programming interface. Routers can be created with very simple configurations to be completed by the user, including single area OSPF and BGP instances. The collaboration with the NOVI project permitted to test the federation capabilities based on SFA.

At the time of writing, notable management and user interfaces for computing elements have been recently developed for the cloud environment (see e.g. the OpenStack software [18]) and their use in the production environment is under evaluation. The procedure to automate the creation and integration of a virtual network connecting virtual resources remains very complex.

The slice contained more than one hundred resources in the case of the PHOSPHORUS project (see [5]), with an average of 20 virtual resources in a slice. FEDERICA has always been able to create the slices as requested. The virtual routers are provided either in the form of system images using standard Linux based software routers, or of virtual routers (with completely separate routing and forwarding tables) in the Juniper equipment.

After testing ends, each user is requested to provide a feedback on the facility and its service using a form available in [2] users' section. The feedback has been quite positive, signalling the usefulness of a facility with services as described above. The users main suggestions were to increase the user support, automation of procedures and NOC responsiveness. The following paragraphs analyses the main users' comments and support experience.

3.1. User interface and access

The ease-of-use and effectiveness of the user interface and the NOC support are of paramount importance to attract and serve users. The interface should allow the user to be operational at the basic level in a very short amount of time, to encourage to continue utilizing and exploiting the facility services. The preferred interface is based on http access and a graphical representation of the slice and the action capabilities. As the users own different ranges of technical expertise, a command line interface is not preferred, it is accepted when it allows a greater level of control and personalization.

3.2. NOC and complexity

The support requested to the NOC is more demanding than in a network production environment. The issues and user's questions are more complex and often refer to innovative and research area. The personnel should then use or rely on computing and networking experts, equivalent to a second level NOC.

3.3. Balance between physical and virtual resources

Although virtualization technologies are very efficient in creating virtual testbeds, in some cases the research is better served by dedicating a physical resource to a slice, which adds to the virtual ones. An example is the research on virtualization capabilities themselves and/or on substrate management and control, like in the case of the NOVI project. Although it is possible in the computing system to have a vir-

tualization substrate on another virtualization substrate, adding a physical resources is preferred, also to permit better reproducibility. Adding physical resources is also useful when virtualization capabilities in hardware are not mature enough. This is the case of the network interface in computing elements. Each V-Node has been equipped with 8 interfaces to dedicate a single interface to each virtual system.

3.4. Reproducibility and quality of service

The users were always urged to perform the experiment without requesting initially QoS on virtual resources, as the resource default configuration level, as detailed in sub-Section 2.2.1, provides adequate level of reproducibility in the majority of cases. These rules, associated to overbooking avoidance in network resources, results in a behavior of virtual resources which was deemed very good and did not interfere, according to feedback, with experiments. Given the initial QoS-aware configuration of virtual resources and over-provisioning in the network circuits, there have not been any requirements to improve the resource quality of service in terms of reproducibility.

4. Challenges

A facility that has as main objective the support of research has to provide an environment which is stable, reliable, supported by a NOC and with a very sophisticated, yet simple to use, user interface. The maintenance, evolution and sustainability of FEDERICA represents an additional, constant challenge. The facility has to provide production level support and at the same time evolve and perform research on itself and new technologies. Support personnel has therefore to include senior networking and computing experts, which are scarce and expensive. Evolution implies almost constant hardware and software upgrades e.g. to support a new technology like Openflow. The user interface needs a constant programming effort especially to attract users and provide them an incentive in respect of building their own testbeds.

All these needs imply that operational expenses are usually higher than capital expenses and that every three to four years there is the need of additional funding to renew hardware. The latter procedure is preferred in respect to a constant renewal process due to lower operational expenses implications. The main challenge is therefore to secure a constant funding flow. The FIRE [9] and GENI [11] initiatives are debating these challenges, without having reached a final conclusion. An evolution towards a federation of facilities may improve the breadth of services available to the user, but leaves open for each facility the aforementioned challenge.

In the case of FEDERICA the current support is being provided as a shared effort between NRENs and the GN3 and GN3plus project, as the service offered is considered and confirmed useful to the research area in Europe.

5. Conclusion

The use of FEDERICA has shown the usefulness and potential of a facility based on virtualization in both network

and computing elements that offers testbeds as a service to researchers. The main challenges are: the constant development of a smart and intuitive user interfaces, the unique mix of production quality environment and research and the funding to maintain FEDERICA up-to-date and in line with research needs.

The FEDERICA facility will continue to be operational. The possibility to extend the service to other communities, e.g. researchers in smart cities, and establish and develop a strong international federation with similar initiatives improve is considered are considered key objectives. Federation and collaboration with the “Cloud” environment is the next natural step to increase scalability and to offer a richer environment to researchers.

Acknowledgments

The FEDERICA infrastructure FP7 is since November 2011 partially supported by the Project GÉANT under the Grant Agreement No. 238875. The authors acknowledges the fundamental contribution of all partners of the FP7 project FEDERICA (RI-213107) Consortium GARR, CESNET, Verein zur Förderung eines Deutschen Forschungsnetzes DFN, Fundação para a Computação Científica Nacional FCCN, Greek Research and Technology Network SA GRNET, Nemzeti Információs Infrastruktúra Fejlesztési Intézet NII-FI, Instytut Chemii Bioorganicznej PAN W Poznaniu PSNC, Entidad Publica Empreserial RED.ES, Teleinformatikdienste für Lehre und Forschung SWITCH, Trans-European research and education networking association TERENA, Internet i Innovacio digital a Catalunya i2CAT, Kungliga Tekniska Högskolan KTH, National Technical University of Athens NTUA, Universitat Politècnica de Catalunya UPC, Juniper Networks Ireland, Martel GmbH, HEAnet Ltd., Delivery of Advanced Networking Technologies to Europe DANTE, Politecnico di Torino NORDUnet.

References

- [1] BonFIRE Project. <<http://www.bonfire-project.eu/>>.
- [2] Federated E-infrastructure Dedicated to Research Innovating in Computing Architectures, EC FP7 Project, Grant Agreement No. RI-213107. <<http://www.fp7-federica.eu/>>.
- [3] P. Szegedi, S. Figuerola, M. Campanella, V. Maglaris, C. Cervell-Pastor, With evolution for revolution: managing FEDERICA for future internet research, *IEEE Commun. Mag.* 47 (7) (2009) 3439.
- [4] M. Campanella, V. Maglaris, M. Potts, Virtual infrastructures in future Internet, in: G. Tselentis et al. (Eds.), *Towards the Future Internet – Emerging Trends from European Research*, IOS press, 2010, pp. 63–73. ISBN 978-1-60750-539-6.
- [5] P. Szegedi, J.F. Riera, J.A. Garcia-Espin, M. Hidell, P. Sjodin, P. Soderman, M. Ruffini, D. O'Mahony, A. Bianco, L. Giraud, M. Ponce de Leon, G. Power, C. Cervello-Pastor, V. Lopez, S. Naegle-Jackson, *Enabling future Internet research: the FEDERICA case*, *IEEE Commun. Mag.* 49 (7) (2011) 54–61.
- [6] M. Campanella, FEDERICA: A Virtualization Based Infrastructure for Future and Present Internet Research, in: *Proceedings of 6th International ICST Conference, TridentCom 2010, Berlin, Germany*, May 18–20, 2010, pp. 123–132, (Revised Selected Papers Series: LNICS, Social-Informatics and Telecommunications Engineering, vol. 46, first ed., vol. XVIII, 768, 2011, pp. 305 illus. ISBN 978-3-642-17851-1).
- [7] The NREN of the Czech Republic. <<http://www.ces.net/>>.
- [8] CONFINE Community Networks Testbed for the Future Internet. <<http://confine-project.eu/>>.

- [9] FIRE Future Internet Research and Experimentation. <<http://cordis.europa.eu/fp7/jict/fire/>>.
- [10] GÉANT, The European Research and Education Backbone. <<http://www.geant.net/>>.
- [11] Mark Berman, Jeffrey S. Chase, Lawrence Landweber, Akihiro Nakao, Max Ott, Dipankar Raychaudhuri, Robert Ricci, Ivan Seskar, GENI: A federated testbed for innovative network experiments, *Computer Networks*.
- [12] R. Iakymchuck, J. Napper, P. Bientinesi, Underutilizing resources for HPC on clouds, in: *Procs. of ACM Symposium on Applied Computing (SAC 2011)*, March 2011, pp. 119–126.
- [13] A. Iosup, N. Yigitbasi, D. Epema, On the performance variability of production cloud services, in: *11th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid 2011)*, May 2011, pp.104–113.
- [14] Open Source Monitoring System. <<http://www.nagios.org/>>.
- [15] NOVI Network Innovation over Virtualized Infrastructures. <<http://www.fp7-novi.eu/>>.
- [16] For a Detailed Analysis of NRENs in Europe and Their Role, See the Documents (in Particular the TERENA Compendium). <<http://www.terena.org/publications/>>.
- [17] Srinji Seetharaman, Masayoshi Kobayashi, Guru Parulkar, Guido Appenzeller, Joseph Little, Johan van Reijendam, Paul Weissmann, Nick McKeown, Maturing of openflow and software defined networking through deployments, *Computer Networks*.
- [18] <http://www.openstack.org/>.
- [19] <http://www.planet-lab.org/>, see also PlanetLab Europe at <<http://www.planet-lab.eu/>>.
- [20] SFA Draft 2.0. <<http://groups.geni.net/geni/wiki/SliceFedArch>>.
- [21] rrdtool: an OpenSource Data Logging and Graphing System for Time Series Data. <<http://oss.oetiker.ch/rrdtool/>>.
- [22] J. Schadt, J. Dittrich, J. Quian-ruiz, Runtime measurements in the cloud: observing, analyzing, and reducing variance, *Proc. VLDB Endowment* 3 (1–2) (2010) 460–471.
- [23] N. Spring, L. Peterson, A. Bavier, V. Pai, Using PlanetLab for network research: myths, realities, and best practices, *ACM SIGOPS Oper. Syst. Rev.* 40 (1) (2006) 17–24.
- [24] Guohui Wang, T.S.E. Ng, The impact of virtualization on network performance of Amazon EC2 data center, *INFOCOM*, in: *2010 Proceedings IEEE*, March 2010, pp. 1–9.
- [25] VMware. <<http://www.vmware.com/>>.



Mauro Campanella is working for the Italian research and education Network (GARR) as senior network engineer, as responsible for research. He started working on computers and networks in 1984. He is one of the author of the Premium IP QoS service and of the AutoBAHN service (Bandwith On Demand), now active in the GÉANT European backbone. He acted as coordinator of the EC FP7 FEDERICA project to support research in current and Future Internet. He is participating to various EC projects and is a member of the EXEC of GÉANT project. He holds a laurea in physics, has spent one year at CISCO systems between the years 1999 and 2000 and lives north of Milano in Italy.



Fabio Farina joined Consortium GARR in August 2010. He received a PhD in Computer Science at the University of Milano-Bicocca, Italy in 2008 working on the Grid infrastructure of the LHC experiment at CERN. He contributed to the FEDERICA, NOVI and GN3 Projects. Currently he participates to the GN3 Plus Project in the actions SA2 “Testbed-as-a-Service” and SA7 “Support to clouds”. His research interests include virtualization technologies and cloud.