

# With Evolution for Revolution: The FEDERICA Approach

---

*Peter Szegedi, Trans-European Research and Education Networking Association*

*Sergi Figuerola, i2CAT Foundation*

*Mauro Campanella, Consortium GARR*

*Vasilis Maglaris, National Technical University of Athens*

*Cristina Cervelló-Pastor, Universitat Politècnica de Catalunya*

## ABSTRACT

Over the last two decades the importance of data networking for human beings and systems has increased beyond any expectation in size, complexity and impact on society. Today, the technology is offering the ubiquitous and constant possibility to be connected to the Internet at a wide range of speeds. The traditional management solutions have up to now followed an evolutionary path, however, the scale of the Internet and emerging novel architectures such as peer-to-peer, ad-hoc networks, as well as virtualization capable network infrastructures require focused and possibly revolutionary changes in management approaches. This article elaborates on challenges posed by the virtualizations' renaissance as experienced in planning, developing and operating the FEDERICA infrastructure. The European Community co-funded project FEDERICA, like other worldwide initiatives, including FIND/GENI in the United States, NWGN in Japan, and FIRE program in Europe, is supporting the development of the future Internet. FEDERICA extends the virtualization capabilities of the current hardware and software to provide a flexible infrastructure to host disrupting testing by networking researchers.

## INTRODUCTION

The number of hosts connected to Internet world-wide is growing exponentially. The Internet population is near to 1.5 billion and a wide variety of services is available along with a vast amount of available information. However, it is clear now that the original Internet architecture needs to be improved at least in the areas of security, management, monitoring and mobility. The evolution of the Internet happened faster than it was possible to develop the basic architectural concepts to a comprehensive design. The management concepts have just been developed to keep with the pace of the evolution. The two main paths to solve the management challenges of the future Internet could be evolutionary (incremental) or revolutionary (clean slate). Either would be successful in the end, the bottom line is that we have to think about the manageability of the future Internet well in advance. It is also fundamental to plan for a migration path which has to maintain Internet operational at all stages.

Research activities on novel network architectures and protocols are anticipated all around the world to address a wide range of innovations. Pioneering clean slate initiatives such as Future Internet Design (FIND), Global Environment for Network Innovations (GENI), and Stanford University's "Clean Slate Design for the Internet" in the U.S. clearly indicate a number of possible areas for innovations that will be of significant importance for the Internet of the future. These areas include: addressing and identification, cross-layer design, network virtualization, routing and traffic engineering, dynamic switching of optical circuits, decoupling of control and data, service discovery and composition, as well as management. The European Commission launched the Future Internet Research and Experimentation (FIRE) initiative in its 7th Framework Programme (FP7). A key element of the research is that the new proposals for Internet architectures, protocols and services should not be defined by paperwork, but they should undergo early experimentation and testing in large-scale environments [1]. To provide efficient support for such innovative research activities the FP7 project titled Federated E-infrastructure Dedicated to European Researchers Innovating in Computing network Architectures (FEDERICA) [2] becomes the enabling platform for all stake-holders involved in "Network of the Future" initiatives to develop highly innovative ideas, both evolutionary and clean slate, tested in a production-like, Pan-European e-Infrastructure used for proof of concepts.

To fulfill its goals, the management capabilities of the FEDERICA infrastructure are being developed to include virtualization and actually control and monitor more than one network at the same time. The infrastructure itself will also provide a useful tool to support research in network management. The success of the project does not rely only on the deployment of the infrastructure itself, but also on the results of the user projects dealing with revolutionary research.

The article is organized as follows; after a brief overview on the infrastructure projects supporting future Internet research we give a sort summary on the major network management strategies. Then, we introduce the European project FEDERICA as a tangible use case for a virtualization capable network infrastructure deployment. We discuss its design objectives, architecture and basic processes includ-

ing challenges in virtualized network management. We describe the FEDERICA vision of future Internet, and finally we conclude our article mentioning some future work.

## INFRASTRUCTURE SUPPORT FOR FUTURE INTERNET RESEARCH

It has been globally realized that the novel concepts of future Internet must be experimented and tested in large-scale environments, even in the early phase of the research. That is why such ambitious research initiatives as National Science Foundation programme GENI in the US, FIRE projects in the EU, New Generation Network (NWGN) forum in Japan, Future Internet Forum (FIF) in South Korea or China Next Generation Internet (CNGI) project are supported by infrastructure development activities.

To design and implement large-scale experimental facilities supporting revolutionary research on future Internet the introduction of network resource virtualization techniques is required. The virtualization in computer science is a mature and widely used technique. PlanetLab [3], a global research network that supports the development of planetary-scale network services, was the first community test-bed started in 2002 using virtualization as main principle in this context. The network slice has been defined as the collection of virtual machines running in a number of computer nodes of PlanetLab connected through the global Internet. All participating members are allowed to use and manage slices independently for their experiments. Site administrators can manage slices only by adding or removing virtual machines.

Although PlanetLab is a popular tool for networked services research, it may not be used for research on lower layer network protocols and architectures. The GENI concept [4] addresses these shortcomings of PlanetLab and extends the aforementioned virtualization principle to the whole diameter of the infrastructure with the notions of substrate, slicing and federation. GENI allows the creation of experiments that create a version of the network that could have its own management protocols and could be incompatible with the Internet.

FEDERICA follows an architecture similar to GENI, using substrate, slicing and federation principles. Currently, finding the proper management concept for virtualization capable networks and its services is a real challenge. The bottom line is that the proposed solution should be simple and easy-to-use by the clients (i.e., re-searchers) without wasting time on managing things and concentrate only to the research objective.

## STRATEGIES IN NETWORK MANAGEMENT

Network and service management in future networks become an important topic involving a whole set of evolutionary and novel approaches.

## Evolutionary Approaches

The functional requirements of the traditional network management are summarized by the well-known term FCAPS (Fault, Configuration, Accounting, Performance and Security) management. During the last few years the importance of configuration management as well as network monitoring and measurement (i.e., fault and performance management) has been increased. Currently, the accounting and authorization (i.e., security management) is also becoming more and more important in a virtual environment.

The classical management frameworks such as the OSI Network Management Framework, the Telecommunication Management Network (TMN), or the Internet Network Management Framework, followed the traditional agent-manager centralized paradigm. With the increasing size, management complexity and service requirement of the emerging virtual network infrastructures, such a management paradigm is no longer adequate, and should be replaced with distributed management concepts. That is clearly illustrated by the evolution of the Simple Network Management Protocol (SNMP) for IP networks. As an early effort in 1992, SNMPv2 introduced the concept of intermediary manager towards distributing management functions. Although SNMPv3 with embedded security features is widely deployed for monitoring network devices, it has some shortcomings in configuration management and managing complex systems [5] including virtualization capable network devices, servers and PCs.

Modern routers and switches are highly programmable devices and can accommodate management functions directly on the device. This design trend leads to new management approaches based on Extensible Markup Language (XML) and Web services, as well as some vendor-specific solutions. We do not call these trends revolutionary because they have been around for several years now and did not really change the network administrators' way of thinking. Since XML can be efficiently used as a textual encoding mechanism for application protocols, object and interface specifications, the XML-based management technologies seemed to be obvious choice for next evolutionary step. In parallel with some proprietary developments like Juniper's JUNOScript [6] the benefits of the combined XML and Web based approaches became clear. Such low cost infrastructure with proven scalability and security features was developed by Distributed Management Task Force (DMTF) called Web-based Enterprise Management (WBEM). The ideas behind JUNOScript and WBEM formed the basis of the IETF's standardization effort towards the Network Configuration (NetConf) protocol [5].

Like SNMP, the WBEM and NetConf also follow the agent-manager paradigm, albeit with important differences regarding the capabilities of the management protocol and the associated information model. NetConf is qualitatively different to all the other approaches, given that it follows a

document-based concept [5] as opposed to individual managed object access approach.

### **Novel Approaches**

Several novel ideas are emerging. Candidates for different network management approaches have been proposed based on highly distributed and sophisticated intelligence. Those kinds of management approaches require revolutionary thinking ‘out-of-the-box’ and may drop the complete set of existing solutions. These approaches must be supported by the large-scale experimental network facilities, like FEDERICA virtual infrastructure, for adequate proof of concepts.

The policy-based network management concept started in early 1990s [7]. The most significant benefit of policy-based network management is that it promotes the automation of establishing management level objectives over various types of network devices. Network administrator would interact with the network by providing high-level abstract policies. Such policies are device independent and human-friendly that is needed in a heterogeneous environment of virtual resources.

The concept of Management by Delegation (Mbd) [7] was introduced in parallel with the policy-based approach pushing more management tasks to the agent side dynamically. Code mobility can be considered as the capability of a management procedure to distribute and relocate its components at run-time. With code mobility, management tasks no longer have to be performed by the managers. They simply generate management objectives and outline task procedures, the execution of tasks are delegated to the agents.

The concept of intelligent agents and active networks is extending the initial idea [7]. An intelligent agent is an independent entity capable of performing complex actions and resolving management problems on its own. Unlike code mobility, an intelligent agent does not need to be given task instructions to function, rather just the high-level objectives. However, the entire network management system is autonomous the problem of managing the intelligent agents also becomes increasingly important. Active networks, combined with code mobility, could present an effective enabling technology for distributing management tasks to virtual network resource level. These kinds of network management approaches are the subject of the current research efforts all over the world and also supported by FEDERICA.

### **USE CASE: THE FP7 PROJECT FEDERICA**

We mentioned earlier the importance of the infrastructure projects and early prototyping. We describe in the following the European project FEDERICA. The project partners include a wide range of stake-holders on network research, National Research and Education Networks (NRENs),

DANTE, TERENA, academic and industrial research groups and vendors.

### **Motivation and Design Objectives**

FEDERICA is a two-and-a-half-year European Commission funded project to implement a versatile experimental network infrastructure for trialing new networking technologies [2]. It exploits the leading edge optical networking environment in campus LANs, NRENs and GÉANT. The FEDERICA virtual infrastructure will be able to support experimental network research in data, control and management planes as well as all protocol layers, over an existing pan-European substrate, but in isolation from the service environment.

FEDERICA is not directly about the future Internet but provides a key infrastructure for such research activities accomplished by hosting virtualized facilities in the substrate nodes. Among the main FEDERICA design principles we point out the stress on experiments reproducibility. We have to provide a controlled environment that is flexible, reliable and allows reproducible running conditions with the advent of simple and easy-to-use management solutions. Hence, the management of FEDERICA does not directly imply how the future Internet will be managed, but gives a good opportunity to understand it. FEDERICA infrastructure is not about performance, but rather about testing new ideas and principles.

### **FEDERICA Architecture and Procedures**

The FEDERICA infrastructure is being built on the top of the GÉANT2 backbone interconnecting several NRENs around Europe (Fig. 1). The topology is composed of 13 physical sites, 4 core and 9 non-core nodes, connected by point-to-point links. The FEDERICA node facilities include:

- Programmable high-end routers/switches (Juniper MX series) enabling logical router/switch instances only at the core nodes;
- Multi-protocol switches (Juniper EX series) extending virtual links to non-core nodes;
- PC-based virtualization capable nodes (V-nodes) running virtualization software, implementing e.g., open source software routers and emulating end-user nodes.

The FEDERICA infrastructure is engineered so as to enable reproducible experimentation of user scenarios when requested. This is accomplished by leveraging virtualization with appropriate hardware and software capabilities of the infrastructure, for instance: provisioning virtualized computing resources; binding separate physical network interfaces to virtual slices; assigning separate logical instances of switching/routing equipment; and partitioning physical communication links. Users are assigned virtual networks, via a centralized decision making procedure. They subsequently access their slice resources via a dedicated proxy

that safeguards user slices and the FEDERICA substrate from unauthorized and/or unscheduled demands, also acting as an access authentication engine.

The main feature of FEDERICA is that the users will be able to fully configure and manage the resources of their own slices, without affecting the physical infrastructure operation. The basic slice provisioning procedure performed by the FEDERICA Network Operation Center (NOC) is illustrated in Fig. 2. Note that initially, almost of the actions performed by the NOC follow a manual procedure. However, the project envisages a dynamic system and a tool-bench able to create slices automatically upon the users' request.

## CHALLENGES IN VIRTUALIZED NETWORK MANAGEMENT

Within the scope of FEDERICA, it is very important to define a management strategy capable to deal with virtualized resources. Currently there are some research projects, in which some partners are involved, such as UCLPv2 [8] and MANTICORE [9], dealing with management approaches that could fit within the FEDERICA scope with some constraints. But, it is important to consider two novel design requirements around FEDERICA: the isolation and the federation. Isolation means that, resources from one slice allocated to a specific user must act independently from resources allocated to other user when both are sharing the same physical infrastructure. Meanwhile, federation means that the FEDERICA's resources in a slice can cooperate with outside resources.

Those requirements need to take into account the virtualization paradigm within a Service Oriented Architecture model. As an example virtualization can be considered in the context of the Infrastructure as a Service (IaaS) [10]. IaaS is a technique which represents a physical device or substrate as a software entity.

The challenge behind network virtualization is with providing infrastructures as a service, which brings some advantages in respect to legacy management system architectures. A particularly important challenge is the capability to federate different infrastructures.

In terms of scalability, the infrastructure can follow the business or projects needs, since resources can scale efficiently, and when the needs or the project is over, network resources can be brought back to the provider.

The IaaS concept is the basis, for instance, of the MANTICORE project, where its virtual router configuration capabilities are going to be implemented within the FEDERICA tool-bench to be delivered by the end of the project.

## FEDERICA'S VISION ON THE FUTURE INTERNET

### *Trends and Requirements*

The future Internet will develop according to two basic requirements: interconnect a constantly growing number of elements and offer ubiquitous, permanent connectivity to the majority of its users. The possibility for any user to connect to any other user is also considered fundamental to allow access to all kind of information, also in sensors and objects. These requirements mandate a robust authorization and authentication infrastructure which must be available at all layers of the Internet.

The Internet will continue to be implemented as a set of inter-connected autonomous domains due to scaling requirements, each with different technologies and capabilities, without a centralized management or control. Mobility, reachability and security requirements will impose a strong inter-domain communication of various types of information. Monitoring capabilities should then extend to mobile users, traversing administrative boundaries and to multiple parallel virtual networks on the same physical infrastructure. Such advances require ad-hoc support in the hardware and the development of new standards for virtual resource representations. In particular the need is for a more rich information system which tracks the relationships between the entities in each domain.

### *Life after FEDERICA*

FEDERICA may act as a bridge between the current Internet and the new one, acting also as a platform to test migration paths. Federating similar infrastructures is of paramount importance. FEDERICA is already engaged in defining a service interface to federate resources with other initiatives, e.g. the European project PanLabII and the US GENI. The virtualization architecture proposed by FEDERICA allows developing the revolutionary schemes on the top of the current Internet architecture isolating the new disruptive protocols and services from the current protocol suite.

FEDERICA is planned to run until June 2010. Its results and research will be used in the European NREN environment and by the pan-European GÉANT backbone.

## CONCLUSION

The Future Internet will need novel control and management capabilities. All related research approaches must be validated on large-scale experimental facilities without affecting production environments. The FEDERICA project uses simple evolutionary solutions to create an almost clean slate infrastructure for revolutionary research on existing NRENs infrastructures, including disruptive reproducible testing.

## ACKNOWLEDGEMENT

The FP7 project FEDERICA is partially supported by the European Commission under the Grant Agreement No.: RI-213107. The authors acknowledge the contribution of all project partners and specifically for their input Radek Krzywania (PSNC), Björn Rhoads (KTH) Jean-Marc Uzé (Juniper Networks) and Jirí Navrátil (CESNET).

## REFERENCES

- [1] Future Internet Research and Experimentation: An overview of the European FIRE Initiative and its projects, 1 September 2008, <http://cordis.europa.eu/fp7/ict/fire/>
- [2] M. Campanella: "Federated E-infrastructure Dedicated to European Researchers Innovating in Computing network Architectures", Future Internet Conference, 2 April 2008, <http://www.fp7-federica.eu/>
- [3] L. Peterson: "PlanetLab: A Blueprint for Introducing Disruptive Technology into the Internet", January 2004, <http://www.planet-lab.org/>
- [4] J. B. Evans, D. E. Ackers: "Overview of GENI Overview of GENI and Future Internet in the US", 22 May 2007, <http://www.geni.net/>
- [5] G. Pavlou: "On the Evolution of Management approaches, Frameworks and Protocols: A Historical Perspective", Springer, 13 October 2007
- [6] J. Schönwalder, A. Pras, J-P. Martin-Flatin: "On the Future of Internet Management Technologies", IEEE Commun. Mag., vol. 41, no. 10, Oct. 2003, pp. 90-97.
- [7] R. Boutaba, J. Xiao: "Network Management: State of the Art", IFIP World Computer Congress, 2002
- [8] E. Grasa, G. Junyent, S. Figuerola, A. Lopez, M. Savoie: "UCLPv2: a network virtualization framework built on web services", IEEE Commun. Mag., vol. 46, no. 3, March 2008, pp. 126-134.
- [9] V.Reijs, E. Grasa: "The MANTICORE Project: Providing Users with a Logical IP Network Service", TERENA Networking Conference, May 2008
- [10] M. Lemay: "An Introduction to IaaS Framework", 5 August 2008, <http://www.iaasframework.com/>

## BIOGRAPHIES

Peter Szegedi (szegedi@terena.org) is currently the Joint Research Activity 2 leader of the FEDERICA project. He received his MSc degree in Electrical Engineering at Budapest University of Technology and Economics (Hungary, 2002). He then worked towards a PhD at the Department of Telecommunications. His main research interests include design and analysis of dynamic optical networks, especially optical Ethernet architectures, network control and management processes. He worked for Magyar Telekom (2003-2007) then he joined TERENA in January 2008.

Sergi Figuerola (sergi.figuerola@i2cat.net) is the Coordinator of the Network Technologies Cluster of the i2CAT Foundation since 2004. He graduated in Telecommunica-

tion Engineering on the Technical University of Catalonia (UPC), in 2002, he holds a Masters in Project Management from La Salle (Universitat Ramón Llull, 2004). He is also pursuing his PhD within the UPC's Optical Communications Group. He is acting as Joint Research Activity 1 leader of the FEDERICA project and work package leader of the PHOSPHORUS project.

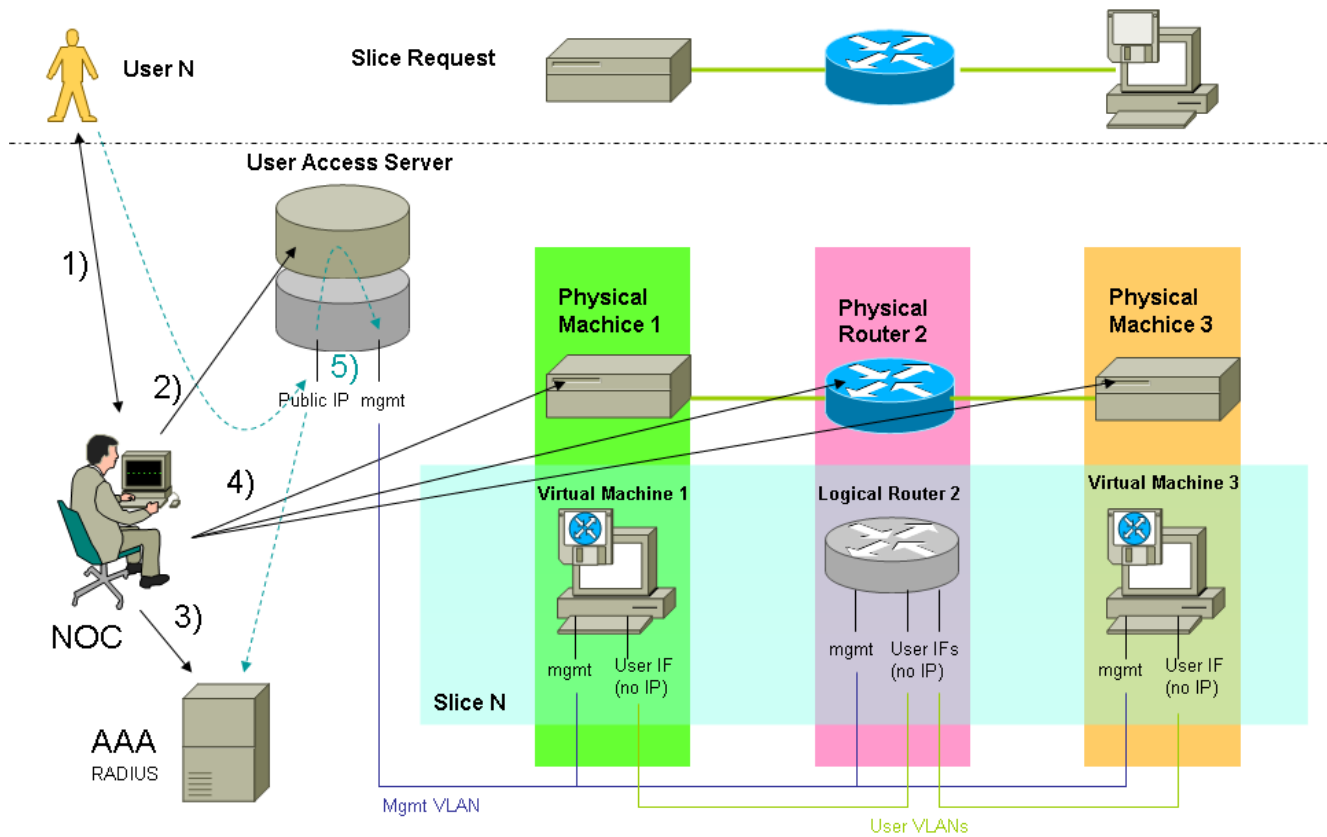
Mauro Campanella (mauro.campanella@garr.it) graduated in physics in 1985. He is at present working for the Italian National Research and Education Network (GARR) and is responsible for research activities. He is one of the creators of the Premium IP QoS service of the European NREN network backbone GÉANT and created the architecture of the Bandwidth on Demand service of GÉANT2 (AutoBAHN). He is acting as the coordinator of the FEDERICA project.

Vasilis Maglaris (maglaris@netmode.ntua.gr) is the Technical Coordinator of FEDERICA. He is a Professor of Electrical and Computer Engineering at the National Technical University of Athens (NTUA) since 1989, directing the Network Management and Optimal Design Laboratory (NETMODE). He holds an Engineering Degree from NTUA (Athens, Greece, 1974) and a PhD degree from Columbia University (New York, USA, 1979). Since October 2004, he has served as the Chairman of the NRENs' Policy Committee, responsible for the Pan-European advanced network platform GÉANT.

Cristina Cervelló-Pastor (cristina@entel.upc.edu) received the Telecommunication Engineering and PhD degrees in 1989 and 1998, respectively, both from the Technical University of Catalonia (UPC), Barcelona, Spain. She is currently an Associate Professor in the Department of Telematics Engineering of UPC, which she joined in 1989. She also belongs to the Technological Committee of the i2CAT Foundation. Her research trajectory has been centered on the field of optical networks. She is currently the Joint Research Activity 1 leader of the FEDERICA project jointly with i2CAT.



Fig 2 Basic slice provisioning procedure



- 1) Negotiation between the User N and the FEDERICA NOC on the specific request
- 2) Creation of public, private and management interfaces on the UAS for the user
- 3) Creation of User N credentials and decision on expiration date and time
- 4) Creation of Slice N: decision on Physical Machines (SUN, Juniper), physical interfaces, creation of VLANs
- 5) User N can access to his/her Slice N