FIRE Portfolio Analysis

Scott Kirkpatrick and Jerker Wilander (editors), Jacques Magen, Dirk Trossen
carried out under Support Action FIREworks

Executive Summary

With the FIRE portfolio the opportunity now exists to support new classes of users and experiments combining heterogeneous technologies, all representing key aspects of the future Internet. The research done already in FIRE and starting in the current Call compares well with the best work in the world, including the highly visible GENI portfolio in the US and national projects in the Far East. However, the FIRE projects to date have not produced a federated testbed for European researchers, overcoming the considerable obstacles of complexity and unfamiliarity that are faced when trying to explore the effects of new applications that bring future users the increasing power of the future Internet, from its edge to its core with computation and storage embedded everywhere. Indeed, federation has not yet been considered as an issue they would all have to deal with. Individual projects, for the most part, have stressed creating improvements in one technology, or two technologies that they can connect directly.

FIRE differs from GENI in that FIRE emphasizes the value as seen by an end-user, at the edge of the Internet, with its applications or services, while GENI is focusing more on basic infrastructure technologies.

The usage of individual FIRE facilities has not been extensive among FIRE call 2 projects. The facilities that have been used to a reasonable level are OneLab's PlanetLab Europe and Federica (not only FIRE internal usage). The reason that Wisebed and PanLab have not been used is that they were not made available outside their consortia. For those facilities now being and becoming available further usage will be discouraged by uncertain sustainability. In addition the requirement for stronger user support and better dissemination of availability is important.

Existing FIRE facilities could have been used by some of the research projects, but the choice was to either use other existing test beds or building a small test bed for internal experimental use. Further dissemination of FIRE facilities should have prevented this.

Several FIRE supported projects show that there is a need for an emulation experimental facility. This is emphasized by the use and planned use of the "IBBT iLab.t" (exemplified by OFELIA, CREW, ECODE and Euler).

A major step forward on the understanding of what can be performed on existing experimental facilities is expected from the FIRE day in Gent in December 2010, when Call 2 projects plan to present their major use cases as showcases.

Some specific actions needed to improve the FIRE experimental facilities are:

- FIRE should address the integration the top-down federation in FIRE projects (mainly through PII/TEAGLE) and the bottom up federation as represented by SFA (in FIRE + GENI) shows approaches with different properties.
- The issue of how the NREN and GEANT infrastructure can be used in FIRE should finally be discussed and concluded. This issue has been discussed since the beginning of the FIRE initiative. Experience in ONELAB, FEDERICA, and OFELIA can serve as a starting point.
- Sustainability of each FIRE experimental facility should be studied in the context of a sustainable FIRE federated facility. It is important that the architecture board during the coming months investigate proper business models using the model of contributing and sharing resources and more commercial models as a starting point. Consideration in how far networking costs must be paid for by the user should also be considered taking into account the FEDERICA experiences.
- FIREStation should take the lead in identifying appropriate levels of user support and ensuring that the best practices are shared across the FIRE portfolio.
- Benchmarking and repeatability of experiments is important when conducting experiments of high quality and scientific impact. This will put requirements on sharing/interconnecting experimental data across experiments and supporting comparable measurement capabilities, not only resources. Standards and shared tools in this area should be organized once the shape of the experiments performed under the open calls is visible.
- Last but not least, we believe that this coordination overall needs to be streamlined in the light of the upcoming PPP and the Calls 7 and 8. Less is more! One opportunity in this context is the Call for a FIRE Federation IP under Call 8.

Introduction

Within the context of the European FIRE initiative and the 7th Framework Program, 12 research and facility projects have now been running for two years[1]. A sketch of how these projects could be extended into a FIRE facility with

[1] This set of projects is referred to as "Call 2 projects" in this document.

continuing impact was developed by an expert working group as part of the inputs to the FIRE portion of a new call. This report is entitled “Towards a collaboration and high-level federation structure for the FIRE facility” – often also referred to as the "FIRE Wise Men Report" (http://www.ict-fireworks.eu/publications/papers.html).

Now that the new call has been completed and another set of projects has been launched[2], it is time to analyze the status and outcomes of the running projects, and to include as well in the analysis the new projects in order to provide an overall view of the landscape and status of the FIRE initiative.

Therefore the objective of this FIREWorks portfolio analysis may be summarized as follow:

- Extend and deepen the first “Wise Men” report mentioned above;
- Provide an analysis of running and just starting projects, and show the coverage of the areas and point out strengths and weaknesses. This would be an input both for the orientation of the upcoming FIRE calls for proposals, and for the regularly meeting FIRE Architecture Board which is coordinated by the FIREStation Support Action, the successor of the FIREworks Support Action under which this report was commissioned. In addition, we also provide some suggestions to the European Commission unit in charge.

The primary purpose of this analysis is to identify the need and support the creation of federated and heterogeneous test beds. It shows the various features that are required in test beds in order to eventually converge into such a federation, as well as the missing domains that need to be covered so that all relevant areas of ICT are dealt with to build up a full-fledged FIRE federated facility. Equally important, it shows what has been already accomplished with past and running FIRE projects, and what is likely to be done in the new Call 5 projects. Then, it provides hints on how projects could collaborate, and focuses on three critical issues: what are the most promising current approaches to federation; how to attract new users; and how to get sustainable experimental facilities.

Methodology for the analysis.

It is worth mentioning that:

- Information on new Call 5 projects is based on telephone interviews and current Descriptions of Work which include the objectives and the activities which were planned, various presentations, and outcomes of interviews conducted by the experts groups. Obviously the actual results of these projects will not be known until they are completed; this is at the same time a risk and an opportunity. It is a risk because we had to consider that the

[2] This set of projects is referred to as “Call 5 projects” in this document.

outcomes that are described will happen; it is an opportunity because potential deviations improving on the original work plans can still be considered.

- Information on past and running Call 2 projects is based on documents that are publicly available, as well as some (limited) additional information that was communicated to the experts group.

Indeed, for the analysis groups, finding out sufficient details to understand each of the IPs in Call 5 was not simple. The project summaries and initial presentations provided at the FIRE workshop in Brussels on March 9 gave us a start. Each IP and most of the STREPs responded to a set of preliminary questions which we addressed to them a couple of weeks before the meeting. We have discussed the projects in greater depth during telcos with partners of the IPs. Finally, all of the Call 5 IP projects have provided us their DOWs in a nearly-final state, and this has proved essential for our understanding. In requesting this information, we have assured each project that the information is to be used only for the purposes of this study, as a part of FIREworks. FIREstation will continue to need planning information from ongoing projects. Such information is confidential, under typical consortium agreements, so the ground rules for handling it in FIREstation should be made as clear as possible to alleviate doubts in the future. The larger Call 2 projects have in some cases been providing similar information; in other cases available information on web-sites and background information from reviews and participation in projects have been used. We summarize our observations on the individual Call 5 IPs in Appendix B.

In most cases STREPs did not provide information until the conclusion of our study. This has of course hampered the analysis. It is suggested that the projects involved in FIRE should provide FIREstation with detailed plans on experimentation and use of experimental facilities in order to both plan and support deployment of resources and present a common view of FIRE to outside research both in other FP7 areas and internationally.

Many Call 2 projects now have project web sites with information almost a year old. Dissemination of the projects' results and news should be improved in order to spread plans and results. Scientific publications do not provide enough guidance for the common use of experimental research facilities.

All projects mentioned herein were provided plenty of opportunity to contribute their comments and suggestions, which were largely taken into account.

The FIRE Facility Projects - coverage, collaboration, federation, sustainability, use

Overview

With the FIRE portfolio that emerges in Call 5, the opportunity now exists to support new classes of users and experiments combining heterogeneous technologies, all representing key aspects of the future Internet. The research done already in FIRE and starting in the current Call compares well with the best work in the world, including the highly visible GENI portfolio in the US and national projects in the Far East. However, the FIRE projects to date have not produced a federated testbed for European researchers, overcoming the considerable obstacles of complexity and unfamiliarity that are faced when trying to explore the effects of new applications that bring future users the increasing power of the future Internet, from its edge to its core with computation and storage embedded everywhere. Indeed, federation had not yet been considered as an issue they would all have to deal with. Individual projects, for the most part, have stressed creating improvements in one technology, or two technologies that they can connect directly. Some of the Call 5 projects add new technologies to the mix, while others bring new attention to the building blocks of real applications, and exposure to end users. All of the Call 5 integrated projects (IPs) have made a commitment to cooperate with other parts of FIRE, devoting one part of their budgets to shared coordinated development in support of this goal, and another part to open calls for use of their testbeds by researchers outside their projects. This was explicitly required in the call. Still, the IPs has very different notions of users, of use cases, and of the range of collaborations that can be expected to augment the value of the technologies they bring to FIRE. Coordination of these efforts through architecture and careful selection of the timeliest areas to emphasize can pay large benefits in the scope of the problems FIRE can address, in the scalability of the solutions, and in sustainability of the FIRE federated testbed. But it is also important not to impose an abstract goal of federation when there are no use cases for combining two or more specific testbeds. We can see that the motivations for the FIRE Call 5 projects to work together are not necessarily always strong as their first priority seems to support their own specific scope. This situation will have to be managed carefully. One purpose of this report is to identify the best opportunities or synergies within FIRE in the next few years.

FIRE differs from GENI in that FIRE emphasizes the value as seen by an end-user, at the edge of the Internet, with its applications or services, while GENI is focusing more on basic infrastructure technologies. FIRE has less internal overlap between major projects (OFELIA does OpenFlow across five European centers in collaboration with the US, CREW explores cognitive radio, offering four modes in a real-life home/office environment, SmartSantander is deploying a city-scale IoT infrastructure for large scale experimentation with novel IoT protocols). The two top-down IPs in Call 5, BonFIRE and TEFIS, have defined a space of services which includes the end user, but makes less contact with details of the networking infrastructure. For example, BonFIRE implements a cloud facility offered as Infrastructure as a Service. They also aim at seeking the high ground above the technical infrastructure that GENI and some of our technology-intense projects are providing. This means they must eventually federate with distributed resources on very large scales, and probably look into supporting the SFA (the Slice-Based Federation Architecture that has grown out of GENI and that is the subject of joint implementation work by OneLab and PlanetLab) protocols underneath them. BonFIRE is adopting the OCCI (Open Cloud Computing Interface) specification in its architecture in order to enable the federation with any kind of Cloud provider that exposes this interface. While this supports and extends the stated objectives of BonFIRE, it does not advance the technology of federation much, only addressing very homogeneous resources. TEFIS, which is more of an architecture project, might address dispatching resources described by the SFA protocols or being able to act as an SFA clearing house with other peers, not being the central source of control. We believe that addressing hierarchical control as done in SFA will reduce the risk of scaling problems, and lead to an opportunity to affect a rapidly advancing international standard.

Short summary of the facility projects and their cross connections.

The facility projects in FIRE are now becoming a large set with different properties and also with co-usage and usage of facilities external to FIRE (the facilities in the middle could). The figure 1 below attempts to describe most of those connections and some of the major external facilities. The dotted lines show that decisions are still not taken but discussed. Some facilities available within the projects are not shown (too much to fit in one page). The federation of several facilities will occur and even though projects like BonFIRE and TEFIS show many similarities they have different focus. OneLab has federated PlanetLab, NITOS and test beds mainly working on measurement aspects. Further federation is planned.

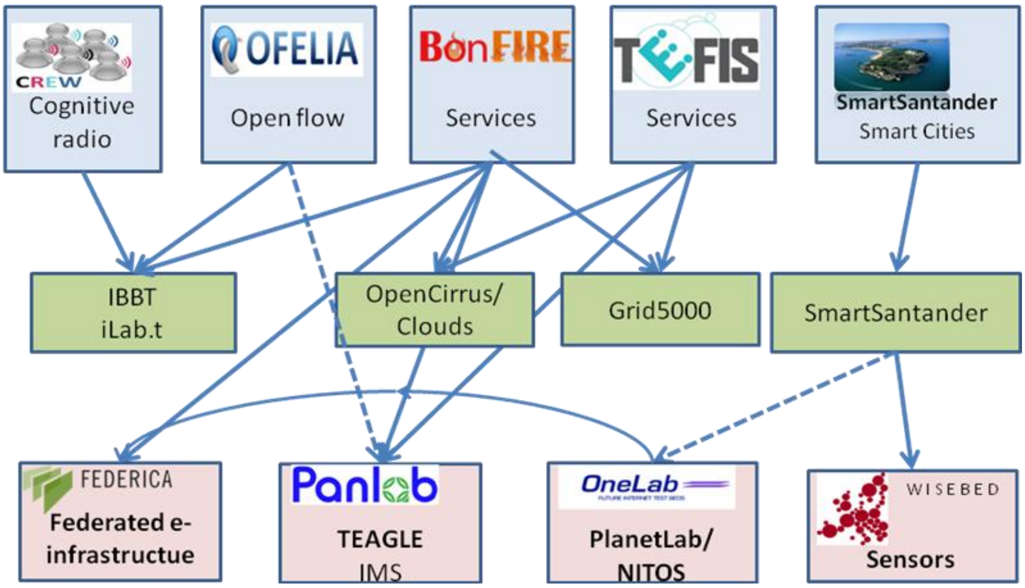

Figure 1. – Experimental facilities in FIRE.

Federation issues

Examples of resource federation already exist, such as PlanetLab, which allows users of two separately managed distributed computing testbeds to have access to each other's facilities under academic fully open ground rules, constrained only by acceptable use policies.

OneLab has extended their PlanetLab Europe (PLE) facilities, already federated with PlanetLab in the United States, by federating PLE with NITOS (a wireless testbed in an instrumented building at the University of Thessaly in Greece) and two active measurement facilities: DIMES (over a thousand virtual "agents" around the world) and ETOMIC (a dozen hardware monitoring sites in Europe). In all of these, user measurements are possible, while in DIMES and ETOMIC there are virtual observatories that provide archived historical data. Additional federation efforts are in progress (e.g., G-Lab), but these four facilities have been open to public users in a combined OneLab experimental facility during the past year. OneLab user documentation exists for PLE on the PlanetLab Europe website, with links to instructional videos that are hosted on Dailymotion and YouTube, and the other OneLab testbeds have documentation on their websites. Project team members have given tutorial workshops both internally to FIRE and at major conferences. An operations team contributes to the development of federation code (e.g., writing SFA components), manages user contracts, runs the servers that host PLE, and responds to the inevitable problems and complaints that arise in a distributed system. As OneLab2's funding ends at YE2010, the facility will remain in service. For the near term it is funded by some smaller grants and national funding.

There is an important difference of approaches between SFA and TEAGLE of TEFIS/PanLab/VITAL++. The latter have centralized the control of the resources, which they federate in a single clearinghouse. In contrast to this approach, SFA defines independent clearinghouses, each of which can allocate and manage resources, and in which identity, authorization and privacy policies are carried out. SFA's clearinghouses can be federated, giving SFA the potential to manage much larger scales. It will thus be possible to include commercial test beds controlled by their own clearinghouses, which can impose localized access and usage policies. SFA is very low level, and does not currently provide the high-level user interfaces offered by TEAGLE in TEFIS/PanLab/VITAL++. Hence, SFA borrows concepts of the current Internet where domains are individually optimized and managed while being interconnected through a distributed mechanism. The Teagle approach, on the other hand, adopts a more centralized telecoms approach with a tight control and governance view of the world. It remains to be seen which approach lends itself to a greater flexibility and acceptance in the research community – only real experiments can shed light on this question! It is also possible that each model will suit different sets of users with different sets of requirements.

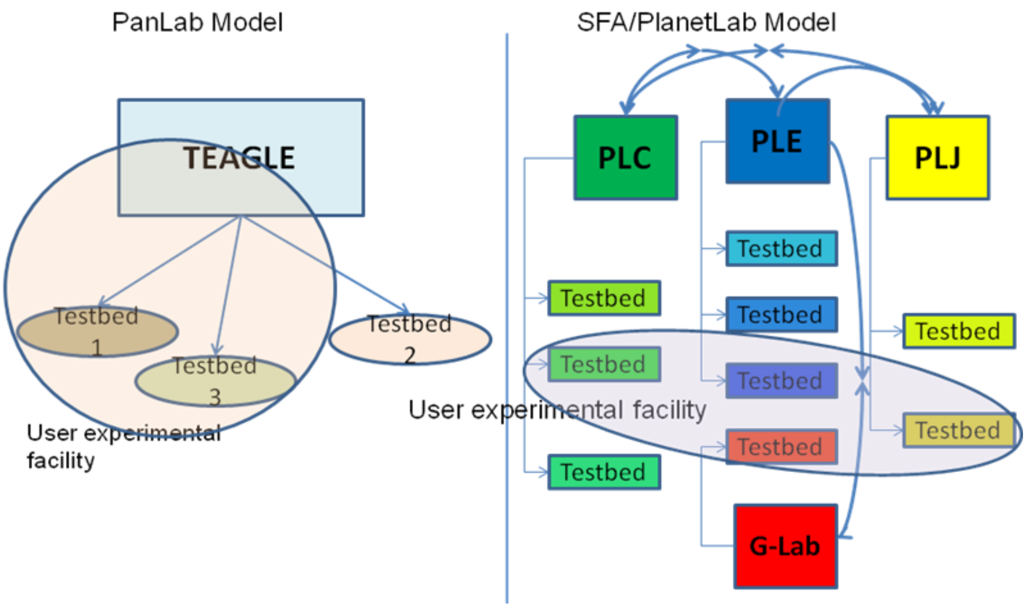

Figure 2 – Teagle approach (centralized) vs. SFA (recursive, hierarchical)

FIRE experimenters in the second round of facility projects will have more complicated requests than earlier. Federating heterogeneous test facilities while introducing new technologies, while keeping it all simple for the sake of outside users, will impose a greater need of collaboration between projects. An experimenter will need support in order to be able to work and this will become more paramount when further federation of facilities will occur. Of the Call 5 IPs, BonFIRE through previous experience of several partners in running open calls in BEinGRID (now represented in BonFIRE) has the most experience with user support and analysis of users' experiments. This knowledge could benefit the whole FIRE portfolio. The remaining IPs will need assistance in managing the open call process, assisting users to overcome usability barriers, and monitoring the experiments for subsequent analysis. OneLab2's experience with user support should also prove relevant to the Call 5 projects.

Test bed issues

General features of testbeds

An effective testbed should be used by both experts in the technology which it exposes and those who just want to integrate that technology with higher or lower levels in an application stack. It must be robust enough for reliable, reproducible tests by the first category of users as well as for durable prototypes needed to expose new functions to end users.

We have constructed a diagram (Figure 3) which captures the set of issues which must be dealt with by a single testbed or (by recursion) by a federation of testbeds. This also provides a simple way of categorizing the issues that must be addressed to support real external users:

- "User facing clearinghouse" is the way a facility is discovered, how you can be authenticated as a user 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.
- "Terms & conditions" includes the cost to use a facility, the acceptable use policy, frequency and duration of use.
- "Security & Privacy" defines both the ability to protect the IPR of the experimenter and the facility provider. It also includes methods to protect privacy of a traffic data.
- "Operational & Research Monitoring" are the functions to start, stop and meter experiments and other operational aspects of experiment control.
- "Define, simulate and control experiments" is the process of creating and supporting the experimental development process.
- "User support to grow the market for test beds" is the process to make public what facilities are available and when, announce federated facilities and promote use of facilities for other user groups.

- "Deployment of resources" is the process of creating the virtual test bed for an experiment where both physical and software resources will be bound to an experiment.

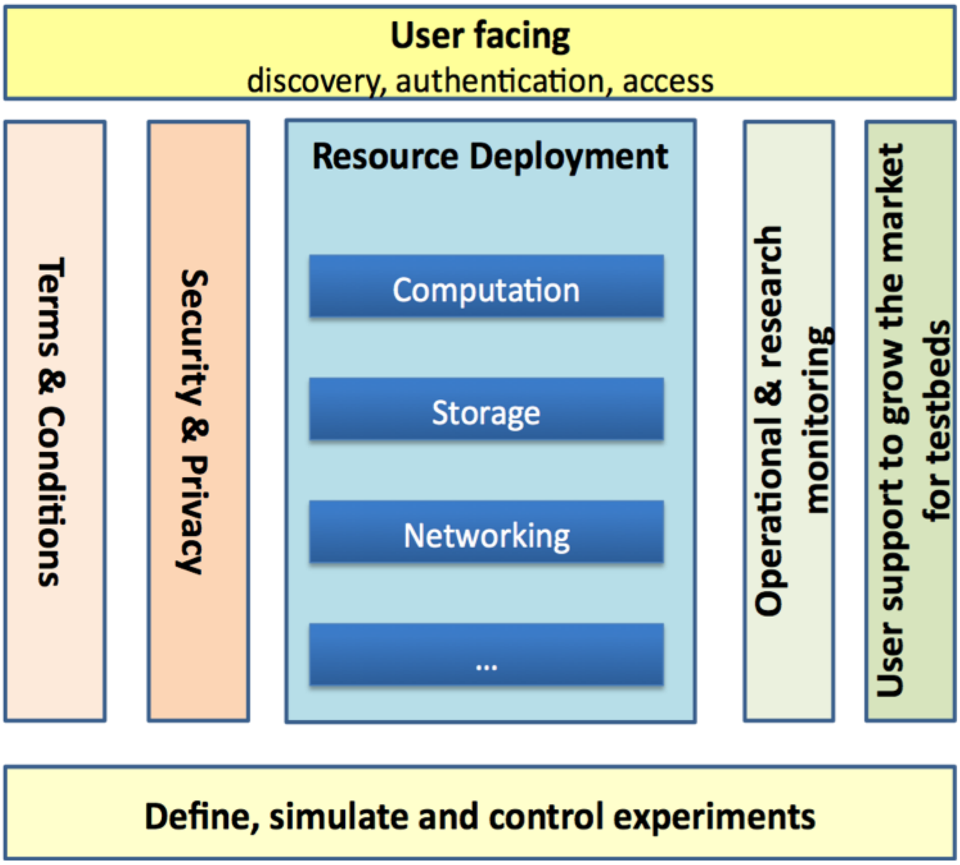

Figure 3 – FIRE experimenter requirements on facilities

Note that this does not mean that it is necessary that all test beds investigate all these features; they should however take advantage of the work which is done in other test beds and/or in common so that all features are available in their test bed. These functions are required for a successful FIRE federated facility.

<u>Current situation</u>

In Call 2 some of these issues have been addressed in the projects, but mostly not deployed to the extent that a user community outside the project itself is able to use the facility. The discovery utility in PanLab and the Teagle brokering tool used for deployment of experiments is still not publically available. The sliced-based facility architecture as defined in PlanetLab and used by OneLab for the definition of its services is a technology used by many today. However, the extensions to further shareable facilities have been in use only a short period of time.

<u>Situation in the future with Call 5 projects</u>

None of the elements in Figure 3 is directly identifiable as the "portal". Almost every project, IP or STREP, has promised to deliver a "portal"; but on closer inspection, most of these are just websites that will allow advertising the project and disseminating its results. A smaller set plans to expose various combinations of the elementary services in Fig. 1 to the external user, but the combinations are each different. For example, the portal in SmartSantander is initially a simple tool, providing information about the facilities offered in the project. The project then intends to extend the interface to support provisioning experiments and obtaining user experience feedback. Bonfire's or TEFIS's public interface (they are considering adopting or extending Teagle), is expected to be a more advanced tool covering most of the areas described in figure 3. 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 would consist of all the boxes in Fig. 2 with the possible exception of the bottom 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 & Conditions and Security (left columns.

With such varying definitions of what a "portal" is and what services it provides, the FIREstation architecture board faces a major challenge and an opportunity in trying to provide some uniformity in external presentation of the FIRE projects as well as encouraging use of common components and reduction of duplicate efforts in portal-building.

The terms and conditions of usage are still complex issues. The testbeds of PanLabII 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 FIREstation support action can contribute through sharing of approaches to user support and best practices.

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 originally developed for the ORBIT wireless testbed and used by OneLab's NITOS testbed) attempt to provide some form of unification of experimental data description (even to the point of describing experiment settings in a unified way). But such work is not receiving the attention that it deserves, despite its crucial importance for truly sharing knowledge in any experimental facility.

Test beds domain coverage

Experimental facilities have different major goals, meeting the needs of diverse and sometimes non-overlapping sets of users. One set of classifications could be:

- Methods for end user involvement and end user experiments as found in user innovation research (this is not fully exploited yet, discussed in PII, planned in STREPs of Call 5, offered as a planned extension in Smart Santander).
- Application deployment and conformance testing (PanLab/PII, Tefis, BonFIRE, and Smart Santander).
- Core networking research (Federica, OneLab, Ofelia).
- Networking component technology (Wisebed sensors, SmartSantander, OneLab and also available in Crew).
- Aspects of networking and/or to create a facility to be used for networking (DTN as in N4C, P2P as Vital++, cognitive radio as in Crew).
- Architecture studies are a special case, where no test facilities have been made open to other experimenters (exemplified by several Call 2 STREPS).

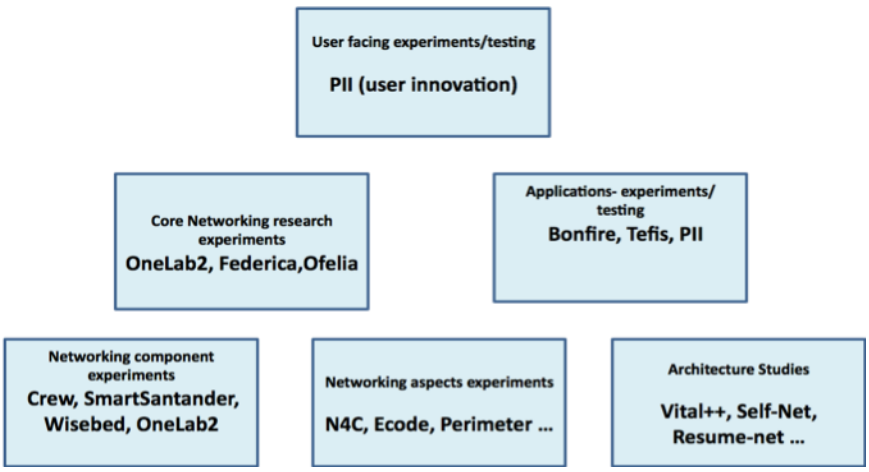

Figure 4 – Domains of experimental facilities

Experiments covering more than one area require collaboration between involved facility owners for its implementation and also for support of the experiment itself. Some Call 5 projects plan to use the open calls to create projects spanning several areas. The open call structure will not encourage such experiments since all proposals will be directed to one facility owner not to a federated facility even though such experiments are a goal in FIRE.

Test beds must share at least some common goals in order to allow profitable federation. The experimental facilities in TEFIS are to be used by services deployed on existing experimental facilities. However, expanding into networking research is not a primary target and is not expected to happen. WISEBED initial experiments are oriented towards verification of networking algorithms and will not have specific support for application development.

Even in the case when a facility is expected to be able to support experiments with a new technology this may not be possible. E.g., data flow experimentation in Ofelia may not be performed on the Federica test bed even though Federica allows lower layer testing, since Federica's routers do not support OpenFlow operations. I.e. novel strategies might not always be supported by the existing facilities.

Federation of some facilities are 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 test bed 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 appear to us to occupy more than 2 boxes in Figure 4. This leads to several problems as we consider how FIRE will complement its portfolio by attracting and supporting experiments that make a 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. Within the open call structure, it is not clear whether the facility owners will select the experiments with an eye to demonstrating the value of their specific technologies, or some not yet determined combination of facility owners, the Architecture Board, outside reviewers, and the program officers in FIRE will select a portfolio of experiments that will achieve the first goal while also moving FIRE closer to a federated testbed facility.

Attracting New Users

The cost and effort to maintain a user community is very high. This will require a user support group that is able to both to provide advice on what experiments can be performed and how the experiment could be adapted to fit the facility.

Few of the FIRE projects launched in Call 2 led to experiments performed by users outside the project, which would have required user support to attract and train new users, to curate their experimental data and metadata, and to create a sustained test facility that will persist beyond the project. WISEBED and ANA, as well as the wireless components are approaching the stage of development at which outside users are first envisioned, and will enter this phase in mid 2010. WISEBED, PII's TEAGLE, and OneLab2 have developed new software components that are being transferred directly to new projects to be launched in Call 5. Several of the Call 5 projects contain testbeds built or started prior to Call 5 with national funding. IBBT, in particular, offers wireless testbeds which should be ready for their first outside user studies early in Call 5's CREW integrated project.

PlanetLab Europe, part of OneLab, is the exception. It is based on a relatively mature technology, servers overseeing hundreds of medium performance clients spread around the world to provide a platform for experiments in distributed computing. It was first developed at Intel Research, Berkeley, and Princeton, and now is developed jointly by teams at Princeton and INRIA, with other international contributions. To support its use, OneLab2 has explicit work packages in operational and educational support of users outside the project.

As a benchmark, we have considered the ways in which users are supported in an established field of science that requires shared use of complex, expensive observational instruments by users at varying levels of experience – radio astronomy. Appendix A contains some details. We believe the astronomers' end-to-end model of supporting outside users as they propose, plan, execute and analyze their experiments is a good one for FIRE to emulate, but it will be costly.

In addition, the issues of federating very different types of networking, ranging from delay-tolerant (DTN) to virtualized optical networks, so that applications which depend on both extremes can be prototyped, may be more severe than the challenges facing astronomy in combining data from different spectral regimes or different instrumental modes. The International Virtual Astronomical Observatory (IVAO) has been less successful and is still considered less critical to good science than effective portals and end-to-end support for new observations at just one telescope. Virtual "networking observatories" are claimed as results of several of the Call 2 and Call 5 projects. These have objectives comparable to astronomy's IVAO – federating access to data from multiple experiments or even from very different types of experiments. The distinction between federating computational and communications resources (which is like federating telescopes) and federating access to archived data resulting from the testbeds must be kept in mind.

Sustainability of Experimental Facilities

Test beds for experimental research on networking solutions and future Internet require long term planning and work. This applies especially when the facility will attract users that have not been involved in the development. A road map of current and some of the future facilities shows that time span and effort is great:

- OneLab project started in 2006, basing its efforts on the work on PlanetLab starting in 2002 (already in 2004 the founders of the European consortium had discussions with the US based PlanetLab group). The first demonstrated federation and collaboration was shown in 2007. Today the number of experiments by European sites is at the order of 100.
- PanLab started as a support action with 4 available test beds in 2006. The work has continued in PII and currently a tool for deploying experiments and test beds into the environment of Panlab is being developed.
- Wisebed has been ongoing since 2008 and has published the availability of its test bed summer 2010.
- The CREW project (Call 5. is not only based on infrastructure from IBBT, but also from TUB (Berlin), TUD (Dresden), and TCD-CTVR (Dublin). The wired portions of the IBBT infrastructure and its emulation environment have come into existence since 2000. The plans for the w-iLab.t (Wireless Lab) started end 2006. The actual deployment was late 2007. The time has been necessary to create a workable environment.

All these efforts show that considerable work has to be performed to create a workable and large scale experimental facility. The case of PlanetLab shows that international collaboration can be of mutual benefit, not only saving cost of development and speeding the ramp-up, but also increasing the level of functionality.

Testbeds, once built, do not last forever. The usability of a facility decreases after a couple of years. This can be exemplified by the number of scientific publications resulting from PlanetLab experiments, which has decreased after the initial 2 years of operation. In order to avoid this situation, continuous development and modernization is required. The initial hype of the facility will create a strong user base, but after a few years many good experiments are already performed leading to less interest. It is also good to remember that a recent study showed that of 800 experiments on PlanetLab approximately 80% are only tryouts ("tire kickers"), which just want to get acquainted with the facility. Only some will return with better experiments, but it is necessary to address this group in order to make it easy for users to reach and understand the facility.

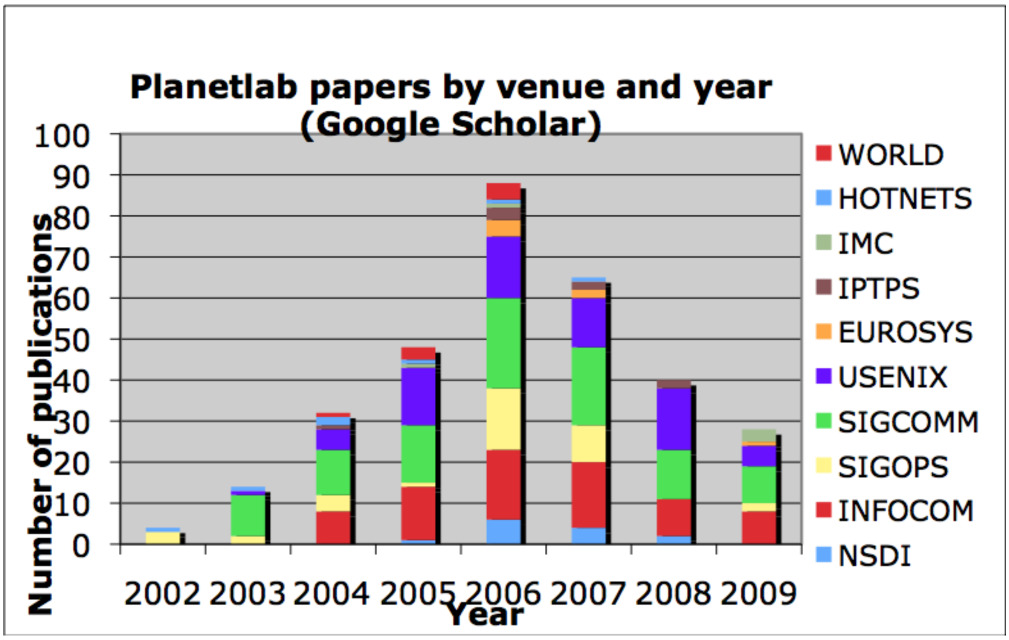

Figure 5 - PlanetLab publication of usage over time.

The sustainability of a facility is today not guaranteed by the facility owner or by the funding organization. This will reduce the trust of the availability of the facility in the potential set of experimenters. This is another cause of non-use by experimenters outside the test bed consortium.

Sustainability of each test bed within a sustainable FIRE federated facility is to be studied. It is however absolutely critical for FIRE as a whole to ensure that its experimental facilities features stay available in future years. This does not mean necessarily that each and every test bed must be kept forever; but this means that at federation level the required features and domains covered will be able to continuously adapt, if not expand, in order to respond to the changing users' demand. A mechanism shall be put in place so that important pillars of the FIRE facilities do not disappear if there is still a demand from the users. This would indeed jeopardize all of FIRE facilities as it would make it much harder to keep a permanent user base.

Support of collaborative effort

In call2 we have seen hardly any collaboration efforts on support tools for experimental facilities; this particular requirement appeared mostly during the course of the projects. This has made sharing of efforts and also sharing and comparing results more difficult. It has been difficult to create a common understanding on what has to be done. In the call 5 list of projects the potential for collaboration is higher, as the requirement has been made clear in the call for proposals; still, there are no established collaboration strategies. Although some 20 % of each IP's funding is reserved for enhancing the effectiveness of FIRE as a whole through collaborative efforts, no project has offered a clear view of what they offer through this channel, or of how they hope to benefit from it. "Wait and see, while using these funds to help our users" seems to be the most prevalent approach.

Much overlap in efforts is found in the proposals. The ambition levels in the different projects are not yet analyzed. This will lead to many issues of coordination when sharing results over project boundaries. This will apply as well to other areas as business policies to resource adaptors (common API). The risk of duplication of work is high, but it is sometimes necessary to pursue more than one approach when alternatives might fulfill different objectives and have different chances of success or satisfy different user sets. Examples of this might be Teagle and SFA. It will take more effort to determine whether they are complementary or should be independent. However, duplication of work should become an exception.

	BONFIRE	CREW	OFELIA	Smart-Santander	TEFIS
Portal/Broker	X	X	X	X	X
Common API		X		X	X
Methodology	X	X	X		
Experiment data storage	X	X	X	X	X
Business/ policies	X	X		X	X

Figure 6 – collaboration opportunities in Call 5 projects

Test bed usage in Call 2 and Call 5

The testbeds developed in Call 2 projects have not yet been used as expected. This may partly be blamed on the non- availability of most of the testbeds, but probably also on the lack of planning on how to attract users to the testbeds developed. In some cases there are testbeds developed to be used by the project only (or at least with unclear goals on how to make the testbed available to external users).

The development of test beds among call 2 projects has been extensive, however very little cross use has been found. In call 5 this is changing but still federation of several test beds is either not there or only outlined in the descriptions of possible work late in the projects.

The 5 IP projects in Call 5 have planned use of many existing testbeds. The major conclusion is that use of the IBBT virtual wall is going to be used by many projects. Much of the use of Call 2 testbed results are planned with a lot of restrictions. Bonfire discusses the use of Federica but a clear description on what this means in actual use is not elaborated. Bonfire and Tefis both plan the use of PanLab results (some portions of or ideas from Teagle) but have not yet planned to work together. SmartSantander sees PlanetLab as an option probably not as a main alternative over standard Internet.

The low usage of PlanetLab in the IPs is improved by STREPS found in FIRE and other areas of Future Internet research. This applies as well to Federica. The use of PanLab facilities is still not possible to evaluate since open use is not available as of today.

FIRE research portfolio

The research projects in FIRE have so far shown little usage of FIRE experimental facilities. The explanation for this is described to some extent in the following report. The analysis is performed using available information on the Call 2 and Call 5 research projects. This might mean that some information has been missed. However, all projects have had the opportunity to provide corrections and extensions to the descriptions provided below.

The research projects in FIRE cover much of the FI area and to do this several experiments are performed. The current FIRE portfolio is described below with a catchword attached to each project.

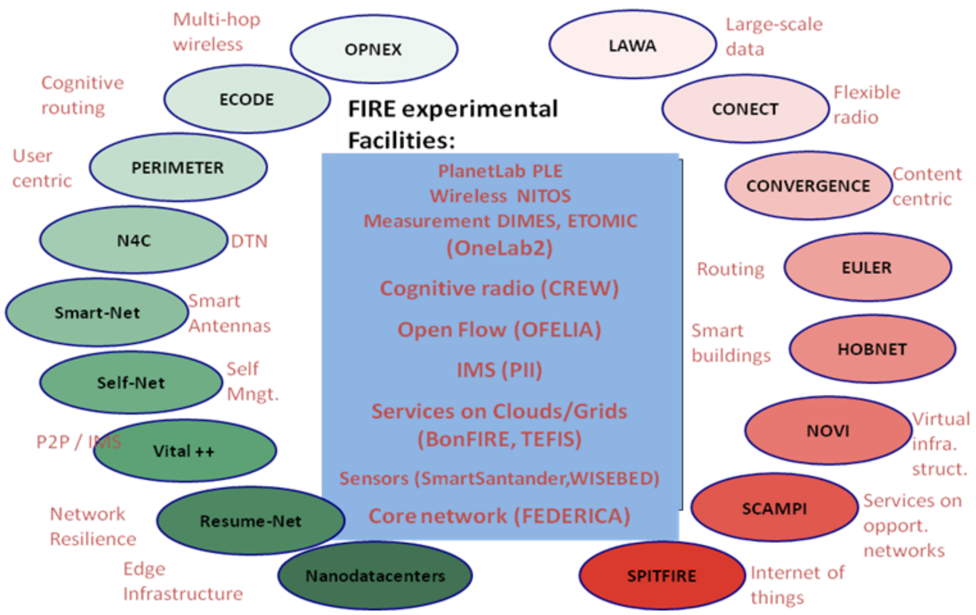

Figure 7 – FIRE portfolio.

The main description of the experimental facility projects are found in part one of this portfolio-analysis. It must be remembered that the actual content of the facilities are not possible to describe in just one word, e.g. OneLab is not only Planetlab Europe but also includes a wireless test bed (NITOS) and two measurement testbeds (ETOMIC and DIMES). PII includes several IMS oriented testbeds, but the main tool, TEAGLE, could be used in more general cases. SmartSantander will integrate not only Wisebed but also other sensor test beds. CREW will integrate several wireless testbeds supporting comparative studies of different wireless technologies.

Since the projects have different maturity the study is divided into two parts; part 1 for the Call 2 projects and part 2 for the Call 5 projects.

Analysis of call 2 STREP projects.

The Call 2 projects are found to the left in figure 7.

The use of FIRE experimental facilities have not expanded as hoped for. This lack of usage has several explanations and this analysis attempts to look at the FIRE call 2 STREPs in order to create an understanding why this is the case.

The current FIRE experimental facility availability is briefly described below:

1. OneLab/PlanetLab facilities have been available before call2 project start. The NITOS wireless test bed as federated with OneLab facilities is working in addition to PlanetLab Europe. The measurements testbeds are also available

since more than a year. OneLab/PlanetLab has at the order of 100 European institutes attached to PlanetLab Europe and has developed use policies and user documentation.
2. The Wisebed sensor facilities will accept external users beginning of September 2010. User documentation is expected to become available at that time. Real use except testing of the facilities has not happened yet.
3. PanLab/PII is still discussing how to create a business offer for users outside the PII consortium, which has lead to that the facilities, are not open for use by all projects. PanLab demonstrated in connection to the Barcelona FIA (July 2010) a set of experiments by PII partners using the TEAGLE tool. This should allow external use outside the consortium in a reasonable future. The research projects Vital++ and Self-Net have shown usage of TEAGLE.
4. Federica has since 2009 a set of users, but will not be able to accept much larger experimental community since the experimental resources are quite limited and the support needed to start experimentation is high. FEDERICA has developed acceptable use polices and user documentation. The number of running experiments is at the order of 20.
5. The IBBT virtual wall and the wireless facilities of IBBT are not part of FIRE facilities, but are used by many projects especially in the call 5 portfolio (all these projects are in some way connected to IBBT Ghent).

The questions this analysis attempts are:

- To what extent has it been possible for the research projects to use the facilities and to what extent have the facilities been used so far by Call 2 research projects?
- Would a higher degree of usage been possible if facilities had been available earlier or if other actions to promote usage had been taken?

This analysis is created as a walkthrough of all call2 projects using available information on their experiments. Due to the lack of available information and difficulties to reach all projects this analysis is certainly still incomplete.

ECODE

ECODE performs research on cognitive routing. One of the use cases in ECODE is directed towards the development of a solution for speeding up the BGP path exploration process and allowing fast network recovery. This can today be experimented using the IBBT Virtual wall (EMULAB copy) supporting programmable control of the network environment. The basic functionality of ECODE is to use a “Machine Learning Engine” that will create the basis for the control of the functionality and of recovery after failures. The IBBT virtual wall allows control of all aspects that may trigger error behavior that ECODE technology wants to manage. However, IBBT virtual wall is not a FIRE facility.

Further experimentation has been planned to use PlanetLab, however this is not documented yet. The use of PlanetLab for cognitive routing might become difficult, due to the fact that it is not possible to control of all routing aspects today in OneLab/PlanetLab. Use of FEDERICA could have been a candidate but this has to be further evaluated. PlanetLab and Emulab (not a FIRE facility) do offer emulation capabilities on an overlay network, but without more detailed knowledge of the experiments planned, it is not possible to say whether they could have been used.

N4C

N4C does experimentation on delay tolerant networking applied in rural areas of Sweden and Slovenia. The goal is both to experiment with the applications and the delay tolerant networking technology. The experiments are based on developed technology within the project and commercially available technology. Initially some of the infrastructure (in northern Sweden) was going to be based on the old analogue infra-structure for mobile telephony (450 MHz) that was going to be converted into a digital system for rural areas. However, this effort went bankrupt and had to be cancelled. The test set up had to be redefined on short notice. This experience illustrates how a combination of DTN with WiMax technology is an option for remote areas.

Existing and planned FIRE facilities would not be able to meet the requirements of the experimenters. An emulation environment like the IBBT Virtual Wall or the NITOS wireless test bed could possibly been used for the initial application development, but not in the real life experimentation (and the real life experimentation is a main goal of the project). PlanetLab would be inappropriate and so would Federica. PanLab/TEAGLE would not help much since the composition of test beds does not seem to be a major problem in N4C. Wisebed facilities would have been out of scope for the project.

SELF-Net

The Self-NET self-management framework is based on the Generic Cognitive Cycle, which consists of the Monitoring, Decision Making and Execution phases. The Network Element Cognitive Manager (NECM) implements the M-D-E cycle at the network element level, whilst the Network Domain Cognitive Manager (NDCM) manages a set of NECMs, implementing sophisticated M-D-E cycle features.

In order to present and test the key functionalities of the Self-NET self-management framework solution, specific network management problems have addressed. Three use cases have been developed and tested, by using the experimentation facilities of the Self-NET network operators (OTE, A1 TA) as well as the facilities of the academia partners:

1) Wireless Networks coverage and capacity optimization
2) Traffic management, re-routing and forwarding to support mobility, QoS and routing adaptation
3) Adaptable routing and mobility management in dynamic self-managed wireless mesh topologies

The goal of one of the use cases is the continuous re-organization of the network

so as to optimally utilise the available network resources under volatile network conditions. In the testbed a heterogeneous wireless network environment, consisting of several IEEE 802.11 Soekris access points and an IEEE 802.16 base station (BS), each embedding an NECM, has been deployed. Moreover, several single-RAT (i.e. WiFi) and multi-RAT (i.e. WiFi, WiMAX) terminals are located in the corresponding area, consuming a video service delivered by VLC-based service provider. The cognitive network manager installed per network element undertakes a) the deductions about its operational status, b) the proactive preparation of solutions to face possible problems, and c) the fast reaction to any problem by enforcing the anticipated reconfiguration actions.
The demonstration experiment shows how to automatically perform (re-)assignment of operating frequencies to wireless network elements and the vertical assisted handover of multi-RAT terminals. The demonstration scenario is divided into:

a) the optimal deployment of a new WiFi AP,
b) the Self-Optimization of the network topology through the assisted vertical handover of terminals from loaded to neighboring less loaded APs or BSs,
c) and the Self-Optimization of the network topology due to high interference situation.

This use case is described in:
http://www.youtube.com/user/ScanLaboratory#p/u/0/EG_iSNrhwkE

The experimentation in Self-NET has recently used TEAGLE for test bed definition. Specifically, the scope of this use case is to experiment on the improvement of QoS features (e.g., packet loss, delay, jitter) using the Self-NET software for self-management over a live network environment (provided by the Panlab Octopus testbed) and exploiting monitoring and configuration capabilities that different administrative domains provide (i.e. access network and service layer). Performance results from the first phase of the PII and Self-NET collaboration have been published here:
http://www.panlab.net/fileadmin/documents/Use-Cases/Testing_end-to-end_Self-Management_in_a_Wireless_Future_Internet_Environment.pdf

Use of PlanetLab is judged not beneficial since Self-NET is a wireless environment. The NITOS wireless test bed in the OneLab federation should have been possible to use. This should have been investigated in the project.

RESUMENET

Resilience in networking is the main target of REUMENET. The goal is to both be able to handle failures and malicious nodes. The testing environment is directed towards wireless networks.

The currently available deliverables do not fully describe the testing environment since they are due during the second year of project life. The first study case investigates the challenge of node misbehavior in a wireless multi-hop network. However, available information indicates that usage of the NITOS test bed should have been investigated. But again, the need of the NITOS test bed was not great since one of the partners has provided an alternative, i.e. ETH Zürich wireless testbed (TIKnet).

This shows that there is an overlap in available test beds, but the possible federation of the NITOS and the Zürich test bed is not discussed and could have been profitable in case a larger scale test would have been desired. Scaling-up of wireless networks in order to create the effect of much larger cells of larger populations of wireless clients served by adjacent cells is an open and important question at present, one which FIRE should put special focus on. It appears that the scale up and federation goals of the existing wireless platforms are not well articulated. This set of issues may also be more complex than the project objectives and funding envision.

PERIMETER

PERIMETER's main objective is to establish a new paradigm for user-centricity in advanced networking architectures. In contrast with network-centric approaches, user-centric strategies could achieve seamless mobility driven by actual user needs rather than simply business considerations. Putting the users at the centre rather than the operator enables them to finely control their identity, preferences and credentials.

PERIMETER implements middleware architecture able to support generic QoE models, signaling-content adaptation, etc. PERIMETER will develop prototype applications and services for obtaining user-centric seamless mobility. User-centric seamless mobility, middleware components and its integrated applications and services will be tested in two large-scale interconnected test-beds on real users. One PERIMETER use case is described in
http://www.youtube.com/watch?v=DZTD9f2h2sY

One important aspect of PERIMETER is to perform research on usability testing. Such tools are not available in FIRE. PII has a work package with the aim to produce tools for UDI. However, these tools are still too rudimentary to be used; however it could have been beneficial if the UDI part of PII and Perimeter had collaborated to mutual benefit. The "Living Labs" have no facilities and tools integrated in the FIRE experimental facilities.

Recently PERIMER has integrated a slice of the FEDERICA test bed to be used in experimentation (however actual experimentation has not yet started). The FEDERICA slice procured by PERIMETER consists of five Virtual Nodes. These nodes will be used to deploy the PERIMETER Support Nodes and will allow for a natural maturation of the of the interconnection from a Layer 3 to a Layer 2 connection between the Waterford Institute of Technology (WIT) and the Technische Universitaet Berlin (TUB) test beds used in PERIMETER. The slice will be used for a variety of purposes within PERIMETER including scalability, performance, network overloading and disruptive testing.

The testing facilities used are part of the ArcLab (WIT Ireland) and TUB test bed (Berlin), which are not part of FIRE. Used facilities are similar to the facilities available in PII/PanLab, which could have been an alternative and a collaboration partner, especially since WIT is partner both in PERIMETER and PII. But at the time of writing of the PERIMETER proposal, WIT was not yet a partner in the PanLab project. Therefore, a potential collaboration with the PanLab facility could not have been anticipated at the outset of both projects. As the projects have progressed, a potential collaboration for PERIMETER with PanLab was investigated, but at the time that the testbed setup was needed in PERIMETER, the Teagle tool was not yet available for use. Furthermore, the particular requirements for the PERIMETER project, regarding Quality of Experience, specific applications and monitoring services, would not have been met by the offerings of the PanLab facility. The PERIMETER testbed facility which has been created, and reused from existing resources, for use in the project belongs to the ArcLab facility in Waterford, and these resources can subsequently be used to enrich the testbed components therein, which can be made accessible to the PanLab federated testbed offering as required. Throughout the PERIMETER project, the focus has been on the re-use and the sustainability of the ArcLab facility in Waterford.

SMART-Net Smart Antenna Multi Mode wireless mesh network

SMART-Net project implements and develops a simulation-based test-bed, providing simulated and real-life facilities . The main objective of the simulation test bed is to simulate large-scale variants of the SMART-Net architecture. SMART-Net has chosen OPNET as the simulator for deploying the simulation test bed. SMART-Net will interconnect its simulation test bed and a real-life test bed, enabling real life evaluation. More precisely, the "real" SMART gateway (WiMAX mesh BS) will be connected, on one hand, to a set of stations located in the real-life test bed and, on the other hand, to additional virtual SMART stations, which belong to the simulation test bed.

One objective is to conduct both stand alone simulations, focusing on different aspects of the protocols, and complex experiments combining the real-life test bed and the simulation test-bed. Some advantages of such an approach are: (a) interoperability (this is a natural step between testing a network or protocol in simulation only and building a real prototype or production device and testing it separately); (b) one can analyze the effects of a simulated network on a real application; (c) the simulation may be used as a traffic generator to load a real network; ((d) one can test the performance of a new protocol being developed by driving real network traffic over a simulated network. An approach as this allows comparisons between the simulation model and the hardware network behaviors.

The interest for other facility projects could be to take up the technology in combining a simulation environment with a real networking test bed for usage e.g. in traffic generation.

SMART-Net Real Life Test-bed is available under request through internet access. Nevertheless, it has not really been used till now by external researcher, certainly due to a lack of dissemination.

Existing FIRE facilities do not include simulation environment thus its connection to the existing FIRE facilities is weak and it would be difficult to reach the goals of this project using FIRE. This is in contrast with the situation in GENI, in which one leading candidate for its control plane is emulation-oriented, and provides this support in a fairly natural way.

VITAL++

VITAL++ develops P2P in a combined IMS and standard P2P environment. The project has been able to integrate some of the features of IMS (as e.g. charging, security) into standard Internet P2P. VITAL++ is sometimes called a facility provider but the actual project should rather be considered a facility user. VITAL++ has recently demonstrated usage of its technology using the PII/Panlab tool TEAGLE to create a virtual test bed on which the algorithms of Vital++ could be tested.

However, large scale testing has not been demonstrated and implementation an environment as e.g. PlanetLab should show important aspects today only shown in simulation.

NanoDataCenters

Nanodatacenters proposes a solution to data hosting and delivery. The current data delivery architecture is "network centric", with content stored in data centers connected directly to Internet backbones. The Nanodatacenters project takes a orthogonal approach through what is called "nano" data centres, which are deployed in boxes at the edge of the network (i.e. in home gateways, set-top-boxes, etc.) and accessed using a new peer-to-peer communication infrastructure.

Nanodatacenters develops communication architecture with security and incentive mechanisms. The full Nanodatacenters architecture will be implemented (i.e. a Nanodatacenters box will be fully specified and implemented). Virtualisation will be used to partition and manage box resources efficiently. Two interactive applications – multiplayer games and VoD – will be designed as a proof of concept.

The Nanodatacenters project is to a high degree an architecture project and less an experimental environment. It is difficult to envisage profitable usage of the FIRE facilities as an experimental resource since the architecture involving mainly local connectivity and the most important problems in the project are architectural issues as security and business aspects (as e.g. do I want to let someone else use my electricity). Possibly PlanetLab could be used in some of the experiments.

OPNEX

Not possible to evaluate since NO public deliverables are available on the web site of OPNEX.

OPNEX, Optimization Driven Multihop Network Design and Experimentation, delivers the first approach of principle to the design of architectures and protocols for multi-hop wireless networks. Systems and optimization theory is used as the foundation for algorithms that provably achieve full transport capacity of wireless systems. Subsequently, a plan for converting the algorithms termed in abstract network models to protocols and architectures in practical wireless systems is given. Finally a validation methodology through experimental protocol evaluation in real network test-beds is proposed.

The project content suggests that use of OneLab/NITOS test bed could be interesting.

Analysis of call 5 STREP projects

The call 5 projects are in some cases similar to the call 2 (see right hand side of figure 7). Examples of new aspects that are researched are smart buildings, flexible radio and large scale data handling. The following overview of the projects is very rudimentary and can only be considered as a discussion on what could be expected in usage of existing and coming facilities. The description of the new facility projects are found in a previous report.

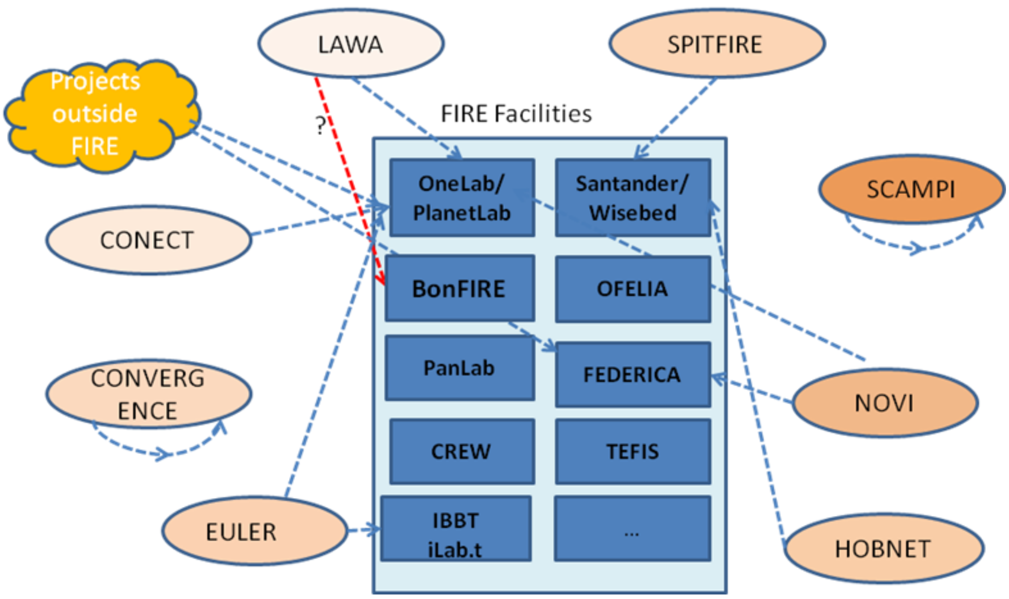

Figure 8 – planned facility usage among call 5 STREPs.

The Figure 8 describes the current plans for the Call 5 research projects (as far as has been possible to deduct from available documentation and personal communication). The arrows in blue are describing planned usage and the arrows having a ?-mark is usage under discussion or to appear in an open call proposal. Further usage may be planned but has not been documented yet.

CONECT

CONECT will deliver a holistic network design approach aiming at enhancement of the performance of wireless networks by unlocking the hidden potential of the broadcast wireless medium.
The fundamental idea is to remove the boundaries among received signals, such that they are not classified as desired or interfering ones. Contrary to existing approaches, different received signals are allowed to mix, superimpose themselves on each other, and interact. This key observation unlocks the potential to exploit the open broadcast wireless medium. Nodes located within the broadcast range of a transmitter are no longer passive listeners trying to distinguish the signal intended for them. Instead, they cooperate among themselves in order to realize efficient and reliable information forwarding by essentially pooling the overheard information.

CONECT experimentation will be based on the OpenAirInterface test beds for flexible radio network design and the NITOS test beds for cooperative wireless access and transport. Experiments will be performed on delay sensitive video multicast and streaming applications.

CONVERGENCE

The Internet is evolving into an Internet of services, an Internet of media, an Internet of people and an Internet of 'things' rather than the classical 'network of hosts',. This leads to a shift from 'host-centric' to 'content-centric' and 'data-centric' networking. Against this background, Convergence proposes to enhance the Internet with a novel, content-centric, publish/subscribe service model, based on the Versatile Digital Item (VDI): a common container for all kind of digital content, derived from the MPEG21 standard.

The experimentation is expected to be run on large scale FIRE facilities. However these are not specified in available documentation but a guess could possibly be PlanetLab.

EULER

The main objective of the EULER exploratory research project is to investigate new routing paradigms so as to design, develop, and validate an experimentally distributed and dynamic routing scheme suitable for the Internet and its evolution. The resulting routing scheme(s) is/are intended to address the fundamental limits of current stretch-1 shortest-path routing in terms of routing table scalability but also topology and policy dynamics. The driving idea of this research project is to make use of the structural and statistical properties of the Internet topology as well as the stability and convergence properties of the Internet policy in order to specialise the design of a distributed routing scheme known to perform efficiently under dynamic network and policy conditions when these properties are met.

EULER is addresses new aspects of routing problems but sharing experimental methodology with the call2 ECODE project. The experimentation will take place in the IBBT emulation facilities and if possible in OneLab/PlanetLab. It seems reasonable to ask whether their proposed alternate routing could be tested at still higher bandwidth by working with OFELIA's programmable OpenFlow routers.

HOBNET

HOBNET is oriented towards smart buildings.
Key objectives of HOBNET are to create:

- A scalable all IPv6/6LoWPAN network architecture to support Future Internet services and applications, particularly for the smart/green building domain.
- A coherent set of novel models and high quality algorithmic solutions that have been implemented, tested and validated along with high-level technical recommendations for smart building scenarios
- An interface layer between the building management system and FIRE experimentation platforms to be used for the rapid development and the evaluation of building management applications
- Contribution to 6lowApp and its standardization towards a new embedded application protocol for building automation

The experimentation will be based on a FIRE/WISEBED test-bed. Also, because of the special focus of the project on IPv6 and smart buildings, the project will also experiment on application-specific test-bed extensions that are created at MANDAT International/Geneva and CTI/Patras.

LAWA

The LAWA project on Longitudinal Analytics of Web Archive data will build an Internet-based experimental test bed for large-scale data analytics. Its focus is on developing a sustainable infrastructure, scalable methods, and easily usable software tools for aggregating, querying, and analysing heterogeneous data at Internet scale. Particular emphasis will be given to longitudinal data analysis along a timeline of Web data that has been trawled over an extended period.

The LAWA project looks at finding possibilities to use the BonFIRE facilities to deploy their web services for data exploration in a cloud environment.
The plan to use PlanetLab looks clearly straightforward and possible.

NOVI

NOVI concentrates on methods, algorithms and information systems that will enable users to compose and manage isolated slices, baskets of virtual resources and services provided by diverse yet federated Future Internet (FI) platforms. NOVI will also enrich the FIRE facility with federated models and methods enabling comprehensive and reproducible experiments.

NOVI plans to create federation between PLE, FEDERICA and GEANT as a main target and extend the virtualization to both facilities.

SCAMPI

The Scampi project envisions a future environment where users will carry personal mobile devices such as cameras, smart phones and PDAs, with a number of resources (several wireless interfaces, a lot of memory, powerful CPUs, components able to generate multimedia content). The resulting networking environment, viewed as a whole, will therefore feature a multitude of heterogeneous resources. The goal of Scampi is to enable each user to benefit not only from the resources available on their own device, but also to opportunistically exploit the other resources of the environment, including those on other users' devices, with confidence and security. Scampi will thus enable users to utilise the functionality of the different resources available in the network, so that users enjoy more than what is solely available on their own device.

The experimentation in SCAMPI has not planned to use any of the existing or coming experimental facilities.

Spitfire

Spitfire works towards the realisation of a stronger connection between the natural and the digital worlds. It will 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. Essentially Spitfire will significantly reduce the effort required for the development of robust and interoperable applications on the 'Internet of Things'. 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. As a result, the architecture developed in Spitfire that unifies the IoT and the Web of Things domains has the potential to increase European competitiveness in ICT technology. Due to the enabling technologies provided by Spitfire, IoT-related technology could permeate private households and enterprises in a way not seen to date.

Spitfire project will rely on the work performed in WISEBED and could be viewed as an additional continuation of that project.

Final conclusions on STREP analysis

This analysis does only include FIRE research projects. It is well known that PlanetLab is widely used in many projects outside FIRE and FEDERICA (not even a FIRE project) has most of its users outside FIRE.

- In order to create greater use of experimental facilities, projects must plan to use facilities already at project definition phases. This has not been possible in Call 2 since knowledge of available facilities was not present. This was still not possible for Call 5 proposers and could still become a

problem for Call 7 and Call 8 proposers, due to the fact that many of the facilities are not yet available. Out of the coming Call 5 facility provider projects expected first external users could be accepted mid 2011.
- A project proposal should in its proposal/work description and in its state of the art analysis include a description of possible usage of existing facilities with clear description how facilities will be used. In case no usage is expected a motivation why not, must be provided. This should be an important aspect when evaluating research proposals.
- Creating awareness of project requirements for experimentation should have been a major part of the effort in FIREWORKS with support from the research projects. This should now become part of FIRESTATION project. Proposers should be supported to find appropriate facilities and to plan extensions to existing facilities that are needed for their experimental work.
- The wireless NITOS test bed in OneLab should be a candidate for many of the projects, which emphasizes the need of efforts to create awareness of existing and planned facilities.
- The use of FEDERICA as a federated test bed in experiments performed e.g. in PERIMETER and PII shows a need for support of heterogeneous federation. This need is also emphasized by the work in integrating the OPNET simulation environment to the physical equipment in the experiments performed by SMART-Net.
- There is an evident overlap in testing facilities in Europe and since many of the research projects use experimental facilities available with partners rather than using FP7 supported facilities. Either these facilities should be federated with the FIRE facilities or the need for further support of existing facilities should be reviewed.
- The end of the OneLab, PII and Federica projects late 2010 will strongly hamper the development of FIRE experimental facilities for European research. This problem should be managed in Call 7 and Call 8.

Short description of research project coverage.

The projects in FIRE cover a large number of areas and are mainly not overlapping. Of course some areas are more often represented but in most cases different aspects are treated. Clustering of activities has not been performed (at least not documented). Below is an indication on what issues can/could create collaboration either on test bed usage or research collaboration. In some cases this is performed and in other the projects are continuation of each other.

Wireless	OPNEX, Smart-NET, Self-Net, N4C,CONECT,RESUMNET OneLab,CREW
Sensors	Hobnet,Wisebed, Spitfire
Routing	ECODE, EULER
Learning technologies for Automation of mgmt	ECODE, Self-Net
DTN	N4C, OneLab
P2P	Nanodatacenter, Vital++
UDI/Usability	Perimeter, PanLab,TEFIS,N4C
Large scale data	Lawa,BonFIRE,TEFIS

Internet of things	Spitfire, SCAMPI
Benchmarking	OneLab, CREW
Measurement	OneLab, BonFIRE, CREW
IMS	PanLab/PII, Vital++

Recommendations

Our recommendations address three areas, namely the immediate steps to be taken by FIREstation and call 5 projects, missing elements for call 7 and 8 preparation, and challenges for the unit in charge of FIRE in achieving the vision of sustainable large-scale experimental facilities. We address these areas in the following.

Immediate Steps for FIREStation – Portals, Federation, and the Open Calls

- A common portal is a required tool to improve usage. However, the different ambition levels of a portal (as shown in the current set of projects and proposals) ranging from a simple catalogue to a full-fledged semantic description of resources, generation, deployment and control of an experiment makes it necessary to find intermediate stages when defining and implementing shared tools. It might also become necessary to support more than one (but not many) competing approaches to the problem.

- The selection on use of facilities when heterogeneous federation is required has to be a decision based on implementation cost and goals/interest of the involved individual facilities. As a result, the combinations (projects, use cases, and users) leading to the most exciting advances are and will be rare. These must be cultivated.

- There has been progress in top-down federation in FIRE projects, while SFA (in FIRE + GENI) has shown potential for scalability in bottom-up federation. FIRE should address the integration of these two approaches.

- It is necessary to clarify how the open calls are to be coordinated. The FIRE Office (i.e. FIREstation) and the IP projects must each have influence on which projects are chosen. If each project, constrained only by commission guidelines on proper procedure, selects those user experiments that best fit its existing resources, we will certainly not end up with much common FIRE activity. There is a tension between choosing user experiments that show off particular benefits of individual facilities versus choosing experiments that strive for demonstrating the need and benefit of federated facilities. This tension must be addressed as the procedures for the open calls are formalized and explained to the IPs, to align the selection process with the overarching vision of a federated experimental facility, to be spelled out at the time of the Calls.

- Besides selecting experiments cutting across facilities through the open Calls of the Call 5 IPs, an additional opportunity arises through Call 8 (closing envisaged for January 2010). Here an opportunity for receiving funding for innovative experimentation projects is provided. This call is likely to be very timely as the Call 5 facilities will be available by the time of closing of this Call.
- Sharing/interconnecting data, not only resources, is crucial since it will enable sharing of data across experiments. Benchmarking and repeatability of experiments is important when conducting 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 once the shape of the experiments performed under the open calls is visible. Overall, such efforts should be coordinated by FIREstation, better being addressed by a dedicated effort with clear goals and milestones to measure progress.
- The issue of how the NREN and GEANT infrastructure can be used in FIRE should finally be discussed and concluded. This issue is under discussion since the beginning of the FIRE initiative. Experience in ONELAB, FEDERICA, and OFELIA can serve as a starting point.
- Sustainability of each FIRE experimental facility shall be studied in the context of a sustainable FIRE federated facility. Here, business models play a crucial role. It is important that in the architecture board during the coming months proper business models are investigated, using the model of contributing and sharing resources of ONELAB2 and the envisaged commercial models of PII as a starting point. Consideration in how far networking costs must be paid for by the user should also be considered taking into account the FEDERICA experiences. Also BONFIRE has first ideas for sustainable business models.
- FIREStation should take the lead in identifying appropriate levels of user support and ensuring that the best practices are shared across the FIRE portfolio.
- Both with Japan and with the US, in recent meetings it was decided to carry out some cross-continental experiments showing the added value of the collaborations beyond networking of the testbed. It is important to present showcases, which prove that experiences can be gained, which would not be possible to be gained internal to FIRE.
- Last but not least, we believe that this coordination overall needs to be

streamlined in the light of the upcoming PPP and the Calls 7 and 8. Less is more! One opportunity in this context sure is the Call for a FIRE Federation IP under Call 8.

Missing Elements –for Call 7 and Call 8

- Much is still missing, for example optical networking and equipment found at the edge (e g hand held terminals) are still underrepresented in the FIRE portfolio. Also the application efforts address computational resources but not storage (possibly "in the network") of the large amounts of data consumed or generated.
- A vision of end-to end support for the FIRE users need to be integrated into the requirements of Calls 7 and 8. Put the FIRE customer at the center of these efforts, with measurements of usage and value to the end-user (both developers and true end-users) as key to success – see below in challenges to the FIRE unit.
- Calls 7 and 8 shall explicitly indicate the relative weight of user support for the projects proposing to build the future FIRE facility.

Challenges to the FIRE Unit

- It is important to establish metrics for facilities with respect to inter-project collaboration and user attraction – these needs to become the main review criteria for technology progress. As a consequence, if collaboration and attracting external users remains modest, projects should be considered to be terminated earlier in order to eliminate wasted efforts that are not used in the wider FIRE environment.
- Experiments in heterogeneous federation must be encouraged and probably become a strong requirement in future calls. As a consequence, there is a need to establish metrics for facility usage that establish strong barriers against unnecessary re-implementation of facilities and tools that already exist. These metrics need to be a major foundation for any experimental research project.
- Together with the projects establish efforts to discuss the creation of a sustainable facility and its business model(s) outside the current project structure since concerns about availability when an experiment is ready, and sustainability after funding terminates will restrict the value of any FIRE facility. Only true sustainability (outside a project timeframe) can change this.

Appendix A: user support in an established field of science

Since user support is clearly a new, untested concept in most of FIRE's space, it could prove useful to describe how it is done in a field of experimental science, astronomy, where this is an accepted part of any world-class effort in an area that new technology has made accessible and important new questions need answers. We examine the 5-year plan of the US National Radio Astronomy Observatory (NRAO) for operations and user support of ALMA, a 50+ antenna millimeter radio array being constructed in Chile and funded equally by ESO in Europe and NSF in the US, supplemented by significant contributions from Canada, Japan and Taiwan, at a total cost approaching a billion Euros.

In planning for user support, operations, and data preservation it helps that astronomers have looking at the sky in disciplined ways since the time of Galileo, but radio astronomy is a new field. Images are obtained only after correlating the radio signals from many receivers, and signals at different frequencies carry different information, so that relatively raw forms of the measurement data must be preserved, along with significant amounts of calibration information. Only a small fraction of astronomers are radio astronomers, so making these tools and this data available to optical, x-ray, and infra-red astronomers is a major objective, requiring specialized training materials. Tools developed for ALMA can be used in support of interferometric observations at other frequencies or for other purposes, such as visualization. Astronomers employed by institutions which lack the resources to do data reduction and analysis, expect these facilities to be made available to them. Finally, astronomy depends for its long-term popularity on keeping an interested public stimulated and informed with exciting images and clear expositions of the science that can be done with these expensive research instruments, so public outreach is a component of work that must be planned from the outset.

The ALMA observatory has three portals, one in Chile, and two which mirror data but manage distinct research and distinct sets of users in Europe and in North America. Common software, developed by NRAO, manages scheduling, monitoring of the observatory, controlling the correlator, and offline data reduction. A pipeline for relatively automatic data reduction combines contributions from the several partners in ALMA. An end-to-end "Observing Tool" (a European contribution) provides a framework that supports users from proposal generation through observation and subsequent data analysis. Large pieces of this exist, or have earlier implementations, but the whole process has never been given this level of business automation before. Staff must be hired and trained to provide support at the submission, schedule validation, and data analysis phases of each experiment. And visiting users will need office space.

In preparing for outside users, NRAO must develop training resources including tutorials, demos, wikis and videos for the Observing Tool, hold workshops to expose this material at various regular professional society meetings, and develop a library of sample observing programs, sample datasets, and a simulator for use in testing and refining proposed observing schemes. All of this must be ready to be tested and validated as part of the early science exercises that occupy about a year of the last stages of the telescope array's construction in Chile. When it operational, help desks will be staffed at all three portals to support the various stages of each user's research.

In all the steady-state operations and support staffing for an instrument of this size and complexity requires about 100 people in North American, and presumably a comparable number in Europe and a smaller number in Chile. The fraction which is required to be added in order to support users outside the organizations which have constructed ALMA is not easily decided, since not supporting such users is not a conceivable option in astronomy. A guess is at least 30 of the 100 represent this extra level of involvement.

Does this model of user support apply to FIRE's experimentally driven exploration of the future Internet? Astronomy has a well-defined set of research paradigms, and these evolve on a comfortable time scale of half a decade or so. New systemic approaches like large scale, long-running surveys have been incorporated into the field without requiring new types of telescopes, only new software to run them. But FIRE and GENI have also developed a time constant of several years, over which strategies of experiment control and data acquisition are being explored and tested through prototypes. We think the differences are only qualitative, and the model of end-to-end user support, with its significant costs, must be considered if FIRE (or GENI) is to succeed.

Appendix B: Integrated Projects in Call 5

All Call 5 IP projects have now provided at least technical portions of their DOWs to the analysis group. We have also arranged telephone conferences between project leaders in each IP and one or more members of the portfolio analysis group. We expect that some aspects of these projects will evolve significantly as the current stage of FIRE proceeds, coordination through the Architecture Board commences (in fall 2010) and the first user proposals are accepted and executed in the subsequent years, so this is a very early view. We sketch our impressions of each project using the following template, supplemented or diverted by additional topics which came up during our discussions.

Q1: How will you engage with users? When, to what extent?

Q2: What are your terms and conditions (Ts and Cs) under which an external user will have access to your testbed?

Q3: Plans for sustainability? How? For how long?

Q4: What efforts are made to ensure security and privacy in the testbed? How will data be handled, and what infrastructure is planned for it?

Q5: What web presence is planned? Is there a portal? What functions are supported there, and how?

Q6: What previous and current projects do you partner with? At what level, and how?

SmartSantander

Q1:

User engagement focuses mainly on the real scale experiment that will be performed in the city of Santander in Spain. The number of IoT devices in Santander alone is expected to be up to 12,000 , increasing to a total of 20,000 when including the other test beds from other locations (Lübeck, Guildford, Belgrade, Århus, and Melbourne), and projects (WISEBED, SENSEI, G-Lab in Germany etc.). SmartSantander claims to look at requirements from "all involved stakeholders: Industries, communities of users, other entities that are willing to use the experimental facility for deploying, validating and assessing new services and applications, and Internet researchers to validate their cutting-edge technologies (protocols, algorithms, radio interfaces, etc.)". They intend to open up the experimental platform to other users than the ones planned in Santander, using the 20% set aside for federation work through open calls, as they consider that "giving money to experimenters will make it much easier and more attractive for them to deploy their experiments on the SmartSantander facility". There is a specific WP looking at user engagement in coordination with FIREWorks and FIREStation.

Tthey seem to have the right intent and associated effort to make a significant contribution on user engagement, even if the focus is on Santander. This may prove to be an asset as this is clearly a real-life experiment with various types of users

Q2: T&C: Any? When?

Terms and Conditions aspects are not yet planned.

Q3: Sustainability: How? For how long?

SmartSantander has a Task within WP5 dedicated to "sustainable exploitation", which focuses primarily on the Santander experiment. It is not clear however how sustainability is considered for other users than the ones involved in Santander. Santander is unusual among the Call 5 IPs in that the city of Santander owns its sensor infrastructure (about 60% of the project's resources), and will maintain it for some time into the future.

Q4: Security & Privacy: Data? Infra? How?

Security, privacy, and trust seem to be important topics for SmartSantander. Although there are not many details, their focus is on "support for testing whether new protocols and algorithms can be used in a secure way and deserve to be trusted by the users". They have a dedicated task within WP2 to make sure that "the algorithms, applications and scenarios uploaded to the experimental facility for testing as well as the results of the experiments and in general the operation of the facility will remain secure and confidential to the users".

Q5: Portal: Any? What function? How?

SmartSantander starts with a web site with all useful information. Otherwise they rely on the common development that will be coordinated by FIREStation. There is an expressed intent to augment the Santander website to support experiments more directly.

Q6: Integration: What projects? What level? How?

SmartSantander is planning to use the results from SENSEI and WISEBED; they will also have access to Telefonica's Ubiquitous Sensor Networks (USN) Platform. Although they intend to "use tools and components which have already been developed within FIRE, and will help in developing new tools and mechanisms which are still missing", there is no detail on which projects such tools and mechanisms could come from (PlanetLab/OneLab2 is mentioned with no further details for interoperability for example); they seem to rely on FIREStation and the FIRE Architecture Board to discuss this. At the same time, they consider

federation through PlanetLab Europe.
They stress several times the concept of "Open Federated Alliance" promoted by WISEBED but seem to admit that they will have to go with what has been proposed by the Wise Men report. The proposal/DoW is quite open to participation in the Architecture Board and in general with collaborating with FIREStation.

Ofelia

Q1: User engagement: How? When? What extent?

The engagement is done largely through the Berlin deployment, which allows access to the student community, i.e., the normal usage of the network through students. This could be extended through the FIRE User call.

There exist letter of intentions from projects to utilize the resources as well as contacts through use cases. A timeframe for the user engagement was not given although T&C and discussions on extent should be done early in the project.

Q2: T&C: Any? When?

T&C discussions started on internal mailing list with task in WP2 picking up the work once project will start. The T&C is supposed to be ready for the ICT 2010 for presentation. The extent seems to cover infrastructure usage. Pricing is unclear due to the NREN involvement. Data usage is unclear to be included. Here we should learn from the use of the NREN infrastructure.

Q3: Sustainability: How? For how long?

Sustainability comes through the island structure of the infrastructure, i.e., island owners commit to sustain infrastructure although it is unclear how and for how long. There is no particular model in place that could ensure the island sustainability through some federated construct. The 'federation' sustainability comes through open sourcing the federation software but there is no particular construct for the federation as an organization (such as PLC).

Q4: Security & Privacy: Data? Infra? How?

Focus on data security and privacy through technology means like virtualization. Data privacy could be part of T&C. However, there exists a mandatory NEC agreement that mandates potential data usage through NEC due to utilizing switch firmware.

Q5: Portal: Any? What function? How?

WP5 will address portal provisioning, leaning towards PII. However, there is a recognition that integration into 'one-stop shop' might be over-ambitious given the heterogeneity of resources in various experimental facilities. Hence, the plans for a portal exist but the required extent of integration is doubted.

Q6: Integration: What projects? What level? How?

FEDERICA is the obvious candidate although OpenFlow approach is seen as superior to FEDERICA. GENI contacts exist, too, in particular through Stanford. CREW is also seen as candidate for collaboration. Unclear how such collaboration and integration could possibly happen. The level of collaboration is largely seen as being on measurement and tool level in general.

TEFIS IP project

1: User set and use cases

TEFIS will develop a portal and a platform for experimental research to support research on software services. More specifically:
The TEFIS platform will support research on future large-scale and resource demanding Internet service technologies by offering a single access point to different testing and experimental facilities for communities of software and business developers to test, experiment and collaboratively elaborate knowledge.

Q2: Have Tx and Cs for use of their facilities been arrived at?

TEFIS has two internal prototype experiments, expects to support a number of external studies through the open call mechanism. The project will plan to allow external use through the open calls at month 18.
However, other aspects of terms and conditions are not elaborated.

Q3: Sustainability?

Not clear. There are defined efforts to develop business models, but the details are still unclear.
TEFIS resembles PII in the way the project will create a business using testing facilities looks similar to that of.

Q4: Approach to security, privacy and data archiving/sharing.

Clouds already offer facilities for configuration and deployment (running the experiment). TEFIS is planning to use Teagle and will experiment with the federation of several existing test beds as PACA GRID and PlanetLab.

Q6: Integration with other IPs and former FIRE projects?

The TEFIS project shares a set of partners with BonFIRE including Fraunhofer

(TEAGLE) which should lead to close collaboration on many aspects.
TEFIS will contribute to the continuation of the development efforts for Teagle (from PanLab) and to a large extent is similar to PII in the way test beds are integrated into the framework.

TEFIS will provide federation of 8 test beds among them Grid5000, but there is no clear discussion on how to collaborate with existing FIRE facilities.

The project is in many respects similar to PII in that it will integrate a number of existing test beds. PII has developed common interface for components used to allow them to. These aspects are only partially addressed in TEFIS.

Bonfire IP project

Q1: User set and use cases

Two themes seem apparent. One is research in services that required combining different standard service or application personalities to create new types of cloud-based services. The second is research on the effectiveness of cloud substrates when the facilities and the data are distributed to some extent. Some proposed controllable interconnect between clusters that participate allows this last issue to be studied as the quality of network interconnect is varied.

Q2: Have Tx and Cs for use of their facilities been arrived at?

Bonfire has internal prototype experiments, expects to support 6-7 external studies through the open call mechanism, and says that is intended that the facilities they create will be open on an unfunded basis to other projects. Bonfire partners have the experience of running open calls in a previous GRID-based project, and developing an evaluation metric as part of each such outside project that they attract and support in this way. They seem to be approaching the FIRE open calls in a similar way. Their experience needs to be shared with the other IPs.

Bonfire is so ready for the open calls that however it is not clear whether they will encourage project submissions that stress their technologies and integrate the BonFIRE facilities with other parts of FIRE. It will be tempting for BonFIRE to focus on exploring their strongest points and staying within their initial definitions of services that they can provide

Q3: Sustainability?

Not clear. These facilities are expensive to operate, need regular technology refreshes to keep the performance competitive and relevant for research. There are mentions in the proposed DOW of developing a business model, but we have heard this before. Bonfire (and CREW) has a better chance of achieving a sustainable facility than the others, but it is still quite uncertain if the most valuable parts of this project will outlive FP7. The "Bonfire store" which is cited as existing at the close of the project seems to be quite undefined at present.

Q4: Approach to security, privacy and data archiving/sharing.

Clouds already offer facilities for configuration and deployment (running the experiment), so BonFIRE only needs to graft above these for project visibility and perhaps some enhanced usability over existing cloud frontends. They agree to evaluate Teagle and have added Fraunhofer as a partner to perform this evaluation and possible integration. They will need a web interface to manage open call proposals, evaluation, and execution/monitoring, and plan to develop this.

Q6: Integration with other IPs and former FIRE projects?

Closest interactions are with Teagle (from PanLab ii) and with IBBT's virtual wall for network prototyping (part of CREW, in a sense). They also present a form of federation which reaches deeper into the programming stack than some of the other FIRE contexts that will be drawn on by the Architecture Board. Yet because of the relatively well-understood and stable test environment that Bonfire provides, these requirements should integrate relatively easily as extensions to the other forms that federation will take in FIRE. Bonfire's views and results on the differing approaches taken to centralized vs distributed control in Teagle and SFA may prove important.

CREW IP project

The objective of the CREW (Cognitive Radio Experimentation World) project is to establish an "open federated platform" 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

Q1: User set and use cases

The project is aiming at demonstrating the CREW functionality through cognitive radio usage scenarios and at establishing demand-driven extensions to support experimentation by new partners. They also intend to "interact with FIRE coordination & support actions to ensure maximal compliance of CREW federation to the overall FIRE vision."

This shows that CREW already intends to develop some federation tools that are required for its own purpose. Such tools, which include a common portal, a common API for linking radio hardware with software components, a common data collection/storage format, and a unified methodology across testbeds for measurements and processing of results, could serve as a basis and/or may be expanded to take into account and serve other FIRE facilities.

In addition, CREW is proposing some "federation modes" that could also be discussed in the Architecture Board and shared with the federation ideas at FIRE level.

Q2: Have Ts and Cs for use of their facilities been arrived at?

The experience from the existing IBBT facilities will give the project a head start, but opening up to a larger user community is still an issue.

Q3: Sustainability?

CREW intends to develop a sustainability business model for the exploitation of the CREW facilities. The CREW roadmap includes the evaluation and monitoring of a sustainable operation of the CREW platform right with definition of a sustainable business model after the end of the project at year 3, and an independent sustainable operation at year 5 with experimenters paying for the use of the facilities.

The CREW project plans to keep the CREW facilities up and running, using the sustainability model developed. However it is not clear in the available documentation who will operate and who will be responsible for maintenance. This needs to be analyzed as part of the business modelling. The operation & maintenance will be a joint activity by different partners (as the expertise, HW & SW resides at different partners), which is quite a challenge.

Q4: Approach to security, privacy and data archiving/sharing.

CREW wants to develop (among other things) a common data collection/storage format, a unified methodology across testbeds for measurements and processing of results. This could serve as a basis and/or may be expanded to take into account and serve other FIRE facilities.

Q5: What web presence is planned? Is there a portal? What functions are supported there, and how?

Tools as a common portal, a common API which could also serve other FIRE facilities are planned.

Q6: Integration with other IPs and former FIRE projects?

The CREW project shows a willingness to collaborate with other projects with common interests. These include portals, API, benchmarking methodology, measurement tools and methodology.

Background material for STREP projects analysis.

Project web sites are indicated below, but when available also deliverables relevant to the analysis have been used and in some cases published scientific reports.

www.onelab.eu

http://nitlab.inf.uth.gr/NITlab/index.php/testbed

www.panlab.net and demos in connection to Barcelona FIA June 2010

www.wisebed.eu and discussions with Stefan Fischer Uni Lübeck.

www.ict-vitalpp.upatras.gr and demos in connection to Barcelona FIA June 2010

www.fp7-federica.eu/

www.nanodatacenters.eu

www.ecode-project.eu

www.n4c.eu

www.opnex.eu

www.ict-perimeter.eu

Blog: http://perimeter.tssg.org/

http://perimeter.tssg.org/2010/07/perimeter_and_federica.html

resumenet.eu

http://www.ict-selfnet.eu/

https://www.ict-smartnet.eu/

Future Internet Research and Experimentation an Overview of the European FIRE Initiative and its projects august 2010 (EU commission Unit 4).

http://www.conect-ict.eu

http://www.ict-convergence.eu

http://www.euler-fire-project.eu/

http://www.hobnet-project.eu

http://www.lawa-project.eu

http://www.fp7-novi.eu

http://www.ict-scampi.eu/

http://www.spitfire-project.eu

www.crew-project.eu