FIRE FUTURE INTERNET RESEARCH AND EXPERIMENTATION

REPORT OF THE EXPERT CONSULTATION MEETING
12 FEBRUARY 2008 - BRUSSELS

EXECUTIVE SUMMARY

This consultation meeting was called to discuss some initial ideas for Objective 1.6 "New Paradigms and Experimental Facilities", which in the future may be more appropriately called "Future Internet Research and Experimentation", under ICT Work Programme 2009/10. As compared to the expert group consulted before Call 2, the new group was expanded in terms of technical expertise in order to include experts on service architectures and infrastructures. The objective of the meeting was to reflect on and discuss the analysis of the Call 2 results as well as some initial ideas for the Objective as tabled by the European Commission. In the form of discussion papers, both inputs were submitted to the invited experts prior to the meeting. Short presentation gave an overview of the findings and ideas.

The Commission suggested that the FIRE Future Internet Research and Experimentation Initiative should continue to have two interconnected dimensions:

1) **Experimentally-driven research** on new architectures and visionary paradigms for the Future Internet: There is an increasing demand from academia and industry to bridge the gap between long-term research and large-scale experimentation, which can be done through *experimentally-driven research* consisting of iterative cycles of research, design and experimentation of new networking and service architectures and paradigms for the Future Internet addressing all levels. FIRE experimentally-driven research should address research which proves and exploits the full value of the FIRE Facility:

- doing truly multidisciplinary experimental research;
- testing new Internet architectures and paradigms including at the service level;
- assessing the socio-economic impacts of future changes to the Internet.

2) Implementing the FIRE **experimental facility** broadly scoped to support research on the Future Internet (technologies and services):

- to include testbeds for different stages of the research and development cycle.
- to support testing the impact of changes to the Internet not only in technical but also in socio-economic terms.
- to cover all levels from fast network connectivity to service architectures.
- to be available to academia, research centres and industry, including the relevant ETPs.
- to allow experimentation with advanced architectures of the Future Internet taking a system view.
- to build on the design principle of "open federation of testbeds".
- to allow for access of broad user communities for experimenting on "user experience".

- to become a sustainable research infrastructure for the Future Internet, serving both industry and academia in their Future Internet related research and to overcome the limited availability of testbeds for the duration of the projects under which they are provided.

There were a number of presentations from the participants, including positions from several ETPs, as detailed in the full report. Some **key statements** made by participants:

- In Call 2, STREPs mainly focused on research topics and IPs mainly on testbed federation. Several proposals planned to integrate testbeds, but without explaining what would be the *advantage* of such federation. Studies on architectures and methodologies *beyond* the basic integration of testbeds were also missing. Many research proposals could have been more innovative or even disruptive in terms of "new paradigms" and "advanced approaches".
- The issue of whether the Future Internet needs a new architecture (clean slate) to fix its problems had also been discussed in the consultation meeting of the Objective "The Network of the Future", where it was noted that an evolutionary trend would certainly preserve more the large investment in legacy equipment but would possibly lead to smaller improvements, and will further increase the complexity.
- Social networking (e.g. MySpace, Facebook) will have an increasing impact and must be considered in any Future Internet scenario.
- The use of the Internet has to evolve from a "fault tolerant network to connect machines" to a "communication medium to exchange content and to build relationships with others".
- One-size does not fit all: Services on the Internet have very different requirements. More in general, one testbed covering at the same time all the four areas of network architecture, service architecture, middleware and system research may be too broad. The same experimental facility cannot serve at the same time for the testing of both advanced Internet architectures and new applications.
- A federation of testbeds, however, would allow testing new and visionary components in one area, while in other areas traditional technologies could be used.
- Keep an open approach to the research: the Future Internet will emerge from work done in research projects, but it is impossible to know from the outset precisely which one(s) will provide the important ideas.
- Research projects should be encouraged to use the experimental facilities for their proof of concept, and to make the data public (build a repository of measurement results). Keywords are “open” and “federated”.
- There is a need for a Security Roadmap, including techniques for security and dependability (e.g. cryptography at gigabit speeds) and considering the convergence of virtual worlds with real worlds.
- IPv6 is the basis for the Internet while FIRE is running. There may be a need for experimentation on the transition from IPv6 to new protocols experimented with in FIRE.
- FIRE should be responsive to the needs of service-oriented architectures in order to ensure that the future Internet delivers an environment (low latency, high reliability, multi-modal, etc) which becomes an integral part of the emerging service architectures on top of which all applications will be built.

- The FIRE facility should be expanded to include the service infrastructure layer, which is prototyped in many FP6 projects in the areas of Grids as well as Software and Service Architectures.
- A NESSI network of testbeds (NEXOF) is planned to host industrial partners, academics and other users, and will provide integration, interoperability, benchmarking and similar services. The NESSI vision on FIRE includes interconnected testbeds as well as the hosting of experiments (pre-services).
- Layers are blurring: the network and connectivity services are only a part of the services delivered to the customer and not the complete picture. Applications are aggregated from raw - and other aggregated – services, from different types of resources. Experimental facilities are necessary to validate service-oriented infrastructure concepts integrating different kinds of technologies.
- Service infrastructures will play a critical role in the evolution of the Internet. Whereas service technologies are maturing, service infrastructures are not yet. Europe should take the lead towards an Internet of services with large-scale open-service infrastructure experimentation.
- FIRE needs to address as well how Internet research can support the next generation of distributed computing.
- There is a lot of work to do on research and validation. The key to success is collaboration across the various EC funding silos.
- Networks must become "resource-aware".
- There was as well the view which is contradicting some of the views above that service architectures could be tested on top of existing facilities, as they do not demand specific facilities.
- Both the eMobility and the NESSI technology platforms support FIRE and want to discuss the role of the ETPs in FIRE with the FIRE stakeholders.

Final discussion:

All participants agreed that the main strength of FIRE is the strong link between experimentation and long-term research. Basic, long-term or advanced research on new architectures and large-scale experimentation must go together. This was substantiated in several statements: "If there is no cooperation between the research projects and the facilities, the facilities may become useless", "vision without implementation is illusion", "research determines which kind of testbed is needed", "experimental facilities do not represent a value per se, only in relation to the concepts to be tested".

When it comes to whether FIRE should be expanded to include service infrastructures and architectures, there were diverging views. Generally, experts from academia focusing on networks intended the facilities only for testing architectures and protocols, whereas experts from the services area wanted facilities for testing service architectures - or even applications - as well. Clearly the facilities can serve multiple purposes at the same time, but the same testbed can probably not serve for the testing of both advanced Internet architectures and new applications. Research on network and service infrastructures is longer term and needs to be treated differently from research on applications, which usually can be developed in shorter time frames.

FIRE FUTURE INTERNET RESEARCH AND EXPERIMENTATION

REPORT OF THE EXPERT CONSULTATION MEETING
12 FEBRUARY 2008 - BRUSSELS

1. INTRODUCTION AND SETTING THE SCENE

1.1. Welcome

Mario Campolargo welcomed everyone to the consultation meeting and remarked that the amount of proposed - and quality of accepted - projects demonstrated that the combination of research and experimentation was of high interest and importance.

The two synergetic parts of:

(i) experimentally driven (longer-term) research (architectures, paradigms, security, ….), and

(ii) building the European experimental facility (not 1 monolithic infrastructure, but one that is gradually evolving, through federation and doing system-testing)

are well represented by the accepted proposals. We are here today to consider what we should do next.

1.2. Objectives and the FIRE portfolio of projects

Max Lemke and Fabrizio Sestini explained that - in accordance with the structure of the Call - the accepted projects can be grouped into 3 areas:

- Co-ordination / Support:
  - FIREWorks (co-ordination action)
  - Paradiso (socio-economic impact on the future Internet)
- Research:
  - OPNEX (design and optimisation of multi-hop wireless networks)
  - ECODE (cognitive routing systems)
  - ECHOS/ NANODATACENTERS (using STBs and home gateways as small data-centers)
  - PERIMETER (user-centric protocols based on QoE)
  - N4C (delay-tolerant networks)
  - Self-Net (self-managed cognitive Internet)
  - SMART-Net (smart antenna multimode wireless mesh network)
  - RESUME-NET (resilience for future networking)
- Testbed-related:
  - PII (follow-on of PanLab)
  - OneLab2 (follow-on of OneLab)
  - Vital++ (integrating P2P into IMS)
  - Wisebed (wireless sensor network testbed)

The testbeds can be used by FIRE research projects and other FP7 projects, including those from the “Network of the Future” Call (Objective 1.1).

1.2.1. Possible next steps (based on proposals received, but unable to be financed in the 1st Call):

- Testbeds:

  The scope of the testbeds could be expanded; connectivity is well addressed, but they could also include the support of new experiments (for example):

  - on service architectures
  - on cross-layer approaches
  - on socio-economics
  - involving larger user communities, academia, research, industry

  Remarks:

  (i) The testbeds should also be sustainable beyond the project duration.

  (ii) Can work be done in parallel on: connectivity, middleware, services?

- Research:

  The focus should remain on (large-scale) experimentally driven (longer-term) research, which can include "clean slate" ideas (and "clean slate" ideas may also turn out to be evolvable.

  Potential topics could be:

  - service architectures and service infrastructures (moving up from the network)
  - security (not as an add-on, but integrated)
  - cross layer approaches
  - socio-economics
  - new Internet paradigms
  - new research into the federation methodology
  - benchmarking and measurements

  Proposals should again focus on highly innovative and revolutionary ideas.

1.2.2. A potential future direction for FIRE

Max Lemke presented a potential way forward for FIRE to stimulate the subsequent discussion. He outlined the motivation for FIRE followed by detailed descriptions of the properties of the FIRE Experimental Facility as well as those of FIRE Research. A paper has been submitted to all participants prior to the meeting. This paper is annexed to this report.

2. A FIRST ANALYSIS OF THE CALL 2 (OBJECTIVE 1.6) RESULTS

Andrea Passarella gave a summary of the results from the evaluation of the proposals received for the Call 2 (Objective 1.6):

- STREPs mainly focussed on research topics
- IPs mainly focused on testbed federation
- a wide range of research topics was covered, but many were low on innovation and not disruptive / clean slate. On the other hand, relevant topics were proposed, but the quality of the proposal was low
- several proposals planned to integrate testbeds, but without explaining what would be the *advantage* of such federation
- studies on architectures and methodologies *beyond* the basic integration of testbeds were also missing

Open issues identified:

- what should be the balance between the amount of research and the experimentation on testbeds? (the process should certainly be iterative)
- what is acceptable as "experimentation" (simulation, emulation, testbeds)? (simulation alone is certainly not sufficient)

Potential new topics could include:

- content and human-centric networking
- co-operation, at the network level
- wireless aspects
- architectures and methodologies (to increase the sustainability of the testbeds)
- added value beyond testbed connectivity
- virtualisation
- flexible composition of real and virtual resources
- heterogeneous networking environments
- flexible parameterisation and configuration
- application-driven proposals

3. FEEDBACK FROM THE CONSULTATION MEETING "INTERNET OF THE FUTURE"

Through the presentation given by Ioannis Stavrakakis, it was clear that many of the topics in accepted proposals were similar to those in FIRE. They also discuss whether the future Internet needs a new architecture (clean slate) to fix its problems, or if (bearing in mind the large investment in legacy equipment) the trend will continue to be more evolutionary. The consequence of the latter scenario is that improvements will probably be smaller, and will further increase the complexity. They also recognise that social networking (MySpace, Facebook) will have an increasing impact. Future Calls for Objective 1.1 will likely be in areas of:

- emerging technologies (Web 2.0)

- new approaches for dealing with content (including user-generated content - the Internet is becoming a large distributed database - and placing users at the centre of the Information Society)
- stronger application / service perspective
- greener Internet (VLSI technology requires much energy, more-energy-efficient transmission modes, disposing of legacy equipment)
- from interconnecting nodes towards interconnecting information
- more emphasis on mobile (including QoS, security,)
- the end-to-end paradigm
- the Internet continues to grow, but revenues do not
- service creation
- competition
- network of global virtualised resources
- wireless should support DSL-like rates (the specific infrastructure should be invisible)
- more spectral efficiency
- optical switching
- network coding

4. INVITED CONTRIBUTIONS

4.1. Serge Fdida: "Experimentally-driven advanced research"

One-size does not fit all. Services on the Internet have very different requirements. Current trends are towards a virtualised backbone, with wireless access at the edges, and an increasing number of sensors.

The future Internet will emerge from work done in research projects, but it is impossible to know from precisely which one(s), since we do not know which will be the important assumptions. We must also take into account projects in other Units.

We should encourage research projects to use our experimental facilities (for their proof of concept), and make the data public (build a repository of measurement results). Keywords are "open" and "federated".

Ideas for the next Call are:

- testbed operation and management
- research on new services
- (maybe) research on middleware
- incorporation of real users (though it is difficult to have real users (needing stability and reliability) on the same facility as disruptive research
- showing how the results can be integrated in commercial environments.

4.2. Martin May: Some comments from the perspective of the ANA project

The FP6 SAC project ANA is in-line with the FIRE philosophy of implementing, experimenting and feeding back the results into the design. It is developing a new communications architecture, which is evolving through this reiteration process. It was not possible to start with the architecture. Other innovations are:

- new software for routers to build federated networks
- a new testing methodology
- a composable stack for devices, which means they can also behave like networks.

Ideas for the next Call are:

- to grow the network and add heterogeneity (this opens up new challenges relating to network interoperability and new devices)
- federate multiple testbeds, each having a specific focus
- focus resources into fewer areas (but which ones?)
- experimental research seems to be the right direction
- mix network architectures and service architectures But what sort of testbed is needed for experimenting with new service architectures? A testbed covering at the same time all the four items of network architecture, service arch, middleware and system research would be too broad.
- the management of large scale facilities
- virtualisation (on devices as well as networks)
- keep the long-term aspect, but consider how to achieve some results in the short term
- communication paradigms beyond client-server
- include socio-economic aspects
- the EC could push for a new frequency band for wireless research

Open issue:

- how to build a pre-commercial Internet?

4.3. Michel Riguidel: A Security View

Michel presented a Security Roadmap and highlighted some topics that will be important and which require further study:

- techniques for security and dependability:
    - how to secure interfaces
    - how to perform cryptography at gigabit speeds (or what is the alternative?)
    - new aspects associated with multiparty and mobility

- the convergence of virtual worlds with real worlds (and how to secure virtual planes)
- identity
- the traceability of swarms of objects
- economics (trust experimentation platforms)
- ethics (privacy and intimacy)
- policy (governance experimentation platforms)

Notes:

- security is different from trust. (The Internet can be made to be secure, but this does not mean that it can be trusted)
- one single system (i.e. in the case of the Internet, IP addresses) fundamentally poses security risks

4.4. Serge Druais: The NESSI (ETP) view

NESSI is working on new services for the Internet to make it more:

- alive
- pervasive
- rich
- invisible

The NESSI Open Framework represents a composition of services on different devices. Demonstrators based on this Open Framework are being developed for:

- health services
- pan-European e-procurement
- European Digital Libraries
- Galileo infrastructures services

A NESSI network of testbeds is planned to host industrial partners, academics and other users, and will provide services such as:

- engineering capabilities
- integration and interoperability platforms (e g. Plug tests)
- benchmarking
- demonstration and visualisation capabilities
- compatibility and compliance (technical, legal, standards)

The Internet is now a strategic resource, because of the services it supports.

Users are key: 1 billion people connected to the Internet, 1 new blog per second.

Trend towards “on-the-move” services.

The NESSI vision on FIRE includes:

- interconnected testbeds
- hosting experiments (pre-services)

4.5. Cees de Laat: Networks "beyond hybrid networking"

NRENs have moved from Telco circuits (expensive) –> dark fibre. This has been a fundamental shift over the last 5 years. Lambdas are getting faster and cheaper. Photonics transmission and switching is cheaper per bit than copper and electronics, and more energy-efficient.

Applications have increasingly stringent requirements, which are stretching the Internet to its limits:

- more predictable behaviour; Ethernet enhanced with circuit properties (PBT) can be a solution
- the data produced by the Large Hadron Collider (LHC) in CERN is too great to store and process everything locally in real time. Huge bandwidths are therefore required to distribute the data throughout Europe
- the SCARIe (radio telescope) project has to process the received data at a rate equivalent to 1000 flops per byte. This also leads to a data flow problem
- The CineGrid project produces film quality pictures of 4000 by 2000 pixels at 60 frames per sec. This is equivalent to 7 Gbit/s for the uncompressed data (in "4K format")
- The Global Lambda Integrated Facility (GLIF) is handling 10Gbit/s lambdas

Ideas for the next Call are:

- a multi-layer network description language to explain how to mix traffic - and at which layer
- attract stakeholders; testbeds are mentioned, but for what purpose will they be used?
- networks will soon be carrying 100Gbit/s lambdas. This traffic needs aggregating and grooming
- programmable L1, L2, L3 objects
- virtualisation
- new types of Internet exchanges
- security

5. SHORT PRESENTATIONS FROM THE PARTICIPANTS

There then followed presentations from (most of) the participants.

Only the topics mentioned as either (i) infrastructure that can be made available for FIRE experimentation, or (ii) ideas for future Calls are listed:

(i) Infrastructure that can be made available for FIRE experimentation

- the federation of the testbeds has to be made using the interfaces that are available externally. The result will be a "network of networks"
- critical mass is important. Federation can bring this, as opposed to individual small-scale testbeds

- significant network infrastructure is available at Uni Trento, and also real-time emulators
- the Living Labs provide a real application environment and end users
- Fraunhofer has a testbed for autonomic communications
- the FEDERICA project offers user-controlled "slices" over GEANT
- the Trilogy project does research on data centres in Unit D1, and invites other projects to build on it
- projects working on the "Internet of Things" (and the SAC projects) are already trying out things through experiments

(ii) Ideas from these presentations to include in future FIRE Calls:

IPv6:

- testbeds to experiment with the transition to IPv6, and conformance testing

Layers:

- 3 layers: Transport, Service Control and Application. Different testbeds could be envisaged for each layer
- 3 layers: Services, Networks, Business
- …. or even layer-less

Services:

- the trend is towards user generated content and user generated ***services***
- flexible services (across all device and network types)
- consider the Internet as an "Internet of Services"
- service infrastructure experimentation
- service architecture research should build on existing research such as on Grid architectures
- the main challenge is to be plug and play. Self-managing complicates the picture. One target could be a toolkit

Content-centric, social-aware Internet:

- the Internet was originally designed for connecting machines in a fault tolerant way. It is now a critical infrastructure and used for content distribution, people to people. Study social relationships and model the network accordingly (pervasive networks
- microblogging (Twitter)
- focus on the problems, not the technologies

Webservices:

- Webservices have received a lot of attention recently, as many are using them to deliver SOAs.
- A lot of work is needed to make them more robust, easier to programme and more compatible with each other.
- FIRE could greatly help with the development of Web Services or related solutions in Europe by showing how the Internet can be designed to support high performance, secure Web-service-based applications.
- "component-oriented Web 2.0"
- towards a Service Oriented Architecture paradigm (Web 2.0 -> Web 3.0)
- machine-machine convergence (devices and interoperability)
- SOA4All: There are currently 30 billion Web pages (10 million are added each day), but only 10'000 "true" Webservices

Energy:

- saving energy consumption on ICT
- carbon-neutral internet
- P2P is greener than centralised Google-like data centers

Socio-economics:

- bear in mind the greater dynamics that are at work, and their impact on traffic, interoperability, management. For example:
  - Open Auth, Open Social, Facebook

Other:

- cloud computing
- virtualisation
- methodologies for developing, verifying and deploying system solutions (not necessarily a common framework, but at least a better formulation of the issues)

Other remarks made:

- aim high (GENI promises 100Mbit/s for 100 people simultaneously and 1Gbit/s to any single user)
- explore the convergence of Telecom, Media, and Software
- bear in mind that FP7 is only one aspect of research funding; there is also national funding and company-internal funding. These are also driving the direction of research

- the EIFFEL group has proposed a “School of Architecture”. This might be relevant for FIRE

- the current Internet infrastructure is not as reliable as we think

- resource-aware networks, dynamically configurable: Bandwidth on demand to minimise costs

- take a re-look at what GRIDs can be used for

- encourage academics and researchers to produce products (rapid prototyping; try it and see)

- social networks might be just a fashion - the topic is very interesting for security, …. but already the number of (eg. Facebook) participants is decreasing (maybe due to recent news about its lack of security)

- the future Internet will be driven by economics

- simulation techniques are becoming unable to reflect the speeds of the latest network technologies

6. THE FOUR QUESTIONS

Four questions were put to the audience in order to stimulate discussion.

1. can the experimental facilities be developed "in isolation" from the research on new architectures which will need testing on such facilities?

Or the reverse: can advanced, experimentally-driven research be carried out without a strong linkage and coordination with the development of experimental facilities?

The meeting provided a consensus on this point. Long-term Research on new architectures and Experimentation must go together.

Yes, testbeds should support the research agenda… but what is the research agenda? We cannot cover everything; for efficiency and impact, we should focus on a few topics, but which ones?. We should identify the borders of different disciplines (access networks, applications, ….) and the cross-layer activities (cross disciplines between services and networks). We can start with some ideas, but be open to change.

Research –> Experimentation is an iterative process. There must be a close co-operation between two groups (Research and Experimentation), or else the testbeds that are built will not be appropriate. Vision without implementation is an illusion. Research determines which kind of testbed is needed. Experimental facilities do not represent a value per se; only in relation to the concepts to be tested.

Collaborate with other Units; the results from Objective 1.1 can be trialled in 1.6.
Especially where platforms are expensive to build (eg. large platforms for RFID and for wireless communications), it makes sense to fund one and then share it.

Europe needs a testbed … and a reason for doing it. There are applications that want Tbit/s. Encourage joint projects between services and networks.

Need real users, or at least realistic traffic.

Need to include Mobile and Wireless issues.

Need new routing, protocols (TCP/IP has well-known limitations that impact on services).

Need fundamental research: We know there will soon be no more IPv4 addresses. IPv6 is the only solution to be available in time, but the longer-term answer may be something completely different.

Service architectures can be tested on top of existing facilities; we do not need to demand specific facilities for them

2: Can the facilities serve multiple purposes at the same time?
i.e. testing of internet protocols/architectures, of service architectures, of middleware for new applications, etc.?

Several views were put forward in this regard:

- It can be done, but implies no interference; else the results will not be consistent.
- The timescales are different for the development of networks, services and applications.
- It is important to see how services interact with the network.
- How deep into the network technology should we go? Do we want to be aware of the low level functions, or rather consider the network to be dumb? Applications should be able to interact with the network at different layers. Photonics is due to explode the bandwidth.
- How high should FIRE go in terms of experimentation at application level. There was strong demand that FIRE could greatly benefit from trying to meet the needs of large-scale SOAs which have only been talked about to date and not deployed in any great numbers.
- A testbed covering all the levels of network architecture (eg. service architecture, middleware and system research) is too broad and unnecessary.

3: What should be the focus on a next call on experimental facilities, given the result of call 2?

Several views were put forward in this regard:

- We need a very close collaboration between research and experimentation. We need also to be able to generate real traffic patterns involving real users in the testbeds.

- Emphasise more *experimentally-driven*, rather than *building experimental facilities*.

- Focus the experimental facilities on a few subjects, as the budget is limited.

- There is a great opportunity to try and bring the networking community closer together with the applications development community. This should be done by calling for large-scale service-based application experimentation which tests all aspects of network provision: bandwidth, latency, reservation, provisioning, multi-modal use, configuration, management etc. Only by doing this it will we ensured that networks resulting from FIRE meet the needs of the applications that will be developed over the coming years.

- Reconsider the word "testbed" – the Internet didn't emerge from a testbed. Call it an "innovation platform". Then it gets a life after the project. The level of end-user participation will be the proof.

4: What should be the focus of a next call on advanced research on future internet, given the results of call 2 (and 1)?

Several views were put forward in this regard:

- Research on network infrastructures is long-term and needs to be treated differently from research on applications, which can be developed in shorter time frames.

- Beware of calling for service/application development proposals that are not linked to networks, or else too many proposals will be received, or it will be too difficult to glue together the resultant projects. There are other groups working on purely services and applications … we should look at where research on services impacts architectures (and vice-versa) … and on the experimentation.

ANNEX I: A DRAFT FIRE VISION PREPARED BY THE EC FOR DISCUSSION

This paper was provided by the EC prior to the Meeting to serve as a basis for discussion at the expert consultation meeting for Objective 1.6 New Paradigms and Experimental Facilities. It describes a potential future vision for the FIRE Future Internet Research and Experimentation Initiative. Due to the limited amount of time for the meeting, not all aspects of this paper were discussed at the meeting. Also, it was not the objective of the meeting to reach consensus on all issues.

Motivation for FIRE

There is an increasing demand from academia and industry to bridge the gap between long-term research and large-scale experimentation, through ***experimentally-driven research*** consisting of iterative cycles of research, design and experimentation of new networking and service architectures and paradigms addressing all levels, including horizontal research on issues such as system complexity and security.

A fundamental need in this approach is the set-up of ***large-scale experimentation facilities***, going beyond individual project testbeds, which are also needed as validation tools, including for interoperability issues. They would help anticipating possible migration paths for technological developments which may be potentially disruptive, discovering new and emerging behaviours and use patterns, as well as assessing at an early stage the socio-economic implications of new technological solutions.

There is a clear need for putting together different research communities in an interdisciplinary approach, which would be stimulated and enabled by flexible experimentally-driven research approaches that cut across layers, in a large systems perspective, for instance from network connectivity and service architectures to security solutions and beyond, i.e. not limited to a few levels of the value chain or to a single objective.

The FIRE Future Internet Research and Experimentation Initiative should continue to have two interconnected dimensions – they cannot and should not be separated:

- integrating existing and emerging testbeds and maturing the concept of federations of testbeds cutting across all layers including service architectures at middleware and application level. This includes setting up and carrying out the management and operations of the resources as a coherent prototype experimental facility.

- do experimentally-driven research projects, which cut across several layers from connectivity via service architectures to applications, thereby addressing the "Future Internet" from a broad system perspective.

In terms of the structure of the ICT Theme under the Cooperation Programme of Framework Programme 7, FIRE should support large parts of Challenge 1 and beyond as illustrated below.

Properties of the FIRE Experimental facility

A European experimental facility for the Future Internet should have the following general properties:

- The experimental facility must be broadly scoped to support research on the Future Internet as defined in the motivation section of this paper thereby supporting all of Challenge 1 and slightly beyond:

    - **from proof-of-concept type testbeds to pre-commercial testbeds:** This explicitly includes testbeds for different stages of the research and development cycle. For example, the facility would support testing of disruptive paradigms for internet and service architectures as developed in an experimentally-driven long term research project. On the other hand the facility would as well support interoperability and validation of technologies, which are in the final development stage.

    - testing the impact of **changes to the internet not only in technical but also in socio-economic terms**, e.g. including experimentation on social acceptability, user experience, economic viability, … .

    - **covering all levels** from fast network connectivity via service architectures to important classes of application services and allowing for heterogeneity on all of these levels, e.g. different network architectures and protocols.

    - being available to **academia, research centers and industry** (e.g. the ETPs under Challenge 1).

- The facility must allow experimentation with advanced architectures of the Future Internet taking a system view aiming at **discovering the technical, societal and economic implications** of changes to the internet and to **identify potential evolutionary transition paths** for concepts and clean slate approaches that on a first glance may appear revolutionary or disruptive.

- The experimental facility must be built on the **design principle of "open federation of testbeds"** to form a facility giving access to a wide spectrum of resources, which may as well in a later stage include some dedicated resources.

- For large scale testing taking a system view and for experimenting on "**user experience**", the facility must allow for access of broad user communities, through adequate access policies, use of open source software, etc..

- The experimental facility must be **open and dynamic**: The experimental facilities must be open and allow for inclusion of testbeds from research projects across all of Challenge 1 and beyond. In concrete terms, testbeds and in particular system level testbeds being part of activities under security, the network of the future, networked media, service architectures, etc. could be federated within the FIRE facilities, following certain rules established by a management board. In addition, projects funded through other objectives or outside ICT, must be able to easily use the experimental facility to validate their research.

- In the long term, the experimental facility must become a **sustainable research infrastructure** for the future internet serving both industry and academia in their Future Internet related research. The limited availability of testbeds for the duration of the projects under which they were provided, must be overcome. **Future projects already at proposal stage must be able to rely on being able to use the facility as a research infrastructure** as they today rely on Geant and in future will rely on the Grid. In for the experimental facility to eventually become self-financed by the stakeholders in the medium term future, in the start up phase there is support needed for

  1. Making the testbeds and federated testbeds that constitute the FIRE facility sustainable beyond the project duration of individual testbed activities, thereby making the experimental facility a reliable research infrastructure for Future Internet research in Europe;

  2. Management and operations of the FIRE experimental facility;

  3. Provisioning of some dedicated resources for the core of the infrastructure.

Properties of FIRE research

FIRE will be open to include testbeds from other projects of Challenge 1 and beyond and for use by the same group of projects. Additionally, broad system-level type projects should proof and exploit the full value of the FIRE Facility, for doing truly multidisciplinary experimental research, for testing new internet architectures and paradigms including the service level, and for doing socio-economic impact assessments of future changes to the internet. The FIRE research projects should be aimed at becoming the most demanding showcases of integrating research for the FIRE facility.

- The Future Internet is both connectivity and services. FIRE research should address the design, prototyping and experimentation with advanced paradigms and architectures for the Future Internet addressing several layers from connectivity to services.

- By making it mandatory for FIRE research projects to cut across layers, FIRE experimentally-driven large scale integrating research projects will address the Future Internet from a broad system perspective, which is needed for Europe to reinforce its position.

- This kind of experimentally-driven yet long-term oriented dimension of the research denotes an open approach and gives strong attention to crucial socio-economic aspects which must be evaluated in a multidisciplinary approach, maintaining a strong link and feedback loop between advanced research on new architectures and their large-scale experimentation.

- The experimentation phase in FIRE research projects may as well include identification of potential evolutionary transition paths for concepts and approaches that on a first glance appeared revolutionary or disruptive.

- There is some need for additonal research on methodologies for benchmarking and measurements.

The specific characteristics of FIRE research under Objective 1.6 as compared to the other Objectives under challenge 1 are:

- Research must be experimentally driven including large scale testing going far beyond demonstrators in the final phase of a project.
- Research must cut across objectives of Challenge 1 and must take a broad system perspective.
- Research must use the FIRE experimental facility - missing testbed components may be provided for federation within FIRE.
- It is essential that FIRE experimentally-driven advanced research is closely linked to the building of the experimental facility as these kinds of projects are the "hard" test and use cases for the facility.

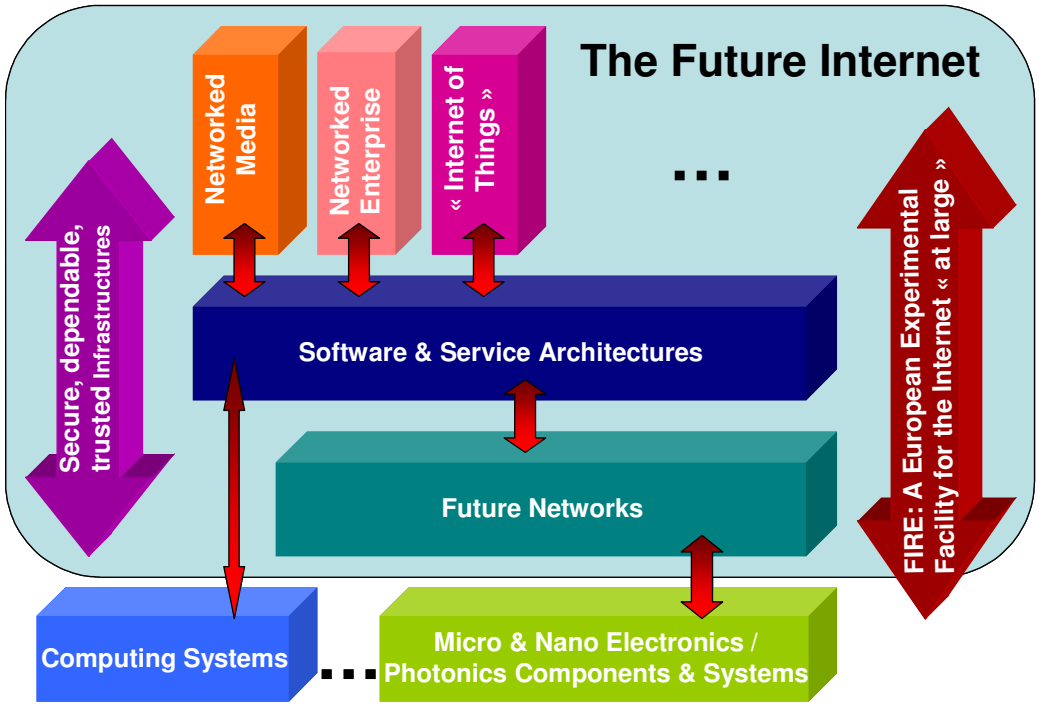

Figure 1: Objectives related to the Future Internet under ICT Work Programme 2007/08

ANNEX II: PARTICIPANTS

ABENI	Luca	University of Trento
ABRAMOVICZ	Henrik	Ericsson, ETP eMobility
AVESSTA	Susanna	Dimes
DE LAAT	Cees	University of Amsterdam
DE LAMA SANCHEZ	Nuria	Atos Research
DIOT	Christophe	Thomson
DRUAIS	Serge	Thales, ETP NESSI
FDIDA	Serge	UPMC
GAVRAS	Anastacius	Eurescom
GIORDANO	Silvia	DTI-SUPSI
HUGHES-JONES	Richard	University of Manchester
LE GALL	Franck	Inno Group
MAGEDANZ	Thomas	Fraunhofer Focus
MAY	Martin	ETHZ
MINERVA	Roberto	Telecom Italia
NIITAMO	Veli-Pekka	Nokia
PARSONS	Mark	University of Edinburgh
PASSARELLA	Andrea	CNR
POTTS	Martin	Martel, Rapporteur
RIGUIDEL	Michel	ENST
STAVRAKAKIS	Ioannis	University of Athens
TROSSEN	Dirk	BT
WESNER	Stefan	University of Stuttgart
WILLMOT	Steven	3Scale
ZAIONTZ	Charles	Create-Net
CAMPOLARGO	Mario	European Commission, INFSO F
BLIXT	Per	European Commission, INFSO F4
LEMKE	Max	European Commission, INFSO F4
SESTINI	Fabrizio	European Commission, INFSO F4
BABOT	Jacques	European Commission, INFSO F4
TSELENTIS	Georgios	European Commission, INFSO F4
Kolodziejski	Marek	European Commission, INFSO F4
De Sousa	Paulo	European Commission, INFSO D1
Paindaveine	Yves	European Commission, INFSO F5

EDITOR OF THIS REPORT: MARTIN POTTS

ANNEX III: MEETING AGENDA

10:00 Welcome: Mario Campolargo, EC, Acting Director Emerging Technologies and Infrastructures

10:15 Objectives - Setting the scene; Max Lemke, EC
10:30 The FIRE portfolio of projects; Fabrizio Sestini, EC

10:45 A first analysis of the Call 2 results; Andrea Passarella
11:00 Feedback from the Consultation Meeting "Internet of the Future"; Ioannis Stavrakakis

11:15 Coffee Break

11:30 Invited presentations:
Serge Fdida, UPMC: "*Experimentally-driven advanced research*"
Martin May, ETH: *title to be announced*
Michel Riguidel, ENST, *title to be announced*
Serge Druais, Thales, Nessi: "*A view from a service perspective*"
Cees de Laat, UVA: "*Beyond Hybrid Networking"*

12:30 Lunch Break

13:30 Position Statements - 5 minutes per participant

15:00 Coffee Break

15:30 Discussion; Chair: Per Blixt, EC

17:30 End of the meeting