CMST SI(1) 19-33 (2010) DOI:10.12921/cmst.2010.SI.01.19-33

COMPUTATIONAL METHODS IN SCIENCE AND TECHNOLOGY **Special Issue 2010,** 19-33

FEDERICA: Federated e-Infrastructure Dedicated to European Researchers Innovating in Computing Network Architectures

R. Krzywania[1]*, L. Dolata[1], P. Szegedi[2]

[1]*Poznań Supercomputing and Networking Center, ul. Noskowskiego 12/14, 61-704 Poznań, Poland*
[2]*TERENA Secretariat, 468D Singel, 1017AW Amsterdam, The Netherlands*
**e-mail: sfrog@man.poznan.pl*

(Received: 25 August 2010; revised: 27 October 2010; published online: 23 November 2010)

Abstract: The FEDERICA project is to create a virtual network environment that is dedicated to investigate and trial novel network protocols and services, as well as Future Internet architectural concepts and emerging applications. With this purpose in mind, a technology agnostic network infrastructure that consists of transmission devices, gigabit connections, and computing resources has been created. The physical FEDERICA network footprint is mapped to the existing GÉANT network links and additional NREN connections and creates its own infrastructure on top that allows resource virtualisation. The project partners do research on novel resource virtualisation techniques as well as share their operational experiences on virtualisation capable infrastructures during the project lifetime. Within the framework of FEDERICA, close cooperation has been established with various standardization bodies such as IETF, ITU-T, OIF, and IPsphere. Although, the main objective of the project is to create a physical environment that allows obtaining dedicated virtual slices of the infrastructure configured by the users, the project also facilitates information exchange, initiates technical discussions among the experts and disseminates scientific results primarily to national research and education network community. The FEDERICA project is partly funded by the European Commission under the INFRA-2007-1.2.2 "Deployment of e-Infrastructures for Scientific Communities" call of the Seventh Framework Program. The FEDERICA project [1] has been devised to provide support and research on current and Future Internet technologies and architectures. The project is linked to the European FIRE initiative [2] and the European Future Internet Assembly [3]. Other similar initiatives exist worldwide, e.g. GENI [4] in the United States in Europe, and AKARI [5] in Japan.

Key words: virtual network infrastructure, virtualisation, pan-European test bed

I. INTRODUCTION

There is a high demand for networking infrastructures providing reliable and efficient interconnectivity between offices, laboratories, data storage areas or grid infrastructures. The technology is developing new standards for transmission speed, protocols, content delivery awareness and many more features that a user is expecting from the network. Users are also willing to do experiments and research on the network, having access to modern technologies and being able to validate their own solution at the same time. But the access to enterprise solution equipment is usually limited by the cost and availability of such equipment on the market. Thus, some solutions that need to be validated are facing issues of lacking test beds with the required throughput, efficiency and functionality, which prevents further research. Often there is high demand on having tests and experiments at the very early stage of development. This may affect the progress of research or innovations in networking and associated industries.

In order to cope with this challenge, the FEDERICA project was created giving access to configurable, scalable, and efficient networking and computing infrastructure spread across Europe. FEDERICA is a European Commission funded project aimed at providing a virtual network infrastructure based on physical backbone links offered by NRENs, GÉANT, and other project partners. The aim is to support future Internet initiatives, yet the project is not limited to this. It is important to stress that FEDERICA is not solving future Internet issues, but rather provides tools for research in this area. The main objective of FEDERICA is to deliver scalable and reliable infrastructure enabling users to create virtual slices on top, in order to perform disruptive experiments at global scale. The vision is that users could validate their research in the

environment that is extremely close to a real physical environment, and will not bear costs of physical equipment at the same time. An answer to this idea was to provide a large, highly distributed pan-European infrastructure built of 14 sites with virtualisation service on-board and a set of configuration and management tools, utilising modern technologies of virtualisation [6]. Since connections between sites are realised via efficient L2 dedicated 1 Gb/s circuits, users gain a unique opportunity to do experiments with network protocols and applications which operate at L2, L3 and above. Most of the user test cases can be modelled in the FEDERICA infrastructure as virtual slices, giving users the opportunity to perform experiments similar to physical infrastructures, but without expenses on specialised hardware.

The FEDERICA consortium consists of experienced partners providing resources and knowledge into the infrastructure construction and research activities within the project. The following organisations are involved in FEDERICA [1]:

1. Consortium GARR (GARR), Italy, as project coordinator,
2. Zajmove Sdruzeni Pravinickych Osob (CESNET), Czech Republic,
3. Verein zur Förderung eines Deutschen Forschungsnetzes e.V. (DFN), Germany,
4. Fundação para a Computação Científica Nacional (FCCN), Portugal,
5. Greek Research and Technology Network S.A. (GRNET), Greece,
6. Nemzeti Információs Infrastruktura Fejlesztési Intézet (NIIF/Hungarnet), Hungary,
7. Instytut Chemii Bioorganicznej PAN (PSNC), Poland,
8. Entitad Publica Empreserial RED.ES (Red.es/RedIRIS), Spain,
9. SWITCH – Teleinformatikdienste für Lehre und Forschung (SWITCH), Switzerland,
10. TERENA Secretariat (TERENA), The Netherlands,
11. Fudacio i2CAT, Internet i Innovacio digital a Catalunya (I2CAT), Spain,
12. The Royal Institute of Technology (Kungliga Tekniska Högskolan) (KTH), Sweden,
13. Institute of Communication and Computer,
14. Systems (ICCS), Greece,
15. Universitat Politècnica de Catalunya (UPC), Spain,
16. Juniper Networks (Juniper), Ireland,
17. Martel GmbH (Martel), Switzerland,
18. HEAnet Ltd. (HEAnet), Ireland,
19. Delivery of Advanced Networking Technologies to Europe Ltd (DANTE), UK,
20. Politecnico di Torino (PoliTo), Italy,
21. NORDUnet A/S (NORDUNET), Denmark.

II. NETWORK INFRASTRUCTURE

The FEDERICA infrastructure [7] consists of 14 virtualisation capable physical sites interconnected by 1 Gb/s capacity GÉANT+ connection services [8] and other tunnelling solutions at pan-European scale. The logical topology is depicted in Fig. 1.

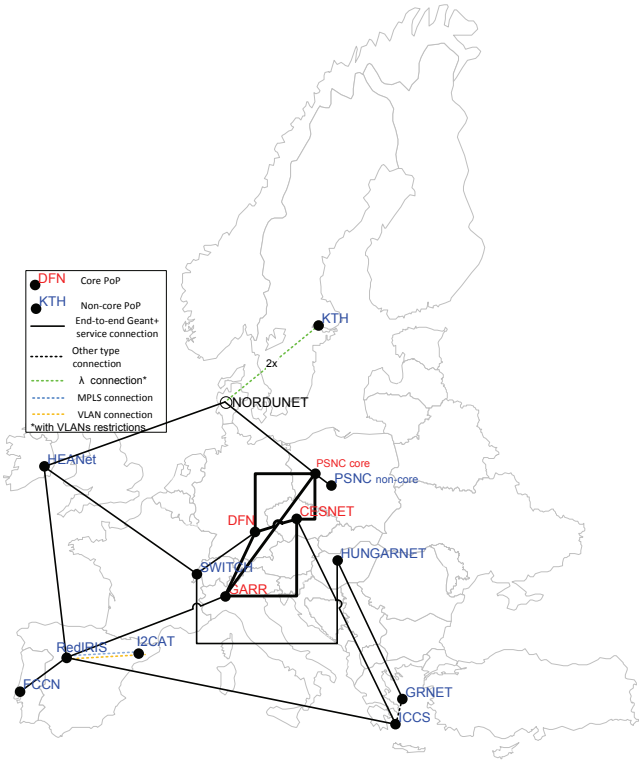

Fig. 1. Logical topology of FEDERICA network

There are 4 core Points of Presence (PoPs), which are fully connected by dedicated optical channels and equipped with more powerful virtualisation capabilities. The other 10 non-core PoPs are connected in less restrictive manner, distributing FEDERICA services a pan-European scale. The designed topology was aimed at reliability, resiliency and load balancing, especially the part on top of the GÉANT2 network which gives the main data transport.

The 4 core PoPs were initialised and delivered for service in the first place. No NREN equipment is placed in the middle of the transmission path, i.e. the optical interfaces from the GÉANT equipment are directly connected to FEDERICA nodes. The GÉANT+ interconnection service guarantees reliable and efficient connectivity between the PoPs, assured by one of the biggest research networks around the world. The circuits are operated and monitored by GÉANT NOC with guaranteed service availability and fast

recovery time in case of failures. The cost of the circuits is a contribution of particular partners to the project. Currently, 18 FEDERICA connections are delivered by the GÉANT network including all core connections (PSNC-CESNET, PSNC-GARR, GARR-DFN, DFN-PSNC, CESNET-DFN, CESNET-GARR) and non-core connections as follows:

1. HEAnet – SWITCH,
2. HUNGARNET – SWITCH,
3. RedIRIS – HEAnet,
4. RedIRIS – FCCN,
5. RedIRIS – ICCS,
6. RedIRIS – GARR,
7. RedIRIS – I2CAT,
8. GRNET – HUNGARNET,
9. ICCS – CESNET,
10. SWTCH – DFN,
11. NORDUNET (KTH) – PSNC,
12. NORDUNET (KTH) – HEAnet.

The GÉANT+ connections are shown in Fig. 2 (core connections) and in Fig. 3 (non-core connections).

Non-core PoP connections are less restrictive in terms of intermediate NREN equipment being placed in between the GÉANT PoP and FEDERICA nodes. It is possible to include additional switches, or pass the traffic through an NREN infrastructure. The only requirement here is to guarantee the 1 Gb/s capacity and allow configuration of VLANs or MPLS tunnels in order to create virtual connections.

In some specific cases, the FEDERICA PoP interconnections are implemented in different ways, which needs more explanation.

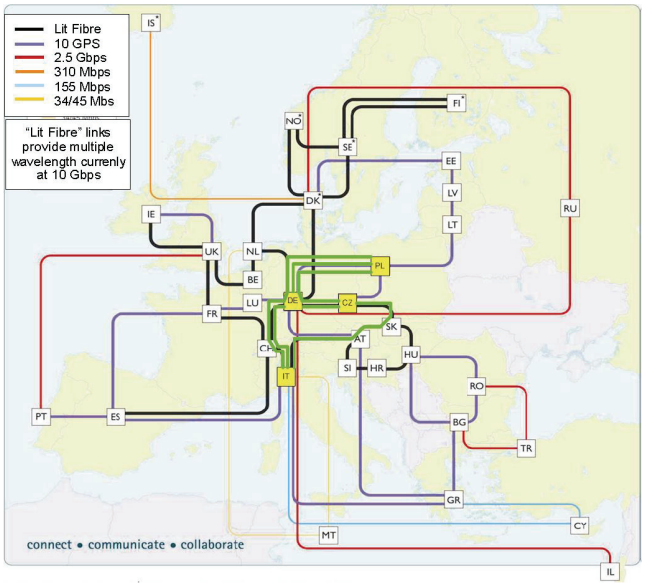

Fig. 2. Core GÉANT+ connection services

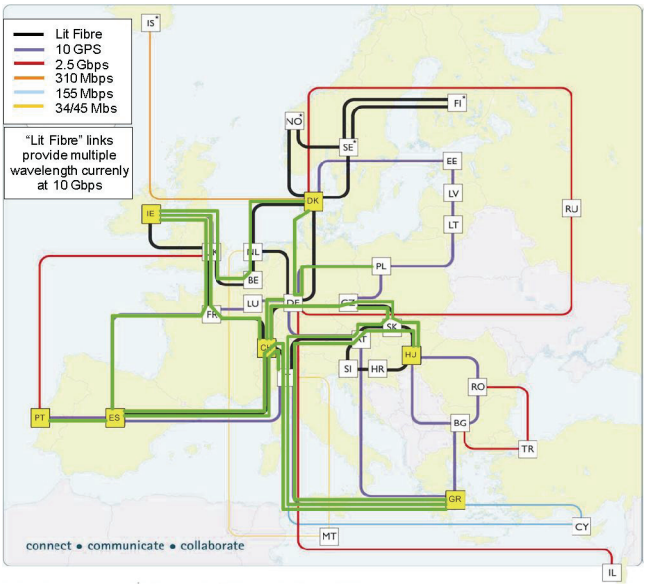

Fig. 3. Non-core GÉANT+ connection services

The first case is the KTH PoP. KTH has no direct access to the GÉANT infrastructure, therefore its connection must be routed via intermediate networks. First, KTH traffic is passed through the UNINETT network using dedicated VLANs (one per connection), and then the NORDUnet network terminates the VLANs in the GÉANT PoP in Denmark. Using VLANs at the level of FEDERICA physical connections in general creates a limitation for the end users. Since virtual links will be defined again as VLANs (using QinQ feature along the path), users will not be able to use full L2 features in their virtual network using these connections.

One of the strategic PoPs of FEDERICA is located in Madrid, Spain, at RedIRIS headquarters. A single PoP usually has the number of connections from one up to three. RedIRIS PoP must be connected to five other sites (FCCN, I2CAT, ICCS, GARR, and HEAnet) of which four connections are implemented via the GÉANT infrastructure except the one to I2CAT. This requirement had an impact on the selection of a more powerful switch that is installed in other non-core location.

Another specific case is Germany where the GÉANT PoP is located at DFN headquarters in Frankfurt, while the FEDERICA node is located at the Friedrich-Alexander University of Erlangen-Nuremberg (FAU). DFN PoP is a core node so it has to fulfil the requirement of having a direct connection between the GÉANT equipment and the FEDERICA node. Therefore, a lambda connection is established between Frankfurt and Erlangen, entirely dedicated to FEDERICA purposes.

Usually there is one PoP per country except three cases such as Spain, Greece and Poland. Connections between two PoPs in the same country are implemented using national or regional network links. In Spain, the connection between RedIRIS and I2CAT uses the NREN infrastructure to reach Barcelona, where I2CAT is located, from Madrid. Due to hardware in the middle, VLAN technology is used to realise this connection. This leads to a similar situation as in case of KTH, where users are not allowed to define L2 VLANs inside their slices at that relation. It is an ongoing work to solve this issue by, for example, using MPLS tunnels on these links. Two PoPs are also presented in Poland, where one of them is a core PoP and the other one is a non-core PoP. In practice, these are installed in the same rack mount. The non-core PoP JUNIPER EX3200 switch is connected directly to the JUNIPER MX480 core switch. The idea of creating an additional PoP in PSNC was to deliver more computational power into the FEDERICA infrastructure for more demanding applications (e.g., routing engines with large number of routing entries). The last case, where two PoPs are present in the same country is Greece. Two Greek research centres are partners of the FEDERICA consortium; GRNET, as the Greek NREN, and the Institute of Communications and Computer Systems (ICCS) that is a non-profit academy research body. The locations are connected to each other with dedicated links.

For operation and maintenance purposes of the FEDERICA infrastructure, a Network Operation Centre (NOC) was created to be led by KTH. The NOC responsibility is twofold; they ensure the sustainable availability of resources, as well as configure the physical resources to create virtual slices.

III. VIRTUALISATION ARCHITECTURE

Each FEDERICA PoP has been built according to the same schema in order to provide virtualisation services for end point machines, software or hardware routers, switches and links. The general PoP concept assumes that each location is equipped with one or more virtualisation servers and a FEDERICA switch that is connected directly to the GÉANT infrastructure or alternatively to partner's infrastructure. As it was mentioned before, there are two types of PoPs defined in FEDERICA; core PoPs and non-core PoPs. The difference between them is mainly in the level of network connectivity and network services being provided to the users for virtualisation purposes.

FEDERICA Core PoPs

A typical core PoP is depicted in Fig. 4. The core FEDERICA switch is JUNIPER MX480 that also includes routing capability. This switch is then directly connected to the Alcatel MCC1678 GÉANT box with 1 Gb/s optical Ethernet interfaces. The number of interfaces determines the number of connections that can be established towards neighbour PoPs. One of the issues during the FEDERICA infrastructure setup was the availability of those interfaces in Alcatel GÉANT boxes to be used by GÉANT+ services. Basically, GÉANT provides the MCC1678 network interfaces in two formats; 1 × 10 Gb/s or 10 × 1 Gb/s. The second format was far simpler to implement connections, as each link consumes only one Alcatel interface and one JUNIPER switch interface. The traffic is simply passed in VC4 containers, remaining transparent for the FEDERICA equipment. The usage of 1 × 10 Gb/s interfaces introduces some complexity in the realisation of connections, as each connection has to use a VLAN tag to be separated from other GÉANT+ services on the same physical interface. This case implies a situation similar to the one with KTH and I2CAT, where users are not able to use VLAN tagging but L3 functionality inside their slices, if those links are used. In collaboration with GÉANT and project partners, the FEDERICA consortium was trying to minimise the presence of 1 × 10 Gb/s interfaces in the infrastructure and its effect on virtualisation capabilities. Eventually, there are only four connections where users cannot apply VLAN tagging inside their slices, and none of them is due to having a 1 × 10 Gb/s interface. The connections without user VLAN tagging capabilities inside slices (e.g., using QinQ) are FCCN – RedIRIS (due FCCN NREN equipment limitations), KTH – PSNC, KTH – HEAnet (due passing UNINETT with VLAN in both cases), and RedIRIS – I2CAT (due to limitations of the RedIRIS equipment).

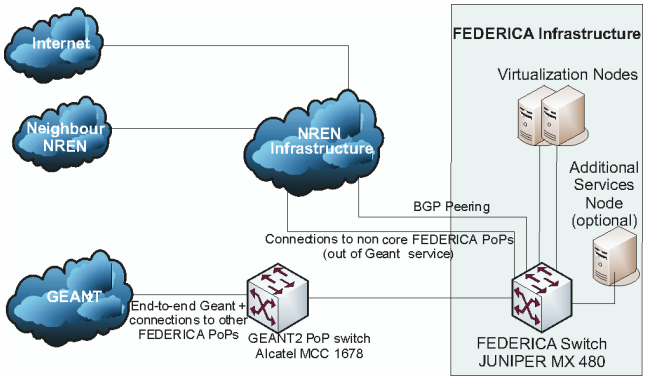

Fig. 4. A typical core PoP architecture

A FEDERICA switch is usually connected to two virtualisation servers with 8 × 1 Gb/s Ethernet interfaces each. The virtualisation servers have to meet several requirements in order to be selected for the FEDERICA infrastructure, as all sites must have the same type of equipment with similar configuration. FEDERICA focuses on network virtualisation, so the servers not need to have much computational power; however, they need to support virtualisation capabilities and have enough memory, network interface, and hard disc space to support multiple instances of operating systems. The best choice, at the time when the FEDERICA infrastructure was created, was the Sun X2200 M2 server that offers enough resources for virtualisation and computations, and has a very good price-value indicator. The servers are equipped with 2 × QuadCore AMD Opteron 2356 processors (2.3 GHz), 16-64 MB of RAM, 1TB of hard disc space, and 7 × 1 Gb/s interfaces plus one 1 Gb/s interface for management purposes. The servers are in size of 1U and fits in 19'' rack mount. In order to simplify the infrastructure management, it was decided to install the same virtualisation operating system on all virtualisation servers. The operating system requirements were to be well documented, including configuration examples, to be efficient and capable to handle 8 physical interfaces per physical machine, to be reliable and to provide an API for automatic configuration. Finally, the VMware software with an ESXi free of charge license has been chosen. The free version has several limitations compared to the commercial one, but its features have been considered to be sufficient for FEDERICA purposes. One of the FEDERICA objectives was to automate slice creation and management of infrastructure, which requires a good quality and well documented API. This requirement was perfectly fulfilled by VMware compared to other alternatives such as OpenVZ, or Xen. The number of features and their practical applications from the point of view of FEDERICA was an additional reason of this choice. VMware, as a well the known operating system, offered a definite way of running virtual machines and a consistent data format to create and transport ISO images of virtual machines. It also supports NFS and mounting remote drivers that helps to organise FEDERICA data repository and management procedures.

FEDERICA is mainly focusing on virtualisation of network resources, so computational power is far less significant than switching and routing capabilities of the infrastructure. Therefore, the main element of the core installations is the Juniper Networks MX480 Ethernet Services Router (Fig. 5) [9].

Fig. 5. Juniper Networks MX480 Ethernet Services Router

It is an Ethernet-optimised edge router that provides both switching and carrier-class Ethernet routing. With a capacity of up to 240 Gb/s, full duplex, the MX480 ESR provides a dense, highly redundant platform primarily targeted for dense dedicated access aggregation and provider edge services in medium and large POPs. Juniper Networks MX480 enables a wide range of business and residential applications and services, including high-speed transport and VPN services, next-generation broadband multi-play services, and high-volume Internet data centre internetworking. The MX480 router is 8U tall. Five routers can be stacked in a single floor-to-ceiling rack, for increased port density per unit of floor space. The router provides eight slots that can be populated with up to six Dense Port Concentrator (DPC) cards and two Switch Control Boards (SCBs) in non redundant fabric configurations. Fully populated, the MX480 provides up to 240 Gigabit Ethernet or up to 24 10 Gigabit Ethernet ports. Six types of DPC cards are available, each of which consists of four Packet Forwarding Engines and enables a throughput of 10 Gb/s.

For FEDERICA purposes, each MX 480 router was equipped with a single DPC card with 40 × 1 Gb/s interfaces and four Packet Forwarding Engines on a single board. The Packet Forwarding Engines on a DPC are equipped with purpose-built application-specific integrated circuits (ASICs) that perform packet processing and forwarding. Each Packet Forwarding Engine consists of one I-chip for Layer 3 processing and one Layer 2 network processor. This configuration guarantees high throughput and efficient data transfer even in large scale installations, which can be virtualized in the FEDERICA infrastructure.

Each MX480 router operates under control of JUNOS® software, which supports the creation of logical routers. Logical routers provide the flexibility that network operators require to consolidate and tune the operation of their networks to accommodate different services without deploying additional physical routers. Logical routers allow partitioning a single physical router into multiple logical routers, where each logical router performs independent routing tasks. Logical routers can be given to the end users within a FEDERICA slice, and seen by users as a physical, fully operational Juniper MX480 router.

The following logical router capabilities are supported by the JUNOS software:

- A maximum of 16 logical routers can be configured on a single physical router.
- Physical and logical interfaces can be assigned to a logical router. After a physical or logical interface is assigned to a logical router, it is considered part of that logical router and cannot be assigned to another logical router. Every logical router takes each input packet it receives on an assigned physical or logical interface, makes a forwarding decision, and then forwards the packet on one of its assigned physical or logical interfaces towards the next hop.
- All physical interface properties (such as encapsulation types and interface related options) are configured using the command line interface (CLI) of the physical router's main router. The main router implements the traditional concept of a router and supports routing across all physical and logical interfaces that have not been assigned to a logical router.
- The Tunnel Services PIC (embedded in MX DPCs) supports the configuration of logical tunnel interfaces that provide point-to-point connectivity between different logical routers configured within the same physical router. This allows the switch fabric to provide inter-logical router connectivity instead of wasting expensive physical interfaces in front of the chassis.
- Unicast routing protocols such as the Routing Information Protocol (RIP), RIP nextgeneration (RIPng), Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (ISIS), and the Border Gateway Protocol (BGP) are supported by each logical router.
- Multicast protocols, such as Protocol Independent Multicast (PIM) and Distance Vector Multicast Routing Protocol (DVMRP) are supported by each logical router, including Rendezvous point (RP) and source designated router (DR) functionality.
- For each logical router, the routing information base (RIB) is maintained separately from the forwarding information base (FIB). The RIB contains all of the routing information that is received from the logical router's peers, including information learned from all routing protocols. The JUNOS software installs active routes from the RIB into the FIB. The FIB is the table that the logical router uses to forward IP datagrams. At a minimum, the FIB contains the network interface identifier and the next hop information for each reachable destination network prefix.
- Multiprotocol Label Switching (MPLS) signalling protocols such as the label distribution protocol (LDP) and resource reservation protocol (RSVP) signalling can be configured on each logical router. Additionally, provider edge and core router functionality such as Layer 2 virtual private networks (VPNs), Layer 3 VPNs, circuit cross-connect (CCC), and virtual private LAN services (VPLS) are supported by each logical router.
- Logical routers are separate from a policy perspective; each logical router can be configured with its own firewall filter policies and routing policies. Firewall filter policies control which packets are allowed to transit each logical router's interfaces to destination networks and which packets are sent and received by the logical router's control plane. Routing policy controls the import and export of routes between the routing protocols and the routing tables, and between the routing tables and the forwarding table for each logical router.

Logical routers provide a flexible model that continues to support subscriber VPNs. A physical router can be configured with multiple software-based logical routers where each logical router is a partition of the physical router's resources. Then, each software-based logical router can be configured with multiple virtual routing and forwarding instances (VRFs) where each VRF supports a different VPN. Note that a virtual router is not the same as a logical router. A virtual router is a simplified routing instance that has a single routing table. A logical router is a partition of a physical router that can contain multiple routing instances and routing tables. Thus, a given logical router can be configured to support multiple VPNs.

In addition to virtualisation and transport capacities, core PoPs operate also as gateways to the Internet through partners' network infrastructure and keeps routing information consistent with dedicated BGP peering. For security reasons, the public Internet is unreachable from inside of a virtual slice; however, there are some practical reasons why this should not be so. One of these reasons is the possibility of virtual operating system updates from the most recent public repositories (e.g., in case of Linux) and the other one is the potential federation with other virtu-

alisation capable infrastructures (e.g., GENI, PlanetLab, OneLab). The connection between the FEDERICA infrastructure and external test beds or other virtualisation infrastructures via the public Internet has limited capabilities in terms of QoS; moreover, it is vulnerable for potential hostile activities. The seeking of proper solutions to this issue is now being extended over the FEDERICA project scope and a dedicated European Comission funded project called NOVI has been established to put more research efforts on this.

At each core PoP, in addition to virtualisation capabilities and resources, an extra server may be installed in order to support the users and the NOC in day to day operations. PSNC PoP is equipped with an additional server implementing the User Access Portal for slice requests and user control, as well as the Data Storage Area for virtual images. Another server dedicated to infrastructure monitoring is installed in CESNET PoP.

FEDERICA Non-core PoPs

There are 10 non-core PoPs offering virtualisation services at a similar level in order to geographically stretch out the core PoPs' functionalities across Europe. The real transmission links between sites allow the users to experience real delays instead of software emulated delays that may be misleading in some network experiments.

Each non-core PoP is built according to the architecture depicted in Fig. 5.

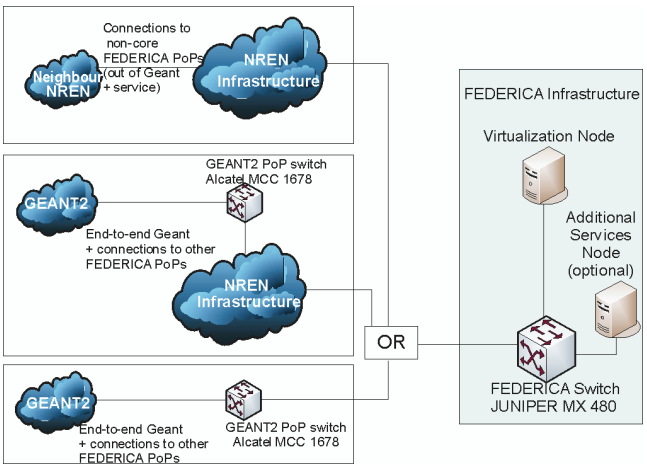

Fig. 5. A typical non-core PoP architecture

At non-core PoPs EX series JUNIPER switches are used connecting remote PoPs and local virtualisation resources. The neighbour PoPs can be connected in three different ways:

- Partners' infrastructure can be used to connect remote PoP directly to JUNIPER EX series switch (e.g., RedIRIS – I2CAT connection).
- Partners' infrastructure can be used to reach the GÉANT PoP equipment, if the location of the FEDERICA PoP is not the same (e.g., FCCN-RedIRIS at FCCN site).
- The FEDERICA node may have direct access to the GÉANT PoP that is similar to the core PoPs connections (e.g., HEAnet-SWITCH, RedIRIS-GARR, etc.)

At non-core PoPs the minimum requirement is to have a single Sun X2200 M2 virtualisation server with the same configuration as in core PoPs. It is envisaged, as the FEDERICA infrastructure is growing, that additional servers will be required in order to provide supporting services for virtualisation such as data storage or secondary User Access Portal. This kind of services can be located not only in the core PoPs but also in the non-core PoPs, according to the needs, development, and utilisation of the infrastructure.

Fig. 6. Juniper Neworks EX3200 switch

All non-core PoPs are equipped with a JUNIPER EX3200 (Fig. 6) switch, except RedIRIS where due to a large number of connections two stacked EX4200 switches have been installed. The functionality of both solutions, from the perspective of the FEDERICA infrastructure, is exactly the same. The JUNIPER EX3200 switch provides 24-port 10/100/1000BaseT (8-ports PoE) + 320W AC PS, 4-Port 1G SFP Uplink Module, and two Small Form Factor Pluggable (SFP) 1000Base-SX Gigabit Ethernet Optics. The switch offers an integrated Routing Engine (RE) that runs the same modular JUNOS software as Juniper Networks router products to ensure the consistent implementation and operation of each control plane feature across an entire Juniper infrastructure. This guarantees full interoperability and compatibility between core and non-core PoPs in the context of technology and communication. The EX 3200 series switches support Q-in-Q tunnelling allowing users to define their own VLAN tags within virtual links realised in VLAN technologies. In addition to VLAN support, EX3200 provides MPLS-based (CCC) links virtualisation, giving more flexible management capabilities and resource control.

FEDERICA Slices

FEDERICA slices can be composed of the following elements:

- Server end points – created from standard OS distributions or delivered from a user as virtual machines on physical servers. It performs as a regular server and may contain any user application, with respect to a software license agreement and the FEDERICA policy.
- Software router – created as a regular OS virtual machine, but with routing daemon software on board (e.g., Quagga). Such slice element may perform routing functions in the slice, according to user's requirements.
- Hardware router – created as a logical router on the JUNIPER MX480 router, available only in core PoPs. It may perform as a regular routing or switching entity in the user slice.
- Software switch – created as a regular OS virtual machine, but with a switching software on board (e.g., bridge Linux package). Such slice element may perform L2 switching functions in the slice, according to user's requirements.

Any of the above elements may be created with different attributes, e.g. defining the number of interfaces, available hard disc space, memory, CPU time, etc. according to user requirements for the slice. The connections between elements can be virtualized by any technology available in the FEDERICA infrastructure that is VLANs with QoS guarantee and MPLS paths. Due to using of 1 Gb/s capacity links between PoPs, the resources configuration process must be very careful, preventing overprovision of resources for intra-slice connectivity. Slice elements can by physically located anywhere in the network that is appropriate for the user. Elements may be virtualised in a single PoP location, if capacity and throughput requirements are more important, or may be geographically distributed in multiple FEDERICA PoPs, if realistic delay conditions are required. FEDERICA NOC is responsible for the proper resource assignment and configuration taking the particular user requirements into account.

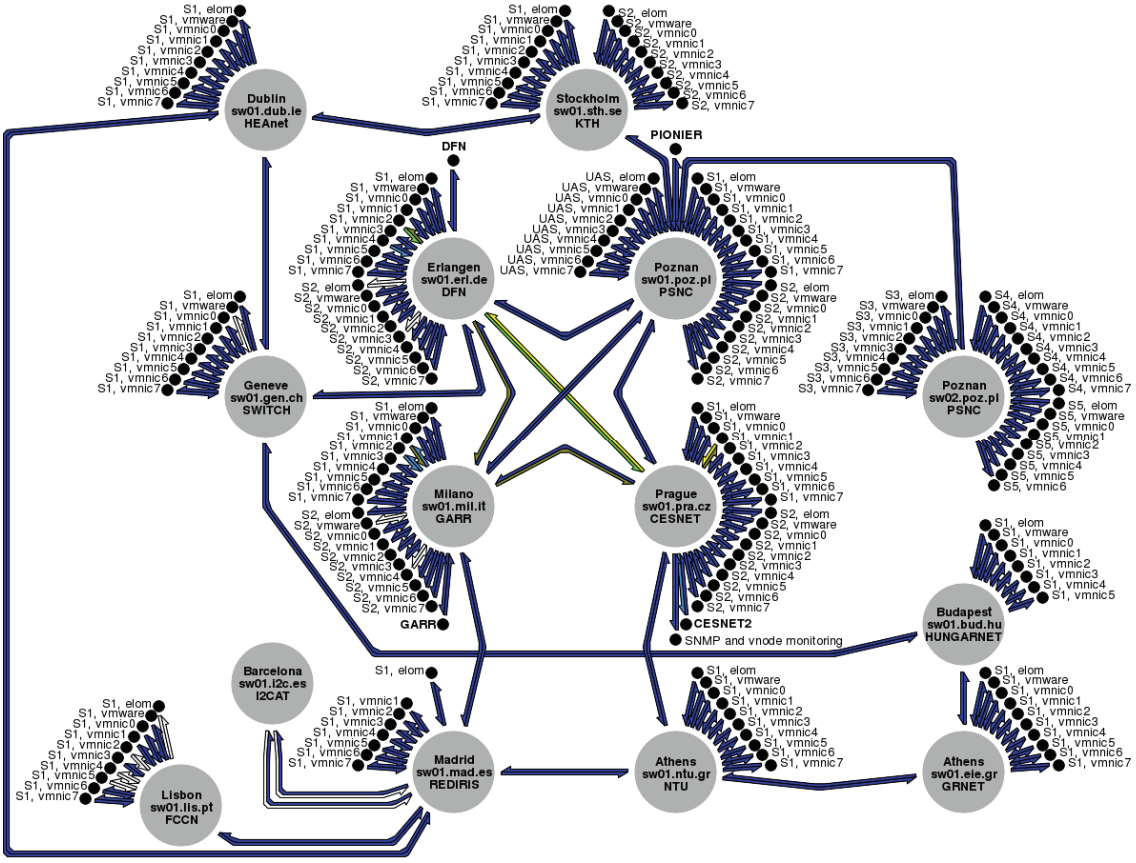

Fig. 7. FEDERICA monitoring map

FEDERICA Monitoring System

In order to keep track of resource utilization and support automated process of resource management, a dedicated monitoring system for the FEDERICA infrastructure was created.

There are four types of data that are being monitored. The first one is the link utilization that is monitored via SNMP on servers, routers and switch interfaces. As an example, Fig. 7 depicts the map of FEDERICA connections and their actual bandwidth utilisation. Another type of data is the utilization of physical resources that is monitored via VMware interfaces and Nagios system, giving information on virtual machines and resources consumed by them. The third type of data is the peformance of virtual resources being monitored within user slices if possible, gathering information on performance of virtual machines inside the VMware system. This kind of information is very useful for tracking efficiency issues and support users in their experiments. The last type of data is the NetFlow traffic analysis, which is a passive monitoring solution with dedicated network interface cards, in order to analyse traffic and provide detailed information on a particular type of router traffics. For NetFlow analysis, additional hardware has to be installed to the PoPs, so currently this function is only available at CESNET PoP.

The monitoring system delivers information to the FEDERICA monitoring web portal in both graphical and numerical formats. The information is stored in a database mainly for statistical and historical purposes. In order to support automation of infrastructure management and to enable a third party application to gain access to monitoring information, a dedicated API has recently been developed and put into the operational status. In the last phase of the FEDERICA project, the main priority is to minimise the number of manual procedures in the NOC and give more rights to the user to create and configure their own slices supported by a set of automated tools. These tools should collect and analyse user slice requests, assign resources to slices, configure equipment, and give users permission to access their slices based on monitoring data, NOC experiences, and dedicated algorithms.

IV. PROVISIONING AND MANAGEMENT OF FEDERICA SLICES

The FEDERICA infrastructure is based on its dedicated physical substrate, therefore it provides a number of powerful features for future network research. Virtual slices of the infrastructure can be assigned to multiple users in parallel allowing them to experiment e.g., various configuration scenarios at the same time. Slices are completely independent of each other so they can even accommodate disruptive research activities which may have fatal influence on the production network environment. Another feature of FEDERICA is the reproducibility of tests. Each experiment can be repeated at any time keeping the same configuration and network conditions.

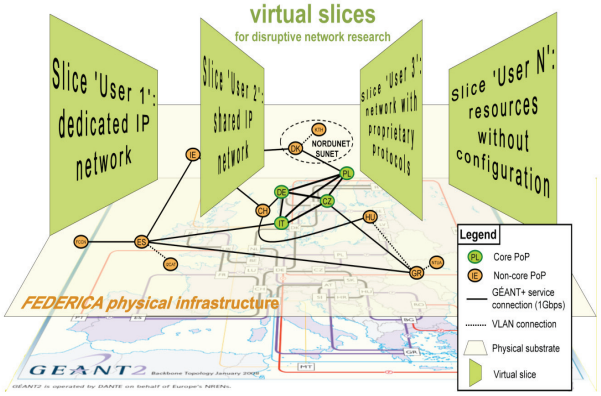

Fig. 8. FEDERICA slicing concept

The overall idea of virtualization in FEDERICA is illustrated in Fig. 8. Taking the project concept and users' requirements into account, FEDERICA engineers have considered two potential approaches for management of physical infrastructure and slice provisioning:

- Providing users with information about the physical network and computational resources in a slice environment. The main disadvantage of this solution is the restriction of slice resources to the number of physical elements in the physical substrate.
- Hiding physical infrastructure information behind virtualization techniques applied for both computing and network resources, in order to provide virtualised entities for a slice environment. This is a much more flexible solution restricted only to the performance of the physical resources.

Virtualization is defined as the capability to create a virtual version of a physical resource, both in the computing and network environment. The virtual resources (e.g., a virtual circuit, a virtual network device, a virtual sever, or a disk partition) are usually created by segmenting a physical resource. This approach brings a higher level of utilizing of physical resources and a higher level of flexibility. In almost all cases virtual resources better fulfil users' requirements and reduce the cost of infrastructure over provisioning. The next important advantage of virtuali-

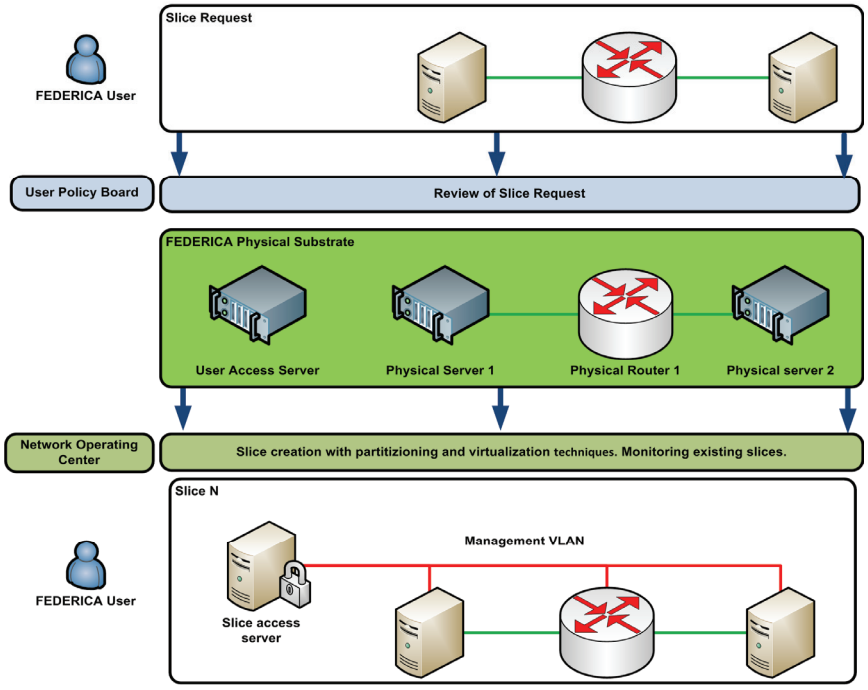

Fig. 9. Slice provisioning process

zation is the possibility to move virtual resources from one platform to another. This brings an opportunity to react better to failures in physical substrate.

Unfortunately, beside advantages, the virtualisation approach increases the complexity of management processes. Operators using virtualisation techniques have to divide their infrastructure substrate into two separate levels such as:

- Physical infrastructure level – the infrastructure part which contains hardware and software providing virtualisation functionality.
- Virtualised infrastructure level – the infrastructure part which contains all the virtualised resources created on top of the physical infrastructure level.

In practice, this division requires the same separation in the management process either. FEDERICA has also defined two levels of infrastructure substrate. The first is the physical substrate of the project that supports virtualisation and contains all devices located in all nodes. The second level contains virtualised resources created during the slice provisioning process. FEDERICA takes challenge to control and manage its own divided architecture and provide a stable and flexible environment for future Internet researchers.

The project has created two working groups [11]: the User Policy Board (UPB) and the Network Operation Centre (NOC), which are responsible for administrative and technical operations, respectively. In the early phase of the project, the FEDERICA UPB was mainly responsible for defining procedures, policies and preparing all the basic documents that are needed to submit a FEDERICA slice request by an external user. Currently, the UPB performs the initial evaluation of all the research projects that wish to use or to be connected to the FEDERICA e-Infrastructure. It ensures essentially that the desired usage is non-commercial, the sources and destinations are reachable, interfaces are compatible and the infrastructure can cope with the requested capacities within the requested time frame. The FEDERICA UPB also ensures the user privacy (if needed) and collects feedback from users on the quality of FEDERICA procedures, level of support, ease of use, as well as the achieved scientific results.

The FEDERICA NOC responsibilities do not differ from the general actions performed by other NOCs, as it was discussed in the previous section. In addition to that, FEDERICA NOC is also responsible for the FEDERICA slice provisioning process. This includes the creation and monitoring of virtual slices required by the users. The basic slice provision process is presented in Fig. 10.

A slice lifetime starts when the user submits the slice request to the UPB including the specification of the planned experiment. The specification should include time constraints, description of work and main objectives,

topology information, and configuration requirements. The UPB analyses the slice request and approves or rejects that. Upon approval, the technical request is passed to the NOC. The NOC administrators prepare the set of physical resources from infrastructure substrate and configure the virtual slice environment. The fully configured slice environment is then provided to the user in order to carry out the experiment within the given time frame. After the end of the slice lifetime all resources are released and can then be used for creating other slices. At the current stage of the project all FEDERICA NOC actions are performed manually. FEDERICA engineers are working on a set of automated tools (called the toolbench) to support the NOC processes. These tools could speed up the whole configuration process, and make the slice provisioning more dynamic. The main parts of the FEDERICA toolbench are described in the next chapter.

The main feature of FEDERICA is that the users are able to fully configure and manage the resources of their own slices, without affecting the operation of the physical infrastructure. In order to get this, the NOC creates a private management network for each user. This network is separated from the public Internet and can be accessed via the dedicated Virtual Access Server (VAS) only. This solution prevents the unauthorized or unscheduled configurations and the VAS also acts as an access authentication engine.

Compared to other projects, FEDERICA appears to be more flexible. The FEDERICA nodes are not restricted to run dedicated operating systems. In practice, a user can upload any appliance which is able to run on the VMware ESXi environment. Nodes can work as common servers, software routers or switches. Compared to PlanetLab or OneLab, FEDERICA is based on dedicated data-links with a higher throughput and higher level of reliability, instead of using best effort public Internet connections. It means that experiments can be done at a lower layer than the IP layer.

FEDERICA toolbench

The architectural and management concept of FEDERICA forces the NOC to use various virtualization and configuration techniques. However, it is very difficult for the administrators to handle different client interfaces for various devices in the physical substrate. The complexity of slice configuration slows down the whole FEDERICA provisioning process. The growing number of slices worsens everything. FEDERICA engineers realised this problems and have started to work on a toolbench designed to support the management process. The toolbench has two major parts; one is to support the FEDERICA users and the other one is to manage the physical substrate. The development has been done in parallel with a manual management approach in order to allow users to use FEDERICA even in the early stage of the project.

FEDERICA User Portal [12]

The fruitful results of the FEDERICA users' experiments may have the biggest impact on the success of the whole FEDERICA project. Their scientific results, validated by the FEDERICA infrastructure, can demonstrate how important is the role that FEDERICA plays in the Future Internet research. In order to facilitate the usage of FEDERICA, the project partners have taken an effort to make FEDERICA resources available in the most convenient and flexible way. All the user recommendations and internal project partners' inputs have been taken into account when the FEDERICA User Portal (FUP) was specified and developed. The FUP is publically available, and allows users to submit and manage FEDERICA slice requests.

Each FEDERICA user has to provide a basic specification for the slice creation. This specification consists of:

- Time constraints of the experiment.
- Description of the purpose and the main subject of the experiment.
- Virtual topology information (required nodes, routers, switches and connections).
- Appliances files structure (virtual appliances files which should be deployed on FEDERICA physical nodes).

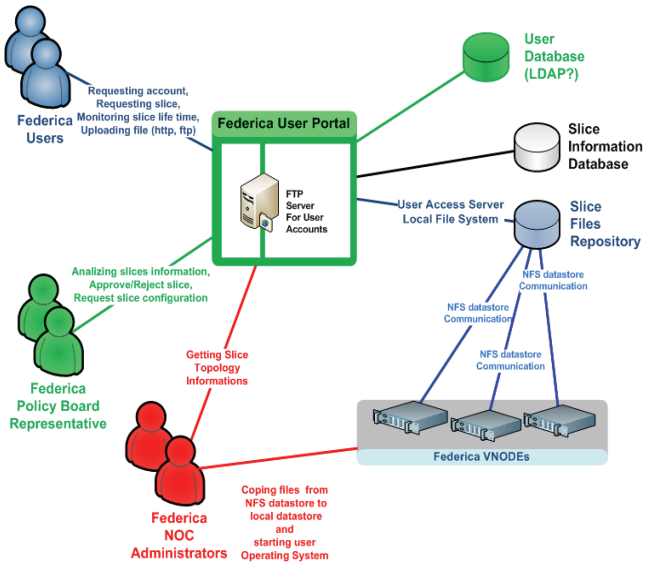

Fig. 10. FEDERICA User Portal architecture

The FUP keeps track of the slice history including the list of people who approved, created, and removed the slice. FPU can also be used to provide feedback on the usage. The results can be used in future experiments. The overall architecture of the FEDERICA User Portal is shown in Fig. 10.

FEDERICA Slicing Tool

The FEDERICA slicing tool is willing to automate some of the NOC procedures, therefore it slightly reduces the slice configuration time. Research projects such as UCLPv2 [15] and MANTICORE [16] are dealing with management approaches that could fit in the FEDERICA scope with some constraints. These projects are following the novel Infrastructure as a Service (IaaS) approach [14].

The IaaS framework is an open-source solution made available by Inocybe Technologies and partners to allow developers to create their own compatible middleware solutions. The framework is a set of software tools, libraries, and applications as well as documentation and best practices to follow in developing P2V (Physical to Virtual) solutions. The framework manages devices from different domains allowing their interconnection. Each type of resource (switch, router, optical device) is managed by a Web Service, and each device is represented as a Resource (using WSRF). These Web Services can be distributed in several machines.

IaaS Framework offers common functionalities that can be used by several types of networks and devices. Manticore and Argia are services that use this framework. Manticore provides virtualisation mechanisms for IP networks while Argia provides virtualisation mechanisms for optical networks. Services can be activated though the GUI or with an application that calls the WS API. Figure 12 shows an overview on the IaaS architecture and the framework-based products and research projects.

The biggest challenge in adopting the IaaS Framework for the FEDERICA project [13] was to deal with the abstraction methodology used in order to virtualise physical resources. Each physical resource or a partition of that needs to be represented by a software instance (called Web Service resource).

The FEDERICA project defined new types of infrastructure resources for the framework adoption:

- Software router: representation of various types of open software routers (such as XORP or QUAGA).
- Virtualization node: representation of servers with installed virtualization hypervisor (such as Xen, VMware or OpenVZ) in order to create and manage virtual machines.
- Ethernet switch: representation of different types of the Ethernet switches.

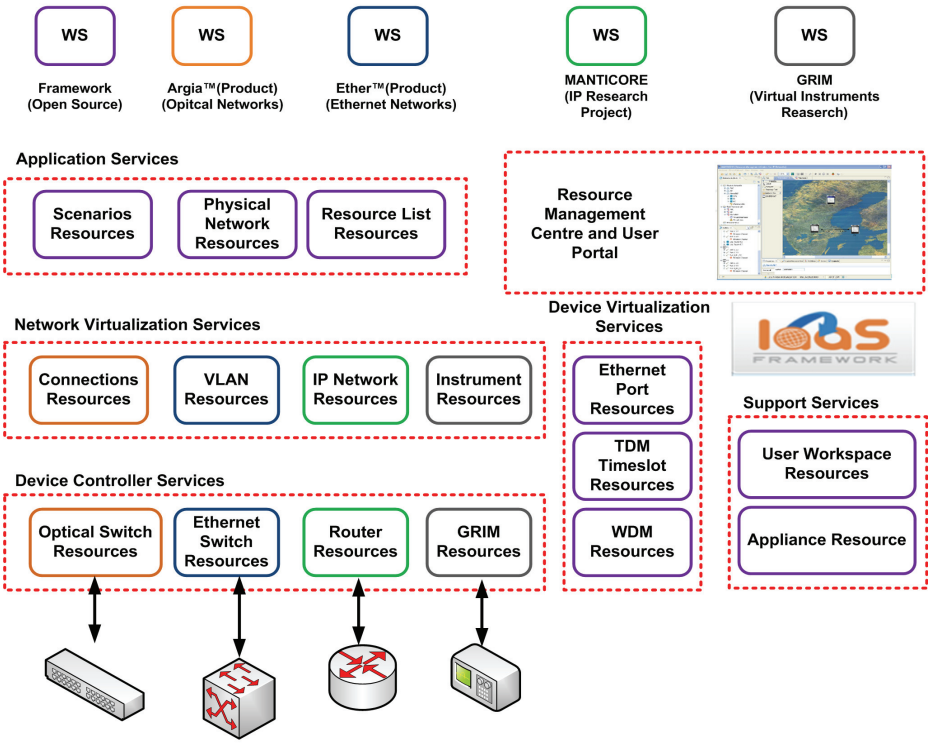

Fig. 11. IaaS Framework architecture

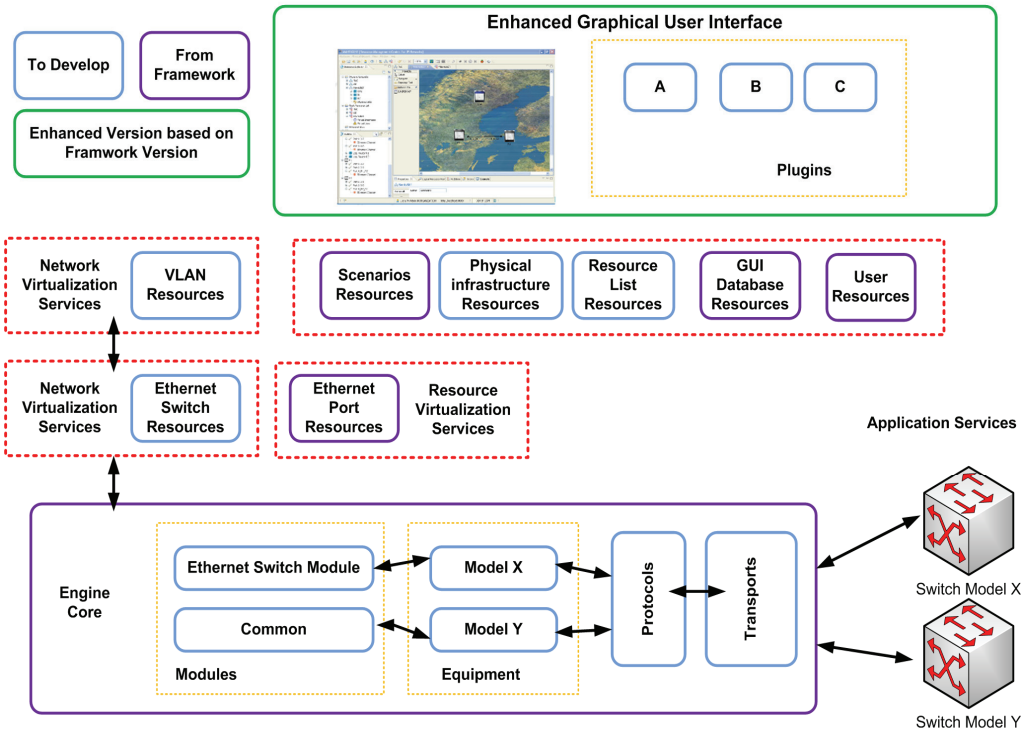

Fig. 12. IaaS Framework adoption for FEDERICA

- VLAN network: representation of a VLAN network.
- MPLS network: representation of an MPLS network.

In order to gain control on all these resources, the FEDERICA project had to design new control WS modules. With the developed Web Services it is possible to fully manage the IP, MPLS, VLAN and Ethernet networks. Figure 13 depicts (blue colour) the components that have been developed initially in the FEDERICA project. An Enhanced Graphical User Interface box is also shown (on the top right), which makes it possible to add some plug-ins in order to enhance the functionalities of the current Manticore GUI to control Ethernet devices.

At the current stage of the project the FEDERICA Slicing Tool, which is adopted in the IaaS Framework, is being tested in operation. The preliminary results of the tests look very promising. The slice configuration process supported by the FEDERICA Slicing Tool is much more efficient and actually speeds up the slice creation that may lead to more FEDERICA users.

V. FEDERICA USAGE

The Networking Activity 2 in FEDERICA is about building and consolidating the user community, led by TERENA. The activity is focused on identifying and establishing a strong relationship with users to the scope of gathering requirements, and facilitating the flow of information and ideas between users, in particular when they are in different research communities. The project is demand-driven from day one. User groups are not considered as external entities to be addressed with questionnaires or to be informed on the FEDERICA concept only, but NA2 partners work together with selected user groups, visit their locations, understand the scope of their research and explain and design to them how FEDERICA can be used.

The FEDERICA project will officially run until the end of October 2010. In this later phase of the project, the NA2 activity is trying to deepen the understanding of the behaviour of current users and the usage patterns, as well as consolidating and further enlarging the user community. NA2 partners have done technical and non-technical segmentations on both the current user basis and a potential future user basis. The current user basis consists of 11 projects partly finished, on-going, or just being requested. The expected user basis includes all the potential user projects or communities (more than 30) have been approached and interviewed by FEDERICA partners. The technical segmentation classifies the followings:

- Level of usage (Control),
- Access requirements (Connection),
- Size of the request (Type of research),
- Lifetime of the request (Length of research).

The non-technical segmentation studies the followings:

- Motivation of research (Community)
- Rationale of research (Business)
- Expected results (Contribution)

Some of the technical and non-technical user segmentation results are depicted in the following graphs.

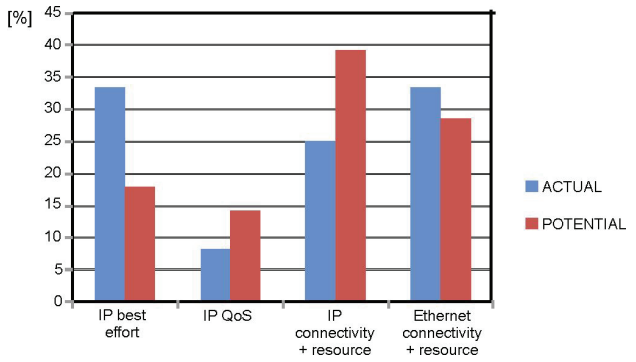

Fig. 13. Level of usage of the infrastructure

Fig. 13 shows that roughly one third of the current and foreseen users want to exploit the unique L2 features of FEDERICA requesting raw resources and pure Ethernet connectivity (last columns).

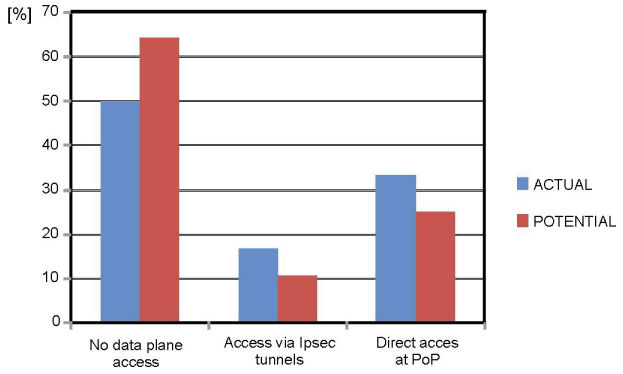

Fig. 14. Access requirements

Fig. 14 illustrates that half of the current users are connecting external resources (i.e., their own test beds) to their virtual slice via IPsec tunnels or direct connections at the FEDERICA PoPs. This feature is not available in PlanetLab or OneLab.

Among the non-technical aspects, Fig. 15 shows the distribution of ICT user communities interested in network research, application and service development, or both. It can be seen that the network researchers will dominate in the future.

Figure 17 explains that most of the users are willing to exploit the unique technical features (e.g., L2 access, reproducibility, real link delays, etc.) provided by FEDERICA,

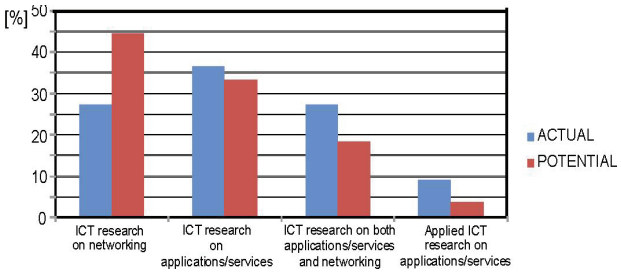

Fig. 15. Motivation of research

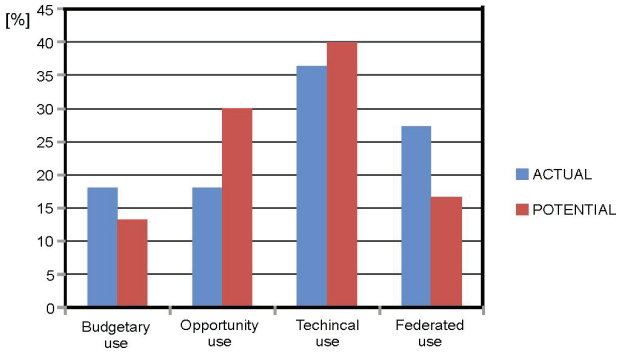

Fig. 16. Rationale of research

although a significant number of users want to gain the opportunity of fast provisioning and large number of virtual resources. Some users (especially Universities) simply go for FEDERICA slices because of budget issues (currently FEDERICA usage is free of charge) as well as some user groups want to federate with other similar infrastructures allowed by FEDERICA.

VI. SUMMARY

The FEDERICA project is constantly evolving according to user needs and expectations. Currently, the main emphasis is placed on dynamic service provisioning and providing users with the expected quality of service. The tools developed as research activities of FEDERICA are now being implemented into the operational environment, eliminating a requirement for manual slice configuration. The dynamic service provisioning is the key issue for improving the virtualisation service availability in FEDERICA, as users may request and use resources at any time.

Increasing awareness of virtualisation techniques and their benefits develop a growing interest of the FEDERICA project. Since more users are willing to use the infrastructure, it is planned to continue to maintain the infrastructure after the end of the project. It is also planned to continue the research work of FEDERICA as a project con-

tinuation, leading to develop a reliable, automated, and efficient virtualisation platform for future Internet experiments.

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 to this work.

References

[1] FEDERICA, *Federated E-infrastructure DEdicated to Researchers Innovating in Computing Architectures*. European Commission co-funded in the 7th Framework Work Programme, project n. RI-213107, http://www.fp7-federica.-eu/ and documents therein.
[2] Future Internet Research and Experimentation, http://cordis.europa.eu/fp7/ict/fire/
[3] Future Internet Assembly, http://www.future-internet.eu/
[4] Global Environment for Network Innovation, http://www.geni.net
[5] AKARI Architecture Design Project, http://akari-project.nict.go.jp/eng/index2.htm
[6] N.M.M.K. Chowdhury, R. Boutaba, *Network Virtualization: State of the Art and Research Challenges*. IEEE Communications Magazine, July 2009, 20-26 (2009).
[7] R. Krzywania at al., Deliverable DSA1.2: FEDERICA Infrastructure, 2009.
[8] GÉANT Network, http://www.geant.net
[9] http://www.juniper.net/customers/support/products/mx480.jsp
[10] http://www.juniper.net/us/en/products-services/switching/-ex-series/ex3200/
[11] P. Szegedi et al., Deliverable NA2.2: FEDERICA User Community and Requirements.
[12] C. Cervelló-Pastor et al., *Deliverable DJRA1.2: Solutions and protocols proposal for the network control, management and monitoring in a virtualized network context.*
[13] A. Berna, Deliverable SA2.2: IP slice service Prototype.
[14] M. Lemay: *An introduction to IaaS Framework*. 5 August 2008, http://iaasframework.com
[15] E. Grasa. G. Junyent, S. Figurerola, A. Lopez, M. Savoie, *UCLPv2: a network virtualization framework built on webservices*. IEEE Commun. Mag. 46 (3), 126-134 (2008).
[16] V. Reijs, E .Grasa, T*he MANTICORE project: Providing Users with Logical IP Network Service*. TERENA Networking Conference (2008).

RADOSŁAW KRZYWANIA received the M.Sc. degree in Computer Science – Software Engineering from the Poznan University of Technology in 2003. He is working in Poznań Supercomputing and Networking Center as a network applications engineer. He is responsible for implementation work in GEANT2 AutoBAHN Bandwidth on Demand system, and is an activity leader for network infrastructure in the FEDERICA project. He is also interested in resources virtualization, highly efficient network utilization and software development.

ŁUKASZ DOLATA is currently working for the Joint Research Activity 1 and is responsible for design and development of the FEDERICA User Portal. He received the M.Sc. degree in Electrical Science from Poznań University of Technology in 2002. He is mainly involved in the network management and provisioning applications designing and developing. He was involved in creating the PIONIER Monitoring and Management System. Since 2004, he was involved in GEANT2 JRA3 activity and took part in implementing Automated Bandwidth Allocation across the Heterogeneous Networks (AutoBAHN) system. He is also participating in the EXPReS project.

PETER SZEGEDI is currently the Networking Activity 2 and Joint Research Activity 2 leader of the FEDERICA project. He received his M.Sc. 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.