1st FIRE Portfolio Update – Concept of Experimentally-Driven Research and its Facilities

Deliverable/Milestone: D2.1

Date: 17th October 2011

Version: 1.0

Editor:	Martin Potts, Martel
Deliverable nature:	Report (R)
Dissemination level: (Confidentiality)	Public (PU)
Contractual Delivery Date:	31/08/2011
Actual Delivery Date	17/10/2011
Suggested Readers:	Network researchers, system architects, FIA and FI-PPP projects' stakeholders, EC
Total number of pages:	107
Keywords:	future networks, experimental facilities, research, experimentation

Authors:	Martin Potts, Martel Monique Calisti, Martel Sandra Wittwer, Martel Jerker Wilander, Dimes
Contributors:	Jacques Magen, Dimes Piet Demeester, Ingrid Moerman, Stefan Bouckaert, IBBT Stefan Fischer, University Lubeck Michael Boniface, IT Innovation Anastasius Gavras, Eurescom Michael Nilsson, CDT

Abstract

This 1st FIRE Portfolio Update is intended for anyone who is interested in the FIRE initiative, the current topics being researched and the status of the testing facilities that are available for use by FIRE STREPs and beyond. It describes the current FIRE landscape of projects, the experiments that have been - and are being - performed on the facilities and also common topics being studied in the various Working Groups comprising representatives of the IPs within the Architecture Board. It supersedes the original "FIRE Portfolio Analysis" document that was produced by the FIREworks project in 2010.

At the time that the original "FIRE Portfolio Analysis" document was written, very little information was available from the Call 5 STREPs and the Working Groups within the

Architecture Board; this version addresses these aspects. A summary of the usage of the FIRE facilities has also been added.

The situation is continually evolving and a further update will be issued in 12 months to provide information about the 3 new IPs from Call 7, the new STREPs from Call 8 and further experiments performed on the facilities.

Table of Contents

1. Introduction

FIRE (Future Internet Research and Experimentation) combines research into new paradigms with comprehensive test facilities upon which the ideas are experimented. Together, this creates a key resource for driving European research into future networks. This environment enables both incremental and disruptive approaches, supports multi-disciplinary research that goes beyond network layers, scholastic dogmas and public-private discussions. It provides a core infrastructure, and also a playground for future discoveries and innovations, combining research with experimentation.

This 1st FIRE Portfolio Update describes the current FIRE landscape of projects, the experiments being performed on the facilities and the common topics being studied in the various Working Groups comprised of representatives of the IPs within the FIRE Architecture Board. It supersedes the original document that was produced by the FIREworks project in 2010. Further information about the FIRE facilities can be found in the FIRE Roadmap document, which can also be found on the FIRE website.

The objectives of this FIRE STATION Portfolio Update may be summarized as follows:

- To describe the current facilities, in terms of:
  - Their offerings (testbed capabilities and support services)
  - The Use Cases (ongoing and planned), categorized as:
    - Use Cases from internal partners
    - Use Cases from FIRE STREPs
    - Use Cases from other FP7 projects
    - Use Cases from outside FP7
    - Use Cases from the Open Calls
  - Joint activities within the framework of the Architecture Board Working Groups
  - Joint developments to extend the facilities, such as a resource brokering portal for federating testbeds
- To describe the research and experimentation projects in terms of:
  - Which research areas they are addressing.
  - How - and to what extent - they made - or will make - use of the FIRE facilities.
  - The core issues that have been identified so far, in particular with respect to the availability and usability of the FIRE facilities.
  - If, and how, end users have been - or are planned to be - involved.

The primary purpose of this report is to present and promote the usage of the facilities and, as a side effect, to determine if there are missing domains that need to be included, in order that all relevant areas of ICT are dealt with to build up a full-fledged FIRE federated facility. Equally importantly, it shows what has been already accomplished with past (Call 2) and running (Call 5) FIRE projects, and in the scope of the new (Call 7) projects. It also describes three critical issues that have been identified by the IPs as being of common interest: Resource Brokering Tools for a federated environment, Measurement and Benchmarking and Sustainability.

Methodology used for the data collection

- Information on the - now completed - Call 2 projects was taken from the previous version of this report, updated with additional information from their Websites, presentations made in FIRE events, answers provided by the projects' consortia to a set of specific questions that have been elaborated for each specific project by the FIRE STATION consortium and the personal knowledge of the editors and contributors.
- Information on the running Call 5 projects is based on information found on their Websites and in their project summaries, telephone interviews, presentations made in FIRE events (especially, the Ghent conference in December, 2010, the Budapest Workshop in May, 2011 and Architecture Board meetings), answers provided by the projects' consortia to a set of specific questions that have been elaborated for each specific project by the FIRE STATION consortium and the personal knowledge of the editors and contributors.
- Since the new Call 7 IP projects (Confine, Experimedia and OpenLab) have not yet started, only very preliminary information about them is available in this version.

Summary of Conclusions

The original version of the Portfolio Analysis produced by the CSA project FIREworks highlighted a rather low usage of the FIRE experimental facilities from the Call 2 STREPs, which to some extent is also re-confirmed by this analysis and for which several reasons have been identified:

- For some of the Call 2 STREPs, the available FIRE facilities were not always able to meet the technical requirements of the planned experiments.
- Easier access to local (not publicly offered facilities) testbeds was available and more specifically tailored to the foreseen experiments.
- Access to FIRE facilities was not always available at the time when the Call 2 STREPs were entering their experimental work phase.

However, this situation seems to have changed for the Call 5 projects, which appear to have planned a more systematic and active uptake and deployment of existing FIRE facilities. This is considered to stem from several main factors:

- The level of maturity of the FIRE facilities has improved as well as their visibility within the overall European R&D landscape.
- Some of the Call 5 STREPs are somehow follow-ups of previous Call 2 STREPs, which means there is now increased experience and know-how of the overall FIRE context, including available experimental means.
- The market-pull for advanced and improved Future Internet technologies has increased over the last few years, while most of the R&D investments have been shrinking. The

European Commission's 2010 "EU Industrial R&D Investments Scoreboard"[4] shows that *"R&D investment by top EU companies fell by 2.9% in 2009... The fall in R&D investment by leading players in the US, at 5.1%, was twice as sharp as in the EU, but the worldwide reduction was at 1.9%."* Therefore, the availability of openly accessible and large-scale experimental testbeds becomes increasingly crucial in today's critical economic situation for both academic and corporate (industrial) R&D labs.

The Open Call mechanism is further encouraging the use of the facilities and making the facilities better known.

The STREPs have highlighted that the FIRE facilities make it possible to:

- Have a greater diversity of technologies and infrastructures.
- Make experiments on a larger scale.
- Obtain more advanced experimentation results, based on multiple metrics.
- Better justify the research results, since the experiments are performed in "close-to-real-life" conditions.
- Gain technical know-how about an increased number of technologies and equipments.
- Reuse the same technologies and resources across several projects, thereby allowing return on investment to be maximized.

They also recognise that testbeds are in principle living beyond a single project lifetime as they remain to be available to the broader audience for follow-up work.

Despite their very different technical foci, the IPs have discovered topics of common interest (resource brokering tools, measurement and benchmarking and sustainability), which are proving beneficial to study together. Furthermore the potential for making joint development work on a common portal for automatically federating testbeds is being investigated.

2. The current FIRE landscape

The FIRE Initiative was launched at the beginning of 2007 (FP7). It was built upon a few projects that were previously assigned to the "Situated and Autonomic Communications" or "Future and Emerging Technologies" Programmes, as well as several Research Networking Testbeds from FP6.

The FIRE portfolio of projects is now becoming increasingly visible, with the testbed facilities being available for the experimentation of a large number of technologies and at European scale. The research that is now possible in FIRE includes many different areas involving, for example, end user communities, large scale distributed software applications and advanced radio technologies. The upcoming call (Call 8) for research projects (with the requirement to use the existing FIRE facilities) and the future Open Calls, will further increase the volume of experimentation in key areas of the future Internet.

However, the expected stronger focus on the federation of facilities has not yet taken place. This is mainly because the research experiments are not sufficiently innovative to require more than an individual testbed, or testbeds within a single facility. Nevertheless, federation has attracted interest and is included as an important aspect within several FIRE facility projects; particularly those that have performed a degree of federation between heterogeneous testbeds within one technology area (as in BonFIRE, WISEBED and CREW). Furthermore, some experiments involving federation across facilities (OneLab and FEDERICA) are found in the project NOVI. Nevertheless, there is still a need to work on creating more advanced experiments involving more than one of the FIRE experimental facilities. We can see that the motivations for the FIRE facility projects to work together are not necessarily always their first priority which is to meet their own contractual commitments. The upcoming call for a single federated facility (in Call 8) will address this.

The experimental work and the use of FIRE facilities by research projects in FIRE has evolved during the previous year, and the availability of some of the facilities have been sustained, even without funding from FIRE. However, there is still need to extend the usage of the experimental facilities in order to obtain value for money for the investment. The Open Calls have shown that there is a geeat need for experimentation of moderate size, not only in academic environments but also in industry. Further support of such experimentation is important.

FIRE projects from Call 2

The FP7 ICT Call 2 represented the first wave of FIRE-specific projects. These started in January 2007 and the last ones ran into 2011.

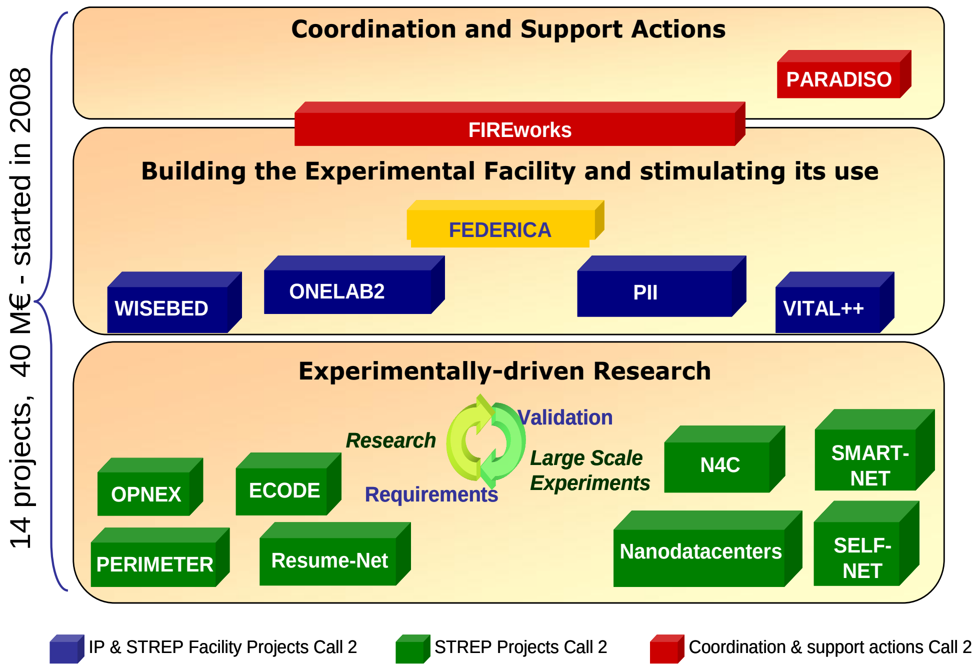

Figure 1 – FIRE projects from Call 2

Four of the projects (WISEBED, ONELAB2, PII, VITAL++) were categorized as "facility projects" (although VITAL++ was a STREP) building experimental platforms for Future Internet researchers, whilst eight projects in the FIRE portfolio (OPNEX, ECODE, N4C, SMART-Net, PERIMETER, Resume-Net, Nanodatacenters, SELF-NET) were research-focused and experimentally-driven. One project (FIREworks) was funded to co-ordinate and support the FIRE Programme and a further Coordination and Support Action project (PARADISO) examined the socio-economic aspects of the Future Internet. The FEDERICA project provided similar testbed services as the IPs funded through FIRE, but was based on the NREN/GÉANT infrastructure and funded by the Research Infrastructures part of the FP7 Programme. Nevertheless, for all practical purposes, it was considered as a FIRE facility.

FIREworks made the FIRE Initiative a globally known, well-recognized brand. The project opened discussion with the GENI Office and NSF, as well as with several GENI projects. These meetings resulted in project-level co-operation and ideas to share efforts in solving common challenges in high-level federation, together with usage stimulation of the facilities. Strong links were also established between FIRE and similar activities in Japan.

- GENI (US): http://www.geni.net
- AKARI (Japan): http://akari-project.nict.go.jp/eng/index2.htm

FIRE projects from Call 5

Figure 2 shows the FIRE projects in Call 5. These projects started in Summer or Autumn 2010. Most have a duration of 2.5 or 3 years. However, the IP project CREW has a duration of 5 years.

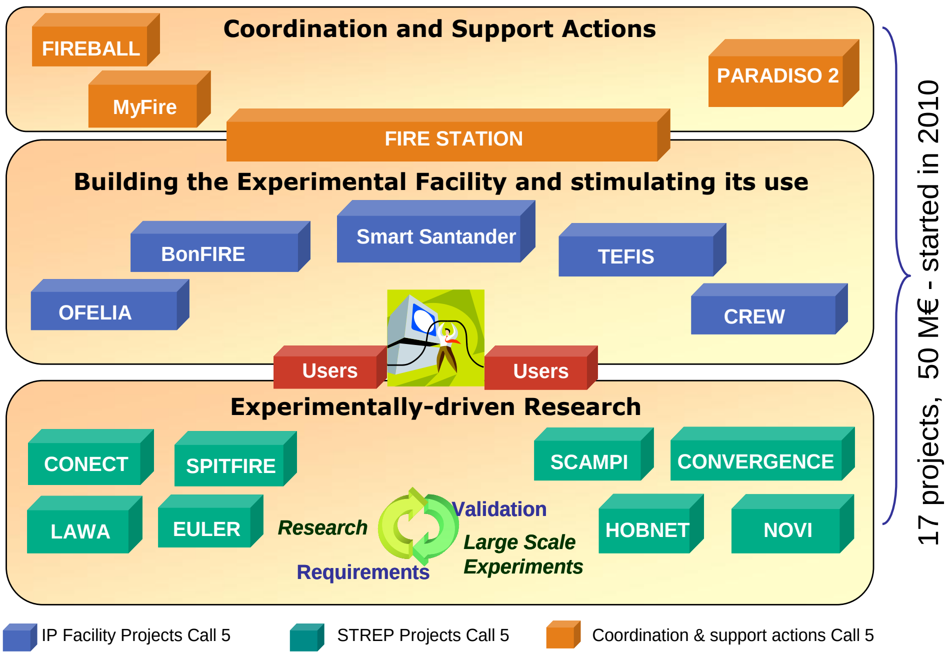

Figure 2 – FIRE projects from Call 5

The five IP projects (OFELIA, BonFIRE, SmartSantander, TEFIS, CREW) are "facility projects", whilst the eight projects (CONECT, SPITFIRE, SCAMPI, CONVERGENCE, LAWA, EULER, HOBNET, NOVI) are research-focused and experimentally-driven. One project (FIRE STATION) was funded to co-ordinate and support the FIRE Programme and three further Coordination and Support Action projects were funded to (i) examine the socio-economic aspects of the Future Internet (PARADISO), (ii) liaise with the LivingLab community

(FIREBALL) and (iii) liaise with the Future Internet activities in Brazil, Russia, India and China, and keep the community aware of important standardization issues.

Whilst OFELIA, BonFIRE, TEFIS and CREW provide facilities in new technological areas, SmartSantander can be considered as a continuation of WISEBED (from Call 2), but on a larger scale and in a real city environment.

International collaboration has been extended to the BRIC countries through the CSA project MyFIRE: http://www.my-fire.eu

PanLab-PII, OneLab2 and FEDERICA are maintaining their facilities beyond the official contractual end date without EC funding, at least until the three new Call 7 IPs start (see below).

FIRE projects from Call 7

Three new IP projects will start in Autumn 2011: Confine, Experimedia and OpenLab. In addition, CREW (additional testbed) and BonFIRE (new Use Case) will extend their facilities.

A specific call for collaboration between Europe and Brazil resulted in one new FIRE project FIBRE-EU:
http://cordis.europa.eu/fetch?CALLER=PROJ_ICT&ACTION=D&CAT=PROJ&RCN=99654

The main goal of the FIBRE project is the design, implementation and validation of a shared Future Internet research facility between Brazil and Europe.

A new experimental facility will be built in Brazil, comprising various technologies (fixed layer 2 and layer 3, wireless, optical) as well as the design and implementation of a control framework to automate the use and operation of the testbed. On the European side, two existing FIRE infrastructures (OFELIA and OneLab) will be enhanced and federated.

Experiments will include pilot public utility applications.

More details will be given about these items in the next version of this report.

3. The testbed offerings

The testbed offerings described below show the scope of the technologies addressed by the FIRE facilities that are currently available. Note that the bfirst 3 of these (PanLab-PII, OneLab2 and FEDERICA) are maintaining their facilities beyond the official contractual end date without EC funding.

All facilities in FIRE may be selected for experimentation by Call 8 research projects, though some may have conditions regarding availability, usage costs, etc. Indeed, every (STREP) proposal in Call 8 must include an agreement with a facility provider to be the platform for their experiments. Projects outside the FIRE research portfolio may also use the facilities - either by responding to Open Calls, or by contacting the facility directly. Requests for a federated use of more than one FIRE facility could be directed either to the involved facilities or to the FIRE Office/FIRE Architecture Board for an analysis of the feasibility of the proposed experiments.

1. Panlab – PII www.panlab.net

The Panlab-PII project manages the interconnection of different distributed testbeds to provide services to customers for various kind of testing. Such testing activities need support from a co-ordination centre instantiated here as the so-called Panlab Office. The main roles in the Panlab concept were:

- Panlab Partner as the provider of infrastructural elements necessary to support the testing services. Partners are connected to the Panlab Office for offering functionality to the customers.
- Panlab Customer who utilizes a service provided by the Panlab office, (e.g. to carry out R&D activities, implement and evaluate new technologies, products, or services). A Panlab Customer takes benefit from the Panlab testing offerings.
- Panlab Office which realizes a brokering service for the test facilities by co-ordinating i) the provision of the testing infrastructures and services, ii) the Panlab Partner test-sites iii) the communication path between them and customers.

The Panlab overall architecture relied on a number of additional architectural components including the Panlab search and composition engine called "Teagle" and a Panlab repository that stores testbed descriptions and testing results. In its most basic working mode, Panlab offered to go manually through the operational steps for the creation and realization of testing projects. Operations were then executed by personnel of the Panlab Office involving partners and customers. Thus, the testbed metadata held in the Panlab repository was entered manually as well as testing configurations, etc. However, in a more elaborated working mode, intending to automate the Panlab related processes, the so-called Teagle tool offered, (among other functionalities) an online form where the testbed representatives could enter the relevant data describing the testbed and its resources, and Panlab customers could then search the Panlab repository to find suitable resources needed for performing their tests.

2. PlanetLab Europe – OneLab – OneLab2 www.onelab.eu

The majority of OneLab's testbed nodes are open to the public Internet, meaning that researchers can experiment with distributed applications in a real-life testing environment. This ability makes the PlanetLab (http://www.planet-lab.org/) environment, for example, an essential complement to controlled experiments in simulation or emulation environments. Unpredictable real-world traffic loads, routing changes, and failures put applications to the test in ways that a controlled environment cannot. Furthermore, researchers can deploy services that are used by regular Internet end users worldwide. For example, one content distribution experiment on PlanetLab offers faster web downloading to thousands of end users in countries across the world. The researchers who deployed this service use it to study application performance and end user behaviour. A key benefit of OneLab's offer to researchers is that it allows them to deploy their experiments at a global scale, exposing their applications to geographic and network topological diversity and allowing them to deploy services in proximity to end users. While the testbeds' own nodes are scattered essentially across Europe, federation with the global PlanetLab system provides European researchers with access to the combined system, which consists of over 1,000 nodes at over 500 sites worldwide.

The OneLab project provided an open federated laboratory, built on PlanetLab Europe (PLE), which supports network research for the future Internet. It provided an open, general-purpose, shared experimental facility, both large-scale and sustainable, which allowed European industry and academia to innovate and assess the performance of their solutions. OneLab also aimed to develop strong international partnerships, in order to explore and experiment with the concept of federation.

The second phase of the OneLab project, OneLab2, started on 1 September, 2008, and ran for 27 months. It built on the original OneLab project's foundations, continuing work on the PLE testbed, increasing its international visibility and extending it in both functionality and scale. PLE extended the PlanetLab service across Europe, federating with other PlanetLab infrastructures worldwide (PLC - Central, PLJ - Japan). In particular, the OneLab2 project enhanced the network monitoring service that supports experiments, and co-operated with potential customers by directly involving pilot projects to test novel ideas under synthetic or real-world situations. OneLab2 built PLE gateways to unusual, cutting-edge networking environments, and pushed forward a federation model, so that PLE could serve as a basis for a future highly heterogeneous communications environment. It now includes the NITOS (wireless) testbed, Etomic (high precision measurements) and Dimes (large scale measurements). By enabling a user to request resources using OMF controllers in a PlanetLab slice, OneLab2 allows a user to run an experiment on PlanetLab. The PlanetLab wireless extensions allow experiments to span PlanetLab and other OMF-controlled wireless testbeds.

3. FEDERICA www.fp7-federica.eu

The third main FIRE facility element in Call 2, even though funded from a different EC Unit (Research Infrastructures), was FEDERICA. FEDERICA is a Europe-wide infrastructure based on computers and network physical resources, both capable of virtualization. The facility can create sets of virtual resources according to users' specification for topology and systems. The user has full control of the resources in the assigned “slice” which can be used (for example) for Future Internet clean-slate architectures, security and distributed protocols, routing protocols and applications. The FEDERICA project supports research experiments on new Internet architectures and protocols by developing a versatile technology-agnostic network infrastructure that runs over the existing production networks from GÉANT and national academic networks. The aim was to allow innovative and/or disruptive technologies to be trialled. The project finished in October 2010, but the infrastructure will be supported until at least mid-2012 through the GN3 project (GÉANT). It currently comprises Gigabit Ethernet circuits, Layer 2 and Layer 3 switching, and servers supporting virtualization. Circuits interconnect Points of Presence of NRENs, hosting virtualized FEDERICA nodes (V-nodes) separate from the GÉANT production equipment.

Virtual slices of FEDERICA's infrastructure are allocated to network researchers to conduct potentially disruptive experiments within a large production substrate. Users can request a virtual infrastructure slice composed of a combination of virtual circuits (up to 1Gbit/s) and V-nodes. These slices can include: Layer2 circuits or Layer3 IP configured circuits (IPv4, IPv6 unicast and multicast) and/or virtual system(s) and/or virtual routers.

4. WISEBED (Call 2) www.wisebed.eu - SmartSantander (Call 5) www.smartsantander.eu

Real-World Internet (RWI) and Internet-of-Things (IoT) technologies are considered to provide important cornerstones and/or extensions for the Future Internet. The two projects WISEBED (Call 2) and SmartSantander (Call 5) aim at providing a large-scale federated experimental facility for these technologies and for applications. While WISEBED concentrated on pure sensor network experimentation at large scale (1000 nodes distributed over 9 European sites), SmartSantander is building an experimental facility, which will be based on a real-life IoT deployment in an urban setting. The core of the facility will be located in the city of Santander and its surroundings. Santander is the capital of the region of Cantabria, situated on the north coast of Spain with a population in 2007 of about 184,000 inhabitants. The facility will encompass IoT deployments, up to 12,000 devices, in different key areas of the city infrastructure, ranging from public transport, key logistics facilities such as harbour and waste management, public places and buildings, work places and residential areas, thus creating the basis for the development of a future Smart City.

There are several reasons for setting the experimental facility into a city context: the first one is the extent to which the necessary infrastructure of a Smart City will rely on technologies of the IoT. The resulting scale and heterogeneity of the environment makes it an ideal environment for enabling a broad range of experimentation needs. Furthermore, a city can serve as an excellent catalyst for IoT research, as it forms a very dense techno-social eco-system. It is expected that the city itself will be a major source of functional and non-functional requirements from a variety of problem and application domains (such as vertical solutions for the environment control and safety, horizontal application to test network layers, content delivery networks). Cities provide the necessary critical mass of experimental businesses and end users that are required for the testing of IoT as well as other Future Internet technologies for user acceptability testing and market adoption.

SmartSantander will specify, design, and implement the necessary building blocks. An initial high-level architecture for the resulting new experimental facility has already been worked out, and is shown in Figure 3. This architecture heavily relies on existing components which will be supplemented by the building blocks that are currently missing.

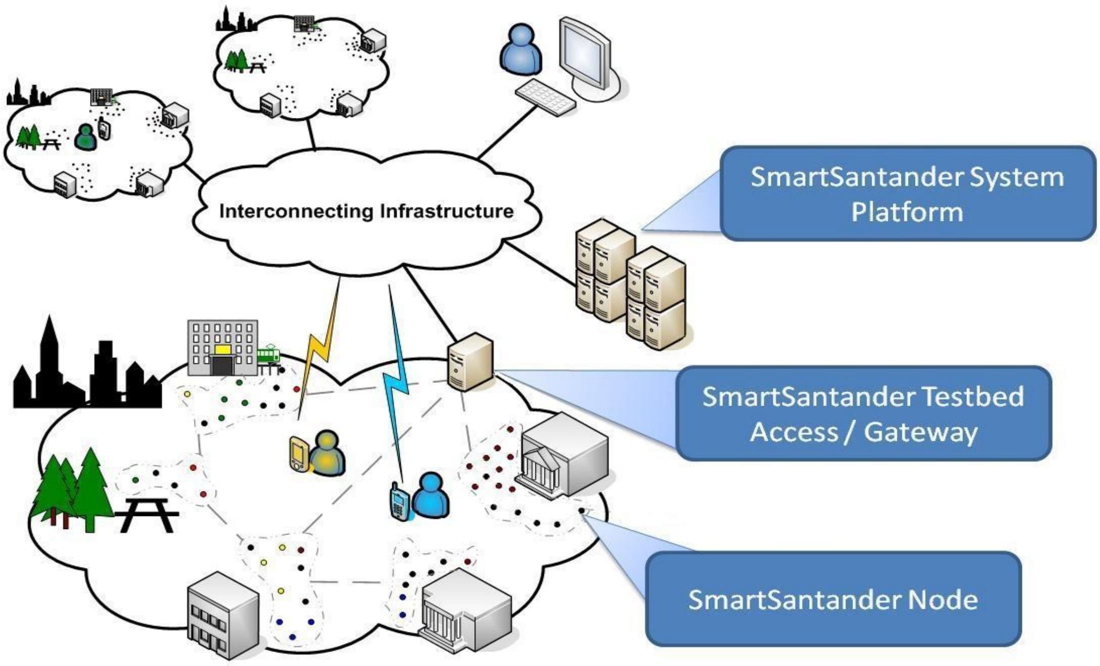

Figure 3 – SmartSantander architecture

SmartSantander will encourage the use of its experimentation facility among the scientific community, end users and service providers in order to reduce the technical and societal barriers that prevent the IoT concept from becoming an everyday reality. To attract the widest interest and demonstrate the usefulness of the SmartSantander platform, a key aspect is the inclusion of a wide set of applications. Application areas will be selected based on their high potential impact on the citizens as well as to exhibit the diversity, dynamics and scale that are essential in advanced protocol solutions, and will be able to be evaluated through the platform. Thus, the platform will be attractive for all involved stakeholders: Industries, communities of users, other entities that are willing to use the experimental facility for deploying, and assessing new services and applications, and Internet researchers to validate their cutting-edge technologies (protocols, algorithms, radio interfaces, etc.).

A first operational and usable prototype will be available within the first year of the project, which is almost over by now.

5. CREW (Call 5) www.crew-project.eu

The main goal of CREW is to establish an open federated testbed, which facilitates experimentally-driven research on advanced spectrum sensing, cognitive radio and cognitive networking strategies in view of horizontal and vertical spectrum sharing in licensed and unlicensed bands.

The CREW platform incorporates 5 testbeds with diverse wireless technologies (heterogeneous ISM, heterogeneous licensed, cellular, wireless indoor sensors, wireless heterogeneous outdoor sensors) augmented with State-of-the-Art cognitive sensing platforms. These will be extended and federated.

CREW Testbeds

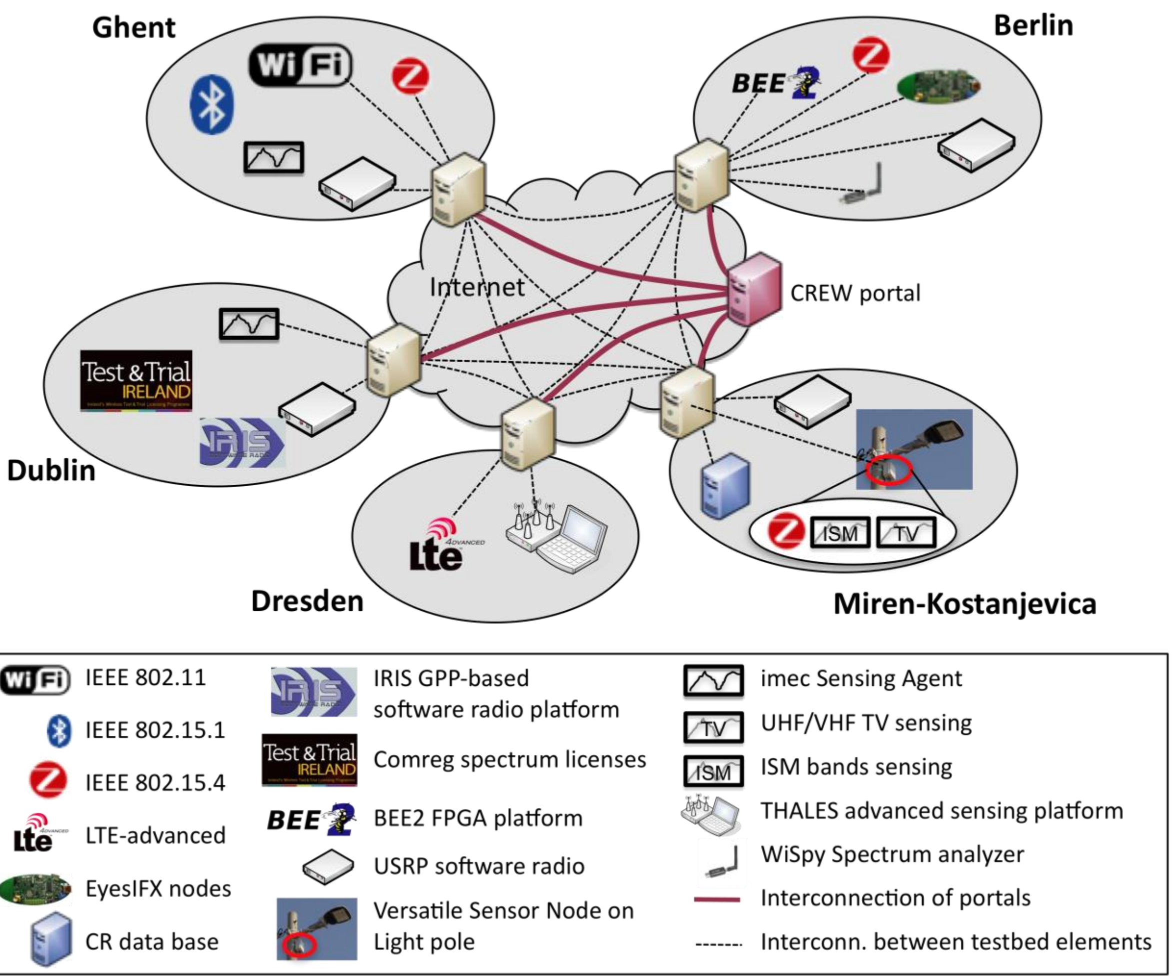

Figure 4 – The CREW testbeds

CREW comprises the following testbeds:

- w-iLab.t run by IBBT

  The w-iLab.t, which is part of the IBBT iLab.t test facilities, is a heterogeneous generic wireless testbed. The w-iLab.t allows flexible testing of the functionality and performance of wireless networking protocols and systems in a time-effective way, by providing hardware and the means to install and configure firmware and software on (a selection of) nodes, schedule automated experiments, and collect, visualize and process results.

  At a first location, the "w-iLab.t Office" consists of a wireless WiFi (IEEE 802.11a/b/g) and sensor network (IEEE 802.15.4) testbed infrastructure, deployed across three 90 m x 18 m floors of the IBBT office building in Ghent, Belgium. At 200 places throughout the office spaces, meeting rooms and corridors, embedded platforms with two WiFi interfaces and one sensor node are mounted to the ceiling. In Zwijnaarde, Belgium, located approximately 5 km away from the "w-iLab.t Office", a second location is equipped with another 60 wireless nodes, each equipped with two IEEE 802.11a/b/g/n interfaces, an IEEE802.15.4 and an IEEE802.15.1 (Bluetooth) interface. This location also hosts cognitive radio platforms such as software defined radio platforms and spectrum sensing engines.

- IRIS testbed run by Trinity College Dublin

  IRIS (Implementing Radio in Software) consists of a general-purpose processor software radio engine and a minimal hardware front-end. IRIS can be used to create software radios that are reconfigurable in real-time.

  The API used in IRIS is used to integrate an IRIS reconfigurable radio application into other applications. Specific functions that manipulate and control the IRIS Framework are abstracted from the core of the IRIS system enabling the sophisticated IRIS features to be accessed using a simple programming interface.

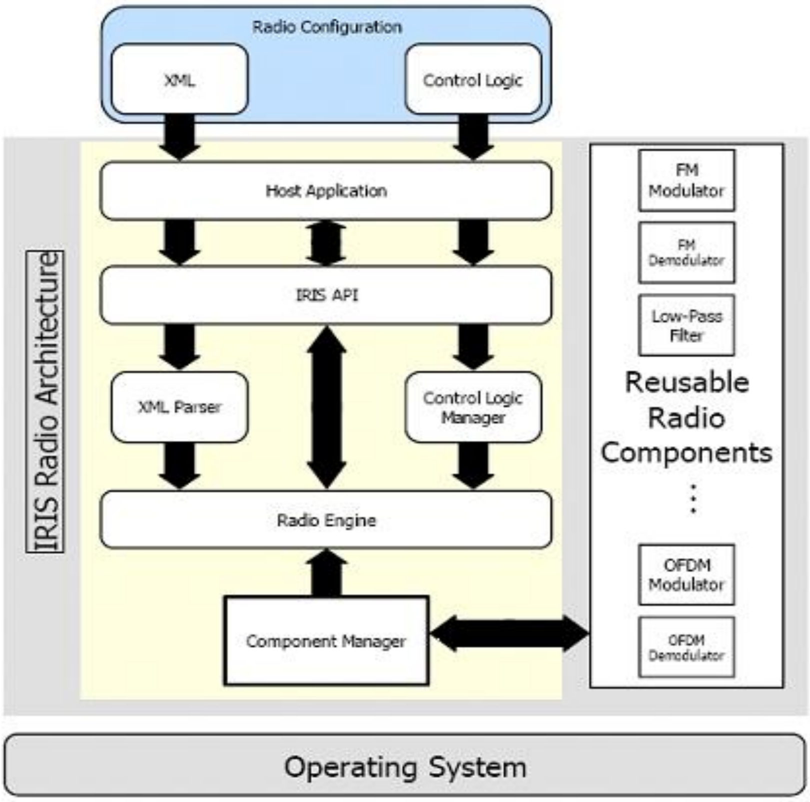

Figure 5 – Overview of the IRIS system

The IRIS testbed facilitates Parametric Reconfiguration (dynamic alteration of individual parameters of signal processing functionality), Structural Reconfiguration (the alteration of the layout of the radio system or the replacement of some aspect of the software of the system while still performing the same overall application) and Application Reconfiguration (completely replacing the software of the software radio with an entirely different software radio configuration).

- LTE / LTE advanced testbed run by the Technical University of Dresden and Vodafone

Two testbeds have been built and operated within the research project EASY-C. One of them is the testbed in downtown Dresden Germany, using existing 2G/3G network sites of operators Vodafone and T-Mobile. Both operators are also involved in the trials. The second testbed, focused on applications enabled through LTE and advanced concepts, is being set up in Berlin. The chosen testbed location in downtown Dresden covers various propagation conditions, which are of special interest for the evaluation of fourth-generation (4G) systems with MIMO links and interference conditions typical in frequency reuse on networks like LTE, and for the development of advanced algorithms such as cooperative MIMO:

The testbed is being built in three phases:

- In the first phase, one site with three cells started operating in April 2008. As shown in the picture below, this central site is located near Dresden's main railway station.

Figure 6 – EASY-C cell structure in downtown Dresden

- The second phase will consist of six sites with a total of 18 cells.
- In the final stage the testbed will comprise ten sites with a total of 25 cells. Additional interferers will surround outer cells in order to emulate the interference intensity and distribution of a network with three tiers of sites.

- TWIST testbed run by the Technical University of Berlin

The TKN Wireless Indoor Sensor Network Testbed (TWIST) is a multi-platform, scalable and flexible testbed architecture for experimenting with wireless sensor network applications in an indoor setting.

The nodes are deployed in a 3D grid spanning 3 floors of an office building at the TUB campus, resulting in more than 1500 m2 of instrumented office space. In small rooms (~14 m²), two nodes of each platform are deployed, while the larger ones (~28 m²) have four nodes. This setup results in a fairly regular grid deployment pattern with intra node distance of 3m.

TWIST relies on COTS hardware and fully leverages the features of the USB 2.0 standard. The SUT nodes are connected via USB hubs, which act as concentrators and provide a power supply management capability. This enables active SUT topology control and node fault injection modelling through selective powering on and off of SUT nodes. The software architecture is designed for easy remote access.

TWIST provides automatic trace collection and centralized time stamping service, as well as raw access to the serial I/O of the SUT nodes. A sample hardware instantiation of the TWIST architecture is depicted below.

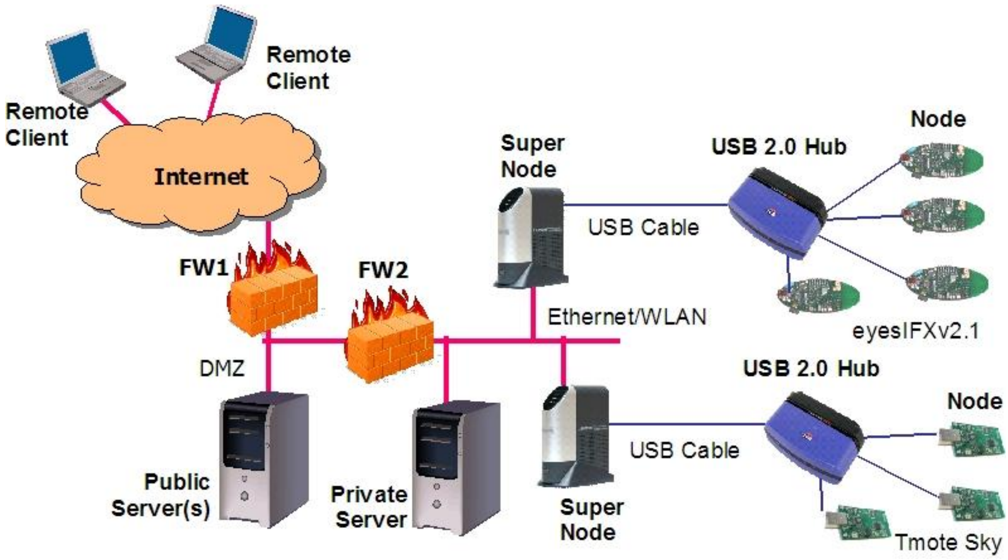

Figure 7 – Instantiation of the TWIST architecture

In addition to the described sensor network SUT components the TWIST infrastructure is complemented by several WiSpy sensing devices: these are low-cost spectrum scanners that monitor activity in the 868 MHz, 2.4 and 5 GHz spectrum, and output the measured RF energy and the quality of the received signals.

- Outdoor heterogeneous ISM/TVWS VSN testbed operated by JSI Slovenia

JSI, together with an industrial partner (Envigence Ltd.) and a public partner (Municipality of Miren-Kostanjevica) deployed an outdoor VSN (Versatile Sensor Node) based testbed for environmental sensing and light control. The first phase of deployment consisted of 5 VSN nodes, 20 VSN nodes are added in the second phase.

The VSN is a WSN platform with high processing capability, long-term autonomy and flexible radio. It supports a broad portfolio of sensors and actuators, while its modular approach allows adaptation to diverse application requirements. In this respect the platform consists of the core module – VSC and a set of special feature modules (radio module – VSR, expansion modules – VSE, power module – VSP) that are used as/if needed. The core module can be powered by batteries, solar panel or external power supply and together with radio module supports wireless sensor networks technologies such as ZigBee, 6LoWPAN and Wireless M-Bus. For the purposes of CREW project dedicated VSE modules were developed, offering low-cost RSSI-based spectrum sensing in the ISM and TV frequency bands. The modules have integrated omnidirectional antennas, but will also allow the connection of additional antennas.

This deployment offers the opportunity to test the VSN hardware in outdoor operating conditions under temperatures around and below zero and at high precipitation levels (both snow and rain). As part of the CREW project, the ISM spectrum sensing VSN nodes, so far tested in laboratory environment at JSI and in JSI campus, will be complemented with TV spectrum sensing capability and moved outdoors in the existing Miren testbed.

What can be tested and how?

The CREW facilities and federation allows the experimental validation of cognitive radio and cognitive networking concepts for the following usage scenarios:

- Radio environment sensing for cognitive radio spectrum sharing: focus of this usage scenario is to match the federated sensing hardware, the XCVR API and novel sensing functionality.
- Horizontal resource sharing between heterogeneous networks in ISM band: this usage scenario will investigate techniques for advanced resource sharing in typical home/office/public building environments, densely populated with various wireless ISM band devices.
- Cooperation in heterogeneous networks in licensed bands: this usage scenario will focus on cooperation in heterogeneous networks in licensed bands.
- Robust Cognitive Sensor Networks: the main objective of this usage scenario is the investigation of the robustness of cognitive radio solutions in Cognitive Sensor Networks (CSNs) in order to achieve a certain QoS, while ensuring non-interference.
- Impact of cognitive networking in primary cellular Systems: in this usage scenario integrate sensing agents use in an LTE cellular environment, which operates as primary system.

To perform these experiments, the CREW platform offers 3 modes of operation:

- Mode 1: Individual CREW testbed usage: this mode offers the common CREW portal (www.crew-project.eu/portal) with clear and uniform information about access and usage of each testbed, usage scenarios, feedback on user experience and results.
- Mode 2: Single CREW heterogeneous testbed usage: this mode allows physically hosting of nodes from one testbed in another, hence creating new nodes from combinations of CREW hardware/software components.
- Mode 3: Multiple sequential CREW testbed usage: this mode enables the capturing of data/behaviours from one CREW testbed and replay for emulation or post-processing purposes on a second CREW testbed.

Usage Information

The CREW project will stimulate research through two Open Calls for experiments. The first of these (for €400K) was issued in September 2011 and the next one will be one year later. Access to the facility will also be open to researchers not funded through Open Calls, depending on the availability and the type of the experiment.

All technical details on the testbeds and components available in the CREW infrastructure, and detailed information on how to access the testbeds (accounts, policies) are available via the CREW portal at www.crew-project.eu/portal.

6. BonFIRE (Call 5) www.bonfire-project.eu

The BonFIRE (Building service testbeds for Future Internet Research and Experimentation) project is designing, building and operating a multi-site Cloud testbed with heterogeneous resources, including computing, storage and networking resources, for large-scale testing of applications, services and systems at all stages of the R&D lifecycle, targeting the Internet of Services community within the Future Internet.

The BonFIRE vision is to give researchers in these areas access to a facility that supports large scale multi-disciplinary experimentation of their systems and applications addressing all aspects of research across all layers. The project will develop and support a framework, which allows service-based computing practitioners to experiment with their latest ideas in service orientation and distributed computing. The overall goal is to encourage new communities of experimenters to take advantage of the opportunities offered by the FIRE infrastructure to guide the development of the Future Internet from a service-based applications standpoint.

How does the infrastructure work?

BonFIRE operates a Cloud facility based on an Infrastructure as a Service delivery model with guidelines, policies and best practices for experimentation. BonFIRE adopts a federated multi-platform approach providing interconnection and interoperation between novel service and networking testbeds. The platform will offer advanced services and tools for research on services, including cloud federation, Virtual Machine management, service modelling, service lifecycle management, service level agreements, Quality of Service monitoring and analytics. Where appropriate, BonFIRE will reuse and adapt existing tools from other projects such as Panlab, FEDERICA and DEISA.

BonFIRE Concept

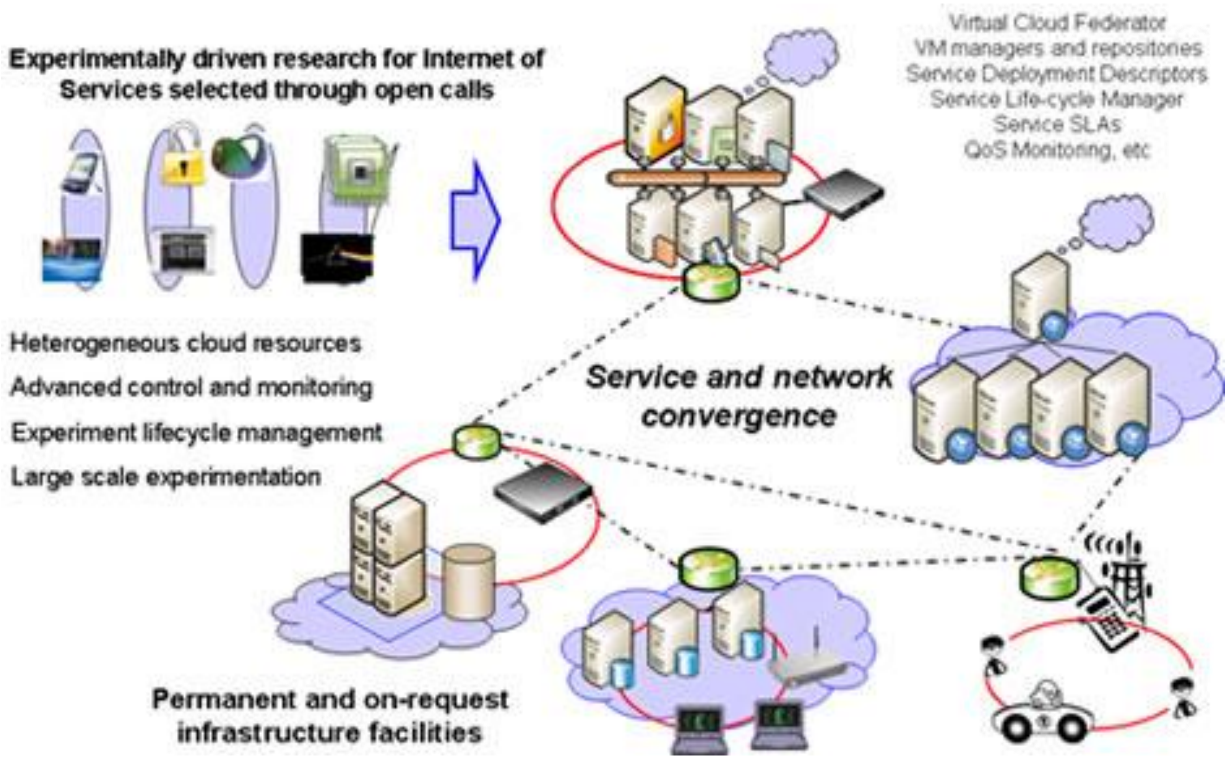

Figure 8 – The BonFIRE concept

BonFIRE comprises an initial set of testbeds at IBBT (Virtual Wall), HP Labs' cloud-computing testbed, University of Edinburgh's world-leading HPC systems (EPCC), INRIA and the University of Stuttgart (USTUTT):

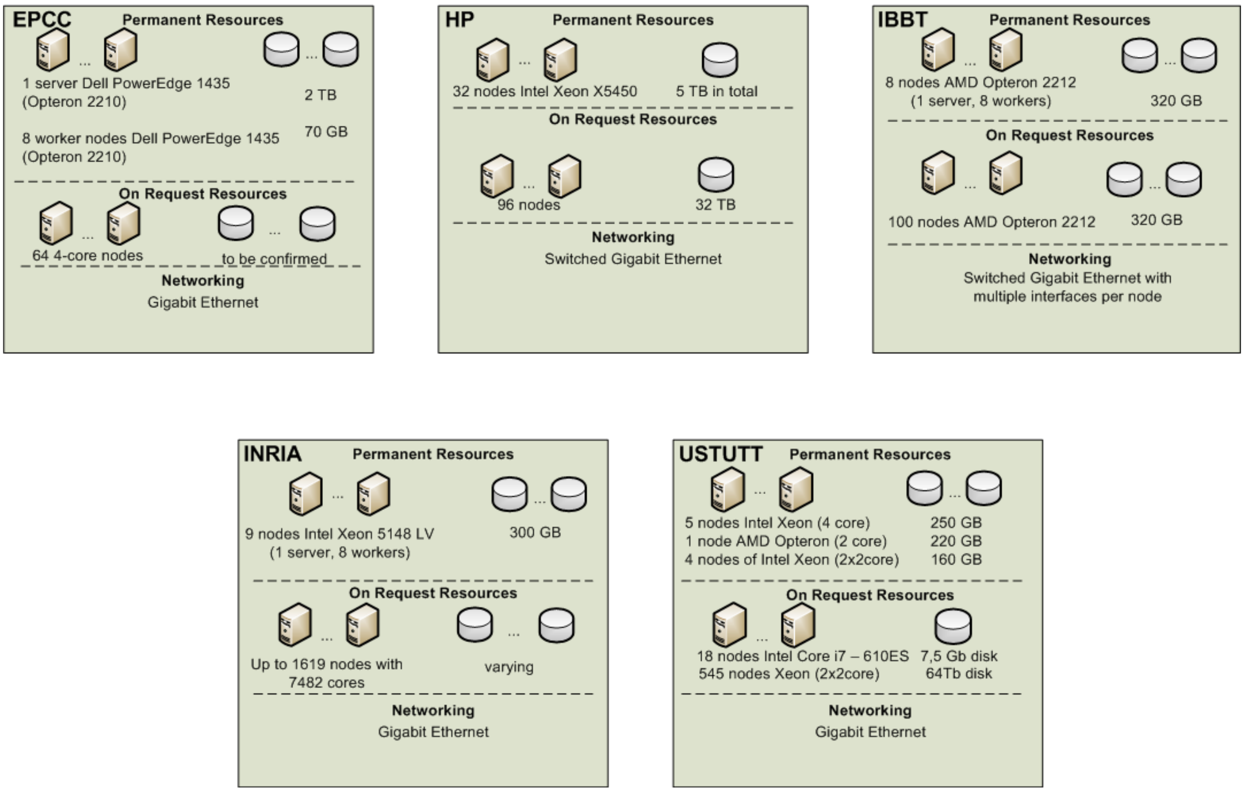

Figure 9 – The BonFIRE testbeds

What can be tested and how?

BonFIRE will support the experimentation and testing of innovative scenarios from the Internet of Services research community specifically focused on the convergence of services and networks. Three scenarios are envisaged:

- Extended cloud: the extension of current cloud offerings towards a federated facility with heterogeneous virtualized resources and best-effort Internet interconnectivity.
- Cloud with emulated network implications: a controlled network environment by providing an experimental network emulation platform to service developers, where topology configuration and resource usage is under full control of the experimental researcher.
- Extended cloud with complex physical network implications: investigates federation mechanisms for an experimental cloud system that interconnects individual BonFIRE sites with FEDERICA, Open Cirrus and Panlab.

BonFIRE will provide innovative methods for describing, deploying, managing, executing, measuring and removing experiments including:

- Uniform test description and deployment descriptors for all the scenarios (including cross-cutting tests)
- Cloud resource federation through the federation of clouds in different administrative domains that provide physical resources to the BonFIRE Project
- User-friendly user interfaces at the facility's entry point with an easy to use portal

Usage Information

The BonFIRE project will stimulate research through two Open Calls for experiments (total of €1.34M funding). The first of these was issued in March 2011 and the next one will be in March 2012. The facility will also be opened to researchers not funded through Open Calls later in the project.

7. TEFIS (Call 5) www.tefisproject.eu

TEFIS supports the Internet of Services community, by providing access to heterogeneous and complementary experimental facilities supporting the entire service lifecycle, including user behaviour, scale, performance and SLA compliance. It brings together 6 test facilities for IT-based testing, including the BOTNIA LivingLab, as follows:

Experimental facilities	Short description	Focus
PACA grid	Addresses parallel, distributed, and multi-threaded computing and cloud applications. Facility description *PACAGrid*	Computing resources for Future Internet experiments.
ETICS	Automatic build, test and quality certification for any distributed software exploiting distributed resources. Facilility description *ETICS*	An e-Infrastructure for Testing, Integration and Configuration of Software.
SQS-IMS	Validation and Testing of Converged Next Generation Services. Emulated and Real IMS networks supporting OMA, SIP, PGM, 3GPP, TISPAN, SS7 and IM standards. Facility description *SQS IMS Testbed*	Infrastructure to validate and test applications over IMS (IP Multimedia Subsystem).
BOTNIA	The BOTNIA LivingLab focuses on the support of human-centric innovation on advanced ICT Services for "Extended Capabilities and Mobility". Facility description *BOTNIA Living Lab*	End user involvement in testing and design.
Kyatera	Resources to develop science, technologies, and applications of the future Internet remotely collaborating via a high capacity optical network in São Paulo State (Brazil). Facility description *Kyatera*	Fiber optic network, web lab, remote collaborative work.
PlanetLab	The facility is a powerful infrastructure for the testing and evaluation of network protocols and distributed systems on a large scale under real conditions. Facility description *PlanetLab*	Facility to develop new technologies for peer-to peer systems, distributed storage, network mapping, distributed hash tables, or query processing.

The combination of testbeds offered by TEFIS allows a broad range of service characteristics including functionality, performance, scalability, usability, maintainability, user experience/ acceptability, and standards compliance.

The platform provides the necessary services that will allow the management of underlying testbed resources. In particular, it handles generic resource management, resource access scheduling, software deployment, matching and identification of resources that can be activated, and measurement services for a variety of testbeds.

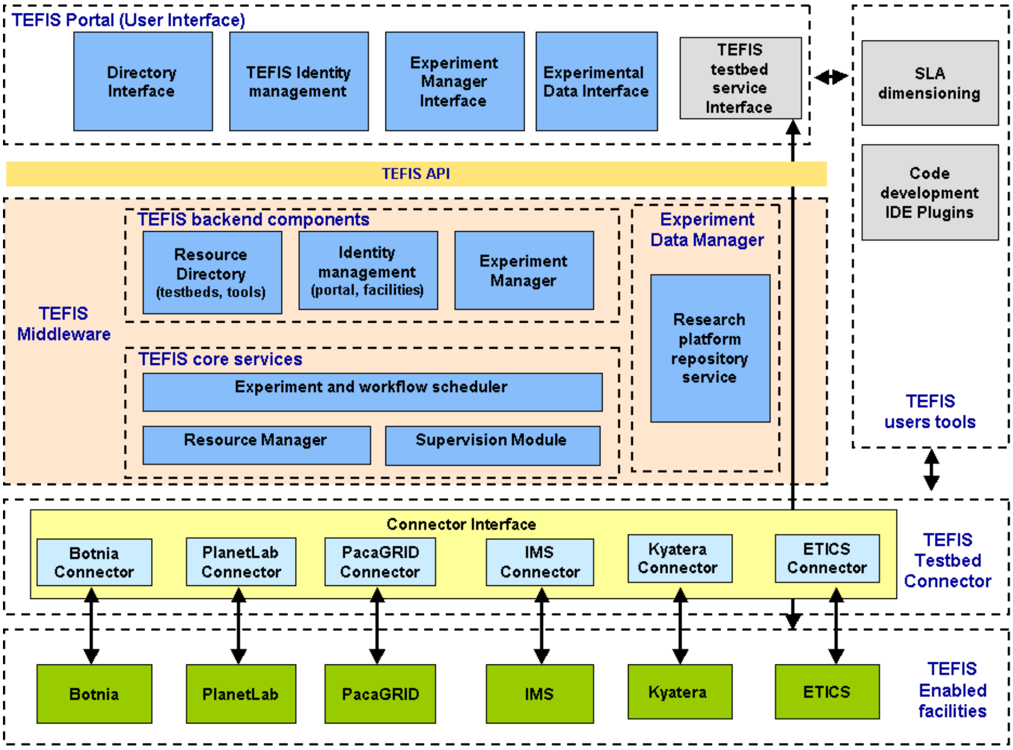

Figure 10 – The TEFIS Functional Architecture

TEFIS offers:

- An open platform to integrate and use heterogeneous testbeds based on a connector models, and exposed as a classical service.
- The integration of 8 complementary experimental facilities, including network and software testing facilities, and user oriented LivingLabs.
- A platform to share expertise and best practices.
- Core services for flexible management of experimental data and underlying testbeds resources during the experiment workflow.
- A single access point (Portal) to testbeds instrumented with a large number of tools to support the users throughout the whole experiment lifecycle (compilation, integration, deployment, dimensioning, user evaluation, monitoring, etc.). Testbed users will be able to create their personalized entry point to the Future Internet, by defining their own environment for their experiments.

What can be tested and how?

TEFIS provides a single access point for all test stages, through to live user trials in a LivingLab environment. It provides the seamless integration and management of test steps, including the automatic transfer of appropriate data from output of one process to input of another. TEFIS can be used to generate simulated load to drive stress testing and validate that SLA values can be met.

The project will contribute to standardize the way network and service facilities may be accessed. TEFIS will establish a connector model that makes it possible to interact with testbeds and their resources in a unified manner using Web services.

The Open Call process will be used to engage new experimentations and to gradually expand TEFIS. Combining the efforts of the software and service industry, the FIRE community and the user-centric LivingLabs, TEFIS will foster research and business communities in collaboratively elaborating knowledge about the provisioning of Future Internet services.

Usage Information

The TEFIS facility will be made available to experimenters from the scientific community and from interested companies, both through the process of Open Calls and through non-funded mechanisms. The first Open Call was issued in February 2011 and the next one will be 12 months later.

The TEFIS portal provides the appropriate support for the design and execution of experiments using multiple facilities (including resource brokering). In addition, TEFIS allows the experimenter to search for and review related experiments that have been run previously on the platform. This is an important feature in support of experiment definition.

The TEFIS facility relies on a distributed set of testbed facilities which each have their own technical requirements, usage protocols and administrative policies. For the duration of the project there is no additional constraint put by the TEFIS project on top of those of pre-existing infrastructures. One of the activities in the project is targeted towards the identification of suitable models for continuing the support (in particular maintenance costs) after the project's end.

8. OFELIA (Call 5) www.fp7-ofelia.eu

OFELIA is creating a unique experimental facility that allows researchers to not only experiment on a test network but also to control the network itself precisely and dynamically. To achieve this, the OFELIA facility is based on OpenFlow, a currently emerging networking technology that allows virtualizing and controlling the network environment through secure and standardized interfaces. In a nutshell, OpenFlow enables experimenters to change the behaviour of the network as part of the experiment rather than, if at all, as part of the experiment setup. OFELIA will provide high-performance OpenFlow equipment to enable experiments at scale and to ensure that the facility is based on mature technology.

The strength of OFELIA is its concept of federated or interconnected islands. A set of five islands creates a diverse OpenFlow infrastructure that allows experimentation on multi-layer and multi-technology networks provided by the different islands. The facility will extend all the way from standard Ethernet to optical and wireless transmission and it will also include an emulation wall for scalability tests comprising thousands of nodes. The facility will grow in three phases to, on the one hand provide an early access to the facility and, on the other hand to evolve during the project lifetime, incorporating the feedback of the user community and extending its reach to other test facilities.

OFELIA Testbeds

- Ghent (IBBT) – central hub, large-scale emulation
- Berlin (TUB) – partial replacement of existing campus network with OF-switches
- Zürich (ETH) – connection to OneLab and GENI
- Barcelona (i2CAT) – L2 (NEC) switches and Optical equipment (ROADM ring)
- Essex (UEssex) – national hub for UK optical community; L2 (Extreme) switches, FPGA testbed

Figure 11 – The OFELIA testbeds

OFELIA comprises 5 OpenFlow testbeds:

- iLab.t Virtual Wall run by IBBT

  The Virtual Wall consists of 300 nodes interconnected via a number of high speed VLAN Ethernet switches, and a display wall (20 monitors) for experiment visualization (see figure: 100 nodes, switch and displays). Each server is connected with 1 up to 6 gigabit Ethernet links to the switches. The nodes can be assigned different functionalities ranging from terminal, server, network node, and impairment node. The nodes can be connected to test boxes for wireless terminals, generic test equipment, simulation nodes (for combined emulation and simulation) etc. A Full Automatic Install feature is provided for fast context switching (e.g. 1 week experiments), as well as remote access.

  The lab also has advanced generic test equipment (SPIRENT Smartbits 2000 - 6000 (packet generator/analyser + L4-L7) and SPIRENT Testcenter; Spirent Avalanche/Reflector 2200 (L4-L7); Agilent N2X (protocol tester); Opticom Opera (voice and audio tester); Fluke Optiview Link analyzer; Agilent Broadband Series Test System (BSTS); 4 Qosmotec shielded wireless environments; several wireless sniffers; ...).

- Berlin testbed run by the Technical University of Berlin

  This testbed is distributed throughout student rooms located in the first wiring-centre in the 5th floor of the TUB Campus Fränklin building. A combination of aggregated and meshed switches is provided, depending on the requirement of the experiments.

  Optical fibres are used for the connection between the switches. This provides a clean electrical network separation and a high-performance connection. The available physical Ethernet links are logically separated in standard configured and OpenFlow-enabled connections. Additionally, the standard links are sub-divided using VLAN techniques in the workgroup specific areas. Additionally, the OpenFlow-VLAN can be dynamically sliced using the FlowVisor tool.

- Zürich testbed run by the Eidgenössische Technische Hochschule Zürich

  This testbed comprises three OpenFlow switches NEC IP8800/S3640-24T2XW. Each of them includes 2 optical 10GB interfaces, and is capable of rewriting MAC, IP, and some of the TCP header fields before forwarding a frame. In order to run the Virtual Machines and workload generators we will deploy additional machines. The machine for the FlowVisor will be equipped with at least a Intel Quad-core processor and 4 GB of RAM.

- Barcelona testbed run by i2CAT

  The i2CAT testbed consists primarly of a set of 5 NEC OpenFlow-enabled switches and 5 XEN-virtualized servers, in which Virtual Machines will be able to be created by experimenters to act both as end-points and OpenFlow controllers. The NEC switches will form a completely meshed topology between them. Each of the servers will be connected to 3 different NEC switches and to the GW switch.

  All of the NEC switches run version 1.0 of the OpenFlow Protocol against the unique island FlowVisor. In addition, Virtual Machines created in the XEN servers will contain a software bundle with all the available1.0-compatible OpenFlow sofware (reference implementation, NOX, SNAC and FlowVisor), as well as other development tools.

- Essex testbed, run by the University of Essex (UEssex)

  UEssex is a national hub for the UK optical community. The UEssex island will allow experiments to be conducted on slices consisting of meshed OpenFlow-enabled NEC Ethernet switches, and End nodes, acting as traffic sources and sinks, residing on Virtual Machines running Linux and XEN virtualization software. ADVA ROADMs will be OpenFlow-enabled in Phase 2 of the project. Users will be provided with a default OpenFlow controller per slice running on a Virtual Machine. It is also possible to run custom images of an OpenFlow controller on the Virtual Machines. FlowVisor will be used for slicing the network topology.

The five islands based on OpenFlow infrastructure will be created and interconnected to allow experimentation on multi-layer and multi-technology networks. The facility will extend all the way from standard Ethernet to optical and wireless transmission and it will also include an emulation wall containing 100 nodes for scalability tests. OFELIA will provide an experimentation space which allows for the flexible integration of test and production traffic by isolating the traffic domains inside the OpenFlow enabled network equipment. This allows for providing realistic test scenarios and for seamless deployment of successfully tested technology.

What can be tested and how?

OFELIA will support the experimentation and testing of OpenFlow in terms of users receiving a network slice consisting of:

- Virtual Machines as end-hosts
- A Virtual Machine to deploy their OpenFlow-capable network controller/application

- Parts (slices) of the network nodes that connect to the user's OpenFlow controller
- Control of a subset of the flowspace in a subset of the switches

OFELIA is intending to establish interoperability with other testbeds as a primary target. The project will start with federation modules developed for the US-based GENI (Global Environment for Network Innovations), but at a later stage extend those by, for example, support for multiple layers and heterogeneous technologies. The automation of managing resources across interconnected OpenFlow networks is another topic of the project.

Usage Information

OFELIA will invite experimenters in Europe to bring their own proposals for Use Cases and scenarios to the OFELIA test facility.

Two Open Calls will be published offering the successful experimenters additional funding from the OFELIA project for conducting their experiments. The first of these Open calls has been held in Spring 2011.

Researchers get remote access to an OpenFlow controller through a SSH connection to a Rack-Server mounted beside the switches. In Phase 2, the data traffic of the OpenFlow workgroups will be separated, forwarded and bridged through a VPN tunnel to the IBBT testbed in Ghent to which all other testbeds will be connected.

The testbed control framework will follow the Slice-based Federation Architecture (SFA). In the first deployment of the control framework, the islands will operate isolated from each other and have at least two aggregate managers: one that manages the network of OpenFlow-enabled devices, and another one that manages the virtualization-enabled servers. A tool called Expedient will be the testbed control framework front-end for both island managers and researchers; the former will use Expedient to assign resources to projects, while the latter will use Expedient to configure, start and stop their slices. The Expedient tool is an ongoing effort lead by a development team at Stanford University that agreed on a joint effort to continue the development of Expedient and its customization for the OFELIA testbed.

In addition, the individual islands have their own particular access conditions. For example:

- In Phase 1, external users will access the testbed by using an OpenVPN L2 connection over the Internet between i2CAT and IBBT OpenVPN. In Phases 2 and 3, i2CAT will have its own user OpenVPN gateway.
- External Access to UEssex is provided via a VPN over a 10GE link from JANET which terminates on Carrier Grade Extreme Black Diamond Switches. 1 or 10 GE GÉANT links will be used to pair with islands already having connectivity to GÉANT. All other islands will be interconnected through VPNs. Extreme Black Diamond switches will be OpenFlow-enabled in Phase 2 of the project.
- One limitation is that the 4K video coder and decoder cannot be shared between experiments, since the Calient Diamond Wave switch is not OpenFlow-enabled, but still can be virtualized.

Testbed timeframes

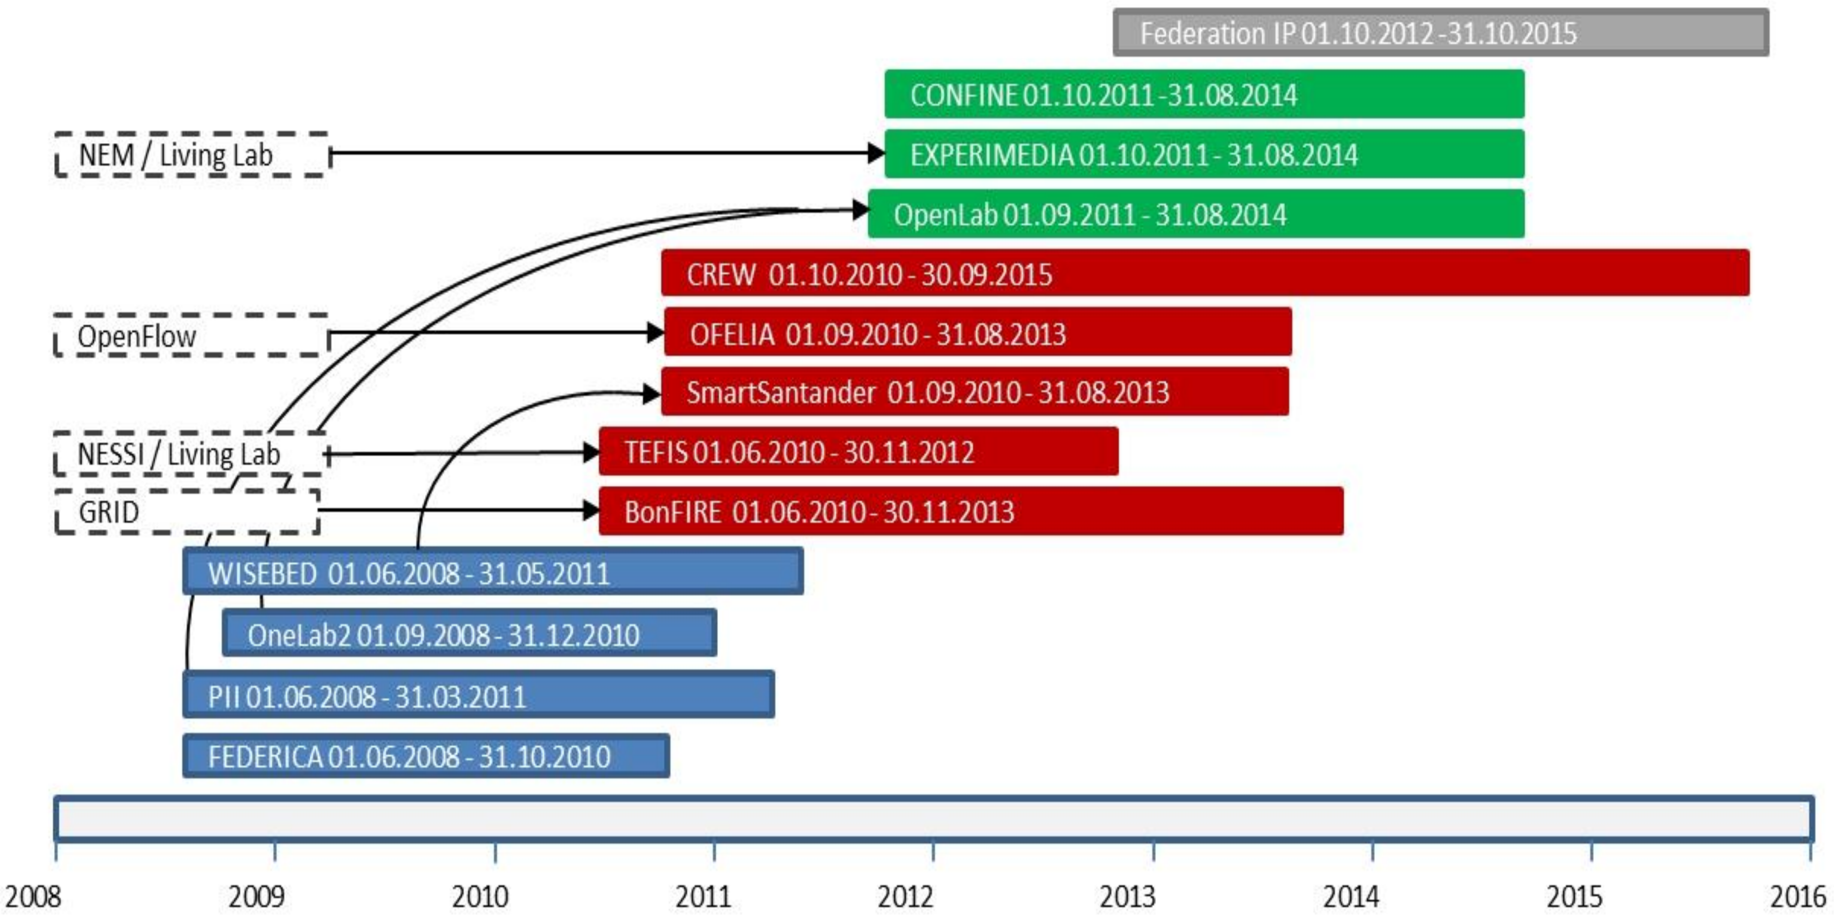

Figure 12 – The Testbed timeframes

4. The usage of the FIRE facilities

Research into new paradigms and the comprehensive test facilities upon which the ideas are experimented together build a key resource for driving European research into future networks. This environment enables both incremental and disruptive approaches, supports multi-disciplinary research that goes beyond network layers, scholastic dogmas and public-private discussions. It provides a core infrastructure, and also a playground for future discoveries and innovations, combining research with experimentation.

The heterogeneous and modular field of Future Internet Research and Experimentation with its national and international stakeholder groups requires community and cohesion building, information sharing:

- **Facilities** need synchronization, resource optimization, and common efforts in order to offer customers the best possible service and ensure their sustainability beyond project life times.
- **Researchers** need correct and timely knowledge about the available resources, easy access, high usability and appropriate tools to run and monitor their experiments.

FIRE STATION provides the FIRE Initiative with an active hub that matches, guides and co-ordinates demand for - and offering of - experimentation facilities in the context of future networks. FIRE STATION implements a FIRE Office and a FIRE Architecture Board that both build on earlier FIRE co-ordination results. The FIRE Office serves as the single contact point and the mediator when either looking for the right experimental resources or new customers for the facilities. The FIRE Architecture Board involves all FIRE facility builders to jointly decide on strategy and means to co-ordinate and facilitate the development of FIRE facility offerings supporting the evolving needs of the customers. FIRE STATION increases global collaboration between relevant stakeholders, promotes experimentally-driven approach in Future Internet research and intensifies the usage experimental facilities, ultimately speeding up the development process of new systems and services.

A federation of testbeds aims at creating a physical and logical interconnection of several independent experimental facilities or testbeds to provide a larger-scale, more diverse and higher performance platform for accomplishing tests and experiments. A collaboration/federation framework is not about having rigid control of all aspects. On the contrary, the aim is to have flexibility and preserve the autonomy and character of the components/projects. In this sense, high-level federation does not at all mean to agree on the same control plane, but rather to allow resource sharing and collaboration towards establishing a sustainable customer-friendly facility. It is important to maintain the major goal of each individual testbed project, which is to create innovative solutions for testing and to support their community of experimenters, while at the same time also contributing to the common goal of collaboration and federation of experimental facilities.

The Use Cases (ongoing and planned), have been categorized as:

- Use Cases from internal partners,
- Use Cases from FIRE STREPs,
- Use Cases from other FP7 projects,

- Use Cases from outside FP7,
- Use Cases from the Open Calls.

Facility	**Use Cases**				
	Internal	**FIRE STREPs**	**Other FP7 projects**	**Experiments from outside FP7**	**Open Calls**
Panlab-PII	• Testing adaptive admission control and resource allocation algorithms • EzWeb application over TID SDPLabs: “PII Message Sender” • Testing enhanced Web TV services over mobile phones • Stress test the Open IMS core	• SELF-NET: Testing self-management in a wireless future internet environment. • The VITAL++ project ported its facilities to Panlab. In particular, access to VITAL++ was enabled via the Teagle tool available in the Panlab portal.		A large number of different and quite heterogeneous Use Cases have been implemented in the PII project. The following list presents the Use Cases which can be retrieved in detail via http://www.panlab.net/use-cases.html • Testing trans-coding video through dynamic cloud allocation • Testing Multicast Streaming on Dynamic Networks • Testing Uncompressed HD Streaming • Testing a VOIP user agent As can be inferred by all Use Cases, typical experiments that can be directly supported by the PII testbeds are Next Generation Network (NGN) services and applications.	N/A
OneLab (the PlanetLab Europe, NITOS, ETOMIC, and DIMES testbeds)	(PLE:) 1. EGOIST: Overlay routing 2. CBG: Geolocation services 3. BitTorrent systems:	OneLab facilities were selected by 6 STREP projects, namely: Resume-Net, EULER, Nanodatacenters, NOVI, SCAMPI and CONECT. CONECT selected OneLab2 for	(PLE:)N/A (NITOS:) 1. FIBRE (Future Internet testbeds/ experimentation between Brazil and Europe)	Since the majority of OneLab’s testbed nodes are open to the public Internet, researchers can experiment with distributed applications in a real-life testing environment. OneLab Use Cases include the	The OneLab testbeds will participate in the 2nd Open Call for FIRE Users, through the OpenLab project.

Facility	Use Cases				
	Internal	**FIRE STREPs**	**Other FP7 projects**	**Experiments from outside FP7**	**Open Calls**

	Content distribution
	4. Online gaming: peer-to-peer scalable alternatives
	5. Radar: Internet topology
	(NITOS:)
	6. Evaluation of network coding implementations
	7. Prototype implementations of user association and frequency selection algorithms for 802.11
	8. Novel routing algorithms for multihop transmission of multimedia content over the wireless medium
	(DIMES:)
	9. Studying routing and delay stability (2010)
	10. PoP level maps (2010-2011)
	11. Accuracy of GeoIP
		federation with its own built facilities.
		
		NOVI plans to federate PlanetLab with FEDERICA: Virtual hosts over a Linux host in PlanetLab interworking with logical routers and VMWare-based virtual nodes in FEDERICA.
		
		The NITOS testbed for network coding experiments within N-CRAVE (FP7 STREP)
			2. COOPLAB
			3. REDUCTION
			(ETOMIC:)
			4. Data sharing from heterogeneous traffic sources (MOMENT FP7 project)
				following:
				• Network topologies
				  ○ Delays
				  ○ Available bandwidths
				• Routing paradigms for wired and wireless networks
				  ○ Real-life routing changes
				  ○ Routing and content sharing topologies: minimizing routing traffic or content queries
				• Real-world Internet traffic measurements
				  ○ Unpredictable traffic loads
				• Applications failure recoveries
				• Concurrent multipath transmissions
				System optimization for multi-hop networks	

Facility	Use Cases				
	Internal	**FIRE STREPs**	**Other FP7 projects**	**Experiments from outside FP7**	**Open Calls**
	database (2011) 12. AS classification using graphlets (2011) (ETOMIC:) 13. Network delay tomography 14. Router fingerprinting 15. Available bandwidth measurements 16. IP geolocation				
FEDERICA		1. PERIMETER experiments have been run on a FEDERICA slice. 2. The FEDERICA facility has also been chosen by the NOVI consortium and it is planned to be federated with PlanetLab: (see above). 3. FEDERICA is discussing with BonFIRE how to enable cross-platform	1. PHOSPHORUS: Experimentation of a multi-domain, multivendor network resource brokering system, called "Harmony" 2. FEDERICA will be supported in GN3 as a Service Activity until April 2012	The FEDERICA consortium has served more than 20 user groups, the paper [P.Szegedi et al, 'Enabling Future Internet Research: The FEDERICA Case' IEEE Communications Magazine - July 2011] details the most interesting ones. Currently two slices, requested from KTH and university of Madrid, are being configured to run new experiments. The experiments typically fall into one of three areas: 1. Validation of Virtual Infrastructure features	N/A

Facility	Use Cases				
	Internal	**FIRE STREPs**	**Other FP7 projects**	**Experiments from outside FP7**	**Open Calls**
		experiments		2. Evaluation of multi-layer network architectures, using the capability of connecting external testbeds or virtual slices provided by other facilities to the FEDERICA user slice 3. Design of novel data and control protocols: this group of experiments aims at designing and validating novel data and control plane protocols as well as architectures.	
WISEBED		WISEBED was chosen by 2 STREPs: 1. HOBNET, primarily because two project partners participate in both projects. However, several distinct features exist, i.e. WISEBED is not focusing on IPv6, so significant extensions and innovations from HOBNET are expected to be ported to WISEBED. 2. SPITFIRE chose WISEBED as it provides a large-scale IoT testbed that is an	WISEBED was extensively used by the European project FRONTS, specifically to experiment with a number of algorithms developed by theoretical computer scientists.	1. Coalesenses: Evaluation of new protocols for sensor nodes like routing, medium access control and time synchronization	N/A

Facility	Use Cases				
	Internal	**FIRE STREPs**	**Other FP7 projects**	**Experiments from outside FP7**	**Open Calls**
		ideal environment to use for exposing services offered by resource-constrained sensor nodes.			
SmartSantander	1. Coding for network transmission on large-scale IoT 2. Service provider scenarios (eg pricing) 3. City services (eg parking, vehicle traffic management, environmental monitoring)	1. SCAMPI has identified SmartsSantander as a possible testbed on which to validate part of its work on socially-driven applications in opportunistic scenarios. 2. SPITFIRE has identified SmartsSantander as a possible testbed on which to validate part of its work on extending the Web into the embedded world to form a Web of Things (WoT).		SmartSantander has identified a considerable number of Use Cases both for scientific experimentation and service execution for the citizens on the platform. In terms of services, Use Cases have been defined in the fields of traffic control (for instance, load/unload area management, disabled parking management, virtual corridor for emergency vehicles, control of vehicles parked in bus stops), supporting of people with disabilities and illnesses, cultural activities, smart metering, environmental monitoring, public transportation, urban waste management, multipurpose physical space interaction, smart schools, precision irrigation management system, and aquifer management-monitoring system. In the first cycle, implementation of these Use Cases concentrates on traffic	

Facility	Use Cases				
	Internal	**FIRE STREPs**	**Other FP7 projects**	**Experiments from outside FP7**	**Open Calls**
				control application area. In terms of scientific experiments, ideas include network coding and rateless coding techniques combined with MOTAP (multi hop over-the-air reprogramming) on top of massive wireless sensor nodes, IPv6 multicast in wireless sensor networks, geocasting routing protocols applied to a WSN environment, geographic routing for wireless sensor networks in presence of network dynamics, IPv6-based interaction of sensor nodes and the Internet, Distributed Game-Theoretic Vertex Colouring for Frequency assignment and energy aware node partitioning, or evaluation of video quality under complex network deployment and different transmission technologies. The platform is open for ideas for further Use Cases.	
CREW				1. Context awareness for cognitive networking: spectrum sensing in unlicensed (ISM) and licensed bands (TV white spaces, cellular systems)	

Facility	Use Cases				
	Internal	**FIRE STREPs**	**Other FP7 projects**	**Experiments from outside FP7**	**Open Calls**
				2. Robust cognitive networks: applications that require robust communications though avoiding harmful interference and using frequency agility to improve communication quality 3. Horizontal resource sharing in the ISM bands: algorithms, protocols and networking architectures for coexistence of and cooperation between independent heterogeneous network technologies 4. Cooperation in heterogeneous networks in TV bands: new techniques for context awareness in unlicensed (ISM) and licensed bands (TV white spaces, cellular systems) 5. Cognitive systems and cellular networks: the impact of dynamic spectrum access by secondary users on LTE cellular primary systems.	
BonFIRE	1. Dynamic Service Landscape Orchestration for Internet of Services 2. QoS-Oriented	1. The LAWA experimental testbed is a platform for analytic tools that will showcase a "Virtual Web Observatory". The initial configuration will			1. TurboCloud *Partners: RedZinc and Cloudium Systems* An experiment which combines 2 complimentary

Facility	**Use Cases**				
	Internal	**FIRE STREPs**	**Other FP7 projects**	**Experiments from outside FP7**	**Open Calls**
	Service Engineering for Federated Clouds 3. Elasticity in cloud based web applications)	be based on data from the LAWA Reference Collection, but the plan is to include a federation with BonFIRE.			technology platforms from 2 SMEs. One technology platform (Cloudium chipset) enables server-based desktop virtualisation. The other technology platform (VPS controller) enables dynamic virtual path slices to deliver a right of way across the Internet without interference from unwanted traffic. 2. VCOC: Virtual Clusters on Federated Cloud Sites *Partner: CESGA* The experiment will investigate the feasibility of using several Cloud environments for the provision of Services which need the allocation of a large pool of CPUs or Virtual Machines to a single user (as High Throughput Computing or High Performance Computing). 3. ExSec: Experimenting Scalability of Continuous

Facility	Use Cases				
	Internal	**FIRE STREPs**	**Other FP7 projects**	**Experiments from outside FP7**	**Open Calls**
					Security Monitoring in BonFIRE *Partner: CETIC* The ExSec experiment aims to determine an empirically validated elasticity function for security monitoring. Besides verifying the scalability of the security monitor on different application loads for a number of Virtual Machines, another important aspect of the experiment is to verify scalability behaviour on different cloud technologies such as different types of hypervisors and different types of cloud environment managers. 4. TEOS: Testing Optimization in Service Ecosystems *Partner: University of Manchester* This experiment aims to determine the conditions for achieving resilient

Facility	Use Cases				
	Internal	**FIRE STREPs**	**Other FP7 projects**	**Experiments from outside FP7**	**Open Calls**
					and optimal service compositions on a distributed cloud infrastructure for the Future Internet. It will deploy and test two service optimization models, characterized as global optimization and local optimization.
TEFIS	The TEFIS partners have included 2 representative Future Internet scenarios as internal use-cases: 1. e-Commerce (eTravel) 2. e-Health These two scenarios provide a powerful demonstration of the benefits of TEFIS, as each one requires multiple federated test facilities to be managed as a single, complete test entity in various stages of the specific tests.			1. The mobile application for content sharing use-case. This external use-case is included as a scenario to empower the development of the TEFIS platform. The use-case is divided into three different phases of the service development life-cycle: concept development, prototype development and Business model definition and involves resources from two different testbeds.	1. Smart Ski resort *Partner: University of Geneva - Switzerland* 2. Dynamic Quality User Experience Enabling Mobile Multimedia Services (QUEENS) *Partner: Institute of Communication and Computer systems (ICCS) - Greece* 3. Augmented Reality Collaborative workspace using future internet videoconferencing platform for remote education and learning (TEFPOL) *Partner: Poznan Supercomputing and Networking Center -*

Facility	Use Cases				
	Internal	**FIRE STREPs**	**Other FP7 projects**	**Experiments from outside FP7**	**Open Calls**
					Poland 4. Experimenting with Quagga Open API and cross-layer Coordinated networks *Partner: Universitat Politecnica de Catalunya - Spain*
OFELIA	1. Performance testing of a programmable flow processing platform, containing processing modules (IP router, IDS, Firewall) on a wide area network setup using OFELIA architecture. 2. Reliability and scalability testing of OpenFlow-enabled networks. 3. Testing methodologies for implementing High performance Datacenters and Access/Aggregation networks using OpenFlow. 4. Energy aware	Initial talks ongoing with CREW in the area of WiFi virtualisation.	1. SPARC is currently testing carrier grade extensions for OpenFlow using the OFELIA facility. Features under investigation include: • Enhanced topology discovery • Virtualization and isolation • Flow OAM • Openness and Extensibility • Resiliency 2. SPARC is implementing a prototype for an OpenFlow based access/aggregation domain. This will document among	Contacts with Japanese OpenFlow testbed (JGN-X), joint demo for summer Olympics planned. (see internal Use Case No.5)	1. VERTIGO (Virtual Topologies Generaliza-tion in OpenFlow networks): is a novel slicing mechanisms to overcome FlowVisor limitations. VERTIGO provides means to adapt the data path model as exposed by the user plane accor-ding to operator/ researcher needs and is an advanced slicing mechanism for OpenFlow based networks. 2. EXOTIC (Extending OpenFlow to Support a future Internet with a Content-Centric model): aims towards extending the Internet with content

Facility	**Use Cases**				
	Internal	**FIRE STREPs**	**Other FP7 projects**	**Experiments from outside FP7**	**Open Calls**
	Virtual Machine migration between distributed Datacenters, to measure the time and performance of VM migration on OpenFlow-enabled networks using OFELIA architecture. 5. UHD (4k, 8k) video transmission over OpenFlow networks to analyse the network performance based on Media aware control applications to be used for the Olympics 2012 demos and games. 6. Evaluate and test the functionality of tools and software developed for virtual network control plane and virtualization mechanisms of layer 2 and layer 1 networks using OFELIA architecture.		others: • Sophisticated layered control plane • Improved namespace management 3. CONVERGENCE will run experiments in 4 trials in cooperation with the OFELIA project. CONVERGENCE project sites have access to GARR/GEANT via which they will reach existing OFELIA islands. The 4 Use Cases that the CONVERGENCE project proposes are: • Management of audiovisual material / semantic and cognitive indexing. • Management and annotation of a large photograph archive. • Customer Relationship Management and		centric networking support. EXOTIC will develop adequate controller architectures for content based routing and forwarding and investigate necessary extensions to the OpenFlow API in order to support content centric networking (CCN).

Facility	Use Cases				
	Internal	**FIRE STREPs**	**Other FP7 projects**	**Experiments from outside FP7**	**Open Calls**
			logistics (shopping environment). An augmented lecture podcast service.		

The table above shows that there was a low usage of the FIRE experimental facilities from the Call 2 STREPs, however, this situation is changing for the Call 5 projects, which appear to have planned a more systematic and active uptake and deployment of existing FIRE facilities.

As seen from Section 7 of this report, the STREPs have highlighted that the FIRE facilities make it possible to:

- Have a greater diversity of technologies and infrastructures.
- Make experiments on a larger scale.
- Obtain more advanced experimentation results, based on multiple metrics.
- Better justify the research results, since the experiments are performed in "close-to-real-life" conditions.
- Gain technical know-how about an increased number of technologies and equipments.
- Reuse the same technologies and resources across several projects, thereby allowing return on investment to be maximized.
- They also recognise that testbeds are in principle living beyond a single STREP project lifetime as they remain to be available to the broader audience for follow-up work.

The Open Call mechanism is further encouraging the use of the facilities and making the facilities better known. In the 1st Open Call (for BonFIRE, OFELIA and TEFIS), 71 proposals were submitted of which about 10 will be retained.

Note that, now the CaLL 5 FIRE facilities have established themselves, they are also making use of the Call 2 FIRE facilities and exploiting testbeds funded from outside the FIRE Programme, as shown below:

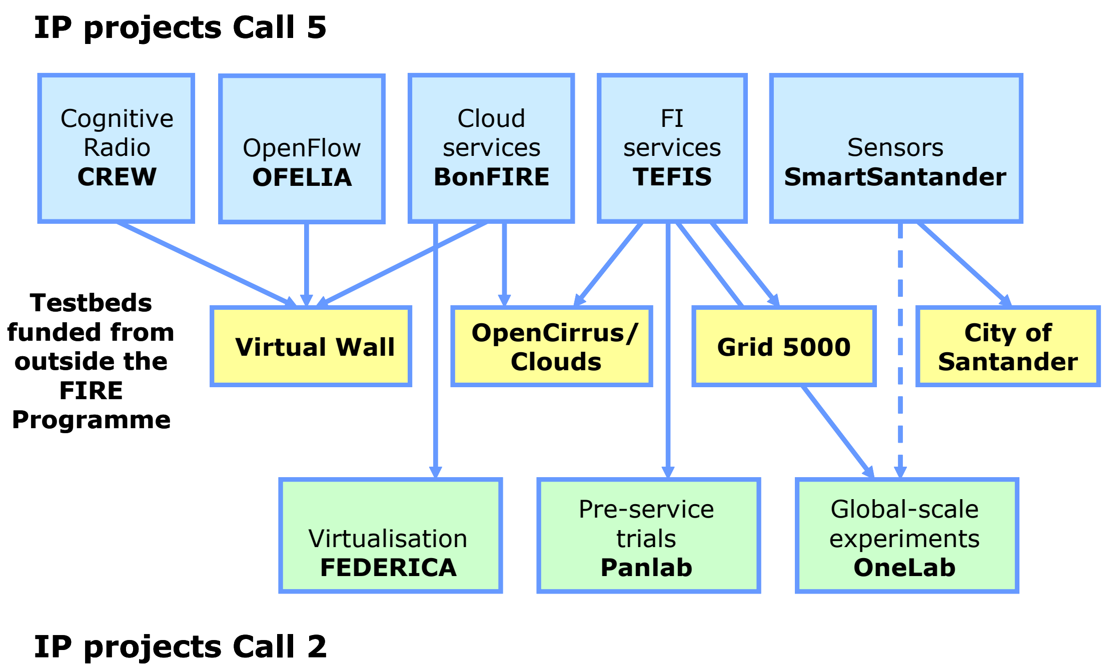

Figure 13 – The mutual usage of the FIRE facilities and external testbeds

5. Joint activities between the facilities

Three Working Groups have been established within the Architecture Board to work on activities that are of general interest to all the facility projects:

- Resource Brokering Tools
- Sustainability
- Measurement and Benchmarking

Resource Brokering Tools

Collaboration between facility building projects is an essential activity to help ensure effective cross-project technical strategy, technical coherence and efficiency in European testbed facility developments. A key aspect of all facility projects is the need to provide methods, tools and services supporting experiment lifecycle management, and how such mechanisms interact with the control and monitoring channels of specific testbed infrastructures. The definition of the experiment lifecycle is not particularly challenging and many have been published. In general there are activities of Experiment Design and Definition, Systems Integration, Resource Provisioning, Execution, Termination, Analysis and Dissemination of Results.

Each of the facility building projects needs to address these activities to some degree either through specific tooling or explicit manual processes. Various solutions have emerged from Call 2 projects and further technologies are being proposed by Call 5 projects. It is noted that the characteristics of specific experiments and underlying testbed infrastructures will tend to have significant influence on how experiment management is implemented. What is appropriate for one facility may not be appropriate for another. It is important to understand the architectural principles and how they are achieved by each project. Example architectural principles (but not limited to) include:

- Common experiment definition language and resource models,
- Automated resource lifecycle management,
- Large-scale experiments,
- Decentralized/centralized experiment management,
- Experiment steering,
- Experimental data as a persistent shared asset.

Some experiments require access to a single testbed provider site whilst others require resources from multiple sites (e.g. for greater scalability or heterogeneity). When a facility is orchestrating experiments at multiple sites under different domains of control, experiment management can be directly related to federation. If experiment management tools address federation concerns, then it is of interest to understand how the federation technologies help contributes to the architectural principles of a facility.

A workshop on Resource Brokering Tools was held on 23 November, 2010 in Brussels, with participation from all facility projects. The workshop was facilitated by the FIRE STATION project.

The topic was identified during the 1st Architecture Board meeting where board members expressed the desire to share experience as the basis for influencing architectural decisions in Call 5 projects. This was especially relevant to CREW, OFELIA and SmartSantander projects which all started in September, 2010.

Workshop Goals

The goals of the workshop were:

- to understand technical details about experiment management tools offered by Call 2 projects and how they are designed to meet desired architectural principles
- to share lessons learnt from Call 2 projects about the success or otherwise of current solutions
- to understand the architectural directions of Call 5 projects, implementation options and constraints

A White Paper has been written: http://www.ict-fire.eu/fileadmin/publications/20101123-Workshop-Experiment-Management-Report-final3.0.pdf.

The differences in approaches are still very great, and this work is still at the level of creating a mutual understanding.

Sustainability

The sustainability of the FIRE experimental facilities is in the interest of all stakeholders. It would ensure that research experiments can be planned and repeated over long periods of time, since nowadays the sustainability of a FIRE facility is not guaranteed by the facility owner or by the funding organization. The current mechanisms do not prevent important pillars of FIRE from disappearing, even when there is still a demand from the users. This decreases dramatically the willingness of potential users/experimenters planning to use such facilities, especially outside the consortium in charge of building and/or operating the facility. Consequently, it is critical for FIRE to ensure that experimental facilities are maintained long enough to allow users to plan usage and perform experiments without risking that facilities will disappear.

For the facilities it is equally important to adjust to new requirements that are in the plans of new experimental projects. This does not mean however that each and every facility needs to remain available forever. In fact, sustainability cannot be addressed as an independent topic without having a broader strategy for FIRE as a whole. Clearly, just asking for more public funding for each facility to last as long as possible is not viable. The suggestion from the projects is to draft (and then implement) a "FIRE Roadmap", with information on what the overall objectives of FIRE are, and how the issue of sustainability (at the global FIRE level and then for each facility, or rather each available feature or set of features) can be addressed to respond to these objectives.

The Roadmap will include aspects relating to the initial development and deployment of the facility, operation and maintenance, advertising for external usage, managing external users including customer support, extension and expansion of the features and the accessibility of the facilities, etc. shall all be considered. Traditional ways shall be investigated (EU programmes such as Research Infrastructure, EUREKA, national R&D programmes and initiatives) but also more innovative/unusual ways of attracting public and also private investment, with a focus on activities that are not easily funded (e.g. operation and maintenance, customer support). Sustainability requires innovation in terms of e.g. revenue, payment, and business and governance principles, for all FIRE facilities.

Such a strategy will likely focus on the medium- and long- term as it will take some time to be developed and agreed upon. Therefore, in parallel, the projects suggest to also address short-term aspects as follows:

- *Strengthen the effort to attract more users in the short-term:*
    - Some organizations are interested to use the FIRE facilities as opposed to building a testbed on their own. This cost-saving option could be better advertised to the Future Networks and Software & Services communities (and beyond),
    - Organizations such as Intune Networks from Ireland and the University of Essex from the UK can help attract users by "adding" their test facilities (which focus respectively on optical burst switching and optical networks) to FIRE. Some national testbeds (e.g. G-Lab, F-Lab, PL-Lab) could also help enhance the FIRE offer and attract more users, though the conditions for such an interaction have to be defined and agreed upon,
- *Sustainability in the short-term.* This could be through funding in Call 8 for the STREPs to pay for using the FIRE facilities. This could be presented in such a way that in this case the

A workshop on *Sustainability* for FIRE facilities was held on 8 February, 2011. Presentations from MyFIRE, FIREBALL, FIRE STATION, BonFIRE during this meeting are available from the FIRE STATION Website: www.ict-fire.eu.

A White Paper on the Sustainability of testbeds is currently a working document only for the Architecture Board members. The next version of the Roadmap document, due at the end of November 2011, is a public deliverable and will include a section on Sustainability, derived from the White Paper and further discussions

Measurement and Benchmarking

By "benchmarking tools", one usually refers to a program or set of programs, used to evaluate the performance of a solution under certain reference conditions, relative to the performance of another solution. Since the 1970s, benchmarking techniques have been used to measure the performance of computers and computer networks.

The goal of a benchmarking process is to enable fair comparison between different solutions, or between subsequent developments of a System Under Test (SUT). These measurements include *primary performance metrics,* collected directly from the SUT (e.g. application throughput, node power consumption), and in case of wireless networks also *secondary performance metrics,* characterizing the environment in which the SUT is operating (e.g. interference characteristics, channel occupancy). The primary and secondary performance metrics may be complemented by *techno-economic metrics*, such as device cost; operational complexity. or traffic traces.

Although benchmarking, in its strictest sense, is limited to measuring performance, several additional aspects are important to make benchmarking a meaningful research activity.

Comparability should be a fundamental property of any benchmark; comparability means that two independently executed benchmarks can be meaningfully compared to each other. One of the factors influencing the comparability is *repeatability*: running an identical benchmark on an identical solution at different moments in time should result in a (close to) identical result. Furthermore, well-defined *experimentation methodologies* are a key factor in achieving *comparability*.

Ideally, benchmarking scores should not only be comparable to other scores obtained using the same testbed, but also with scores obtained from different testbeds with similar capabilities but potentially running different operating systems, or based on different types of hardware. The success of a specific benchmark may very well depend on whether this *interoperability* aspect is satisfied or not.

Test infrastructures may be equipped with benchmarking functionality. In this case, the *configurability* of the testbed environment is crucial: in wired networks, a benchmark may require a specific topology and links of specific quality; in wireless networks, a trace containing reference background traffic may need to be played back during the execution of a benchmark.

Benchmarks therefore have to be defined, executed and shared within the research community.

A session in the FIRE Research Workshop at the Future Internet Week in Budapest was dedicated to *Measurement and Benchmarking*, with contributions from Fraunhofer, BonFIRE, IBBT and OneLab.

A White Paper on *Measurement and Benchmarking* has been developed, available at: http://www.ict-fire.eu/fileadmin/publications/Whitepaperonbenchmarking_V2.pdf.

A need for common strategies/standardization has been identified and further work is ongoing.

6. Joint developments between the facilities

Work is starting for joint development work on a resource brokering portal for federated FIRE testbeds. Many of the projects with multiple testbeds already have mechanisms (even if manual) for setting up experiments across multiple sites. However, the intention is to have a scheme that would allow experiments to be set up automatically across testbeds from different projects.

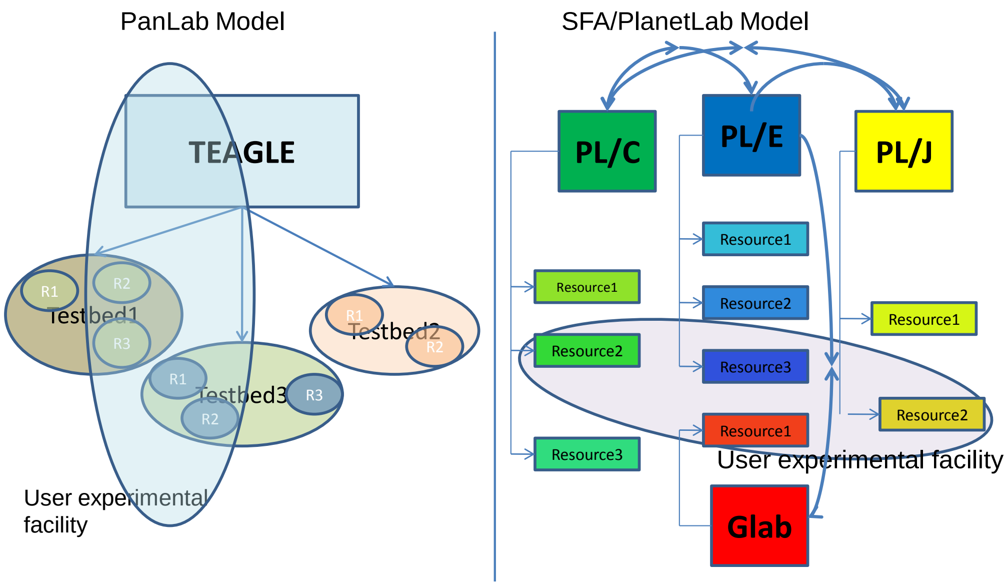

Figure 14 – Teagle approach (centralized) vs. SFA (recursive, hierarchical)

In Call 2 the main two projects OneLab2 and PII were using different approaches for managing of resources; one with more central control as with Teagle and one more distributed as implemented in the approach of SFA. In Call 5 new technologies as Cloud computing, sensors in Smart Cities, large scale distributed software, cognitive radio and OpenFlow were incorporated. The standards in the Cloud community indicate the use of OCCI (Open Cloud Computing Interface) which is neither strongly related to SFA nor Teagle. The federation of the cognitive radio experimental testbeds in CREW will initially also not benefit from these tools in order to support experiments and the sharing of experimental data.

These kinds of issues shows that the picture is from Call 5 and even more so from Call 7 projects (social networks and large scale user experiments) has become more complicated. A practical approach is being taken in order to find commonalities and joint technologies to be used by several projects. However, it is important to try to adhere to internationally accepted standards and to proceed to develop methods and tools that either enhance existing standards or find new methods through international collaboration.

The projects BonFIRE, TEFIS, SmartSantander and CREW are discussing together to expose various combinations of the elementary services in Figure 15 to the external user, though the requirements are different in each case. For example, the portal in SmartSantander is likely to be a simple tool, whereas that of BonFIRE will inherit Teagle and is expected to become an advanced tool covering most of the areas as described in Figure 15. In a "cloud-like" testbed (e.g. BonFIRE) login, service and resource deployment, and running of the experiment may be treated as a single transaction. Its portal will consist of all the boxes in

Figure 15 with the possible exception of the lowest layer, for controlling experiments. In a distributed computing testbed, such as PlanetLab Europe, it is more common to authorize access for an extended period, months or years, during which many experiments can be developed and conducted, each with a different set of resources. The portal in this approach addresses the federation issues in the top bar (authentication), resource deployment and the particular solutions to Terms and Conditions and Security (left columns).

The Terms and Conditions of usage are still complex issues. The testbeds of PanLab and Vital++ were proprietary; customers would not want their results shared nor did they intend scientific publication. In OneLab/PlanetLab, the opposite norms apply, but users still are subject to appropriate use guidelines which are contractually enforced. Across the range of possible use scenarios, we expect to have restrictions on how to use these facilities when we want to support IPR protection and privacy of certain data.

User support will grow in importance as outside usage of the FIRE facility increases. This is an area in which the FIRE STATION support action can contribute through sharing of approaches to user support and best practices.

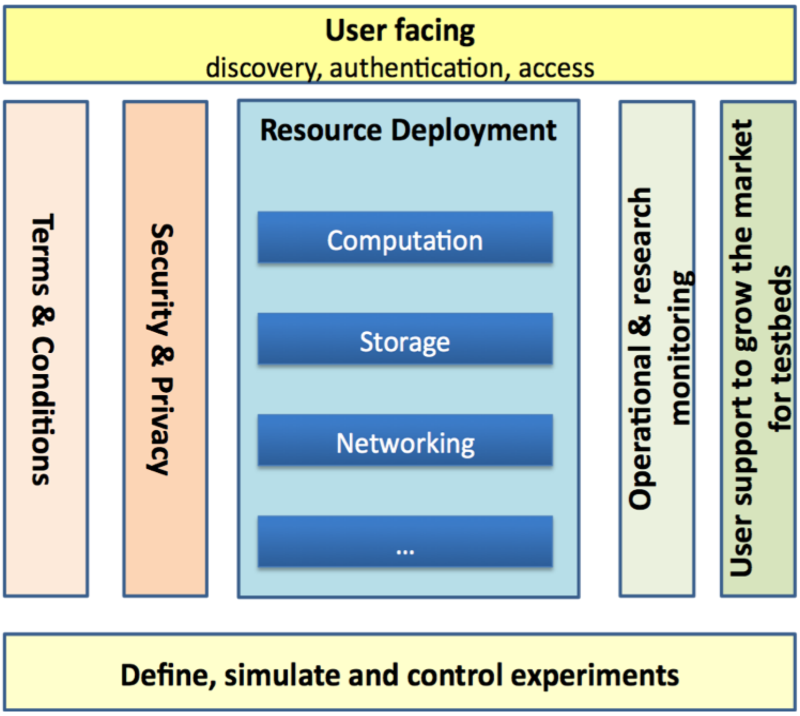

Figure 15 – FIRE experimenter requirements on facilities

Experimental monitoring and experiment definition have an important side effect. If these steps are done properly, the data from one experiment on one testbed has a greater chance of being compared effectively with data from other testbeds. This sort of data interconnection (a form of federation, in fact) is still missing in work within the FIRE portfolio and this field of research more generally, yet may be very important. However, works like OMF (a control plane for wireless testbeds used in Orbit and OneLab2) attempts to provide some form of unification of experimental data description (even to the point of describing experiment settings in a unified way). However, such work is not receiving the attention that it deserves, despite its crucial importance for truly sharing knowledge in any experimental facility.

The purpose of the joint development work is to avoid that all testbeds investigate all features and rather take advantage of the work which is done in other testbeds and/or in common so that all features are available in their testbed. These functions are required for a successful FIRE federated facility.

Experimentation topics

Experimental facilities have different major goals, meeting the needs of diverse and sometimes non-overlapping sets of users. The diagram below (reproduced from the FIRE Website shows one way in which the projects can be classified:

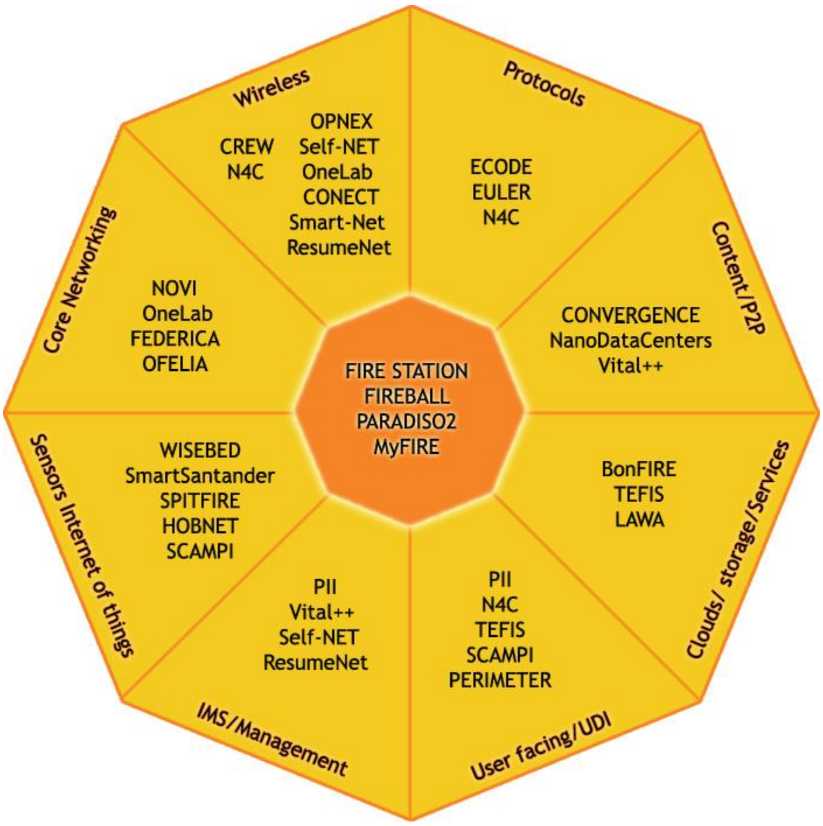

Figure 16 – Domains of experimental facilities

Experiments covering more than one area require collaboration between testbeds belonging to the same facility or involve multiple facility owners for the implementation and also for the support of the experiment itself.

The federation of some facilities is easy to implement, but the control of all aspects is hidden from the experimenter and thus it will be difficult to compare results and judge if it is the design of the federated testbed or the experiment that is the major contributor to the result found. Important aspects in some research as, e.g. interference between radios might have a strong influence on an experiment and if this cannot be controlled or understood this might invalidate experiments.

None of the projects in Call 2 and in Call 5 occupies more than 2 segments in Figure 16. This could be considered a limitation in terms of how FIRE can attract and support experiments that will make a fundamental difference in the future Internet.

Experiments cannot be forced to use the facilities available if they don't fit their needs but facilities are hard to design if experiments' needs are not clearly expressed by the community, either (a vicious circle that needs to be broken!). More effort on the understanding of experiment repeatability and benchmarking of results is required. Also there has to be expert evaluations of the possibility, cost and validity of proposed experiments.

Early indications are that the Open Calls will enhance the facilities with new capabilities and introduce experiments in new fields, while also moving FIRE closer to a federated testbed facility.

7. Research and Experimentation projects

In this section, we focus on analyzing the various FIRE STREPs also called *FIRE research projects*. In line with the previous FIREworks Portoflio Analysis [1], we distinguish between two main groups of STREPs with different levels of maturity (at the time of this analysis):

- The Call 2 STREPs, which are either just finished or about to finish.
- The Call 5 STREPs, which are basically in the second half of their first year.

In this perspective, while for the Call 2 projects more advanced and extensive results are (shall be) available, for what concerns the more recent projects the analysis is rather focusing on their plans, initial work, but most of all foreseen steps.

The main aim of this analysis is to identify:

- Which research areas the FIRE STREPs have covered or will cover.
- How and to what extent they made or will make use of the FIRE facilities.
- What are the core issues that have been identified so far, in particular with respect to availability and usability of the FIRE facilities.
- If and how end users have been or are planned to be involved.

This work shall enable to formulate recommendations for ongoing STREP projects (namely the Call 5 ones) to better tailor their experimental efforts, but also for ongoing FIRE IPs in order to make sure that their facilities can be better accessed and more effectively deployed and tested by a broader audience.

Finally, this work is also expected to provide precious feedback to the European Commission to more effectively steer R&D efforts in the FIRE domain, including refined and better targeted formulation of upcoming FIRE (or FIRE related) Call(s).

7.1. The Call 2 STREP Projects

The information reported in the following about the various STREPs projects has been elaborated on the basis of:

- Information already collected and presented in the FIREworks Portfolio Analysis [1].
- Publicly available web pages, technical reports, project deliverables and recent presentations, like in particular the ones given at the FIRE research workshop recently held in Budapest[1].
- Specific answers provided by the projects' consortia to a set of specific questions that have been elaborated for each specific project by the FIRE STATION consortium.

The ECODE Project -http://www.ecode-project.eu/

ECODE, *Experimental COgnitive Distributed Engine*, is a 3-year STREP project (terminating at the end of August 2011) that aims at developing, implementing, and validating experimentally a *cognitive routing system* that can meet the challenges experienced by the Internet in terms of manageability and security, availability and accountability, as well as routing system scalability and quality. By combining both networking and machine learning research fields, the resulting cognitive routing system aims at revisiting the capabilities of the Internet networking layer so as to address such major challenges.

Research focus:

Machine learning mechanisms and their application to the control of communication processes, including routing, forwarding, and monitoring.

The project has investigated and elaborated semi-supervised, on line, and distributed machine learning techniques as the kernel of the cognitive routing system. During the building phase (first project period), the cognitive routing system has been designed and initially prototyped. In particular, some ECODE demonstrations were given at the FIRE/ServiceWave 2010 event during the Future Internet Conference Week last December 2010.

In the second and current project phase, several Use Cases are being experimented to evaluate the benefits of the developed machine learning techniques.

Deployment of FIRE Facilities:

Although in the project pages collaboration with OneLab2/PlanetLab is claimed, experiments have been conducted at iLab.t experimental facility, located at IBBT in Ghent, Belgium.

The main reason of not having used PlanetLab is, according to the ECODE consortium, that better fit with the ECODE needs and purposes was found in the iLab.t experimental facility. In particular, as PlanetLab it is an emulated open/shared facility, using it would have required emulating a virtual router on top of PlanetLab nodes, instead of running directly the ECODE learning module on Linux router blades taken from the iLab.t facility. Moreover, the deployment of the PlanetLab facility would not have allowed full control of experimental

[1] More information about the FIRE research workshop aims and agenda can be found at: http://wiki.ict-fire.eu/index.php/Research_Workshop

parameters to produce repeatable results, which is a sine qua non conditions for model validation that ECODE aims at. In addition, the number of nodes available on iLab.t platform was largely sufficient for conducting the ECODE experiments.

Experimental work:

The selected experimental scenarios have been defined and presented (i.e. deliverables D4.1 Experimental scenarios including evaluation criteria and methodology, D3.3 Experimental Evaluation of TO1 and D3.5 Experimental Evaluation of TO2) and some concrete and recent results have been discussed and presented in Budapest at the FIRE workshop. In particular, they presented the two core problems that have been studied so far:

- *Packet loss during fast re-routing*:
  - An improvement of up to 20% has been calculated from simulations
- *Shared Risk Link Group* (predicting failures in back-up links from the first failure). Experiments have been made using Emulab and IBBT's iLab.t with 100 nodes:
  - For failures of SRLG containing at least one adjacent link, the recovery time was almost equal to the local failure detection time (~100ms), instead of 1s in standard OSPF.
  - For failures of SRLG containing only non-adjacent links, once the first failing link was known, all paths can be recovered at the same time by the "inferring" node.

The main lessons learned from the packet loss reduction during fast re-routing experiments can be summarized as follows:

- No traces were available in Europe from experimental facilities.
- An experimental validation needs a formal model or a simulation – to test against.
- Moving directly to an experimental facility would have been too time-consuming.

Regarding the Shared Risk Link Group, experiments performed on the IBBT iLab.t platform enabled (i) to control with a fine granularity the routing protocol and machine-learning operations according to the network running conditions, (ii) to produce repeatable and reliable results, and (iii) to verify the scaling properties (so far, only 3% of the number of available of machines have been used).

Unfortunately, three deliverables which should give more details on this part of the work, namely D2.3 Cognitive Engine low-level network and system specification, D4.3 Experimental evaluation of the cognitive engine, and DF.2 Federating Requirements are not yet available, although expected to come up by beginning of September 2011.

Collaborations:

- The ECODE consortium organized a joint workshop with RESUME-NET in order to exchange results for what concerns possible cooperation on network self-healing/self-repair and self-protection. This joint workshop was held on Feb.4th 2011 at ULg which hosted a RESUME-NET and an ECODE project consortium meeting during the same week. The goal of this workshop was to exchange on two common topics the two projects are investigating, namely security and resiliency, to assess possibility of cooperation on either of these. More details are planned to be made available on the public ECODE Website.

- The ECODE consortium invited the Self-NET project participants to participate to the IRTF *Learning-Capable Communication Networks* (LCCN) activity[2]. Even though Self-NET is mostly dedicated to self-management of wireless networks and ECODE focuses on self-adaptive control of core networks, similar learning capabilities and mechanisms could indeed be applied to different control paradigms.

The N4C Project - http://www.n4c.eu/

The N4C, Networking for Communications Challenged Communities, project is a 3-years STREP that ended in April 2011. By recognizing that remote and otherwise communications challenged areas often have specific infrastructural constraints and other obstacles for communications, the N4C main aim was to extend Internet access to people, businesses and authorities operating in remote locations. For this purpose N4C has focused on challenging scenarios with populations and local industry in remote and topographically complex areas, exploring models for deployment and business in the context of current debates on knowledge society inclusion.

Research focus:
Internet access for all by evolving Delay and Disruption Tolerant Networking (DTN) technology combined with wireless technology, methods for power harvesting and power management, off-the-shelf hardware.

The central idea of N4C has been to drive the evolving Delay- and Disruption-Tolerant Networking technology (DTN) towards practical usage. The strategy has been to combine this emerging future internet protocol set with mobile computers used as data mules and off-the-shelf wireless technology to provide carriers, including upgrading of WiMAX solutions for challenging scenarios. As lead concept, nomadic solutions were the aim. At the peripheries data transport would be dependent on data mules but it would be robust and independent of conventional infrastructure. The idea was to deliver technology that does not require heavy investments in infrastructure but, can operate opportunistically, with smaller, movable masts, portable data mules and, generally light and portable equipment.

DTN makes the exchange of data between the source and destination tolerant of time delays. The core idea is that in case there is not a continuous path from data source to destination, the data can be held at some intermediate nodes (which can physically be any computer or server) along the path until there is an appropriate one available along which data can be forwarded with some delay. Most importantly, and unlike the conventional Internet, a DTN 'network' does not expect its connections to be fixed in time. Instead it can cope with mobile nodes that are exposes to highly changing connections to other nodes. People, cars, helicopters and similar entities moving “data mules” that travel in the remote

[2] LCCN is the name of the research group that the ECODE consortium tries to launch at the Internet Research Task Force (IRTF), the research arm of the IETF, more details can be found at: http://www.ietf.org/mail-archive/web/irtf-discuss/current/msg00012.html
and http://tools.ietf.org/html/draft-tavernier-irtf-lccn-problem-statement-01

area actually take the data with them and deliver it to the next node in the network "cloud". Movable data mules together with nodes at permanent locations form the so-called DTN cloud. The connections between nodes are "opportunistic": nodes communicate when they meet up; data is exchanged if it appears that this will bring the data closer to delivery to its intended destination.

Contributions were made to nine areas of experimental IRTF3 standards including the PRoPHET Routing Protocol, the Endpoint Discovery Protocol for DTN, the Bundle Security Protocol, and more. The project has included enhancement of WiMAX solutions for challenging scenarios.

Experimental work:

The N4C project has experimentally validated the developments of appropriate mechanisms and technologies in real-world testbeds. The testbeds were facilities in areas (namely Swedish Lapland and Slovenia) that were likely venues for "future DTN for remote regions" on a small scale, but operating under real conditions in scenarios that approximate ordinary life in such areas.

In particular, the N4C-based DTN access solutions and applications have proven stable and run autonomously for several months in real situations.

- Environmental measurement nodes using DTN for transfer of data in critical situations for the nuclear industry and, results in the WiMAX area are already available on the market. DTN testbeds in remote locations were built for the project, and as a sustainable result the Slovenian Kočevje testbed has been suggested for the FIRE federation.
- Communication technologies we use daily are dependant on electrical power. In an urban European setting this is not an issue. Setting up networks beyond the conventional infrastructure and electrical power grid however, implies limited or no access to this commodity. To handle such situations in an efficient and environmentally friendly way, N4C developed strategies for power harvesting, and made power management intrinsic in software as well as hardware solutions and testbeds have been run in Staloluokta, Swedish Lapland.

Deployment of FIRE Facilities:

The N4C project built its own facilities as the FIRE available ones, like for instance the NITOS wireless testbed in OneLab2, could not effectively serve the purpose of reaching the remote identified non-urban areas they have targeted by deploying the underlying DTN mechanisms.

Towards the end of the project, other facilities (external to FIRE) have also been used; in particular, a Spanish setting for the near-to-commercial deployment of new WiMAX techniques, and the Galway facility in Ireland testing the combination of WiMAX and DTN – more information on this is given in N4C Deliverable 6.5 - Update of previous work based on field experiences.

[3] Internet Research Task Force, IRTF, http://irtf.org/

The N4C experimental facility, which the N4C consortium defined and realized, namely the Slovenian Kočevje testbed run by N4C partner MEIS[4], has been offered for federation to the FIRE Facilities based on the OneLab2 strategy. This testbed offers pretty interesting features. As real traces of DTN traffic are generally quite rare to obtain and study, MEIS, upon request, can take care of creating a dedicated testing infrastructure – namely moving data mules – as a concrete facility for remote transfer of DTN data. This can be useful both for testing near-to-commercial stage products and/or for testing routing, naming and other schemes, and protocols in a more R&D-driven perspective.

For the testbed to be offered to users in a broader FIRE context an agreement should be put in place with MEIS. More details on this and on the testbed can be found in the publicly accessible N4C deliverable "Integration Plan for N4C Test Beds with other Future Internet Test Beds, Including offer to Federation by Slovenian Test Bed"[5].

Other forms of relevant collaborations include:

- Several international R&D connections have been established through the experimental standards work in the DTN Research Group in IRTF.
- The N4C partner Trinity College Dublin continues working on similar issues in the context of the FP7 project SAIL[6], which is investigating scalable and adaptive Internet solution at a wide EU-scale, where substantial input is transferred from N4C.
- Direct link with the former ICT HAGGLE project[7] through common roots in a past nationally funded Swedish project named "Sami Network Connectivity" (2004-2006).

The Self-NET Project – http://www.ict-selfnet.eu

The Self-NET project, which recently terminated, aimed at exploiting the autonomic capabilities and cognitive features in networks, according to the feedback control cycle: Monitoring, Decision-making, Execution (MDE).

The core concept is based on Future Internet cognitive managers, located on various network devices.

Research focus:

The main areas that Self-NET project has covered are the following:

- Empowerment of cognition into network nodes for network Management:
    - Knowledge modelling (Information Model. Ontology)
    - Situation awareness
    - Distributed decision making
    - Learning techniques

[4] Contact to MEIS: Dr. Marija Zlata Boznar (Director), *Phone:* +386 1 3663226 http://www.meis.si/en/

[5] http://www.n4c.eu/Download/n4c-wp2-043-D2-3-Offer-to-federation-10.pdf

[6] http://www.sail-project.eu/

[7] http://www.haggleproject.org/

- Break-through holistic architectural framework for cognitive self-Management.
- Modularization of network stack through dynamic re-composition of functional modules.
- Coverage and Capacity Optimization in Wireless Network Environment.
- Self-configuration and routing adaptation for wireless mesh networks.
- Probabilistic reasoning by using the real network management data have been collected and provided by Vodafone MNO in order to experiment on the probabilistic deduction and analysis.
- Experimentation (small scale trial, simulations) for feasibility study and performance assessment of network self-management.

Experimental work:

Self-NET used the PanLab-PII facilities and specifically the OCTOPUS testbed. The goal of this collaboration was the measurement as well as evaluation of the performance improvement (e.g. QoS features, PER, jitter, delay) that the Self-NET self-management framework could provide in an end-to-end scenario (user-side and service domain, core network and wireless domain) having the capability to use a broadband wireless technology i.e. WIMAX. The availability of the WIMAX Base station and the WiMAX clients as well as the management interfaces that OCTOPUS testbed could provide for the remote monitoring and configuration of the WIMAX Base station are important reasons for choosing OCTOPUS testbed.

The experiments included service adaptation and network reconfiguration based on multi-objective optimization (QoS, packet loss, fault, interference, etc.). The OCTOPUS testbed allowed to change the prioritization of a certain number of flows and to move flows onto different VoIP codecs. It was found that changing the prioritization at the WiMAX Base Station and changing the VoIP codec between the service provider and the end user, reduced the number of dropped packets.

A video of the demonstration is available online at: http://scan.di.uoa.gr/prototypes/self-net-and-panlab-demo-12-2010

Issues with the use of the FIRE facilities was discussed and presented in Budapest:

- Facilities should have as many capabilities as possible to reconfigure their components.
- Facilities should have appropriate interfaces for interacting with the experimental resources.
- There is an overhead in terms of effort from both sides for the experimentation and Use Case deployment. This should be minimized as much as possible.
- Currently there is no testbed that specifically supports experiments on Autonomic Communications.

The experience matured through the Self-NET project also highlighted that it was difficult to understand the timescale of the FIRE facilities; ideally, they should be supported for a longer time.

On the other side, the added value of using a FIRE facility can be summarized as follows:

- To have a greater diversity of technologies and infrastructures.

- To make experiments on a larger scale.
- To obtain more advanced experimentation results, based on multiple metrics.
- To better justify the research results, since the experiments are performed in "close-to-real-life" conditions.
- Researchers can gain technical know-how about an increased number of technologies and equipment.

Collaborations:

Apart from the collaboration with the Panlab-PII FIRE IP project (see experimental work sub-section above), Self-NET did not directly collaborate with any other FIRE STREP project, mainly due to the diversity of objectives and underlying technologies, with the exception of ECODE, which is also exploring learning techniques for network management.

Self-NET and ECODE have participated several times to common FIRE events, and the Self-NET partners stated it would be interesting to set-up a dedicated inter-project collaborative session (e.g. with a duration of 3-4 hours), in the context of any upcoming FIA or FIRE related event, so as to facilitate exchange of ideas, specific problems/solutions discussions as well as preparation of common publications.

The RESUME-NET project - http://resumenet.eu/

The RESUME-NET, *Resilience and Survivability for Future Networking: Framework, Mechanisms, and Experimental Evaluation*, project (that will terminate at the end of August 2011) focused on achieving systemic and systematic resilient connectivity in Future Internet environments at both the network and service levels.

The consortium developed a so called resilience framework enabling to build resilient networks and services "by design". This framework relies upon the idea of defining *at modelling time* what are the main challenges to control the network (and thereby service) behaviour and make use of policies to define detection and remediation strategies through the specification and management of specific resilience metrics. Based on such a framework, the consortium focused on identifying and defining mechanisms and algorithms to enact and ensure resilience *at run time*. *Network resilience* relies upon core mechanisms such as redundancy of resources, diversity in routing, and transport incentives for collaboration and challenge detection techniques. *Service Resilience* is enabled by creating "supervisors" operating in a peer-to-peer (P2P) modality and coordinating with each other in order to minimize the probability of failures in combination with "virtualization" techniques. Virtual Machines can be used to guarantee connectivity and thereby service continuity.

Research focus:

The RESUME-NET project is investigating a framework, architectures and mechanisms for network resilience, with specific attention to resilience metrics, multi-level information sharing approaches, policy-based network management for resilience, and various defense mechanisms, such as multi-path routing structures. These aspects are evaluated using different experimental case studies: challenges, including selfish node behaviour, in opportunistic and wireless mesh networks; resilient Voice over IP services; and ensuring the resilience of a publish-subscribe platform that supports an Internet of Things deployment.

Experimental work:

The RESUME-NET approach has been validated by experiments run on different testbed platforms in which Virtual Machine migration has involved different sites in OneLab2 (PlanetLab) and G-Lab (German facility).

- **PlanetLab Europe** was chosen because it enables to run the RESUME-NET overlay and P2P experiments, which require to scale to the order of magnitude of the hosts on PlanetLab, and greater. PlanetLab Europe is open for members of an institution that contribute to the PlanetLab network. However, as restricted access was provided to the users, not all kind of experiments, which could be run on a local testbed, were possible. For example, experiments with Virtual Machine migration were not possible. Also, users running experiments on the same machine could not run services on the same port number. For example, if someone else runs a web server on port 80, then it is not possible to use port 80 for any other purpose. Thus, the concept of the "slice" in PlanetLab does not offer sufficient isolation, as virtualization techniques would allow for.
- **The G-Lab testbed** was used because it provides full control over the employed virtualization solution, allowing us to develop an individual Virtual Machine management. G-Lab access is provided only to G-Lab partners – a total of 29 partners coming from academia, research institutes, SMEs, and large companies. Therefore, it is not open for the public. As a matter of fact fine grained control of virtualization mechanisms requires exclusive control of the underlying hardware. Since several parties use the G-Lab network concurrently, this requires exclusive reservation of hardware well ahead of time. This situation makes development of a prototype difficult, as implementers can only triage bugs and crashes in predefined time windows.

The experience matured through experimentation highlighted the fact that a critical part of deploying FIRE facilities concerns availability of information about data traffic (as discussed at the FIRE workshop in Budapest). In particular, even though BGP information is available, there are few packet level measurements. Also, in terms of packet traces, it is important to have the right data (i.e. traffic data from PlanetLab may not be consistent with that from other environments). It is also necessary to know where the data was measured, with what tool and the timeframe (minutes, days, etc). This should be considered for future work.

In addition, to collect the results of our experiments on PlanetLab was quite a time consuming task. Therefore, the RESUME-NET consortium expressed a recommendation: to have a "dropbox-like" service on PlanetLab that would allow collecting the results as easy as on a local testbed.

Collaboration:

A part from the collaboration with OneLab2 for the deployment of the PlanetLab facility, interactions of the RESUME-NET project with other FIRE STREPS have largely focused on sharing knowledge and results. For example, a half-day common meeting was held in Liège (Belgium) with the ECODE project (www.ecode-project.eu), during which partners from each consortium presented their respective work.

The PERIMETER project - http://www.ict-perimeter.eu/

The PERIMETER project is a three years STREP that terminated at the end of April 2011. The main objective of the project was to establish a new paradigm of user-centric strategies for advanced context centric networking by focusing on Quality of Experience (QoE) retrieval and measurement. QoE is defined in terms of how a user perceives the usability of a service when it is in use, i.e. degree of satisfaction with a service.

This concept embraces and goes beyond Quality of Service (QoS) that is defined as the ability of the network to provide a service with an assured service level. As a matter of fact, perfect QoS may not guarantee a satisfied user. In this perspective, PERIMETER aimed at defining a user-centered networking paradigm where users are Always Best Connected (ABC). This paradigm is based upon the following core principles/mechanisms:

- *Data collection and adaptation* relying upon the definition, monitoring and tuning of QoE metrics relating to the different available network technologies and based on network performance, user's context information and feedback, etc.
- *Data processing* that elaborates the data collected and adapted through the previous step.
- *Awareness and decision.*

Research focus:

- "User-Centric Seamless Mobility", namely to let the user control the QoE and service cost, using:
    - User centric network selection based on QoE and user preferences and high-level rules or policies instead of naming individual technology and technical methods.
    - Mobility based on fast authentication, authorization and accounting (AAA) that utilize privacy-preserving cryptographic digital identity, trust and reputation frameworks.
    - Security and privacy-protection regardless of the type of network and service the user is currently engaged on.
- Advanced "Distributed A3M" (authentication, authorization, accounting and mobility) protocols suite for the "user-centric seamless mobility" paradigm, designed to cope with increased scale, complexity, mobility and requirements for security, resilience and transparency of the Future Internet.
- Experimental tests, validations and evaluations of the technological, behavioural and economic aspects and usability, privacy, security, interoperability and performance

parameters of the "user-centric seamless mobility" paradigm on interconnected large-scale testbeds with real users.

- Novel distributed and reconfigurable protocol and service architectures, and the advanced overlay security, trust and identity management architectures and technologies of PERIMETER with these testbeds.

Experimental work:

The PERIMETER middleware components and its integrated applications and services were tested in two large-scale interconnected testbeds from WIT[8] (Ireland) and TUB[9] (Germany) by enforcing the Living Labs approach, which allows actual users to directly provide valuable information. When the project was planning its experimental work, none of the existing FIRE facilities was suitable for testing the diverse needs of the PERIMETER middleware on a range of heterogeneous devices.

Different experimental scenarios have been identified, namely user-centric, agnostic ubiquitous communication, emergency situation/health care, community networks, and more details on concrete testbed setup and FIRE facilities deployments and/or extensions can be found in the deliverables produced by the PERIMETER work package 6 "Proof of Concept - Integration, Validation and Testing", in particular in Deliverable 6.4 "Report on Phase 2 Integration and Testbed Testing".

The PERIMETER testbed partners have agreed to leave the testbed infrastructure (including the secure IPSEC tunnel linking the two sites, namely WIT and TUB) in place for the foreseeable future for any interested parties to make use of ensuring its future sustainability.

In addition to this, PERIMETER also made use of the FEDERICA facilities (namely 5 virtual nodes connected via the WIT and TUB access points) to extend its federation and test PERIMETER on a larger infrastructure.

About end users' involvement:

The PERIMETER Living Lab study involved 30 users who completed a task list that allowed the assessment of the projects results. What emerged is that it is particularly hard to conduct a living lab study within a limited time-frame.

A more significant involvement of end users could be realized with a field test. End users could be equipped with a PERIMETER enabled smart phone that they could use over a time-frame of several weeks to experience PERIMETER under real life conditions.

If one could provide the users with the necessary tools to directly affect the behaviour of the system (e.g. by tuning some user-intuitive parameter settings provided to them) and observe the performance results over the testing period, that would give more confidence on the results of the study, by making the users a more aware and active part of the innovation cycle. However, the side development of this kind of tools was out of the scope of PERIMETER.

[8] http://www.wit.ie/
[9] http://www.tu-berlin.de/

Main challenges and future work:

- PERIMETER contributions to building of new and flexible business models.
- New modelling approaches and practices including worldwide stakeholders.
- The move to user-centric methodology that solves the seamless mobility needs to create a market for new devices, applications, and services that embed the user-centric seamless communication capabilities as developed by PERIMETER.
- Adaptation of the European legislation for fostering the application of user-centric approaches. This to allow the industry and the Telco market to evolve towards a different perspective in which users can more freely choose depending on their needs and operators can operate by relying upon more flexible business models.

Collaborations:

Collaboration with the FEDERICA consortium for the deployment of their facilities has taken place.

The Smart-Net project – https://www.ict-smartnet.eu/

The Smart-Net project was a 3-year STREP, which ended in May 2011. It had the goal of defining a novel architecture for wireless access networks and analyzing advanced protocols and procedures for this new paradigm.

Research focus:

Smart antennas, traditionally deployed in base stations, are provided for use with each subscriber unit. The ability to combine efficiently these new capabilities with mesh networking allows the system to provide scalability, reliability and security. This approach is strengthened with advanced routing and scheduling protocols specifically designed of heterogeneous Wireless Mesh Networks.

The SMART-Net architecture includes the needed functionalities to enable efficient cross-layer operations and advanced security. In particular, the project developed an attack-resilient security framework that takes benefit from the capabilities of smart antennas.

Experimental work:

The SMART-Net consortium built its own facility that is composed of two parts:

- Real Life Testbed composed of smart nodes of two types
    - Multi radio nodes (WiMAX and WiFi) equipped with Smart antennas.
    - WiMAX relay stations.
- OPNET Simulation Testbed integrating WiMAX, WiFi and UWB simulation model taking advantage of smart antennas.

These two testbeds are interconnected together through Internet connection.

These facilities allow to perform measurements based on real life equipments and to extend the concept in large-scale simulation. The real-life testbed is composed of 5 fixed nodes and a nomadic node. It enables the development and evaluation of mechanisms that manage smart antennas with better network resiliency and that take advantage of smart antennas combined with WiMAX technology. Studies on energy efficiency have also been performed and some mobility management mechanisms defined.

The SMART-Net testbed Internet access is open to external researchers until at least May 2012 *(send an e-mail to: bruno.selva@thalesgroup.com)*.

Note that the SMART-Net testbed would ideally benefit from being extended with connectivity to other radio access technologies, such as LTE.

The VITAL++ project – http://www.ict-vitalpp.upatras.gr/

The VITAL++ project is a 3-years STREP that ended in April 2011 and that besides building an experimental facility based on P2P and IMS technologies, developed new P2P algorithms for live streaming and monitoring as well as improved QoS-based service offering through resource optimization.

Research focus: VITAL++ covered the areas of P2P content distribution in a Telecommunications environment based on the IMS (IP Multimedia Subsystem) technology. VITAL++ has designed and built an integrated communications environment that aimed at combining the best of these two worlds.

Experimental work:

VITAL++ mainly used partner facilities and created its own experimental platform, which combines P2P and IMS technologies[10] and that is realized by interconnecting a number of geographically distributed test sites that are maintained by the project partners. The test sites range from single machines that can have access to the integrated platform to networks with complete IMS installations.

The test-sites are interconnected over a VPN tunnel that is established via an OpenVPN Server operated at TID premises. There is an accompanying DNS server at TID that aids the proper network address resolution. The IMS equipped test-sites are configured to treat users in roaming mode. The home network of all the users is represented by the Fokus OpenIMS Core installation.

Scattered across the test-sites the Vital++ Application Server subcomponents realize the operation concepts of the platform. All are accessible via SIP Instant Messages without requiring any modifications of existing protocols and methods.

Instances of the Vital++ Client assemblies are launched in every test-site so as to allow for evaluation and validation of the designed and implemented methodology.

One of the primary purposes of their experiments has been towards developing a proof of concept of the VITAL++ architecture, its innovative functionality and the corresponding algorithms. Another, objective of the experiments has been to test the scalability of the

[10] For an overview see http://www.ict-vitalpp.upatras.gr/infrastructure.html, while for details about the VITAL++ facilities see D5.1 that can be downloaded from http://www.ict-vitalpp.upatras.gr/pdf/deliverables/VITAL++_D5.1_final.pdf

overall VITAL++ approach[11].

The VITAL++ project ported its facilities and functionality to Panlab. In particular, access to VITAL++ is enabled via the Teagle tool available in the Panlab portal.

About end users' involvement:

Involving users in experimentation in VITAL++ has been necessary for two main reasons: a) versatility of the experiments introduced by various users' behaviour and user's equipment capabilities environment combined with b) large scale experimentation that wanted to be close to real-world scenarios.

Motivating a significant number of regular users in this kind of experiments is difficult as the test environment must offer a number of features like usability, stability, and incentives. However, building these features is not an easy task as it requires extensive testing and debugging of the environment to support users with different needs, profiles and interests.

VITAL++ wanted to verify their simulation results of new P2P algorithms through large scale experimentation scenarios involving real user behaviour and heterogeneous facilities. This was not implementable in practice using PlanetLab or the IBBT Virtual Wall. Using Panlab was much easier for them – mainly due to their familiarity with the framework – though the testing had to be limited to a moderate number of clients.

Main challenges and future work:

In this perspective, the VITAL++ partners believe that existing and/or future FIRE facilities need to be transformed into a more stable and stimulating environment wherein users will be attracted to use interesting applications and services.

Similarly such environments need to be build as plug and play systems wherein components representing new or updated functionality can be easily integrated leaving the rest of the experimental environment intact and operational. In this way, it is expected that specific FIRE facilities will eventually take the form of a production like system customized for seamless experimentation[12].

Collaborations:

With the exception of Panlab as mentioned above, no collaboration with other FIRE STREPs or IPs took place.

The NanoDataCenters project – http://www.nanodatacenters.eu

The Nanodatacenters project, which ended in April 2011, aimed at combining the power of data centres with the scalability of P2P methods, while not threatening the robustness and the stability of the Internet. The core idea behind the Nanodatacenters approach was to enable the next generation of interactive services and applications to flourish, complementing existing data centres and reaching a massive number of users in a much more efficient manner, by leaving the network untouched and locating the content on customer premises. In particular, the project concerned managed peer-to-peer services

[11] More details about the experimental work and its main results can be found at http://www.ict-vitalpp.upatras.gr/pdf/deliverables/VITAL++_D5.3_final.pdf

[12] For more details see Deliverable 6.6, Contribution to Fireworks.

(e.g. gaming), in which (for the example given) the game engine and all of the players are connected via regular Internet links. Nanodatacenters monitoring and measurement tools take care of optimizing storage and paths and circumventing failures. Of all the IP facilities, the OneLab2/PlanetLab Europe environments matched best this kind of situation that has been used to experimentally validate the Nanodatacenters approach.

Research focus:

- Providing a new distributed hosting edge infrastructure that supports interactive applications.
- Integrating a large number of geographically dispersed nano datacenters to complement existing datacenters.
- Exploiting the use of under-utilized resources existing at the edge.
- Supplying a new secure P2P paradigm.
- Supporting system scalability.
- Optimizing the performance from the user's viewpoint.
- Incentive mechanisms for those who provide access to their edge resources.
- Advanced monitoring capabilities.
- Providing service guarantees over the regular internet.

Experimental work:

Experiments were made with 4 applications: P2P VoD, on-line gaming, backup service and P2P video streaming. In addition, experiments were made with network services such as monitoring, measurement and caching. Some users, and also the sources of some of the applications (e.g. P2P video streaming), were located at remote nodes attached to the OneLab2 / Planet Europe network.

The testbed comprised 12 Gateways, running on Intel Atom platforms. Each Gateway had 1 GByte of RAM and multiple network interfaces: one for the connection with the Wire Area Network, one for the in-home connection (1 Gbit/s Ethernet interface), as well as a wireless interface. Each of these Gateways was connected to a real DSLAM using the aforementioned WAN port. As part of the Nanodatacenters infrastructure, there were also two network proxies (caches) that were used to speed up the connectivity between the Gateways and the core network infrastructure.

The backend infrastructure hosted the Nanodatacenters management systems. On this backend server, OMF (with own extensions) was operated. This server was also the repository for the application and service images that were deployed on the Gateways when requested.

In addition to the Nanodatacenters backend, each service provider had to provide its own service platform. In the testbed, the multiple services used in the final demonstration were all hosted on the same server.

Connected to the physical network infrastructure was a Linux router capable of emulating a more complex network infrastructure, i.e. able to add packet losses, delay and jitter to the packet forwarding process.

Main challenges and future work:

The Nanodatacenters consortium reported to have faced the usual issues related to integrating pieces of software for the platform from different partners, which turned out to be more time-consuming than planned, though the integration of the applications from various partners was straightforward. Federating Nanodatacenters with OneLab2/PlanetLab Europe also did not cause problems, but modifications had to be made to the commercial streaming application ("Kangaroo") to comply with Nanodatacenters' P2P environment.

The system will now be productized, requiring the software to be re-written in accordance with Technicolor[13] practices, tools and development environment, and integrated with the rest of their Set-Top Box software.

The OPNEX project – http://www.opnex.eu

The core objectives of this 3 years STREP project can be summarized as follows:

- Rethinking and redefining the protocol stack for multi-hop wireless networks by designing advanced systems optimization and control theory-driven algorithms.
- Delivering a plan for converting the algorithms termed in abstract models to protocols and architectures that extract the full transport capacity in real dynamic multi-hop wireless environments while being amenable to decentralized low-complexity and low-overhead implementation.
- Implement the theory-driven protocols and experimentally demonstrate performance improvement in realistic wireless testbeds over currently used techniques. Two different platforms, an 802.11-based and a sensor-based one will be used to assess the appropriateness of the methods in different application domains.

Research focus:

OPNEX covered the area of distributed coordination and control of networks and in particular that of wireless networks. Due to wireless link volatility and lack of a centralized entity for control, OPNEX attacked the issue of distributed control. In particular, many theoretical covered topics led to substantial performance improvements for wireless network operation. Moreover, the benefits of OPNEX where not only theoretically investigated, but practically shown as well. Extensive algorithm and protocol development and experimentation took place in the OPNEX testbeds, covering a long range of interesting Use Cases and depicting the actual benefits in each one.

Experimental facilities:

OPNEX works on the following experimental facilities:

- DES-testbed (location: Berlin (FUB), http://www.des-testbed.net/). DES-testbed is a deployement with 110 nodes (indoors and outdoors) featuring both wireless IEEE 802.11 network cards as well as wireless sensor nodes. Several tools aiding the experiments have been developed enabling remote operation of the testbed.

[13] Technicolor was the Nanodatacenters coordinator.

- NITOS (location: Volos (CERTH), http://nitlab.inf.uth.gr/NITlab/index.php/testbed). NITOS is a testbed deployement comprised of 50 outdoors nodes, featuring IEEE 802.11 network cards as well as software defined radios. It is equipped with heterogeneous technologies such as cameras, robots for mobility, MIMO cards, energy metering devices and others. The testbed is running OMF language for experimentation and provides several tools like remote reboot, connectivity measuring tools, user scheduler with slicing, etc.
- wnPUT-testbed (location: Poznan (PUT)). This testbed targets mainly home networking scenarios, which are characterized by a limited number of nodes and a small number of hops (mainly one or two) between them. Moreover, the testbed is oriented on experiments concerning queue-level-based control protocols developed and implemented by the research group working in OPNEX at PUT.
- Technicolor wireless testbed (location: Paris (Technicolor)). The Technicolor wireless testbed currently consists of 16 dedicated wireless nodes deployed in the offices at the Technicolor premises. The testbed spans three buildings and one partly covered parking garage, and nodes are deployed across four different floors. This deployment allows for constructing a multitude of multi-hop network configurations with up to six hops. In addition to the standard off-the-shelf wireless hardware, the Technicolor testbed includes programmable WiFi cards and custom-made sectorized antennas.
- INRIA testbed (location: Paris (INRIA)). The testbed at INRIA is composed of a flat topology of 50 wireless sensor nodes, randomly scattered on a field for each new experiment.

The OPNEX testbeds allowed the development of distributed network control algorithms for wireless networks. They were strategically designed to serve this purpose as well as to allow wireless network experimentation in diverse environments and technologies. For this reason, the testbeds complement each other in a perfect way.

Both the NITOS and the DES-testbed (the two largest testbeds developed in OPNEX) are publicly available for experimentation and federated with other FIRE activities and projects (e.g. OneLab).

More information can be found in the public OPNEX deliverables D4.1, D4.2 and D4.3.

Experimental work:

OPNEX developed and deployed several algorithms and protocols spanning several layers of network operation (MAC layer approaches, scheduling, routing, flow control, etc) including cross-layer approaches. Experiments took place to show the performance of the proposed OPNEX algorithms/protocols that utilized the network queue sizes. Further experiments where run including sector antennas, MIMO systems, network node mobility, channel assignment and association, etc. Moreover, multiple experiments took place for the testbed management, connectivity tools, interference avoidance, slicing tools, power consumption measurements, to name a few. Finally, the NITOS testbed was used for the demonstration of the OPNEX approaches in a video streaming scenario.

Main challenges and future work:

The OPNEX testbeds, although having reached a level of maturity, are live and constantly developing. The OPNEX partners consider them as a strategic activity, central to their efforts of improving communications and ICT and for this reason they maintain dedicated personnel working on them. They try to enhance them with more capabilities and improve their operations. At the same time, many experiments run on these testbeds supporting several research efforts.

Collaborations:

The OPNEX project had a strong connection to the WISEBED project. The wireless sensor network of the DES-TESTBED has been used for development and experimentation in both projects. OPNEX profited by using algorithms from the WISELib, a software library for wireless sensor networks developed in WISEBED. In particular, algorithms regarding the energy consumption and routing were used in OPNEX experiments. The experiment results were also published back to the WISEBED consortium.

7.2. The Call 5 STREP Projects

As already mentioned earlier in this document, most of the Call 5 STREPs have now run for about one year. While for most of them it was possible to obtain interesting information from the answers they gave to the specific questions that were sent by the FIRE STATION consortium, most of their Websites are still quite scarcely populated.

Most of the various projects' Websites repeat some of the Description of Work key objectives, but they do not elaborate much more beyond that on their ongoing and planned work, including intended use of the FIRE facilities and planned or ongoing collaborations with other FIRE projects.

The CONECT project - http://www.conect-ict.eu/

The CONECT (COoperative NEtworking for high Capacity Transport architectures) project (that has been running over a year now) aims at proposing a new holistic approach for wireless network design that exploits the inherently broadcast property of the wireless medium towards achieving superior network performance (via constructive use of signal interference, opportunistic overhearing and node cooperation) in terms of throughput, delay and power expenditure. Special emphasis is placed on currently under-developed information transfer modes, such as one-to-many and many-to-many, which are expected to prevail in the Future Internet.

Research focus:

Information and communication theoretic principles, coupled with optimization theory techniques for open broadcast wireless medium. Signal interaction and node cooperation, including architectural design choices and protocols for end-to-end information transport. Optimization of video transfer in one-to-many and many-to-many modes.

Within CONECT, different mechanisms of signal interaction have been investigated and fundamental performance limits for wireless network performance have been derived, while distributed cooperative coding schemes that perform close to these limits have been proposed. This has also motivated the development of packet-level cooperation algorithms and the investigation of their impact on end-to-end information transfer, especially video streaming. Apart from the above theoretical contributions, two experimental testbeds have been constructed and some of the proposed principles and algorithms have been already implemented with encouraging results. Since CONECT is currently in progress, additional experimental results are expected to be produced in the future.

Experimental facilities:

Under CONECT, the following facilities have been constructed and/or upgraded from existing infrastructure:

- ***NITOS*** (http://nitlab.inf.uth.gr/NITlab). The NITOS testbed is located in Volos, Greece and maintained by the University of Thessaly, the latter being part of CERTH. It consists of 50 outdoors nodes (including Orbit, Commell and diskless nodes) equipped with commercial IEEE 802.11 network cards with open source *madwifi* drivers, as well as software defined (GNU) radios. NITOS incorporates multiple heterogeneous technologies such as cameras, temperature/humidity sensors, robots for implementing node mobility, MIMO cards, energy metering devices and others. The testbed offers remote access and control to registered members and supports the OMF framework for setting up an experiment and collecting its results; in addition, it offers tools such as remote hard reboot capability, connectivity measuring tools, user scheduler with slicing etc.
- ***OpenAirInterface*** (http://www.openairinterface.org). OpenAirInterface has been developed by EURECOM and is actually an open-source hardware/software development platform (rather than a specific "actual" testbed) for rapid experimentation in radio communications using software-defined radios. The platform is geared towards cellular and mesh networking and offers a subset of the LTE Rel. 8/9 standard. Both computer simulation and actual real-time implementation via a Linux-like RTAI and suitable hardware are supported and all developed tools are available to collaborating partners. Regarding CONECT, a dedicated outdoor testbed has been constructed in Sophia-Antipolis, France, consisting of 2 transmitters (eNb in LTE terminology), 2 relay stations and 2 user terminals (UE), and will be used to demonstrate the efficiency of the proposed signal-level cooperation algorithms. An emulator is also under construction for evaluating end-to-end video transport over adhoc networks and can be accessed at https://emu.openairinterface.org

Both testbeds are developed in an open-ended fashion using commercial hardware and open-source software, which allows for rapid modifications and upgrades. To exploit the strengths and capabilities of both testbeds, different experiments are planned for each of them, with NITOS focusing on the evaluation of packet level cooperation in ad-hoc networks and OpenAirInterface being used for signal level cooperation in a cellular setting (the final proof-of-concept scenario of video multicast over adhoc networks will also be tested on OpenAirInterface). Thus, the testbeds complement each other nicely and cover the full scope of CONECT. More information regarding the testbeds can be found in the public deliverable D4.1.

The NITOS and OpenAirInterface testbeds are currently available for remote experimentation through web interfaces and will also be federated with the Onelab/OpenLab FIRE infrastructure.

Experimental work:

Given that CONECT has just completed its first year of operation, the experiments performed so far have aimed in simply verifying the basic principles of the project and providing the necessary components for the more realistic experiments to be performed in the future. Specifically, a simple quantize-and-forward coding scheme, to be used as a baseline benchmark for new algorithms, has been developed and integrated into the OpenAirInterface testbed. A simple cooperative multicasting protocol exploiting channel variations has also been developed for OpenAirInterface, while NITOS has undergone the necessary modifications to allow for packet-level cooperation in the context of the IEEE 802.11 MAC.

Main challenges and future work:

The two testbeds implemented in the scope of CONECT are under continuous development (with new capabilities being added in regular intervals) to address the project objectives and provide real feedback to concurrent theoretical investigations. Different topologies will be examined and improved signal and packet level cooperation algorithms will be experimentally evaluated as they become available during the project's course. The expertise gained from this process, along with the developed infrastructure, will provide valuable assistance in future research efforts.

The CONVERGENCE project - http://www.ict-convergence.eu/

The CONVERGENCE project proposes to enhance the Internet through a content-centric, publish-subscribe service model. In this scheme, users would be able to query the network for content directly. In particular, the main project goal is to deliver a full specification and Open Source, implementation of the CONVERGENCE framework (middleware allowing third party developers to develop and deploy CONVERGENCE applications), a set of prototype applications demonstrating the functionality of the middleware in different business scenarios, reports demonstrating the scalability, robustness and security of the middleware, and studies from trials showing its applicability in real-life business scenarios.

The core concept is based on Versatile Digital Items (VDIs): a common container for any kind of digital content/data, derived from the MPEG21 standard. This approach in-line with the communications trend from circuit switching -> packet switching -> content switching, in which content is addressed with names, rather than hosts with addresses. Routing is also performed by name and the system is fundamentally an anycast network.

Research focus:

Content/centric networking, Middleware, MPEG Extensible Middleware (MPEG-M), Publish/Subscribe, Semantic searches, Security/Privacy.

The main identified research challenges in these domains include:

- To find the best naming scheme.
- Routing scalability (name-based routing tables would have to handle billions of names, which cannot be easily aggregated, though it should be possible to exploit caching and the fact that a lot of information is common locally).

- Content delivery (a content-centric network should deliver content to requesting devices, which do not have anymore a "native" address).
- Data-centric security (protecting information at the source with recipients not known a-priori).
- Continuing the support for existing Internet features (services, domain names).
- Supporting push services (e.g. streaming).
- Supporting per-content QoS differentiation.

Experimental work:

Experimentation will be made in 4 trials in cooperation with the OFELIA project. As a matter of fact, the CONVERGENCE project sites have access to GARR/GEANT via which they will reach existing OFELIA islands, which provide the necessary network infrastructure that can be configured/programmed so as to support the planned experimental work. The 4 Use Cases that the CONVERGENCE project proposes are:

- Management of audiovisual material / semantic and cognitive indexing.
- Management and annotation of a large photograph archive.
- Customer Relationship Management and logistics (shopping environment).
- An augmented lecture podcast service.

The types of experiments and measurements, which are planned to include real end users, will include:

- Throughput measurements:
  - To compare results achieved by using CONET, content-centric networking, with results achieved by using current protocols.
- Routing-by-name measurements
  - To assess the possibility to reduce the routing table size by using the lookup-and-cache approach implemented by CONET and to assess the scalability of CCNs.
- In-network caching measurements:
  - To assess the ability of CONET to reduce the time needed to retrieve content.
- Measurements related to routing-by-name for real-world-objects:
  - To verify the ability of CONET to address real-world-objects identified by their tag ID.

About FIRE facilities extensions:

One of the expected outcomes out of the experimental work of the CONVERGENCE project is the extension of the OFELIA facility through improvements of the OpenFlow functionality and extensions to the OFELIA network. In particular, through CNIT (coordinator of the CONVERGENCE project) it will be possible to connect to OFELIA the Rome and Catania islands through IP tunnels. The CNIT islands are connected with 1Gbps connections with GARR, the Italian National Research and Education Network. GARR is integrated with the broader GEANT initiative and also interconnected to several commercial networks.

Collaborations:

Apart from the collaboration with the OFELIA FIRE IP project, as discussed above, the CONVERGENCE consortium plans to collaborate with the COMET[14] and the PURSUIT[15] projects.

The EULER Project – http://www.euler-fire-project.eu

The EULER project's main focus is into the Internet routing system and routing engine resource consumption, using data collected from operational networks. A core aspect consists of understanding the dynamics of the Internet routing system to ensuring its stability and improving the mechanisms of the BGP routing protocol, as current BGP-based solutions often suffer from instability problems (policy-induced and protocol-induced) determining quite negative effects (non-deterministic/unstable BGP states and long convergence time after topology changes/failures).

Some of the reasons for this situation can be found in the increasing load on BGP:

- Growth in Internet traffic is 50% per year.
- Growth in routing table size is 15-25% per year (345'000 entries in September 2010).
- Growth in the number of Autonomous Systems is 10% per year (35'000 in September 2010).
- The Internet is becoming denser (the number of ASes is increasing, but the path length (3.7) remains steady.
- The BGP routing tables are updated about 1 million times per day (peak of 1'000 times per second).

Although current routing engines could potentially support up to 1 million routing table entries, instabilities resulting from (i) routing protocol behaviour, (ii) routing protocol information exchanges, and (iii) changes in network topology, may adversely affect the network's ability to remain in a useable state for extended periods of time.

In this perspective, the goal of the EULER project is to:

- Develop a methodology to process and interpret routing table data.
- Define a set of stability metrics and a means to apply them.
- Define ways of improving the stability of current and future routing systems.

Research focus:

The key research areas of the project sit at the intersection between applied graph theory, Internet topology/dynamics modelling and distributed routing algorithmic.

Past research on compact routing algorithms has focused on centralized algorithms in a static context. The EULER project investigates the design of distributed compact routing algorithms and greedy routing algorithms able to update quickly the routing tables with the evolution of the topology. These routings schemes aim to be more scalable than BGP.

[14] http://www.comet-project.org/
[15] http://www.fp7-pursuit.eu/

EULER makes use of measurement-based modelling/tomography, meaning that their models and algorithms rely on knowledge of the actual Internet topology obtained through measurements. Experiment data has been obtained from RouteViews, which is a project that comprises archives containing BGP feeds from a set of worldwide distributed Linux PCs running Zebra. RouteViews routers use multi-hop BGP sessions to peer with transit providers and other ASes at selected locations around the globe. The data comprises:

- RIBs (collected every two hours).
- Update files are rotated every 15 minutes.

RouteView was chosen as it offers BGP dataset (taken out of several locations around the Internet) that no experimental facility could offer (or accessible to experimenters). The RouteViews data were used for one of our measurement-based experiments on path-vector routing stability, i.e. experiments in the context of the EULER project will not exclusively rely on RouteViews data.

Experimental work:

To verify scalability properties of the protocol components implementing the designed (and validated) routing algorithms, the use of PlanetLab experimental facility is foreseen. It is important to underline the difference between large scale facility and scalability experiment: the scale of a facility (the number of resource units and their distribution) does not determine the scalability properties of the corpus of the experiment. The scalability properties of the experimented corpus (in the EULER case routing protocol component) determine the number of resource units that are locally required for the corpus to be executed up to a certain scale. Thus, such experiment can be performed to i) verify a pre-determined level of scaling of the experimented corpus and/or ii) iteratively determine the scaling properties of the corpus with the risk that the dependency on the global properties could never be found (hidden dependencies, correlations, non-linearities, etc.). Hence, only the former type of experiment leads to verifiable results. In other words, a large-scale facility can help verifying known scaling properties of an experimental corpus but not identify its actual scaling properties.

The EULER project also relies on measurements of the actual Internet topology, for which PlanetLab is a useful resource that provides a large number of monitors distributed in the Internet. Henceforth, integrated experiments where topology measurements are exploited by their routing algorithms are foreseen on PlanetLab.

About the type of experiments:

So far, the project has defined a set of stability metrics that characterize BGP's instabilities. Initial results suggest that they can be used to both identify and quantify the impact of instability events on local routing tables.

Ongoing work includes determining if their stability function can be integrated within the BGP decision process and what would be the possible trade-offs.

In particular, to test scalability properties of the EULER algorithms (validated by means of simulation) and in particular of some of the critical components of the associated routing protocols, experiments on PlanetLab are foreseen (feasibility study is expected by mid-2012). As mentioned above, scaling verification experiments require identifying a priori and by other means the scaling properties of the designed routing algorithms.

It is essential for the EULER Internet degree distribution measurement method, as well as for radar measurements. Such work cannot be conducted using emulation facilities, as their aim actually is to explore accordance between current hypotheses and actual behaviour of the Internet. PlanetLab provides a unique platform embedded in the actual Internet that is absolutely needed.

Moreover, experimental evaluation and comparison of routing protocol components will be realized by using the iLab.t controlled emulation environment. As of today, a single physical node of this experimental infrastructure can emulate about 20 routers leading to a total of number of 2000 routers (if the whole experimental facility would be mobilized).

Collaborations:

Formal collaboration with other project has not started yet. As a first step, an analysis of possible collaborations has been performed. Among the available facilities, OFELIA, with its OpenFlow technology, may offer extended/enhanced features with respect to the Virtual Wall, which is the EULER default experimental facility (notice that the Virtual Wall is one of the five islands which will be available in OFELIA). OFELIA is currently considered as experimental emulation platform within WP4 of EULER. Experiments that could be conducted on facilities such as OFELIA, include measuring performance of local exchanges between routing and forwarding components (upon routing table updates resulting from topological changes).

The EULER experimental approach assumes that simulation and emulation both in closed and open experimental environments provide complementary information (some experiments are complex to realize by simulation while others are simply impossible by means of emulation). In this context, validation of a new routing algorithm, for instance, is better conducted on a simulation platform first (after theoretical verification) not only because their resulting cost is lower but because associated tools enable producing results verifying repeatability, reproducibility (assuming multiple simulators are available), and reliability criteria. On the other hand, emulation experiments can lead to reproducible and repeatable results but only if their "conditions" and their "executions" can be controlled. Realism can thus be improved when compared to simulation (in particular for time-controlled executions of protocol components on real operation system). Nevertheless, such experiments are more complex and time consuming to configure and execute; performance evaluation is possible if the experimental platform comprises a "sufficient number" of machines (representative of the experiment to conduct). Emulation still requires synthetic network conditions (models) if executed in controlled environment and either injecting real traffic or replay traffic traces (not that even when available "spatial distribution" of traffic is available remains problematic to emulate because the spatial aggregation of address prefixes is not necessarily as the routing tables are often not provided together with traffic traces).

Emulation-based experiments can be realized using other experimental facilities than PlanetLab. PlanetLab provides an open distributed environment for emulation-based experiments. Whereas emulation is "enough" in the EULER case, PlanetLab does not provide better capabilities when compared to other emulation platforms that allow control of experimental parameters to ensure repeatability and reliability of results (like the use of

IBBT and VINI[16]). Nevertheless the combination with topology measurements exploitable by the EULER routing algorithms makes such environment (i.e. PlanetLab) suited for integrated experiments that can be foreseen towards the end of the project.

The HOBNET project – http://www.hobnet-project.eu

The HOBNET project attempts a holistic approach addressing several aspects of Future Internet Green/Smart Buildings, at different networking layers and system levels.

Research focus:

In particular, the proposed research addresses algorithmic, networking and application development aspects of Future Internet systems of tiny embedded devices:

- An all IPv6/6LoWPAN infrastructure of buildings and how IPv6 can integrate heterogeneous technology (sensors, actuators, mobile devices etc).
- 6lowApp and its standardization towards a new embedded application protocol for building automation.
- Novel algorithmic models and scalable solutions for energy efficiency and radiation-awareness, data dissemination, localization and mobility.
- Rapid development and integration of building management applications.
- Support for the deployment and monitoring of resulting applications on FIRE testbeds.

Experimental work:

For their experimental work two facilities were chosen: IBBT and WISEBED. WISEBED was primarily chosen because two HOBNET partners participate in it and some of the know-how gained can be exploited in HOBNET. Still, several distinct features exist (e.g. WISEBED is not focusing on IPv6), so usage of existing knowledge will be partial and significant extensions and innovations are expected and needed. During the evolution of the project, additional facilities that could be of interest to HOBNET have been identifies, such as IBBT; actually there is ongoing cooperation between HOBNET and IBBT partners.

It is important however to underline that the HOBNET approach is, by its very initial design, application-specific, focusing on the particular application domain (building automation and energy efficiency), so the plan is to mainly use facilities that HOBNET itself will create: an IPv6 sensor network test-bed in CTI-Patras, a test-bed (using but significantly extending the local WISEBED installation) in University of Geneva (UNIGE) and a new HOBNET test-bed at the MANDAT building in Geneva (MANDAT is a UN-related, public utility foundation participating in HOBNET as an end user[17]).

Each of the 3 HOBNET test-beds is planned to include a rich variety of wireless devices (like sensors of various types, actuators for ambient control, robotic elements, RFID etc). The CTI test-bed will include 200 devices, the UNIGE test-bed 150 while the MANDAT test-bed will include 150 devices in total. In all 3 test-beds a variety of smart/green building scenario will be implemented and validated; characteristic scenarios include Local adaptation to

[16] VINI is an extension of PlanetLab. More details can be found in [3].
[16] More information about MANDAT can be found at: http://www.mandint.org/

presence, Emergency management, Electricity monitoring, CO2 monitoring, Maintenance control, Customization to personal profile, Building, 3D visualization and monitoring, Mobile phone, User awareness, Early warning system and oil tank monitoring, Garden watering, Resource tracking and monitoring[18]. It is worth noting that the main novelty is meant to lie not in the scenario themselves but in the methodology that is followed: the HOBNET architecture chosen, the networking protocols used, the algorithms designed and implemented, the building monitoring tool used (BMF) and the integration of these in a holistic manner.

User involvement:

As mentioned above, MANDAT International is participating in HOBNET as an end user; its buildings (both the current and planned ones) will serve as the main users of the HOBNET results. However, the consortium is also exploring additional end users to be involved by including in the experiments other scenarios such as schools and public authorities.

Collaboration:

It was recently decided to explore the possibility to make use of additional facilities such as the ones offered by OFELIA. However, the HOBNET approach is quite application-specific, focusing on the particular application domain that is building automation and energy efficiency, and in this perspective the consortium plans to mainly use facilities that HOBNET itself will create.

The LAWA project - http://www.lawa-project.eu/

LAWA is a FIRE research project in scalable data analytics with emphasis on the temporal dimension or content evolution over time. The focus is on domains with large, perhaps distributed, archives of Web data. They are exploring issues of suitable and optimized algorithms as well as trying to remove limitations of the typical "cloud" infrastructure and prepare for a "distributed cloud" in the future. The aim is creating a test environment, the "Virtual Web Observatory", with significant (Terabytes of) amounts of test data to analyze with LAWA tools and on the LAWA infrastructure.

In the first year of the project, the consortium is consulting with other researchers who share their interests to obtain requirements for the testbed. In the second year of the project, the aim is to identify and involve a few external users of the LAWA testbed. In year 3, the aim is to invite more extensive outsider users to validate their methods and enable more scientific studies on the data they will provide/render accessible.

Research focus:

- **Web Scale Data Provision**. The main objective in this area is to ensure that Web data is provisioned on the LAWA infrastructure at the right scale and in an optimal structure for further processing. This encompasses efficiency and scalability of the infrastructure at "Web scale", and also the development of an on-demand crawling service which allows researchers to craft focused content collections.

[18] For a detailed description see D1.1 at http://www.hobnet-project.eu/files/D1.1.PDF

- **Distributed Access to Large Scale Data Sets**. This area focuses on making large data sets accessible to researchers around the world. Their primary efforts will be to extend the Hadoop system so that data can be read and filtered at its home location and then distributed out to processing nodes anywhere in the Internet. This approach will require extensions to Hadoop for wide area operation where LAWA makes extensive use of distributed index structures to capture and access the location and contents of the data stores.
- **Web Analytics**. This part of the work is dedicated to methods and tools for the discovery of latent knowledge from aggregated Web data. The work includes investigations on the dynamics of Web sources and dependencies among them. The aim is to develop algorithms and tools for systematically aggregating, querying, mining, and analyzing statistical patterns, cross-data correlations, and temporal variability, in order to distil and reveal knowledge latently embedded in Web sources.
- **Virtual Web Observatory**. The Virtual Web Observatory will implement a collection of core analysis solutions for the application stakeholders by integrating the methods developed in the project into a FIRE infrastructure application. The Virtual Web Observatory will support experimentally driven research by enabling cycles of experiments on large-scale Internet Data from collection to analysis. Moreover, the results of experiments will be used to refine the Virtual Web Observatory and its component. Experiments will include two highly innovative steps in Web analytics: data-quality based classification and detection of temporal patterns in Web data.

Experimental work:

The LAWA experimental testbed is a platform for analytic tools and will showcase a “Virtual Web Observatory”. The initial configuration will be based on data from the LAWA Reference Collection, but the plan is to include a federation with the FIRE facilities and to incorporate supplementary data from FIRE.

The BonFIRE and OpenLab facilities are currently considered. BonFIRE because it provides a cloud environment made up of multiple small and separated clusters, which fits with the LAWA model of the future heterogeneous large data store – computational facilities that exist where the data is being created or first aggregated. The LAWA consortium is currently working on how to structure queries and how to schedule the computational assets to minimize the obvious difficulties that such a computing/communications environment will introduce. One of the LAWA partners (HUJI) has already submitted an OpenCall proposal to BonFIRE addressing some of these points, and will elaborate on it with the help of additional LAWA participants. In addition, through participation in the PlanetLab-Europe consortium, as PlanetLab begins to offer standard logins as part of its federation activities, LAWA will consider to eventually make use of their facilities as well.

About FIRE facilities extensions:

By providing access to Web scale data with powerful analytical tools, LAWA aims to enable the FIRE experimental facility to explore a new range of issues related to the Future of Internet that are out of reach at the moment for European Researchers. To this end, LAWA will utilize Web archive data contributed by its partners European Archives and Hanzo Archives. Several reference corpora will be made freely available to the research community at large.

The LAWA aim is to compile reference corpora for research on the Future Internet content and make them freely available to the research community. These will include large-scale, citable datasets that can be harnessed, for instance, for classification, link analysis, spam filtering, entity tracking, and other Web-analytics purposes. They anticipated, via early feedback from potential users (e.g. from the LAWA User Workshop), that this data will be of high interest for a significant number of Future Internet Research projects by creating unique view and research opportunities on the data layer of the Internet. This encompasses but is not limited to Information Retrieval, Web Mining, Network Analysis, Social Sciences, etc. Although this, including users' base, is planned to be delivered by the project, it will be of fundamental importance for their success to get the help of the FIRE community to extend their exploration of the infrastructural issues and enhance their outreach.

About users' involvement:

LAWA stages annual workshops for those interested in using the services developed by their consortium (the next upcoming in November 2011). These events are meant to establish a target user group for the services of LAWA, including a research community of library and archiving organizations as well as other stakeholders interested in large scale Web analytics. LAWA monitors the requirements of these researchers in order to develop tailored services for comprehensive time-travel Web archive access. The idea is to develop demonstrators to allow target user groups to apply, evaluate, and comment on LAWA services for Web data analytics. The formats used for documenting these requirements will be of different type. At the beginning of the project, LAWA has set up a target user community comprised of active Web archiving institutions and researchers working in the area of Web analytics. These users will be contacted in consecutive stages throughout the project to ensure LAWA will meet their requirements. User requirements are gathered on a Web-based wiki. In addition, it is planned to provide access to citizens at large by an introductory user front-end as part of the Virtual Web Observatory. The demonstrator will have limited scope and customized/restricted functionality to ensure robustness for public Web access.

Collaboration:

Besides foreseen collaboration with BonFIRE and PlanetLab Europe in the perspective of their experimental work and with several EU projects in the area of Web technologies, there are no ongoing collaborations with other FIRE STREPs.

The NOVI Project – http://www.fp7-novi.eu

NOVI targets innovative research in a bottom-up approach, addressing a critical area in Future Internet (FI) services: How FI users (including Data Center Managers and Cloud Service Providers) securely share a multi-domain networking substrate. NOVI aims to empower end-users and FI service providers with algorithms and tools to compose and manage isolated slices, baskets of virtual resources and services provided by diverse yet federated Future Internet (FI) platforms.

Research focus:

NOVI's research objectives can be classified within four closely related areas within the *NOVI Innovation Cloud*:

- Federated Virtualization Technologies: open architectures enabling the formation, control and management of isolated slices across local and wide area multi-domain networking environments.
- Virtual Resource Brokerage (virtual resource allocation algorithms to meet service requirements of user-slices, e.g. reproducibility of experiments).
- Semantic Resource Description (formal description of virtualized network and cloud computing objects in a complex environment, assisted by semantic methods and ontologies).
- Monitoring Architectures: passive and active monitoring of virtualized resources in a federated environment, accessible to a specific user-slice and/or the federated infrastructure administrator.

The NOVI Information Model provides abstractions and semantics of federated virtualized resources, enabling ontology based tools and algorithms of the NOVI Innovation Cloud.

NOVI's Control and Management Federation Architecture is based on the well established GENI Slice Federation Architecture (SFA), complementing it with a higher-layer Control and Management Plane offering the following services:

- An Intelligent Resource Mapping Service that provides the functionality to support embedding user requests within the physical substrate.
- A Resource Information Service that includes the capability of efficiently finding resources based on their context and for storing the objects representing entities defined according to the NOVI Information Model.
- A Monitoring Service that enables NOVI users and administrators to retrieve temporal information behaviour of resources of the network via active and passive network measurements.
- A Policy Management Service that governs the intra and the inter domain behaviour within or amongst member domains of a NOVI federation.
- A Request Handler Service that delivers requests of services and resources appropriately formatted to the NOVI Service Layer.

The NSwitch (NOVI virtual switch) distributed software switch is designed by NOVI researchers to extend virtual connectivity of user slices across federated platforms.

NOVI monitoring addresses novel monitoring techniques capable of operating in virtualized infrastructures and in the underlying federated network substrate to make sure that measurement processes in different heterogeneous networks can be synchronized and all available measurement resources can be used efficiently.

Experimental work:

Research prototypes will be developed according to the spiral experimentally-driven methodology. In NOVI’s first phase (Spiral 1) early prototypes are being designed and implemented over two existing FIRE facilities, notably PlanetLab Europe and FEDERICA. Validations of an integrated prototype and refinements of Spiral 1 prototypes will follow in the second phase (Spiral 2).

The NOVI testbed, configured in Spiral 1, consists of a private PlanetLab installation (where nodes run PlanetLab's VServer virtualization technology) with a MyPLC operated by our partner PNSC and sites within the premises of NOVI partners in Warsaw - Hungary (ELTE), Poznan – Poland (PSNC) and Athens – Greece (NTUA). The testbed is complemented with three core FEDERICA routers at PSNC (Poland), DFN (Germany) and GARR (Italy) and with several FEDERICA VNodes running the VMWare ESX virtualization software. This is a testing environment that captures the essence of some key issues of the project: Management and control of federations of heterogeneous virtualized infrastructures, data plane connectivity among virtual resources (e.g. between one PlanetLab sliver and one FEDERICA Logical Router), role based access control, inter-domain policy models, combination of measurements from multiple domains, sub-optimal embedding of virtual resources within the federated substrate. To support these key research issues, NOVI proposes a semantic-aware Information Model that enables ontology-based resource descriptions of virtual resources and services within heterogeneous domains that may have different (or no) resource description formats.

Further experiments may be jointly planned with related initiatives in Europe (FIRE projects, National Research and Education Networks – NRENs and GÉANT) and FI researchers in North America and Asia – Pacific regions.

About end users' involvement:

NOVI aims to develop and implement a set of tools for enabling federation of heterogeneous virtualized infrastructures. During the experimentally driven software development in NOVI, users of the experimental NOVI federation are drawn from the NOVI consortium.

Should the NOVI framework be finalized and deployed in existing virtualized infrastructures, users of the NOVI federation would be the users of the virtualized infrastructures. Additional domain-specific policies should be specified by domain administrators to allow some of their users to have access to resources across the NOVI federation. User involvement is being dealt with as usage scenarios during NOVI's experimental research; nevertheless, sustained operational user support is beyond the scope of this STREP project.

About FIRE facilities extensions:

NOVI strives to serve as an enabler technology to aid establishment and management of federations of FI virtualized infrastructures. The project will offer the community a semantic-aware NOVI Information Model, an SFA-based Control and Management Federation Architecture with advanced services (Resource Information Service, Intelligent Resource Mapping Service, Policy Service, Monitoring Service) and the NSwitch distributed virtual switch. The overall goal is that NOVI's research outcomes will enrich FIRE facilities with federation mechanisms and tools, enabling them to collaborate offering their users federated baskets of virtual resources and services.

Collaborations:

A concerted effort is carried out in NOVI by academic researchers and operators of advanced public e-Infrastructures (NRENs and GÉANT, FEDERICA, PlanetLab Europe) in partnership with a pioneer vendor in Internet technologies (Cisco). As such, it is expected to contribute in evolving standardization efforts and provide pre-normative research and proofs of concept on federated virtualized solutions to vendors and service providers. NOVI partners interact with the OGF, IPFIX and DMTF standardization bodies, aiming at contributing with resource description formats for federations of heterogeneous virtualized infrastructures (NOVI Information and Data models) and measurement formats for federations of virtualized infrastructures which employ domain-specific measurement tools. NOVI disseminates its findings in scientific events and global FI initiatives; it assists on developing the FIRE facility at the FIRE Architecture Board.

The SCAMPI project – http://www.ict-scampi.eu

SCAMPI, *Service platform for social Aware Mobile and Pervasive Computing*, is a three years STREP project that started last October 2010. In the context of a Future Internet characterized by an increasing diffusion of devices with heterogeneous capabilities and resources, SCAMPI aims to empower end users through a combination of pervasive and opportunistic networking capabilities as to allow a more effective access to a variety of services in a secure manner.

Research focus:

Pervasive and opportunistic networking, reliable and secure service access for nomadic users in opportunistic scenarios.

SCAMPI focuses on opportunistic networking environments, where (i) devices are spread in the environment, (ii) events such as long disconnections and partitions are the rule, and (iii) no simultaneous multi-hop paths can be guaranteed. Thus, SCAMPI generalizes the pure opportunistic networking concept, and investigates the novel concept of opportunistic resource usage in challenged networks.

Experimental work:

Optimized service access is enabled by putting in place social and context awareness mechanisms that bridge resources at various layers, including human social, application, service, network and physical layers. To do so, SCAMPI takes an experimentally-driven approach that relies upon:

- Developed solutions distributed to real world users.
- Data from usage as further input to research.

As very few large opportunistic testbeds are available (versus fixed ones), initial experiments are led in scenarios that are expected to grow in size later on and that in a first phase rely on fixed Internet settings, e.g. campus-sized trials with low power devices (typically smart phones) with socially-driven applications in the Internet.

However, SCAMPI aims to extensively use the testbeds provided by the FIRE facilities as it follows:

- Running SCAMPI solutions on testbeds.
- Extending testbeds with ubiquitously connected devices that interact with testbeds.

The three main testbeds identified so far by the SCAMPI consortium are the one under construction within the SmartSantander project, the NITOS within the OneLab2 project and iLab.t (IBBT). The main criteria for choosing these specific FIRE facilities are:

- Ease of access and use,
- Easy integration with simulation tools,
- Additional features (e.g. sensors, broadband access, fixed WiFi infrastructure).
- Large scale.

About end users' involvement:

The current plan is to experiment with real users and real applications with trials of increasing scale. Examples of experiments thus far consisted of trace collection using sensors and smartphones and smartphone apps relying on opportunistic communications. Collected measurements (mobility, traffic, and sensor) will further extend the SCAMPI models, influence the design of solutions, and will be used as inputs to simulations.

The core idea beyond involvement of end users is correlating mobility with the social dimensions of humans-beings and active users interacting with their devices are of prominent interest to us. In a few cases though, for more theoretical works, the involvement of users might not be required. An example could be to reproduce a very dense topology of static communicating devices (sensors and/or smartphones) to simulate opportunistic communications in a very crowded place.

SCAMPI plans therefore to leverage FIRE facilities to evaluate and validate the SCAMPI results at a larger scale that what can be provided by individual partner testbeds (or even cross-partner testbeds) and with more features (e.g. sensors, WiFi hotspots with storage for caching, 3G/LTE broadband access).

Collaborations:

To optimize the experimental work through the selected FIRE facilities initial collaboration has been initiated with SmartSantander and OneLab2. In addition the plan is to collaborate with other FIRE facility projects to federate the SCAMPI platform and testbeds to extend their reach to the outside world.

The SPITFIRE project – http://www.SPITFIRE-project.eu

The SPITFIRE project works towards the realization of a stronger connection between the natural and the digital worlds. The main idea is to investigate unified concepts, methods, and software infrastructures that allow the efficient development of robust applications that span and integrate the Internet and the embedded world. Thus, SPITFIRE aims at lowering the effort required for developing robust, interoperable, and scalable applications in the Internet of Things (IoT).

This will facilitate building new kinds of applications and services that were not possible before thus having an impact on research, industry, and private households. Industry will be able to evaluate new solutions and pick those that operate satisfactorily under realistic conditions.

Research focus:

Internet of Things, Web of Things.

The starting point for SPITFIRE is the observation that only the development of the Web has opened the Internet to the greater public by providing open standards that resulted in a usable and interoperable software infrastructure and well-understood methodologies for design, implementation, and evaluation.

The approach is to extend the Web into the embedded world to form a Web of Things (WoT), where Web representations of real-world entities offer services to access and modify their physical state and to mash up these real-world services with traditional services and data available in the Web to create novel applications.

Experimental work:

The embedded component of the IoT is largely affected by its surrounding real-world environment. Therefore, experimentation on real platforms in realistic environments is a key for successful IoT research. WISEBED was selected by the SPITFIRE consortium as it provides a large-scale IoT experimental facility that is an ideal environment to use for exposing services offered by resource-constraint sensor nodes. The idea is to deploy software on the sensor nodes so that they expose their functionality and data (e.g. sensor readings, topology data, etc.) as CoAP services over 6LoWPAN that are interoperable with RESTful web services in the Internet. In particular, these experiments shall demonstrate the feasibility and effectiveness of the SPITFIRE approach and techniques in real-world settings.

Moreover, once the SmartSantander platform will be available and compatible with the WISEBED APIs, for SPITFIRE it will be possible to extend their experiments in a smart-city context and obtain even more realistic data, provided by a larger-scale and a wide range of devices and sensor data.

About FIRE facilities extensions:

The SPITFIRE project is currently developing a platform-independent implementation of the CoAP IETF standard draft which enables RESTful web services on extremely resource-constraint devices. The implementation is based on the Wiselib algorithm library (www.wiselib.org) which already supports virtually all IoT platforms from sensor nodes to smartphones and tablets. This implementation could be used to as a service-oriented basis for smart-city applications.

About users' involvement:

Within the WISEBED experimental facility, no users are planned to be involved. A potential extension of the experiments to the SmartSantander facility (which should be compatible to WISEBED) could involve users (e.g. citizens of Santander).

Collaboration:

Besides foreseen collaboration with WISEBED and SmartSantander FIRE IPs in an experimental perspective (as discussed above), SPITFIRE conducted a joint school on semantics and the Wiselib with WISEBED and the FP7-project FRONTS.

7.3. Analyzing the FIRE STREPs

All seventeen FIRE STREPS answered the specific questionnaires that the FIRE STATION consortium prepared.

In the following Table, the core information that has been collected is summarized. The first nine projects listed in the Table are the Call 2 ones and the following eight are the more recent Call 5 projects.

In line with the main objectives of this analysis, the information summarized hereby reports on:

- The key research areas the FIRE STREPs have covered or are covering.
- The kind of facilities they deployed, whether existing FIRE facilities or proprietary ones or otherwise developed within the project context.
- Whether or not end users have been involved in the projects' experiments.
- The collaborations established in the FIRE context with other projects, if any.

FIRE STREP Name	*Research key area(s)*	*Deployed FIRE Facilities*	*Deployed proprietary facilities*	*Developed their own facilities*	*Involved users*	*Collaborations*
ECODE	Cognitive routing systems, Machine learning mechanisms		Deployed the iLab.t (IBBT) - better suited than PlanetLab			Collaborations: -RESUME-NET: via a dedicated workshop - Self-NET: potential synergies have been envisaged
N4C	Internet access for all, Delay and Disruption Tolerant Networking (DTN), wireless technology			Built its own facilities, namely the Slovenian Kočevje testbed run by N4C partner MEIS	Yes	Collaborations on a general level with other FIRE projects and more specific ones with other EU and not only projects and initiatives
Self-NET	Cognitive mechanisms for network management and optimization in autonomic and wireless domains, routing adaptation for wireless networks	PanLab-PII facilities and specifically the OCTOPUS testbed (WIMAX)				Collaboration with PanLab-PII through the usage of the OCTOPUS testbed. Potential closer synergies with ECODE have been suggested but not explored
RESUME-NET	Network and service resilience in opportunistic ad wireless mesh networks - Internet of Things, BGP,	OneLab2 (PlanetLab)	G-Lab (German facility) – because not all planned experiments were possible on PlanetLab			Collaboration with OneLab2 through the usage of their facility

FIRE STREP Name	***Research key area(s)***	***Deployed FIRE Facilities***	***Deployed proprietary facilities***	***Developed their own facilities***	***Involved users***	***Collaborations***
	Publish-subscribe		facilities			
PERIMETER	QoE, always best connected	Some experiments have been run on a FEDERICA slice	Deployed proprietary facilities, at WIT and TUB (project partners)		Yes	Collaboration with FEDERICA through the use of their facilities
SMART-Net	Wireless Mesh Networks and smart antennas			Developed their own facilities that can be accessed to the broad audience upon request to the consortium		
VITAL ++	P2P content distribution and IMS technology			Developed their own facilities that have been offered through the Teagle tool available in Panlab	Yes	Collaboration with Panlab
Nanodatacenters	Peer-to-peer services in Internet-based networks, Infrastructure for interactive and secure applications, Monitoring services	OneLab2 (PlanetLab)		Part of the back-end system was hosted by Nanodatacenters partners	Yes	Collaboration with OneLab2 through the use of their facilities
OPNEX	Distributed network			Developed and		Close collaboration with

FIRE STREP Name	***Research key area(s)***	***Deployed FIRE Facilities***	***Deployed proprietary facilities***	***Developed their own facilities***	***Involved users***	***Collaborations***
	coordination and control, Multi-hop wireless networks by designing advanced systems optimization and control theory-driven algorithms			deployed several testbeds (DES-testbed, NITOS, wnPUT-testbed, Technicolor wireless testbed, INRIA testbed), NITOS and DES-testbed are offered for federation through OneLab		WISEBED as the same wireless sensor network of DES-testbed has been used by both projects. Some of the WISELib algorithms have been deployed by OPNEX
CONECT	Information and communication theoretic principles, coupled with optimization theory techniques for open broadcast wireless medium			The plan is to develop and deploy real life experimental testbeds within NITOS and OpenAirInterface. Federation with OneLab2/Openlab is foreseen		
CONVERGENCE	Content centric networking, Middleware, MPEG Extensible Middleware (MPEG-M), Publish/Subscribe, Semantic searches, Security/Privacy	OFELIA facilities		One of the expected outcomes out of the experimental work is also the extension of the OFELIA facility through improvements of the OpenFlow functionality and	Yes	Collaboration with OFELIA through the usage of their facilities

FIRE STREP Name	***Research key area(s)***	***Deployed FIRE Facilities***	***Deployed proprietary facilities***	***Developed their own facilities***	***Involved users***	***Collaborations***
				extensions to the OFELIA network		
EULER	Graph theory, Internet topology/dynamics modelling and distributed routing algorithmic, BGP routing	The use of PlanetLab experimental facility is foreseen	Evaluation and comparison of routing protocol components also planned on iLab.t controlled emulation environment			Collaboration with OneLab2 through the usage of PlanetLab, and investigation of potential collaboration with OFELIA is ongoing (use of OpenFlow technology)
HOBNET	Algorithmic, networking and application development aspects of Future Internet systems of tiny embedded devices, all IPv6/6LoWPAN infrastructure	WISEBED facilities	IBBT facilities		Yes	WISEBED and currently exploring further collaboration opportunities with OFELIA
LAWA	Web Scale Data Provision, Distributed Access to Large Scale Data Sets, Web Analytics, Virtual Web Observatory	Plans to federate with BonFIRE and OpenLab2 (PlanetLab)		They will build their own facilities, but plans to federate with BonFIRE and OpenLab2 (PlanetLab)	Yes	Foreseen with BonFIRE and OpenLab2
NOVI	Federation of resources in virtualized	OneLab2 (PlanetLab) and FEDERICA			Yes (users intended as	With OneLab2 and FEDERICA through

FIRE STREP Name	***Research key area(s)***	***Deployed FIRE Facilities***	***Deployed proprietary facilities***	***Developed their own facilities***	***Involved users***	***Collaborations***
	e-Infrastructures, Formal description virtualized network and cloud objects with semantic methods, Allocation of resources with QoS attributes				network and service providers)	deployment of their facilities
SCAMPI	Pervasive and opportunistic networking, reliable and secure service access for nomadic users in opportunistic scenarios	SmartSantander facilities and the NITOS testbed within the OneLab2		Initially, their own defined scenarios and testbeds, later on federation with SmartSantander and OneLab	Yes	SmartSantander, OneLab2
SPITFIRE	Internet of Things, Web of Things	WISEBED was selected as a FIRE facility for SPITFIRE as it provides a large-scale IoT experimental facility, Later on exploration and potential deployment of SmartSantander facilities is also planned			Not initially, maybe later on through deployment of SmartSantan der facilities	WISEBED, and SmartSantander

7.3.1. Research/technological areas coverage

The main research/technological areas covered by the past and ongoing FIRE STREPs are listed hereby. Notice that some of these areas are intrinsically linked and sometimes overlapping and that some of the STREPs cover therefore several research/technological areas.

- **Cognitive and learning mechanisms**:
  - ***ECODE*** explores and deploys cognitive and learning mechanisms to define a new routing approach for Future Internet networking scenarios which improves the Internet in terms of manageability, security, availability and accountability.
  - ***Self-NET*** explores and deploys cognitive and learning mechanisms to define and develop an innovative approach for network self-management relying upon autonomic monitoring, decision making and execution focusing on wireless networks.
- **Wireless technology / Opportunistic networking**:
  - ***Self-NET*** research/technological work, as just mentioned above, covers improved autonomic management mechanisms for wireless networks.
  - ***N4C*** focus is on the definition and evolution of delay- and disruption-tolerant (DTN) networking technology to improve opportunistic wireless connectivity in challenging network scenarios.
  - ***SMART-Net*** focus is on wireless access networks and smart antennas for use with each subscriber unit. They investigated this approach for advanced routing and scheduling protocols specifically designed of heterogeneous Wireless Mesh Networks.
  - ***RESUME-NET*** focuses on network and service resilience through the definition of policy-based network management and defense mechanisms to serve complex networking scenarios, such as opportunistic and wireless mesh networks.
  - ***OPNEX*** focus is on multi-hop wireless networks by designing advanced systems optimization and control theory-driven algorithms.
  - ***CONECT*** studies digital radio communications, namely signal- and packet-level cooperation and end-to-end information transport in the form of multicasting, in wireless networks.
  - ***SCAMPI's*** main research/technological focus is on opportunistic resource usage in challenged networks. The core idea is to put in place optimized service access by putting in place social and context awareness mechanisms.
- **Networking Protocols**:
  - ***N4C**,* as mentioned above, focuses its research activity on improving data transport, routing and delivery for turbulent DTN network scenarios.
  - ***SMART-Net*** investigates the deployment of smart antennas and the definition of advanced routing and scheduling protocols for heterogeneous wireless mesh networks.

  - ***OPNEX*** has also developed and delivered algorithms at the routing level within their wireless network scenarios.
  - ***ECODE's*** main research efforts are in the area of routing system scalability and quality by deploying cognitive and machine learning algorithms.
  - ***EULER*** concentrates its R&D efforts in the area of dynamic modelling and distributed routing mechanisms that aim at evolving more traditional BGP routing technologies and solutions.
- **Sensors / Internet of Things / Web of Things**:
  - ***HOBNET***'s main R&D focus lies on the addresses algorithmic, networking and application development aspects of Future Internet systems of tiny embedded devices empowered by sensors able to provide the necessary information about their surrounding environment.
  - ***SCAMPI*** plans to bridge resources at various levels for improved opportunistic networking also by making use of sensors captured trace collections correlating mobility with social dimensions of human-being and active users interacting with devices.
  - ***RESUME-NE****T* also focuses on resilience of a publish-subscribe platform that supports the Internet of Things deployment.
  - ***SPITFIRE*** is aiming at extending the Web into the embedded world to form a Web of Things and thereby facilitate the development of robust and scalable applications in the Internet of Things dimension.
- **Content-centric networking / P2P**:
  - ***RESUME-NET*** enables service resilience by creating "supervisors" operating in a P2P modality and coordinating with each other in order to minimize probability of failures. PlanetLab was chosen as facility to run experiments also because of the possibility to runP2P experiments.
  - ***CONVERGENCE*** proposes to enhance the Internet through a content-centric, publish-subscribe service model. The core idea is to enable users to directly query the network for content delivery.
  - ***Nanodatacenters*** focus is on enabling the next generation of interactive services and applications by combining the power of data centers with the scalability of P2P methods.
  - ***VITAL++*** main activity and work is on P2P techniques for content distribution and control, namely live streaming and monitoring, as well as improved QoS-based service offering in IMS-enabled networks
- **Data/service management / Clouds**:
  - ***LAWA*** main research focus is on ensuring improved distributed access to large scale data sets by paving the way to the "distributed cloud" in the Future Internet.

- ***NOVI*** focuses on federation of different kinds of resources in virtualized e-Infrastructure by providing mechanisms for the description, monitoring and management of virtualized object clouds.

- **Network Management/ Resource Optimization**:
  - Thorough technical work at various levels (either at the routing or at the service level, either in wireless or opportunistic scenarios, either by deploying cognitive and learning mechanisms) the following projects are contributing to advance the know-how in the areas of network management and resources optimization: **RESUME-NET**, **Self-NET**, **OPNEX**, and **CONECT**.

7.3.2. STREPs usage of the FIRE Facilities

Call 2:

- Four projects, meaning the 45% of the Call 2 ones, namely Self-NET, RESUME-NET, PERIMETER and Nanodatacenters, are making use of some FIRE facilities. PERIMETER, Nanodatacenters and RESUME-NET also run in parallel experiments on proprietary testbeds.
- N4C built and deployed its own facilities as the FIRE available ones could not serve the project's main purpose of reaching the remote identified non-urban areas they targeted by deploying DTN mechanisms. Important to notice is that ***the N4C facilities are offered for federation to the FIRE facilities based on the OneLab2 strategy***.
- Also the Nanodatacenters project built and hosted part of the back-end system that was used in combination with OneLab2 (PlanetLab) facilities.
- ECODE only run experiments on iLab.t (from IBBT) as it was better suited that the available FIRE Facilities ones.
- SMART-Net developed their own facilities (smart nodes - WiFi, WiMAX – equipped with smart antennas) that can ***still be accessed upon request*** to: bruno.selva@thalesgroup.com.
- VITAL++ built its own facilities that combine P2P techniques and IMS technologies that have been made available through Panlab to the broader audience.
- OPNEX developed and deployed various testbeds (some of which are proprietary). However, the two largest testbeds they developed, namely NITOS and DES-testbed, are ***publicly available for experimentation*** and federated with other FIRE facilities like for instance OneLab2.

Call 5:

- Six Call 5 projects, meaning 75% overall, are planning to make use of available FIRE facilities.
- LAWA will build its own facilities, but they plan to federate them with BonFIRE and OpenLab and make it available to the broader audience.
- CONECT is building its real-life experiments within NITOS and the OpenAirInterface, testbeds and plans for federation with the OneLab2/Openlab FIRE infrastructure.

- SCAMPI will also initially build its own facility, but in the medium-long term the plan is to federate and deploy with SmartSantander and the NITOS testbed within OneLab.

Summary of facility usage:

- ***OneLab2/PlanetLab/NITOS***: selected by 7 STREP projects, namely RESUME-NET, Nanodatacenters, EULER, NOVI, SCAMPI and CONECT, which has selected OneLab2 for federation with its own built facilities. In particular, SCAMPI selected the NITOS platform.
- ***FEDERICA***: selected by 2 STREP projects, namely PERIMETER and NOVI.
- ***WISEBED***: selected by3 STREP projects, namely HOBNET,OPNEX and SPITFIRE.
- ***SmartSantander***: selected by 2 STREP projects, namely SCAMPI and SPITFIRE.
- ***PanLab2-PII***: selected by Self-NET and VITAL++, which federated through the Teagle tool its own facilities.
- ***BonFIRE***: selected by LAWA.
- ***OFELIA***: selected by CONVERGENCE.

7.3.3. Involvement of end users

The involvement of end users in the experimental part of the projects' work has been adopted quite systematically, according to the answers provided to the FIRE STATION questionnaires. As a matter of fact, ***from the seventeen projects that provided feedback, nine of them confirmed direct involvement of users***.

However, the involvement of end users, especially in big numbers and outside of the usual campus/universities/Labs communities can become quite challenging as the access to the facilities (either technical or geographical) is often difficult and support instruments, like tutorial, examples, GUIs, etc. are in large part missing or rather basic.

The PERIMETER project adopted the LivingLab approach for a study that involved 30 users who completed a task list that allowed the assessment of the projects results. What emerged is that it is particularly hard to conduct a living lab study within a limited time-frame.

The various projects that did not directly involve end users are the ones that deal with more technical back-end aspects mostly in the networking area, e.g. BGP algorithms, routing, network resource optimization etc. for which validation of the work has not required the direct intervention of end users in the value chain.

7.3.4. Collaborations

Most of the collaborations that took place are:

- Between FIRE STREPs and specific FIRE IPs through the selection and deployment of the facilities the IPs offer(ed) and the STREPs adopted (or plan to adopt) for their experimental work.

- Mostly "exclusive", meaning one STREP interacted in most of the cases basically only with one and only one IP. This can be related to the fact that STREPs usually have rather limited amount of resources and rather short duration to explore more than one experimental context effectively. However, ***for the Call 5 projects one can notice that almost all STREPs have already started exploring synergies and deployment of facilities of two distinct IPs***.
- Almost never between FIRE STREPs: this could be related to the lack of proper communication means, but also because the different FIRE STREPs covered basically different areas of expertise and work. Nevertheless, very few synergies have been set up between complementary projects, like for instance between ECODE and Self-NET, where it is felt that concrete collaborations would have been very interesting and possibly productive.
    - As mentioned, potential synergies between ECODE and Self-NET have been identified, but not explored.
    - Potential synergies between Self-NET and RESUME-NET have been discussed through a dedicated workshop.
- Some of the FIRE STREPs established collaborations with other EC-funded projects that do not necessarily belong to the FIRE bulk, but they are related from an R&D point of view. In several cases this liaison seems to have been established by project partners that belong to different consortia.
- STREP projects from outside the FIRE are beginning to explore the opportunity of exploiting the FIRE facilities. An example is given by the ICT project COAST that contacted the FIRE office for information about appropriate FIRE testbeds for their content management experiments.

7.4. Drawing some conclusions about the FIRE STREPs analysis

The previous version of the FIREworks Portfolio Analysis [1] ***highlighted a rather poor usage of the FIRE experimental facilities from the Call 2 STREPs, which to some extent is also re-confirmed by this more recent analysis*** and for which several reasons have been identified:

- For some of the STREPs, in particular the Call 2 ones, the available FIRE facilities were not always able to meet the technical requirements of the planned experiments.
- Easier access to local (not publicly offered facilities) testbeds was available and more specifically tailored to the foreseen experiments.
- Access to FIRE facilities was not always available at the time when the Call 2 STREPs in particular were entering their experimental work phase.

However, ***this situation seems to have changed for the Call 5 projects, which appear to have planned a more systematic and active uptake and deployment of existing FIRE facilities***. This is considered to stem from several main factors:

- The level of maturity of the FIRE facilities has improved as well as their visibility within the overall European R&D landscape.

- Some of the Call 5 projects are somehow follow-ups of previous Call 2 STREPs, which means there is now increased experience and know-how of the overall FIRE context, including available experimental means.
- The market-pull for advanced and improved Future Internet technologies has increased over the last few years, while most of the R&D investments have been shrinking [4]. Therefore, the availability of openly accessible and large-scale experimental testbeds becomes more and more crucial in today's critical economic situation for both academic and corporate (industrial) R&D labs.
- It has been noticed that for Call 8, STREP proposals must include a commitment from at least one IP facility to support their experiments. This shall further promote the uptake and usage of the FIRE facilities.

7.4.1. Lessons learnt about FIRE Facilities deployment

On the negative side:

- Difficult access to the facilities (either technical or geographical) and therefore difficult to involve end users, also because of poor support instruments (service support, tutorial, examples, GUIs, etc).
- Not always possible to ensure full control of experiments and in particular:
    - The reproducibility of repeatable results.
    - Reconfiguration of components.
    - Poor interfaces.
- Information about data traffic is poor.
- To collect results of experiments can become quite complicated and time-consuming.
- Lack of support for experiments on autonomic communications.
- Unclear timeline on the availability and maintenance of the FIRE facilities (sustainability issue).

On the positive side:

The FIRE facilities make it possible to:

- Have a greater diversity of technologies and infrastructures.
- Make experiments on a larger scale.
- Obtain more advanced experimentation results, based on multiple metrics.
- Better justify the research results, since the experiments are performed in "close-to-real-life" conditions.
- Gain technical know-how about an increased number of technologies and equipments.
- Reuse the same technologies and resources across several projects allows return on investment to be maximized.

- They also recognise that testbeds are in principle living beyond a single STREP project lifetime as they remain to be available to the broader audience for follow-up work.

Recommendations:

- To have a "drop box-like" service on PlanetLab that would allow collecting the results as easy as on a local testbed.
- Make sure future calls in the FIRE direction better target the need of providing more intuitive and user-friendly means to final end-users for effective deployment and uptake of services.
- Documentation and support of the facilities is guaranteed on a more stable and continuous way.
- At proposal submission time, FIRE STREPs should show evidence that they already know and have considered which FIRE facilities better fit their project objectives.
- The EC shall continue to promote at a broad and international level the availability of experimental FIRE facilities to improve their visibility and facilitate their uptake. This could of course be done through dedicated Support Action projects.

8. The FIRE Coordination and Support Actions (CSAs)

Within the FIRE landscape, besides the IP and STREP projects, there are also dedicated Coordination and Support Action projects (CSAs), aiming at facilitating coordination across the R&D projects and relevant pan-European initiatives.

Below, we provide an overview of the FIRE CSAs.

FIRE STATION – http://www.ict-firestation.eu

FIRE STATION is a 3-year project which started in June 2010. It provides the globally known FIRE initiative with an active support hub that guides and coordinates the demands and requirements of experimentation in the context of future networks and services.

The heterogeneous and modular field of the Future Internet Research and Experimentation (FIRE) Initiative, with its national and international stakeholder groups requires information sharing, cohesion, community building and a single point of contact to coordinate and promote the approach with respect to the following requirements:

- Facilities need synchronisation, resource optimisation, and collaboration with one another in order to offer customers the best possible service and ensure their sustainability beyond project lifetimes.
- Researchers need correct and timely knowledge about the available resources, along with easy access and highly usable and appropriate tools to run and monitor their experiments.

The purpose of FIRE STATION is therefore:

- To connect stakeholders to allow for the most efficient bilateral (and multilateral when and if appropriate) collaboration, reduce duplication of work, share experiences and best practices and work for the future of experimental research.
- To move the FIRE Facility towards a more customer-driven, dynamic, effective, sustainable, easy-to-access and easy-to-use experimental platform.
- To intensify the collaboration amongst the FIRE Community, for example, all FIRE projects and stakeholders, such as other testing facilities and their customers (projects, researchers from academy and industry).

Actions and implementations

The FIRE STATION project (and its predecessor, the FIREworks project[19]) is responsible for the FIRE Web portal that aims at making the benefits and achievements of the FIRE projects known to all relevant target groups (FP7 IPs, STREPs, SAs, KICs, EIT, industry, Celtic, Eureka), including policymakers, NRENs, end-users and other actors in and around the European and international Future Internet environment.

FIRE STATION also operates a FIRE Office and FIRE Architecture Board. The FIRE Office serves as the single point of contact and as mediator when it comes to looking for the right experimental resources or new customers for the facilities.

The FIRE Architecture Board involves all facility builders in the FIRE initiative to jointly decide on the strategy and means to coordinate and facilitate the development of FIRE facility offerings to support the evolving needs of the customers. FIRE STATION increases the global collaboration between relevant stakeholders, promotes an experimentally driven approach in Future Internet research and intensifies the usage of experimental facilities, ultimately speeding up the development process of new systems and services. FIRE STATION also actively promotes experimentally driven research and the usage of experimental test facilities in Europe and beyond.

PARADISO 2 - http://www.paradiso-fp7.eu

Designing the Internet of the future is not merely a question of technological developments, but one of societal issues. Even though it is impossible to predict the future, it does not mean that we cannot prepare for it.

The PARADISO project, defined in early 2007 and selected at ICT Call2, has already advocated and explored this probable paradigm shift in global societal developments through the 'PARADISO reference document'. It thus appeared quite visionary and timely when the economic and financial crisis expanded during the year 2008, which is probably one of the reasons for the significant project impact.

The PARADISO-2 project, selected through the ICT Call 5, is an 18-month project which started in April 2010, and which contributes to build on the strengths and achievements of the PARADISO project by aiming to be:

- More focused: while 'PARADISO1' was encompassing the ICT sector at large, PARADISO-2 focuses specifically on the Future Internet, and

[19] http://www.ict-fireworks.eu/

- More concrete: the two facets of the problem ('how should societies evolve' and 'which FI can be suited to these societies') will be further explored, including functional and technological specification of the envisioned Future Internet, and recommendations concerning research to be developed in the framework of FP8.

The PARADISO-2 project counts on the involvement of a multidisciplinary high-level expert panel composed of around 25 representatives of leading research institutions, companies and NGOs from Europe and all over the world.

In this perspective PARADISO-2 offers a great opportunity to influence and direct FIRE in accordance with the potential social and economic effects new developments in Future Internet may have, particularly with the launch of the PPP.

All project activities are open to any organisations interested in the project approach and expected outputs. External contributions are welcomed, particularly through the project website.

The main findings, based on various inputs including those from any interested stakeholders contributing through the PARADISO social networks, and their participation in project events, are detailed in successive versions of the "PARADISO Reference Document". This document includes recommendations to the European Commission on research topics to be addressed in future EU-funded programmes and suggests in particular to "explore the Internet at its limits". The latest version has been release in July 2011 and presented at the conference organized by PARADISO at the European Commission in Brussels on September 7-9, 2011 on the theme "Internet and societies: new innovation paths". This document version is now available for public consultation until end of November 2011. Its final version will be released end of December 2011.

A PARADISO scientific workshop[20] was held on November 23, 2010 at the European Commission in Brussels on the theme "Understanding the interaction between Internet and societal developments". Over 120 delegates from Europe and other regions of the world (access the final attendee list here) have participated in this event offering an intense programme, including over 30 presentations from high-level experts in societal and Internet developments, and breaks for discussions and networking.

Other dissemination activities included substantial contributions to publications and external events as reported in more details in the PARADISO web pages.

MYFIRE - http://www.my-fire.eu

The challenge for the MyFIRE project, a 2 years project started in June 2010, is to develop the use of experimental facilities in Europe specifically by increasing awareness of testing-related best practices.

It will seek to ensure there is balance between stakeholders' expectations and the collaborative capacities of researchers in order to achieve better experimentation and develop sustainable testing methodologies while contributing to European standards development.

[20] More details about this workshop can be found at: http://paradiso-fp7.eu/events/2010-workshop/

The framework will be developed through an open dialogue between the ICT networking research communities and experts from the key fields in sociology, public policy and economics.

In order to succeed in these challenges, MyFIRE's mission is centred on four key activities.

- To identify the main issues and needs in the testbeds approach:
  - Researchers' and users' needs for the experimental facility.
  - Standardisation.
  - Exploitation process of Future Internet research.
  - Economic data regarding testbed sustainability.
- To define the testing methodologies used by the projects in Europe and also by international testing facilities so that the best practices model can be analysed and documented, leading to improved design, set-up and use of the experimental facilities and standards.
- To build common tools and develop a roadmap to increase the effectiveness of the testing approach.
- To disseminate the results and create a network through a series of workshops in Europe and advanced emerging countries such as Brazil, China, India and Russia. For instance, during the MyFIRE workshop held in Sao Paulo on 13th September 2011, a small contest was organised where Brazilian researchers could present the usage they would like to do of the FIRE facility. The price, that two representatives from two distinct research teams won, is a trip to Poznan to attend the Future Internet week (24th-28th October 2011) to meet European researchers and increase concrete chances of collaborations.

The MyFIRE project contributes towards the methodologies related to efficient testbed design and set-up with an approach that crosses multidisciplinary techniques and research areas.

Making use of known standardised approaches, together with socioeconomic analysis, MyFIRE will provide tools to optimise the design, set-up and usage of FIRE testbeds. This will target the optimisation of investments in FIRE testbeds and he further improvement of well-established testbeds. The efforts made on international collaborations with third countries will allow further expansion of the FIRE community including those countries less involved currently.

The MyFIRE project organised its work and activities within the following four support activities.

- Research and technology: the real needs for testbeds — focuses on the identification of user and industry needs regarding testbeds.
- Standardisation — focuses on the testing methodologies used by the projects in Europe and by international testbeds so that the 'best practices' model can be analysed and documented aiming to improve the design, set up and use of the experimental facility.
- Innovation process and exploitation - covers the process between the scientific proof of concept to the industrial proof of concept. The objective is to optimise the results of the

process research by using best practices and evaluate the socio/economic/environmental impacts.

- Networking with stakeholders and dissemination - aims to increase awareness collaboration with multidisciplinary stakeholders by sharing the roadmap, tools and best practices (organisation of workshops amongst the FIRE stakeholders to present the results and define a common roadmap); and by networking with other similar initiatives (organisation of a series of workshops at European and international level to networking the European FIRE initiative).

FIREBALL - http://www.fireball4smartcities.eu

The main objective of the FIREBALL Coordinating Action, a 2 years project started in May 2010, is to coordinate and align activities in the FIRE and Living Labs communities into a sustainable network of European cities that pave the way for Smart Cities by utilising both facilities and people. This is done by bringing three communities and assets together, the FIRE community, the User- Driven Open Innovation (Living Labs) community, and users in city environments, thus creating a sustainable city-centred network of open user-driven innovation.

To this end the FIREBALL project works on:

- Achieving a European-wide coordination of methodologies, approaches and activities in the domains of FIRE and User-Driven Open Innovation to benefit innovation towards Smart Cities. The coordination is driven by a network of Smart Cities and includes the key constituencies involved in Future Internet innovation (FIRE, User-Driven Open Innovation) to benefit open, sustained and user-driven Future Internet innovation in cities and urban areas, to align and accelerate innovation activities and to exchange knowhow, experiences, information, innovation plans and activities.
- Leveraging European-wide available assets (scientific excellence, technologies, methodologies, tools, experimental facilities, Living Labs, user communities) of the constituencies involved to enable Smart Cities across Europe to explore and exploit the opportunities of the Future Internet in future showcases.
- Ensuring a coordinated development and sharing of best practices and showcases of Future Internet innovation across pilot cities, covering different themes relating to Smart Cities innovation. Within FIREBALL, Smart Cities are considered as the drivers of Future Internet innovation. A core network of cities will engage in practical collaboration to explore the opportunities of the Future Internet and user-driven open innovation environments, underpinning a roadmap and action plan for cities towards Future Internet innovation.

In doing so, FIREBALL provides the opportunity to combine FIRE and Living Labs research communities and assets, and open up a fresh new approach to the coordination of experimental research and open user-driven innovation activities in collaboration with Future Internet research, experimenting in large city environments involving actual citizens.

Beneficiaries will be the FIRE and Living Labs communities, especially projects from the CIP and PPP call for autumn 2010, cities and the public authorities responsible for strategic planning, infrastructure, service delivery, etc., as well as national agencies and actors responsible for developing new R&D programmes, and industry involved in discovering new market opportunities by observing user needs.

Policy makers developing strategies to explore the Future Internet and user-driven innovation for the benefit of social and economic development will fit this group. The FIREBALL concept will make it significantly easier for both individuals and organisations in the private and public sectors to initiate, test and evaluate new and innovative smart services.

To date, Future Internet, Living Labs and cities have always been considered as separate domains of activity. Bringing them together opens up the opportunity to achieve the following results:

- Creation of a European-wide community of Future Internet Innovation constituencies (FIRE, Living Labs, Smart Cities).
- Creation of a common vision and shared agenda by the Future Internet innovation constituencies mentioned.
- Development of showcases to represent innovative uses and future needs of Future Internet in Smart Cities.
- Organisation of a workshop/Conference on Future Internet and User-Driven Open Innovation; identification of common concepts, methodologies, tools and processes to enable the constituencies to work together for Future Internet innovation.
- Definition of processes and arrangements to enable the three constituencies to access, share and use common assets.
- Ensure coordinated development and sharing of best practices of Future Internet innovation in pilot cities and sectors.
- Creation of a Smart City network for Future Internet innovation, based on a core group of advanced cities.
- Development of a roadmap and action plan for exploring Future Internet innovation; Identification of Future Internet pilot areas, and sharing of practices across Cities.

These actions will be ensured through the organisation of a workshop/Conference on Future Internet and User-Driven Open Innovation Identification of common concepts, methodologies, tools and processes to enable the constituencies to work together for Future Internet innovation.

9. Recommendations

The general recommendations in the previous portfolio analysis have been addressed to a large extent in Call 7, since OpenLab has included the University of Essex with its optics lab, Experimedia will create large scale Use Cases involving many users with handheld devices (World Cup media experiment) and Confine will facilitate large scale social networking experiments.

Other specific recommendations have also been addressed:

- The common portal for testbed federation was a recommendation that has evolved into a concrete investigation by several of the IPs, with the intention that they will be able to define a common approach and share the development effort.
- The Open Calls are proving to be a success. The mechanism is a relatively lightweight way to enable even a single user to be funded to make an innovative experiment and to become part of the existing project consortium. Through the Open calls, the facilities can also be enhanced beyond what was foreseen in the original proposal, thereby keeping them attractive for new experiments. The 1st Information Day (for the Open calls from BonFIRE, OFELIA and TEFIS) was co-ordinated by FIRE STATION and proved to be very popular by both potential experimenters and facility providers. 71 proposals were submitted of which about 10 will be retained.
- Sustainability was raised as an important issue in the previous report and has become one of the Architecture Board Working Group topics. Both short-term and longer-term ideas have been identified, taking into account the experience from the GRID community, through a specific Workshop session on the topic in the Ghent FI Week.

New recommendations from this report are:

- To have a "drop box-like" service on PlanetLab that would allow collecting the results as easy as on a local testbed.
- Make sure future calls in the FIRE direction better target the need of providing more intuitive and user-friendly means to final end users for effective deployment and uptake of services.
- Documentation and support of the facilities should be guaranteed on a more stable and continuous way.
- Proposals for FIRE STREPs should show evidence that they already know and have considered which FIRE facilities better fit their project objectives. This ensures the continued use of the facilities and, by including the costs for the usage of the facility in the proposal, can also help towards their sustainability.
- To continue the Open Call mechanism, and even expand it as an excellent way of attracting more innovative experiments to testbeds, throughout their lifetime. By using

these Calls also as a way to fund the facilities for their support, it can also be a way to bring sustainability to the most popular testbeds.

- The EC should find a way to satisfy the growing international interest in Europe's Future Internet activities. FIRE STATION has limited resources to address this rapidly evolving need, but the opportunity should not be missed, whilst Europe is perceived as having a lead. There is a demand for both strategic and technical collaboration, through a flexible scheme, such as joint calls or specific calls for tender. This could be done through a dedicated Support Action project.
- To share experimentation results data. Benchmarking and repeatability of experiments is key to experiments of high quality and scientific impact. This will put requirements on format, storage, and access to experimental input data and results. It also will require comparable methods for measurement. Data archiving will become increasingly important as user projects evolve. Standards and shared tools in this area should be organized. Overall, such efforts should be coordinated by FIRE STATION, better being addressed by a dedicated effort with clear goals and milestones to measure progress.
- FIRE STATION should take the lead in identifying appropriate levels of user support and ensuring that the best practices are shared across the FIRE portfolio.

Missing Elements – for Call 8 and beyond

The projects selected in Call 8 should constitute a further consolidation and integration of the FIRE facilities both in methods of use and allowing experiments that utilize more than one experimental facility.

The smaller experiments supported by the Open Calls, international collaboration and experimenters from other Challenges can help to create a large experimental usage base of the FIRE facilities and a proof of the their usefulness. The ongoing work taking place in the FIRE STATION Architecture Board concerning the incorporation of end users and measurements under real-life conditions, etc. are still in the early stages and should be further elaborated.

Challenges for the FIRE Unit

- Evolution of the FIRE concept into Horizon 2020
- Sustainability of the FIRE facilities
- Collaboration with similar initiatives worldwide and within Europe
- Relationship with the FI-PPP

10. Conclusions

The previous version of the FIRE Portfolio Analysis [1] ***highlighted a rather poor usage of the FIRE experimental facilities from the Call 2 STREPs, which to some extent is also re-confirmed by this more recent analysis*** and for which several reasons have been identified:

- For some of the STREPs, in particular the Call 2 ones, the available FIRE facilities were not always able to meet the technical requirements of the planned experiments.
- Easier access to local (not publicly offered facilities) testbeds was available and more specifically tailored to the foreseen experiments.
- Access to FIRE facilities was not always available at the time when the Call 2 STREPs in particular were entering their experimental work phase.

However, ***this situation seems to have changed for the Call 5 projects, which appear to have planned a more systematic and active uptake and deployment of existing FIRE facilities***. This is considered to stem from several main factors:

- The level of maturity of the FIRE facilities has improved as well as their visibility within the overall European R&D landscape.
- Some of the Call 5 projects are somehow follow-ups of previous Call 2 STREPs, which means there is now increased experience and know-how of the overall FIRE context, including available experimental means.
- The market-pull for advanced and improved Future Internet technologies has increased over the last few years, while most of the R&D investments have been shrinking. The European Commission's 2010 "EU Industrial R&D Investments Scoreboard"[4] shows that *"R&D investment by top EU companies fell by 2.9% in 2009... The fall in R&D investment by leading players in the US, at 5.1%, was twice as sharp as in the EU, but the worldwide reduction was at 1.9%."* Therefore, the availability of openly accessible and large-scale experimental testbeds becomes more and more crucial in today's critical economic situation for both academic and corporate (industrial) R&D labs.
- It has been noticed that for Call 8, STREP proposals must include a commitment from at least one IP facility to support their experiments. This shall further promote the uptake and usage of the FIRE facilities.

The FIRE portfolio of projects is becoming increasingly visible and the facilities are available to experimenters to a much larger extent than was the case at the time of the previous analysis report. The upcoming call for research projects, with the requirement to use the existing FIRE facilities and the future Open Calls, will further increase the opportunity to experience relatively large volumes of experimentation, representing key aspects of the future Internet.

The research that is now possible in FIRE has been extended through Call 5 to many different areas involving, for example, end user communities, large scale distributed software applications and advanced radio technologies.

The expected stronger focus on the federation of facilities has not yet taken place, this is mainly because the research experiments are not sufficiently innovative to require more than an individual testbed, or testbeds within a single IP project. Nevertheless, federation has attracted interest and is included as an important aspect within several FIRE facility projects; particularly those that have performed a degree of federation between heterogeneous testbeds within one technology area (as in BonFIRE, WISEBED and CREW). Some experiments involving federation across facilities (OneLab and FEDERICA) are found in a project like NOVI. However, there is still a need to work on creating more advanced experiments involving more than one of the FIRE experimental facilities. We can see that the motivations for the FIRE facility projects to work together are not necessarily always their first priority which is to meet their own contractual commitments. This situation will have to be managed carefully also in the future.

The experimental work and the use of FIRE facilities by research projects in FIRE has evolved during the year, and the availability of some of the facilities have been sustained, even without funding from FIRE. However, there is still need to extend the usage of the experimental facilities in order to obtain value for money for the investment. The Open Calls have shown that the need for experimentation of moderate size is great not only in academic environments but also in industry. Further support of such experimentation is important.

The previous portfolio analysis identified several issues that were necessary to work on. The FIRE facilities are now dealing with these issues within the context of the FIRE Architecture Board, however the following issues require further attention:

- To implement a "User-facing Clearinghouse" as a way that a facility can be discovered, how a user can be authenticated and how the access to the facility can be defined. This is not only a list of facilities it is also the ability to create an understanding in what way a facility may be used.
- To document the "Terms and Conditions" for using a facility, including the cost, the acceptable use policy, frequency and duration of use.
- To clarify "Security and Privacy" issues, in terms of the ability to protect the IPR of both the experimenter and the facility provider. It also includes methods to protect the privacy of the traffic.
- "Operational and Research Monitoring" functions and tools to start, stop and meter experiments and other operational aspects of experiment control.
- Develop a process to "Define, Simulate and Control Experiments", in which experiments can be created and supported.
- Stimulate the "User Support to Grow the Market for Testbeds" through making public what facilities are available and when, announce federated facilities and promote the use of facilities for user groups outside of FIRE.
- "Deployment of Resources" in order to create a virtual testbed for an experiment where both physical and software resources will be bound to an experiment.

All these efforts show that considerable work still has to be performed to create a workable and large scale experimental facility.

11. References

[1] FIRE Portfolio Analysis, Scott Kirkpatrick and Jerker Wilander (editors, Jaques Magen, Dirk Trossen – FIREworks Support Action

[2] Experimentally driven research - white paper Version 1, April 2010 - On the existence of experimentally-driven research methodology. Anastasius Gavras - FIREworks Support Action

[3] A. Bavier, N. Feamster, M. Huang, L. Peterson, and J. Rexford. In vini veritas: Realistic and controlled network experimentation. In ACM Sigcomm, 2006

[4] The 2010 EU Industrial R&D Investment Scoreboard, Joint Research Centre European Union 2010, http://iri.jrc.ec.europa.eu/research/docs/2010/SB2010_final_report.pdf