The European FIRE
Future Internet Research and Experimentation Activities

Recommendations prepared by
the FIRE Scientific and Industrial Preparatory Expert Groups

D R A F T

Draft Version 14 June 2007

1. CONTEXT – WHY FIRE ?

1.1. Evolution of the Internet – growth and problems - patches

The architecture of the Internet is approaching 30 years of age. In that time, it has succeeded beyond the wildest dreams of its creators in changing the way people throughout the world live, work, and play. The EU research and industrial communities have played an essential part in the international efforts that have led to this success, not least under the successive ICT R&D programmes aimed at building the foundations of the Information Society.

This success has been achieved while the scale of the Internet has increased by 7 orders of magnitude. In order to cope with such growth, the simple, original Internet architecture (16 specification documents) has accreted several hundred additional protocols and extensions. Networks based upon this significantly more complex architecture are increasingly difficult to manage in a way that enables the qualities of service delivered to meet the needs of the over 1 billion users.

The increasing, and implicit, reliance on the Internet has stimulated a major debate amongst experts [CLARK et al.] as to whether the current architecture and protocol can continue to be patched, or whether it will collapse under the demands of future applications. There are signs that the current suite of protocols and solutions are becoming inadequate to cope with some common Internet trends: mobility of users and devices, unusual but legitimate traffic load (e.g. flash crowds), large heterogeneity in terms of devices capabilities and service features, delivery of real-time high-bandwidth video services, requirements for episodic connectivity, scalability in terms of number of nodes and users, complexity related to network, service and security management.

Additionally, the original Internet was designed and built in an era of mutual trust, probably due to the small size of the original community addressed by the DARPA project. Many of the protocol additions/extensions have been to retrofit protection mechanisms that are required in the current Internet environment, which does not merit mutual trust. The volume and types of attempts to subvert the Internet will continue to increase, further stressing the current architecture. Current solutions for security are added a posteriori as a patch to overcome the limitations encountered, instead of being embedded in the system functionality.

Furthermore, mobile network hosts are rapidly becoming the norm for the devices with which users access the Internet. An increasing number of the protocol additions/extensions have been needed to retrofit support for mobility into the (initially wireline-focussed) Internet architecture. The growing use of mobile sensors will continue to drive the need for solid mobility support in the architecture (and the efficient transfer of small data units).

IPv6 has emerged as the solution for resolving the shortage of IPv4 addresses, and restores the original end-to-end concept of the Internet. It has stimulated the development of many associated protocol additions/extensions, which bring benefits such as auto-configuration, security, and simpler handling of mobility. However, it addresses only one specific layer of the Internet and, especially in a long-term perspective, cannot be seen as *the* solution for all the above mentioned problems.

1.2. The starting point - EU research in the field

In Europe since several years industrial and academic researchers and the Commission have recognised the need to investigate possible means for overcoming the intrinsic limitations of

the current Internet architecture. Within the framework of FET ("Future and Emerging Technologies") under FP6, the SAC ("Situated and Autonomic Communications", [SAC]) activities and other future Internet related projects with an overall budget of ~€30m were launched in 2004. These projects address several decentralised (i.e. autonomic or distributed) approaches for new Internet architectures and services, featuring self-organising properties. Under the IST Strategic Objectives "Broadband for all" and "Mobile and Wireless Systems and Platforms beyond 3G" several projects also address important aspects of the future Internet. A more complete list can be found in Annex C.

In parallel, under the Strategic Objective on Research Networking Testbeds [RN] under FP6, several testbed projects were started to pave the ground for a future generation of e-infrastructures for research (~€30m funding). These testbeds address issues such as convergence and quality of service. They include two projects, started in 2006, which are envisaged to be instrumental in bootstrapping the experimental facility planned within FIRE: ONELAB which addresses proof-of-concept testbeds for researchers, by extending and deepening the current Planetlab [Planetlab] approach to Europe (Planetlab is a worldwide interconnection of testbeds for networking technologies and is managed by Princeton University), and PANLAB which is deriving a framework for the federation of testbeds on a broad scale.

Indeed, there is a multitude of projects in Europe running under various programmes, which carry out research on networks and services in a broad sense and which address aspects related to the evolution of the Internet. Several of them include testbeds, such as many relevant RTD projects in the above mentioned Strategic Objectives of the EU framework programmes, as well as projects from European regional, national, or multinational research initiatives such as EUREKA – CELTIC. Several of these testbeds may be candidates for federation within FIRE. Collectively these testbeds form a *European testbed ecosystem*, which FIRE aims to structure so as to maximise its utilisation.

A large number of European experts (called the EIFFEL "Evolved Internet Future For European Leadership" think tank) have recently published their first white paper on "The future networked society", which addresses many aspects of research into future networks in a European context [EIFFEL].

1.3. Need for experimentally-driven research on the Future Internet – definition and methodology

Many networking researchers around the world have identified the above mentioned emerging limitations of the current Internet architecture and agree that it is time for research to take a long term view and to reconsider the basic architecture of the Internet, to see if any improvement can be identified, even if it does not appear to be backward-compatible at a first glance.

To be effective and to produce applicable results, this long-term, fundamental research in new communication and networking paradigms has to be tested, at least as a proof-of-concept, in large scale environments, so as to assess the feasibility of the new concepts, verify their large scale effects (not only at technological level, but also as for their foreseeable implications on users, society and economy) and derive further requirements, orientations and inputs for the long-term research. This kind of experimentally-driven approach avoids that the long-term research will remain at the level of paperwork and will allow exploring dramatic improvements over the current Internet.

Researchers need an experimental facility for validating innovative – and potentially disruptive – architectures, technologies and protocols. History has shown that many occurrences are only discovered when systems are deployed in "real-life" situations. For instance, experimentally-driven security research includes experimentation with intentionally and unintentionally misbehaving programs and machines in a large, heterogeneous, real-world-like testbed environment, which nevertheless is isolated from the outside world. Such experiments need to be conducted in a planned, controlled, responsible and legal manner. The FIRE experimental facilities would enable such testing, in a controlled way, within a federation of more generic testbeds, reducing the dependency on specialised security-oriented testbeds such as the US testbed [DETER].

The experimental facility on future Internet technologies should also broadly support medium- to long-term research on networks and services so as to compare current and future approaches. Practical experiments are needed to give credibility and raise the level of confidence in the research findings; furthermore, the experimentation must be performed on a large scale to be representative, convincing, and to prove scalability. Experimental facilities based on federating testbeds at different levels of maturity, from proof-of-concept to validation, are needed to test compatibility, interoperability and to derive potential migration paths for innovative technologies.

Large scale experimental facilities cannot of course be developed anew and repeatedly in every single research project. Therefore there is a strong need to ensure that synergies are achieved, through a sustainable and dynamic evolution of the federation of testbeds.

The rationale for deploying an experimental facility is to provide the research activities with a tool for assessing the characteristics of the solutions derived, or an instrument to speed up the derivation of the results themselves. The value of the experimental facilities is expected to significantly increase over time if data produced by the facility is archived, together with the corresponding metadata duly describing the testbed environment and the tests performed, allowing for better experimenting and comparing at scale or with diversity. This would also allow researchers not having direct access to the federated testbed to perform some simulation/emulation studies based on real-world traces, adding significance to their own research. To base such an archive on solid grounds and to make results comparable, there is a need for additional research on methodologies for testing and benchmarking.

1.4. Why should the Commission fund FIRE?

The FIRE Research and Experimental facilities are expected to strengthen European industry by giving them a means for the development of advanced networking technologies and experimentation at network and service levels so as to reinforce their competitiveness in a global market, which in this sector is dominated by US patents and companies. Being successful, FIRE can give a significant contribution to redesigning the map of industrial competitiveness in Internet technologies, exploiting also the more favourable European situation regarding mobile and wireless technologies. In so doing, the FIRE vision is fully in line with the objectives set in Lisbon [LISBON].

FIRE is expected to create an environment promoting strategic, long-term research on new Internet concepts, while at the same time giving the opportunity to Europeans for carrying out large scale experimentation of such new paradigms, in order to measure and compare results, give credibility to what would otherwise be only theoretical studies, prove concepts, learn from the experiments (and not just at the technological level).

FIRE will create a credible and sustainable environment for European research into the future Internet. It is a mechanism to mobilise European resources around a common theme, creating a critical mass of expertise and encouraging Europeans to work together. Under this umbrella, European experimentally-driven research will achieve credibility and be able to compete with similar activities worldwide.

In summary, funding the FIRE activities will be strategic for Europe, given its interest both for industry and academia. The ultimate benefit expected from FIRE is that EU industry is better positioned on future Internet technologies and services than it was for the first generation, where the US currently holds a clear advantage.

2. THE FIRE VISION

FIRE aims at addressing the emerging expectations which are being put on the Internet by providing a research environment for investigating and experimentally validating highly innovative and revolutionary ideas.

New technological solutions may follow either an incremental approach or a “clean slate” approach. While the first evolves a system from one state to another by implementing incremental patches, the latter demands a radical redesign to offer new abstractions and improved performance, and has no constraint of “backward compatibility”. The latter approach might even prove to be disruptive in terms of technologies, services or business models. Whereas the FIRE experimental facilities are obviously meant to be open to both types of approaches, the research carried out under FIRE is addressing only more visionary approaches, more risky but at the same time more likely to yield significant improvements.

It has to be stressed here that incremental and "clean slate" approaches are not competing, but complementary. Incremental solutions are obviously necessary fixes in the short term, whereas in the medium or long term we have also to consider the possibility of fundamentally changing the Internet architecture or some of the underlying paradigms. Testing them in large scale experimental facilities will be essential for proving their feasibility, for identifying potential migration paths and for assessing their possible technological and socio-economic impact.

FIRE should at the same time investigate new and possibly unexpected Internet concepts, and gradually federate existing and future testbeds to become an integrated, sustainable, dynamic, large scale experimental facility on future Internet technologies, as a major support instrument for both medium- and long-term research on networks and services by industry and academia.

The vision is one of a broad range of advanced and interconnected testbeds, spanning the technology chain from the network connectivity to the service architecture and the development chain from proof-of-concept to operational testbeds, which can be used by both industry and academia.

By addressing future challenges for the Internet such as mobility, scalability, security and privacy, this new experimentally-driven approach is challenging the mainstream perceptions for future Internet development.

3. HOW TO IMPLEMENT IT

The FIRE Future Internet Research and Experimentation activities have two related dimensions:

- Experimentally-driven long term research on new paradigms and networking approaches for the future Internet;
- Building a sustainable, dynamic, large scale experimentation facility by gradually federating existing and new testbeds for emerging or future Internet technologies.

3.1. Experimentally-driven long-term research related to the Future Internet

Long-term, visionary research on Internet protocols and architectures is expected, as a key and distinguishing core part of these activities. The main characters of such a research are:

3.1.1. Openness of the approach

There should be no boundaries for the research, but rather the freedom to address *any* emerging or radically new but promising concepts to address the fundamental limitations of the current Internet. This can span all layers of the communications protocol – and, of course, explore different paradigms which might not be based on layered models. There should be no backward-compatibility constraints from the outset.

3.1.2. Multidisciplinarity

Real innovation often comes at the intersection of different disciplines. Moreover, the Internet is a complex system, depending on a delicate equilibrium between technologies, users, services and applications. Evaluating carefully these interrelations will be key to harnessing and exploiting the full potential of the future Internet for economy and society at large.

Non exclusive examples of such cross-fertilisation are:

- To apply bio-inspired principles to network design, which for instance can exploit the presence of redundancy, random patterns and noise in the system to obtain extreme robustness and reliability, even in unpredictable environments.
- To explore means by which networks can change, learn and develop spontaneously.
- To apply concepts from neural development in the paradigmatic exploration of concepts for the new Internet. Unlike evolution, which is a long term process that may include items that are no longer of any use, development is a dynamic process of optimisation.
- To apply findings from social psychology experiments to capture or predict human interaction. Taking into account the interdependency of the users is crucial to optimise future communication systems.

- To apply traditional market or economic rules, such as try, buy, change, borrow future Internet services as is done today with material commodities.
- …

3.1.3. Experimentally-driven Research

Large or small research projects in this area should base their research theories and results on testbeds and experiments. Building dedicated testbeds can be in the scope of the research projects themselves, or the projects can envisage exploiting existing or planned testbeds, which are part of the FIRE experimental facilities under construction. This should be seen as part of a new research methodology in Internet concepts, tightly coupling the research and experimentation of long-term, potentially disruptive Internet concepts. New Internet paradigms require testing - and eventually validation - in large scale environments in order to adequately assess their potentially far-reaching implications. The results of such testing will then be fed back again as concrete requirements for the long-term research.

3.2. Experimental testbeds and their federation into FIRE

The other dimension of the FIRE activities is an experimental facility for future Internet technologies which supports both medium-term and long-term research on networks and services. This facility should be built up gradually by federating existing and new testbeds, in order that they become an integrated, sustainable, dynamic, large scale experimental facility for use by both industry and academia.

A testbed here is to be understood as an environment allowing rigorous, transparent and replicable testing and experimentation in research and development projects.

A federation of testbeds is to be considered as an interconnection of two or more independent testbeds for the creation of a richer environment for testing and experimentation, and for the increased multilateral benefits of the users of the individual testbeds. In a federation, testbeds would normally be geographically dispersed and owned by different organisations. They would however be considered as being part of a single resource, in so far as they are operated in a common management framework under a common management authority. Federations are dynamic and evolve over time based on the requirements of the users.

The federation of testbeds should address issues from early proof-of-concept to validation aspects, thereby allowing industry and academia to collaborate, to exploit synergies, to identify migration paths for new concepts, and in particular to support the exploitation of research results.

The FIRE experimental facility should be built by gradually federating existing and new testbeds supporting research on networks and services. In its full implementation, the FIRE integrated, sustainable, dynamic, large scale experimental facility shall allow for integrated system level testing of future Internet technologies across technology layers from the network connectivity to the service architecture and across stakeholders at different levels of the value chain from technology providers via service providers to users at all levels.

3.2.1. Research on methodologies and tools for testing

Research in testbeds for future Internet research is an important area in its own right. To convert networking from art into an exact science, testing methodologies are needed that can help increase the scientific value of research results. A scientific approach to testing is needed, including repeatability of the tests in a dedicated and controllable environment (such as simulators and emulators) and their interaction with the experimental testbeds.

There are several issues that need to be measured in a quantifiable and scientific way, including the following (out of which the last one is the most challenging):

- Performance (including also robustness and resilience).
- User experience.
- Quantitative impact on societal goals.

Methodology for the three areas can build on established methods from computer science, HCI (Human Computer Interaction), and economics, but the scale and innovativeness of the FIRE questions will require new ways of combining and extending these methods.

3.2.2. Operation and management of federated testbeds

The operation and management of federated testbeds over multiple networks and multiple administrative domains is difficult and requires professional support from both industry and organisations with experience of operating large networks. The following issues need to be addressed:

- **Openness:** Federation implies openness at all levels, including provision, implementation and use. Concerning provision, the FIRE experimental facility must be open for any testbed of the *European testbeds ecosystem*. There are many related issues which make the implementation of openness a challenge. These include Intellectual Property Rights (IPR), how to exchange results and what may be shared, the standardisation of interfaces, common procedures, usage policies, etc. The implementation could be done through open source tools and open hardware design details which will enable easy replication, interfacing and inter-operation among various testbed components, and consequently, would increase the overall FIRE effectiveness and impact. Use will be open to any relevant European and worldwide initiative under fair access terms.

- **Excellence:** The federation principle promoted by FIRE aims for "best of breed" and must assure excellence of the capabilities of the federation from various aspects. These aspects include among others (i) degree of innovation, (ii) measurable quality or compliance to essential standards, (iii) diversity, (iv) scale or geographic coverage. Criteria must be developed that allow a transparent evaluation of excellence, and also the re-classification of a testbed as a commodity and thus its removal from the federation.

- **Management:** Managing federated testbeds is complex, but necessary for achieving scale, diversity, cost-efficiency and, last but not least, to improve the sustainability and quality of the individual testbeds and the federation as a

whole. The traditional network management objectives apply for the management of federated testbeds as well.

- **Governance**: Interconnecting different test-beds belonging to different administrative domains means granting access to remote resources that are owned by different stakeholders. This has the consequence that a legal framework must be in place that governs this relation and which must include among others: handling of IPR, definition of confidentiality and trust, process for resolution of conflicts and handling of the mis-use of rights or neglecting obligations, etc. In order to lay the foundation for establishing a long-term sustainable testbed federation, a legal entity might prove necessary. This legal entity must manage the different relations, possibly manifested as contracts, between the stakeholders that interconnect their resources in the scope of the federation.

3.2.3. Building the FIRE experimental facilities

It is not possible to build a single large infrastructure from the outset, not only due to the funding limitations, but also because it is not feasible to select the most relevant research topics at this early stage. The FIRE facility should therefore be deployed incrementally, responsive to the evolution of the outside stimulations. Several facilities and research topics need to be deployed and studied in parallel.

Moreover, the construction does not start from scratch, but builds upon - and federates - existing connectivity, testbeds and facilities which can be shared. Examples are:

- Advanced testbeds already existing which include the experimental proof-of-concept testbeds emerging from the projects of the SAC activities.

- More mature testbed facilities for testing interoperability, compliance, performance, and quality of service which have been launched under Research Networking in recent years.

- Testing and validation components of RTD projects in the FP6 Strategic Objectives related to networking and services.

- The PANLAB and ONELAB projects launched in 2006 and which are investigating and trialling the concept of testbed federation.

- The Full Scale European Demonstrator currently under development by the NESSI ETP which in the long term could provide the link from ICT infrastructures to service architectures and to business processes in enterprises.

- For FIRE to have maximal economic impact, resulting in new successful European products, services, and companies, it is important that in the long term it can cover the whole chain from basic research to broad pilot projects with real customers. Several of the LivingLabs currently under establishment in Europe [LIVINGLABS] have the potential to provide the link to these user communities and bring the user in the loop at all stages of research and development thereby improving the innovation process.

- Network connectivity: For the federation of geographically dispersed testbeds in general and in order to assure the replicability of results in particular, FIRE

could profit from and build as much as possible upon GÉANT, the Pan-European Gigabit Research Network. In order to do so, however, a change of policy would be necessary to give researchers access directly to the fibre on a European scale, which is currently not foreseen in GÉANT, but which is already possible in several NRENs (National Research and Education Networks).

Building upon existing activities as described above, a logical step-by-step approach towards establishing the integrated, sustainable, dynamic large scale experimental facility would start with the following activities:

(1) **Pilot projects of interconnected testbeds, including R&D on test methodologies**: Federating first groups of testbeds also at system level; validating the methodology and procedures, and demonstrating the value added through federation. This should also result in guidelines/best practices for the federation of future testbeds into FIRE. This pilot group of federated testbeds should already be open to relevant external projects.

(2) **One or several closely co-operating projects which provide a framework and an umbrella or clustering mechanism for federation**: In the medium term these projects should facilitate the establishment of an organisation governing and coordinating the experimental facilities.

In subsequent phases the following steps should be envisaged. These however are subject to a critical review as soon as first results become available.

(1) The existing federated testbeds should be developed further and extended both in scale and functionality by introducing new testbeds. It should be envisaged that the testbeds will develop from being focused on aspects of a system (devices, application, control, communication protocol, physical network, measurement, monitoring, security, etc.) into more complete systems, combining the elements previously mentioned, and further including issues like interoperability between heterogeneous systems or bridging to different user communities.

(2) Under a single umbrella, the pilots of federated testbeds should evolve into the targeted European integrated, sustainable, dynamic large scale experimental facility.

4. SOCIO-ECONOMIC ISSUES

Nowadays, the Internet is not just a technology, but a complex system deeply integrated in the fabric of our society; consequently, radical technological changes in its architecture could have unexpected consequences at economic and social level, and even possibly carry some ethical concerns, of which everybody, and researchers in the first place, should try and be aware. For instance, an architecture different from the current one could have more

"intelligence into the core", which could be used to guarantee a better quality of video streams or improved security. At the same time, such an architecture would be less flexible than the current "dumb" one, and current P2P applications, used to exchange files, music and video, sometimes illegally, and which represent more than 80% of the current Internet traffic, could not be possible, or be more strictly controlled. Would that be for the good or for the bad? Dominant content providers are clearly in favour of a stricter control, but others argue that for many people P2P could represent the only means to access culture. As widely recognised, the Internet has so far flourished thanks to the end-to-end principle, the use of open standards and open source, the lack of generalised censorship and its role in ensuring fair use in the circulation and use of knowledge. Promoting and maintaining an "Internet for Everyone" is therefore a key mission for future generations.

In general, the success of the current Internet is evident by how it has influenced our society. Yet at the same time, the society is facing the Internet with continuously new challenges. A specific and unique strength of FIRE is that it will provide the basis for a scientifically rigorous impact assessment of network architecture proposals, at both technological and social levels. One such core related question is the impact on Network Neutrality (NN), a network design principle that states that a maximally useful public information network aspires to treat all content, sites, and platforms equally. NN is perceived by many as being endangered by policies such as differential quality of service, yet others see NN as a deterrent to innovation and investment in network infrastructure. Opinions are split about how changes to the fundamentals of network architecture will affect it, in the absence of measured data. It is clear that any intelligence in the network – as opposed to today's "dumb pipes" – involves decisions (e.g., what is legitimate traffic and what is an attack) that may affect NN.

The FIRE sustainable, dynamic, large scale experimental facility on future Internet technologies is an excellent place for doing a rigorous and quantitative research on the broader impact of technological alternatives. Specifically, the aim is to treat socio-economic requirements and effects as much as possible in parallel to technical requirements and effects: to specify desired outcomes, to define measures and metrics of goal achievement, to perform analytical and experimental performance measurements, and to derive recommendations.

In addition, FIRE's heterogeneous, federated testbeds can become realistic models of heterogeneous socio-economic and end-user situations, negotiations and other developments, and experiments in these models can build on currently emerging research in social and economic sciences and in HCI. FIRE will enable interdisciplinary teams to perform rigorous scientific experiments under controlled conditions in order to determine effects on and interactions with end users and communities.

5. STRATEGIC ISSUES

5.1. Sustainability of the testbeds, beyond the lifetime of the projects

Sustaining the operation of a testbed of advanced technologies beyond the duration of a research project has proven very difficult to achieve in the past. This is unfortunate, not only

because of the loss of a wealth of experience; but also because it shows that the business models behind the design of the testbed were fundamentally flawed. Ensuring that the testbeds can be maintained and exploited also beyond the lifetime of projects (whenever it makes sense) is a key issue for the FIRE activities, as well as finding mechanisms to ensure that results from past and current projects can be effectively exchanged and compared.

A Cluster of future Internet research and testbed projects could be formed that is also open to membership from non EU project organisations. This group could jointly define the requirements for testbeds and services provided by them. The sites could be part of national research initiatives, not bounded by EU project timescales.

Sustainability of the measurement data and its metadata is an important issue. The huge amounts of information, their heterogeneity and the expected privacy concerns are thorny questions, and solutions for their management (within or outside the FIRE activities) should be found to set up a common data repository.

Incentives for joining the federation

Motivating research projects and existing activities to join the FIRE federated experimental facility is a challenge in itself. As one possible means, EU calls in this area should positively consider future Internet research projects that are committed to validating their results on the federated FIRE testbed, and to testbeds that will actively federate and develop their services.

A transparent classification and rating of the testing sites would improve for everyone the overall quality of the offered services. By bringing the individual test sites on the same line, a competitive situation is induced between the sites. This alone is an incentive for the sites to stay up-to-date in their offering, and to strive for improving their services.

Risks

The main risk is that the facility does not meet the requirements of its user's community. Besides technical risks, issues related to management and governance must be considered at an earlier stage in order to understand the model that will provide sustainability and easy access to the facility. The organisation should be visible in order to discuss with similar initiatives worldwide.

5.2. European strengths and weaknesses - competitiveness issues (SWOT-type)

A number of European strengths and weaknesses at large must be taken into account, in order to build on the strengths and avoid the pitfalls rooted in the weaknesses.

Strengths	*Weaknesses*
• Ability and experience in conducting long-term research in a collaborative manner • Head start in long-term research in communications	• Technology transfer, spin-offs, exploitation of results (no start-up culture) • Limited interaction between industry and research organisations • Heavy administrative procedures

• Deployed large scale broadband infrastructures (GÉANT, NRENs, …) • Expertise in mobile communications	• Marginal European impact on the evolution of the Internet • No major equipment vendors (except wireless technologies)
Opportunities • Wide adoption of the open source model (to speed up development of needed tools) • Trend towards wireless access to the Internet - given Europe's strengths in mobile communications • FIRE as a means to spur new technological areas and business models • "Disruptiveness" of the long term research leading to innovation • Possibility of developing European standards • Collaborate and influence FIND/GENI and other international initiatives • Consideration of socio-economic aspects	*Threats* • Leading applications over the Internet might change faster than the FIRE vision • Following dead end approaches • Other initiatives across the globe (e.g. FIND/GENI) dominate the future design

5.3. Analysis of the constituency and their interests

The multitude of the stakeholders having an interest in the FIRE activities is large, yet can be categorised in several groups:

The research and academic communities have a valuable head start when it comes to capturing the complexity of requirements for future networks. With its roots in the research community, the Internet now links more than a billion people and extends deeply into the commercial world. One of the most important topics in the future will be the linking of research of multiple different disciplines.

The commercial deployment of new services is a risk if they are not previously tested with real networks and users. Telecommunications operators need to integrate and validate new "telecommunication operator's concepts" within a research networking infrastructures and they must be enabled to test interoperability on a large scale. They must be part of the future Internet right from the beginning, since they can provide valuable insights in building large operational networks. Through their involvement they can influence the design of the future networks.

Service and content providers at all levels and SMEs have to be aware about the future platform for delivering services and content. Similarly they can influence its design through their involvement. Content will drive the design of the future networks.

Equipment vendors will benefit by having the opportunity to trial and evaluate technologies and system solutions in a large scale with real users. FIRE provides significant market opportunities for new technologies based on new architectures that overcome current limitations of deployed equipment.

Additional stakeholders are the individual users, who are involved through a network of communities focused on how people use the technology. The fundamental principle is to give users early access to technology and learn from their needs and values in order to improve the technology.

Finally, there is an ecosystem of stakeholders responsible for supporting the activity through technology testing at all levels, such as interoperability, usability, accessibility, etc.

Due to the convergence between networking, computing and content, in future it will be more difficult to distinguish between telecommunications operators, service providers, content providers and users themselves.

5.4. Positioning of FIRE in the European and international contexts

ICT Programme context

Under Challenge 1 ("pervasive and trusted network and service infrastructures") of the ICT theme of FP7, research into the future Internet is addressed in all the objectives under Challenge 1 of the ICT WP 2007-08. Two focal points under these objectives specifically target this field:

1. "***Technologies and systems architectures for the Future Internet***" (focal point (c) of objective 1.1 on "The Network of the Future" with a pre-allocated budget of 200 M€).

2. "***Advanced networking approaches to architectures and protocols***" (focal point (a) of objective 1.6 on "New Paradigms and Experimental Facilities" with a pre-allocated budget of 40 M€).

A second focal point under 1.6 aims to establish interconnected test beds open to all innovative technologies and research results from across the objectives.

European general context

With respect to its relation with industry, FIRE lives in the complex ecosystem of the European ICT research, whose main constituencies are the established Framework Programmes like the EUREKA ICT Clusters and the IST European Technology Platforms (see Annexes C and D). These Programmes consolidate and cover almost the full scope of the ICT industry in Europe, since they have been established and running for several years. Please also refer to the EIFFEL report [EIFFEL].

With respect to regional or national initiatives in the Member States of the European Union, it is necessary to take stock of similar or complementary activities in order to exploit synergies in the context of the European policy towards a European Research Area (ERA).

International Context

In an international context, clearly Europe is not doing research in isolation. The European activities under FP7 as mentioned above are building on research in FP6 under the FET Situated and Autonomic Communications activities, under the FET Open scheme and under Strategic Objectives such as "Broadband for all", "Mobile and wireless systems and platforms beyond 3G", and "Towards a global dependability and security framework". These European activities are also running in parallel with activities across the globe, for example:

NSF has launched in 2005 a research initiative called FIND (Future Internet Network Design) to study what are the requirements for the global network of 10 or 15 years from now, and what should that network look like (see http://cfp.mit.edu/events/slides/jan06/Dave-Clark.pdf). It has so far received some US$ 25m.

Recently NSF has started planning for an infrastructural initiative called GENI (Global Environment for Network Investigation), aiming to build an experimental facility to enable the research community to invent and demonstrate a global communications network and related services that will be qualitatively better than today's Internet.

Elsewhere, Japan and Korea have publicised their ambitious u-Japan and u-Korea initiatives, and China also has an ambitious and integrated industrial policy. These initiatives are not all tackling the issue of the Internet evolution as part of their core objectives, but are certainly related to technological and socio-economic scenarios that are needed for the Internet of tomorrow.

There is a close dialogue between the US and Europe. The Commission has participated to the OECD workshops on the economic aspects of the future Internet. The last one was organised in by OECD and the US National Science Foundation and contacts have been established and reinforced.

In order to promote collaboration and the exchange of ideas with US colleagues on views and experiences on concepts for the future Internet, both on research and experimentation, a **EU-NSF workshop on the future Internet** took place in April 2007 (http://www.net.t-labs.tu-berlin.de/ARCADIA/index.shtml). The results were presented at the second FIRE preparatory Group Meeting.

ANNEXES:

Annex A List of participants in the FIRE Preparatory Groups

Sahin	Albayrak	Germany	DAI Labor
Susanna	Avessta	Finland	DIMES
Stephan	Baucke	Germany	Ericsson GmbH
Bettina	Berendt	Germany	Humboldt University Berlin
Wolfgang	Brandstaetter	Austria	Telekom Austria
Heinz	Brüggemann	Int. Org.	CELTIC
Marco	Conti	Italy	CNR
Costas	Courcoubetis	Greece	Athens University of Economics and Business
Philippe	Cousin	Int. Org.	ETSI
Jon	Crowcroft	UK	University of Cambridge
Walid	Dabbous	France	Inria Sophia Antipolis
Peter	Deussen	Germany	Fraunhofer Fokus
Christophe	Diot	France	Thomson
Serge	Druais	France	Thales, ETP NESSI
José	Enríquez Gabeiras	Spain	Telefonica I+D
Serge	Fdida	France	LIP6
Anja	Feldmann	Germany	TU Berlin / Deutsche Telekom
Afonso	Ferreira	Int. Org.	COST
Timur	Friedman	France	Université Pierre et Marie Curie
Anastasius	Gavras	Germany	Eurescom
Silvia	Giordano	Switzerland	DTI-SUPSI
Andreas	Heinemann	Germany	TU Darmstadt
Arto	Karila	Finland	Helsinki Inst. of Inf. Techn.
David	Kennedy	Germany	Eurescom, ETP NEM
Peter	Key	UK	Microsoft Research Ltd
Scott	Kirkpatrick	Israel	Univ. of Jerusalem
Michael	Kleis	Germany	Fraunhofer FOKUS
Dina	Kronhaus	UK	St. Catharine's College
Vasilis	Maglaris		GEANT, NRENS
Antonio	Manzalini	Italy	Telecom Italia
Martin	May	Switzerland	ETHZ
Cezary	Mazurek		
Ioannis	Mertzanis	Greece	Space Hellas, ETP ISI
Daniele	Miorandi	Italy	CREATE-NET
Edmundo	Monteiro	Portugal	University of Coimbra

Andrea	Nicolai	Italy	T-6
Veli-Pekka	Niitamo	Finland	Nokia
Yoram	Ofek	Italy	University of Trento
George	Polyzos	Greece	Athens University of Economics and Business
Radu	Popescu-Zeletin	Germany	Fraunhofer Fokus
Martin	Potts	Switzerland	Martel
Bart	Preneel	Belgium	Katholieke Universiteit Leuven
Michael	Riguidel	France	ENST
Eric	Robert	France	Thales
Pablo	Rodriguez	Spain	Telefonica
Juha	Saarnio	Finland	Nokia
Ioannis	Stavrakakis	Greece	University of Athens
Joe	Sventek	UK	Univ. of
Leandros	Tassiulas	Greece	University of Thessaly
Christian	Tschudin	Switzerland	University of Basel
Fiona	Williams	Sweden	Ericcson, ETP e-Mobility
Andreas	Willig	Germany	Technical University of Berlin
Tanja	Zseby	Germany	Fraunhofer Fokus

This report was edited by Serge Fdida, Anastasius Gavras, Arto Karila, Martin May, and Martin Potts.

The work of the FIRE Preparatory Groups was stimulated by the European Commission Directorate General Information Society and Media, Unit F4 "New Infrastructure Paradigms and Experimental Facilities".

Annex B Extract from FP7 Work Programme

Objective ICT-2007.1.6: New Paradigms and Experimental Facilities

Target outcome

a) **Advanced networking approaches to architectures and protocols**, designed to cope with increased scale, complexity, mobility and requirements for security, resilience and transparency of the Future Internet coupled with their **validation in large scale testing environments** based on a combination of physical and 'virtual' infrastructures.

b) **Interconnected test beds** addressing novel distributed and reconfigurable protocol architectures; novel distributed service architectures, infrastructures and software platforms; and advanced embedded or overlay security, trust and identity management architectures and technologies. Test beds for systems that provide trusted access to e-Services with users requiring no administration and security skills.

c) Coordination and support actions for: i) standardisation and conference support; ii) coordination with related national or regional programmes or initiatives.

Expected Impact

- Strengthened European position in the development of the future Internet.
- Wider take-up of technological developments in networks and service infrastructure facilitated by a comprehensive validation of the technological and service choices.
- Global consensus towards standards and strengthened international co-operation through interconnected test beds and interconnection capabilities offered to third countries.
- Higher confidence in the secure use of the Internet through test beds enabling trusted access to e-Services.

Funding schemes

CP, NoE, CSA

Indicative budget distribution

40 M€:

- CP 36 M€ of which a minimum of 12 M€ to IPs and a minimum of 15 M€ to STREP
- NoE 3M€
- CSA 1M€

Text to be added.

Annex C Short descriptions of existing projects in the area

Within its multi-annual framework programmes the European Commission have been supporting numerous projects and initiatives addressing the intrinsic limitations of the current Internet architecture. This annex provides an incomplete list of exemplificative projects considered relevant to the scope of this report. Further information on the projects funded by the European Commission can be found at: http://cordis.europa.eu/ist/

FET SAC and Open projects

Eight IST-FET projects are considered as very relevant for future Internet research. ANA, BIONETS, CASCADAS and HAGGLE are Integrated Projects which have been selected in the 4th IST call under the SAC FET initiative (Situated Autonomic Communications). The EVERGROW IP had been originally selected in the 1st IST call under the Complex System FET initiative. CATNETS, COOPCOM and Net-Refound are STREPs which have been selected in FP6 under the FET OPEN call.

ANA (Autonomic Network Architectures) aims at developing a novel network architecture (beyond IP) that can enable a flexible and autonomic formation of network nodes according to working, economic and social needs. Its objectives are mainly focused on studying adaptation and reorganisation of the network. In particular, the project aims at analyzing fundamental principles around functional scalability, both horizontally (adding more functionality) and vertically (integrating the functionality).

BIONETS (BIOlogically-inspired autonomic NETworks and Services) looks at nature and society for introducing novel networking/service provisioning paradigms tailored to pervasive computing environments. BIONETS exploits opportunistic communications as a mean to provide a localised peer-to-peer support to autonomic services, which embody evolutionary capabilities in order to be able to adapt and self-organise without requiring any explicit human intervention.

CASCADAS (Componentware for Autonomic, Situation-aware Communications and Dynamically Adaptable Services) aims at defining a new generation of highly distributed, pervasive, situation-aware, semantically self-organising communication-intensive services, by developing a common abstraction level for autonomic communication services. The focus is on situation awareness, semantic self-organisation, self-similarity, and autonomic component-ware.

HAGGLE (An innovative Paradigm for Autonomic Opportunistic Communication) aims at developing a new (cross-layer) network architecture exploiting intermittent connectivity. In this new architecture, message forwarding is driven by the application rather than by only control information of the network layer. This way, the architecture is able to support an opportunistic networking paradigm in which the delivery of messages is based on store and forward exploiting situated information.

EVERGROW (ever-growing global scale-free networks, their provisioning, repair and unique functions) studies Internet measurement technologies and their effects on distributed applications. Its goal is to invent methods and systems, and build infrastructure and peer to peer overlay networks and algorithms, for measurement, mock-up and analysis of network traffic, topology and logical structure, so as to start addressing now the opportunities presented by the future Internet, also by exploiting complex systems research principles and physics based approaches.

CATNETS (Evaluation of the Catallaxy paradigm for decentralised operation of dynamic application networks) proposes an alternative approach for realising resource allocation in dynamic application layer networks, by applying a decentralised economic self-organisation mechanism called "Catallaxy". The project is studying the applicability and implementation possibilities in P2P middleware, eventually providing a prototype, to obtain performance results both from a technical and economic point of view.

COOPCOM (Cooperative and Opportunistic Communications in Wireless Networks) studies the combined use of cooperation and opportunism, two concepts that have recently revolutionised the way engineers think about wireless system design and which target the maximisation of the spectral and power efficiency at the system level. It will explore performance limits of cooperative and opportunistic schemes and develop efficient strategies also for efficient resource allocation with limited feedback, including decentralised resource allocation, eventually implementing selected schemes on a test-bed.

NET-REFOUND (Network research foundations and trends) aims to develop the theory, methods and algorithms suitable for the modelling, analysis and design of future telecommunication networks. The long-term goal is the theoretical understanding of the collective interaction of a multiplicity of communicating nodes beyond the boundaries posed by specific telecommunication standards. This will lead to a quantitative characterisation of the fundamental performance limits of these systems and eventually to algorithms for achieving them.

IST Projects under networking-related Strategic Objectives

Under the Strategic Objectives "Broadband for all" and "Mobile and Wireless Systems and Platforms beyond 3G" several projects address important aspects of the future Internet. EURO-NGI and WIP are the most important examples:

EURO-NGI (Design and Engineering of the Next Generation Internet): The main objective of the EURO-NGI network of excellence is to develop and maintain a most prominent European centre of excellence in next generation Internet design and engineering, acting as a "collective intelligence think tank", representing a major support for the European industry and leading towards a European leadership in this domain.

WIP (An all-wireless mobile network architecture): to design, implement, and experimentally validate an all-wireless interconnection architecture based on advanced wireless transmission techniques, mesh networking, cross-layer optimisation, and mechanisms for seamless mobility.

Research Networking Testbeds

Under the Strategic Objective on **Research Networking Testbeds** (6^{th} Framework Programme), several testbed projects were started to pave the ground for a future generation of e-infrastructures for research (~€30m funding). The goals of these projects are integrating, testing, validating and demonstrating new fixed and wireless networking technologies and services in both real-world settings and production environments.

Two of the projects, started in 2006, are experimenting with the concept of federating or interconnecting testbeds on networking and services:

ONELAB (An Open Networking Laboratory Supporting Communication Network Research Across Heterogeneous Environments) addresses proof-of-concept testbeds for researchers, by

extending and deepening the current Planetlab approach to Europe (Planetlab is a worldwide interconnection of testbeds for networking technologies managed by Princeton University);

PANLAB (Open Interconnect for the Internet Community) is deriving a framework for a federation of testbeds on a broad scale.

Three of the projects deal with **Quality of Services**:

OPENNET (Open Interconnect for the Internet Community) brings solutions to the main barrier for a fuller-scale IP deployment: the lack of predictable support for QoS, when packets have to cross many domains. The project brings together the major Internet router manufacturers worldwide and identifies QoS parameter values for Premium IP services;

NETQoS (Policy Based Management of Heterogeneous Networks for Guaranteed QoS) is addressing the issue of management of QoS in an heterogenous environment with the objective of adapting the QoS system according to the changing physical networks taking into account QoS management policies;

EuQoS (End-to-End Quality of Service Support over Heterogeneous Networks) resolves the outstanding design issues presently associated with the delivery of end-to-end QoS service across heterogeneous networks. The key objective of EuQoS is to research, integrate, test, validate and demonstrate and-to-end QoS technologies to support the infrastructure upgrade for advanced QoS-aware applications.

Two of the projects deal with the problems of developing and implementing the **IPv6** protocol in Europe:

IPv6 TF-SC (IPv6 Task Force Steering Committee) The aim of the project is to assure the success of IPv6 deployment. IPv6 TF-SC project, in its role as the facilitator of the Task Force, which is invited to create strategic Roadmaps, will continuously monitor the academic, market and industrial activities, and provide guidance where appropriate to avoid duplication of work;

RING (Routing in Next Generation) investigates the current routing protocols and infrastructure and their ability to support the new services in a scalable manner; multihoming situations, the Inter-domain traffic engineering, the security of the routing infrastructure and utilisation of IP in airplanes, cars, trains, ships and other transport.

Two projects take an application-oriented view:

ANEMONE (Advanced Next Generation Mobile Open Network) will realise a large scale testbed providing support of mobile users and devices and enhanced services by integrating cutting edge IPv6 mobility and multihoming initiatives together with the majority of current and future wireless access technologies.

VITAL (Enabling Convergence of IP Multimedia Services over Next Generation Networks Technology).

Building on the results of past and ongoing IST projects related to the research on IMS technology, VITAL aims at consolidating the technological framework that will enable the smooth transition of multimedia communications, including voice, from circuit to packet switched domain of the communications. With respect to this goal, the project is setting-up an experimental IMS distributed platform, elaborate solutions to the aforementioned

problems and test, validate and assess the resultant integrated IMS functionality in the context of advanced traffic experiments and theoretical simulations.

Annex D Short descriptions of relevant ETPs

The objectives of the European Technology Platforms (ETPs) is to help industrial and academic research communities in specific technology fields to co-ordinate their research and tailor it to a common "Strategic Research Agenda" (SRA), which sets out research and development goals, time frames and action plans for technological advances that are relevant to industry and society. Each ETP represents all major stakeholders, including small and medium-sized enterprises, provided that they are willing to pool their knowledge and resources and possess the requisite expertise. ETP stakeholders agree to support their SRA financially and to monitor its implementation. A research project that supports ETP aims may be submitted for EU research funding in just the same way as any other. The main relevant ETPs in the context of FIRE are:

eMobility (Mobile and Wireless Communications) The eMobility technology platform aims at the improvement of the individual's quality of life, achieved through the availability of an environment for instant provision and access to meaningful, multi-sensory information and content. The eMobility vision is shifting from the current concept of "anywhere, anytime" to a new paradigm of "any network, any device, with relevant content and context in a secure and trustworthy manner".

NEM (Networked and Electronic Media) For a competitive networked and electronic media industry, the NEM initiative focuses on fostering the development and introduction of novel audiovisual and multimedia broadband services and applications to benefit European citizens and enterprises. The NEM vision is promoting the convergence of existing and new technologies, including broadband, mobile and new media across all ICT sectors, to create a new and exciting era of advanced personalised services.

NESSI (Networked European Software and Services Initiative) NESSI aims at the provision of new approaches enabling the transformation of the European economy through service-oriented business models. The NESSI vision is transforming the European economy into a knowledge-based economy and enabling the European Software and IT services industry to attain a stronger global position.

ARTEMIS (Advanced Research and Technology for Embedded Intelligence and Systems) The ARTEMIS approach is to break down the barriers between application sectors, stimulating creativity and yielding multi-domain reusable results. The ARTEMIS reference architecture can support product development in a diversity of application domains, for example automotive, aerospace and nomadic.

ISI (Integral Satcom Initiative) The purpose of the ISI technology platform is to bring together in a converged way, all research and technology aspects related to satellite communications, including mobile, broadband, and broadcasting applications, in order to foster and develop the entire industrial sector, maximise the value of European research and technology development, and contribute to EU and ESA policies.

Annex E Related European Activities

EUREKA is an inter-governmental organisation, created in 1985, which aims to enhance European competitiveness through its support to businesses, research centres and universities who carry out pan-European projects to develop innovative products, processes and services. EUREKA projects are co-financed though national public funding and private investments. Through a EUREKA project, partners develop new technologies for which they agree the Intellectual Property Rights and build partnerships to penetrate new markets. By encouraging and assisting businesses to innovate, the EUREKA Initiative complements the European Union's Framework Program in working actively towards the common European objective of raising investment in R&D to 3% of GDP by 2010. The main relevant EUREKA clusters in the context of FIRE are:

CELTIC is an initiative to overcome societal, economical and technical challenges to put Europe in a leading position with respect to the R&D in the telecommunications domain.. One of the main added value of CELTIC can be identified in the development of pre-competitive comprehensive “Integrated Communication System Solution” being eventually at the core of a CELTIC Pan-European Laboratory. The major technical domains covered by CELTIC include: i) Services and Applications, ii) Broadband Infrastructures, iii) Security.

ITEA2 with the overall objective to develop a European leadership in embedded Software Intensive Systems. The rationale for ITEA's launch was a response to the fact that the digital age is imminent and the digital transition is proceeding rapidly. 'These changes, the most significant since the Industrial Revolution, are far-reaching and global. They're not just about technology. They will affect everyone, everywhere. Software is the key to this revolution. Never before has European industry mounted a concerted strategic R&D initiative to address this challenge.

MEDEA+ is a co-operative research and development aimed at increasing the performance and competitiveness of the European microelectronics industry. Achieving superiority in microelectronics, which is the key enabling technology for the entire industry, starting with the Information and Communication Technologies (ICT) Industry but also encompassing the automotive, professional and leisure sectors, is the most effective way for creating higher added value, thus stimulating new innovative markets and leading to growth and additional highly skilled jobs in Europe.

Annex F Related US and Far-East Initiatives

In the **US**, the GENI facility (Global Environment for Networking Innovations) will be unique among experimental platforms in that it will be designed to support both research and deployment, effectively filling the gap between small-scale experiments in the lab, and mature technology that is ready for commercial deployment. It is envisaged that researchers will use GENI to evaluate new network systems on large scale, learning from real workloads as users begin to exercise those systems, and hardening the successful systems as user adoption grows. GENI foresees also interconnection on the global level.

In the Far-East both **Japan** and **South Korea** have relevant telecommunications initiatives, with a focus on mobile communications, such as the mobile IT forum for Japan and Next Generation Mobile Communication Forum, NGMC Forum, for Korea. These Asian countries have founded several test-oriented technology fora, consortia, and platforms, namely Yokosuka Research Park (YRP), Asia-Pacific Information Infrastructure (APII), Cyber Kansai Project and the Japan Gigabit Network II (JGN II). In particular the "u-Japan" is a policy, striving for the realisation of a ubiquitous network society in 2010. It is a core part of "e-Japan Strategy," which was launched by the Cabinet as Japan's comprehensive IT Strategy.

China is currently deploying the CNGI (China Next Generation Internet) which connects several nation-wide core networks covering more than 40 major cities in the country. It is supporting development and experiments on next generation network technologies, middleware and applications. The CERNET is the largest of the core networks in CNGI and supports science and education. It provides services for studies, development and test of applications for the Future Internet. The CERNET is available for international collaboration.

References

[CLARK et al]:

“The Internet is Broken” in Technology Review, December 2005/January 2006 by David D, Clark, MIT;

"Overcoming the Internet Impasse through Virtualization" by T. Anderson, L. Peterson, S. Shenker, J. Turner, IEEE Computer, April 2005;

"New Architectures and Disruptive Technologies for the Future Internet: The Wireless, Mobile and Sensor Network Perspective", NSF WMPG Planning Project Interim Report, August 2005;

"Overcoming Barriers to Disruptive Innovation in Networking", Report of NSF Workshop, January 2005.

[SAC] http://cordis.europa.eu/ist/fet/comms.htm

[RN] cordis.europa.eu/ist/rn/

[ONELAB] www.fp6-ist-onelab.eu

[PANLAB] www.panlab.net

[LIVINGLABS] http://www.livinglabs-europe.com

[Eiffel]: www.future-internet.eu

[EUQOS]: www.euqos.eu

[EUROLABS]: www.ist-eurolabs.org

[6DISS]: www.6diss.org

[DETER] http://www.isi.edu/deter/docs/testbed.overview.htm

[CELTIC] www.celtic-initiative.org

[EUREKA] http://www.eureka.be/

ETPs http://cordis.europa.eu/technology-platforms/home_en.html/

[ARTEMIS] http://www.artemis-office.org/

[ISI] http://www.isi-initiative.eu.org/

[eMobility] http://www.emobility.eu.org/

[NEM] http://www.nem-initiative.org/

[NESSI] http://www.nessi-europe.com/

GENI/FIND: http://www.geni.net/

Planetlab: https://www.planet-lab.org/

This consultation activity: cordis.europa.eu/fp7/ict/fire